

Three Classes of Deep Learning Architectures and Their Applications: A Tutorial Survey

Li Deng

Microsoft Research, Redmond, WA 98052, USA

E-mail: deng@microsoft.com, Tel: 425-706-2719

Abstract— In this invited paper, my overview material on the same topic as presented in the plenary overview session of APSIPA-2011 and the tutorial material presented in the same conference (Deng, 2011) are expanded and updated to include more recent developments in deep learning. The previous and the updated materials cover both theory and applications, and analyze its future directions. The goal of this tutorial survey is to introduce the emerging area of deep learning or hierarchical learning to the APSIPA community. Deep learning refers to a class of machine learning techniques, developed largely since 2006, where many stages of nonlinear information processing in hierarchical architectures are exploited for pattern classification and for feature learning. In the more recent literature, it is also connected to representation learning, which involves a hierarchy of features or concepts where higher-level concepts are defined from lower-level ones and where the same lower-level concepts help to define higher-level ones. In this tutorial, a brief history of deep learning research is discussed first. Then, a classificatory scheme is developed to analyze and summarize major work reported in the deep learning literature. Using this scheme, I provide a taxonomy-oriented survey on the existing deep architectures and algorithms in the literature, and categorize them into three classes: generative, discriminative, and hybrid. Three representative deep architectures --- deep auto-encoder, deep stacking network, and deep neural network (pre-trained with deep belief network) --- one in each of the three classes, are presented in more detail. Next, selected applications of deep learning are reviewed in broad areas of signal and information processing including audio/speech, image/vision, multimodality, language modeling, natural language processing, and information retrieval. Finally, future directions of deep learning are discussed and analyzed.

I. INTRODUCTION

Signal processing research nowadays has a significantly widened scope compared with just a few years ago. It has encompassed many broad areas of information processing from low-level signals to higher-level, human-centric semantic information (Deng, 2008). Since 2006, deep structured learning, or more commonly called deep learning or hierarchical learning, has emerged as a new area of machine learning research (Hinton et al., 2006; Bengio, 2009). Within the past few years, the techniques developed from deep learning research have already been impacting a wide range of signal and information processing work within the traditional and the new, widened scopes including machine learning and artificial intelligence (Hinton et al., 2012; Deng, 2011; Yu and Deng, 2011; Arel et al., 2010; Bengio et al., 2013); see the recent New York Times media coverage of this

progress in (Markoff, 2012). A series of workshops, tutorials, and special issues or conference special sessions have been devoted exclusively to deep learning and its applications to various classical and expanded signal processing areas. These include: the 2013 ICASSP's special session on New Types of Deep Neural Network Learning for Speech Recognition and Related Applications, the 2013 ICML Workshop for Audio, Speech, and Language Processing, the 2012, 2011, and 2010 NIPS Workshops on Deep Learning and Unsupervised Feature Learning, the 2012 ICML Workshop on Representation Learning, the 2011 ICML Workshop on Learning Architectures, Representations, and Optimization for Speech and Visual Information Processing, the 2009 ICML Workshop on Learning Feature Hierarchies, the 2009 NIPS Workshop on Deep Learning for Speech Recognition and Related Applications, the 2008 NIPS Deep Learning Workshop, the 2012 ICASSP deep learning tutorial, the special section on Deep Learning for Speech and Language Processing in IEEE Transactions on Audio, Speech, and Language Processing (January 2012), and the special issue on Learning Deep Architectures in IEEE Transactions on Pattern Analysis and Machine Intelligence (2013). The author has been directly involved in the research and in organizing several of the events and editorials above, and has seen the emerging nature of the field; hence the need for providing a tutorial survey article here.

Deep learning refers to a class of machine learning techniques, where many layers of information processing stages in hierarchical architectures are exploited for pattern classification and for feature or representation learning. It is in the intersections among the research areas of neural network, graphical modeling, optimization, pattern recognition, and signal processing. Three important reasons for the popularity of deep learning today are drastically increased chip processing abilities (e.g., GPU units), the significantly lowered cost of computing hardware, and recent advances in machine learning and signal/information processing research. Active researchers in this area include those at University of Toronto, New York University, University of Montreal, Microsoft Research, Google, IBM Research, Stanford University, University of Michigan, Massachusetts Institute of Technology, University of Washington, and numerous other places. These researchers have demonstrated successes of deep learning in diverse applications of computer vision, phonetic recognition, voice search, conversational speech recognition, speech and image

feature coding, semantic utterance classification, hand-writing recognition, audio processing, visual object recognition, information retrieval, and even in the analysis of molecules that may lead to discovering new drugs as reported recently in (Markoff, 2012).

This paper expands my recent overview material on the same topic as presented in the plenary overview session of APSIPA-ASC2011 as well as the tutorial material presented in the same conference (Deng, 2011). It is aimed to introduce the APSIPA Transactions' readers to the emerging technologies enabled by deep learning. I attempt to provide a tutorial review on the research work conducted in this exciting area since the birth of deep learning in 2006 that has direct relevance to signal and information processing. Future research directions will be discussed to attract interests from more APSIPA researchers, students, and practitioners for advancing signal and information processing technology as the core mission of the APSIPA community.

The remainder of this paper is organized as follows:

- Section II: A brief historical account of deep learning is provided from the perspective of signal and information processing.
- Sections III: A three-way classification scheme for a large body of the work in deep learning is developed. A growing number of deep architectures are classified into 1) generative, 2) discriminative, and 3) hybrid categories, and high-level descriptions are provided for each category a literature survey.
- Sections IV, V, VI: For each of the three categories, a tutorial example is chosen to provide more detailed treatment. The examples chosen are: 1) deep autoencoders for the generative category (Section IV); 2) DNNs pre-trained with DBN for the hybrid category (Section V); and 3) Deep stacking networks for the discriminative category (Section VI).
- Sections VII: A set of typical and successful applications of deep learning in diverse areas of signal and information processing are reviewed.
- Section VIII: A summary and future directions are given.

II. A BRIEF HISTORICAL ACCOUNT OF DEEP LEARNING

Until recently, most machine learning and signal processing techniques had exploited shallow-structured architectures. These architectures typically contain a single layer of nonlinear feature transformations and they lack multiple layers of adaptive non-linear features. Examples of the shallow architectures are conventional, commonly used Gaussian mixture models (GMMs) and hidden Markov models (HMMs), linear or nonlinear dynamical systems, conditional random fields (CRFs), maximum entropy (MaxEnt) models, support vector machines (SVMs), logistic regression, kernel regression, and multi-layer perceptron (MLP) neural network with a single hidden layer including extreme learning machine. A property common to these

shallow learning models is the relatively simple architecture that consists of only one layer responsible for transforming the raw input signals or features into a problem-specific feature space, which may be unobservable. Take the example of a SVM and other conventional kernel methods. They use a shallow linear pattern separation layer with one or zero feature transformation layer when kernel trick is used or otherwise. (Notable exceptions are the recent kernel methods that have been inspired by and integrated with deep learning; e.g. Cho and Saul, 2009; Deng et al., 2012; Vinyals et al., 2012). Shallow architectures have been shown effective in solving many simple or well-constrained problems, but their limited modeling and representational power can cause difficulties when dealing with more complicated real-world applications involving natural signals such as human speech, natural sound and language, and natural image and visual scenes.

Human information processing mechanisms (e.g., vision and speech), however, suggest the need of deep architectures for extracting complex structure and building internal representation from rich sensory inputs. For example, human speech production and perception systems are both equipped with clearly layered hierarchical structures in transforming the information from the waveform level to the linguistic level (Baker et al., 2009, 2009a; Deng, 1999, 2003). In a similar vein, human visual system is also hierarchical in nature, most in the perception side but interestingly also in the "generative" side (George, 2008; Bouvrie, 2009; Poggio, 2007). It is natural to believe that the state-of-the-art can be advanced in processing these types of natural signals if efficient and effective deep learning algorithms are developed. Information processing and learning systems with deep architectures are composed of many layers of nonlinear processing stages, where each lower layer's outputs are fed to its immediate higher layer as the input. The successful deep learning techniques developed so far share two additional key properties: the generative nature of the model, which typically requires adding an additional top layer to perform discriminative tasks, and an *unsupervised* pre-training step that makes an effective use of large amounts of unlabeled training data for extracting structures and regularities in the input features.

Historically, the concept of deep learning was originated from artificial neural network research. (Hence, one may occasionally hear the discussion of "new-generation neural networks".) Feed-forward neural networks or MLPs with many hidden layers are indeed a good example of the models with a deep architecture. Back-propagation, popularized in 1980's, has been a well-known algorithm for learning the weights of these networks. Unfortunately back-propagation alone did not work well in practice for learning networks with more than a small number of hidden layers (see a review and analysis in (Bengio, 2009; Glorot and Bengio, 2010)). The pervasive presence of local optima in the non-convex objective function of the deep networks is the main source of

difficulties in the learning. Back-propagation is based on local gradient descent, and starts usually at some random initial points. It often gets trapped in poor local optima, and the severity increases significantly as the depth of the networks increases. This difficulty is partially responsible for steering away most of the machine learning and signal processing research from neural networks to shallow models that have convex loss functions (e.g., SVMs, CRFs, and MaxEnt models), for which global optimum can be efficiently obtained at the cost of less powerful models.

The optimization difficulty associated with the deep models was empirically alleviated when a reasonably efficient, unsupervised learning algorithm was introduced in the two papers of (Hinton et al., 2006; Hinton and Salakhutdinov, 2006). In these papers, a class of deep generative models was introduced, called deep belief network (DBN), which is composed of a stack of Restricted Boltzmann Machines (RBMs). A core component of the DBN is a greedy, layer-by-layer learning algorithm which optimizes DBN weights at time complexity linear to the size and depth of the networks. Separately and with some surprise, initializing the weights of an MLP with a correspondingly configured DBN often produces much better results than that with the random weights. As such, MLPs with many hidden layers, or deep neural networks (DNN), which are learned with unsupervised DBN pre-training followed by back-propagation fine-tuning is sometimes also called DBNs in the literature (e.g., Dahl et al., 2011; Mohamed et al., 2010, 2012). More recently, researchers have been more careful in distinguishing DNN from DBN (Dahl et al., 2012; Hinton et al., 2012), and when DBN is used to initialize the training of a DNN, the resulting network is called DBN-DNN (Hinton et al., 2012).

In addition to the supply of good initialization points, DBN comes with additional attractive features. First, the learning algorithm makes effective use of unlabeled data. Second, it can be interpreted as Bayesian probabilistic generative model. Third, the values of the hidden variables in the deepest layer are efficient to compute. And fourth, the over-fitting problem, which is often observed in the models with millions of parameters such as DBNs, and the under-fitting problem, which occurs often in deep networks, can be effectively addressed by the generative pre-training step. An insightful analysis on what speech information DBNs can capture is provided in (Mohamed et al. 2012a).

The DBN training procedure is not the only one that makes effective training of DNNs possible. Since the publication of the seminal work in (Hinton et al., 2006; Hinton and Salakhutdinov, 2006), a number of other researchers have been improving and applying the deep learning techniques with success. For example, one can alternatively pre-train DNNs layer by layer by considering each pair of layers as a de-noising auto-encoder regularized by setting a subset of the inputs to zero (Bengio, 2009; Vincent et al., 2010). Also, “contractive” autoencoders can be used for the same purpose

by regularizing via penalizing the gradient of the activities of the hidden units with respect to the inputs (Rifai et al., 2011). Further, Ranzato et al. (2007a) developed the Sparse Encoding Symmetric Machine (SESM), which has a very similar architecture to RBMs as building blocks of a DBN. In principle, SESM may also be used to effectively initialize the DNN training.

Historically, the use of the generative model of DBN to facilitate the training of DNNs plays an important role in igniting the interest of deep learning for speech feature coding and for speech recognition (Deng et al., 2010; Dahl et al., 2011, 2012; Hinton et al., 2012). After this effectiveness was demonstrated, further research showed many alternative but simpler ways of doing pre-training. With a large amount of training data, we now know how to learn a DNN by starting with a shallow neural network (i.e., with one hidden layer). After this shallow network has been trained discriminatively, a new hidden layer is inserted between the previous hidden layer and the softmax output layer and the full network is again discriminatively trained. One can continue this process until the desired number of hidden layers is reached in the DNN. And finally, full backpropagation fine-tuning is carried out to complete the DNN training.

In the next section, an overview is provided on the various architectures of deep learning, including and beyond the original DBN published in (Hinton et al. 2006).

III. THREE BROAD CLASSES OF DEEP ARCHITECTURES: AN OVERVIEW

As described earlier, deep learning refers to a rather wide class of machine learning techniques and architectures, with the hallmark of using many layers of non-linear information processing stages that are hierarchical in nature. Depending on how the architectures and techniques are intended for use, e.g., synthesis/generation or recognition/classification, one can broadly categorize most of the work in this area into three classes:

- 1) Generative deep architectures, which are intended to characterize the high-order correlation properties of the observed or visible data for pattern analysis or synthesis purposes, and/or characterize the joint statistical distributions of the visible data and their associated classes. In the latter case, the use of Bayes rule can turn this type of architecture into a discriminative one.
- 2) Discriminative deep architectures, which are intended to directly provide discriminative power for pattern classification, often by characterizing the posterior distributions of classes conditioned on the visible data; and
- 3) Hybrid deep architectures, where the goal is discrimination but is assisted (often in a significant way) with the outcomes of generative architectures via better optimization or/and regularization, or discriminative criteria are used to learn the

parameters in any of the deep generative models in category 1) above.

Note the use of “hybrid” in 3) above is different from that used sometimes in the literature, which refers to the hybrid pipeline systems for speech recognition feeding the output probabilities of a neural network into an HMM (Bengio et al., 1991; Bourlard and Morgan, 1993; Morgan, 2012).

By the machine learning tradition (e.g., Deng and Li, 2013), it may be natural to use a two-way classification scheme according to discriminative learning (e.g., neural networks) vs. deep probabilistic generative learning (e.g., DBN, DBM, etc.). This classification scheme, however, misses a key insight gained in deep learning research about how generative models can greatly improve DNNs and other deep discriminative models via better optimization and better regularization. Also, deep generative models may not necessarily need to be probabilistic, e.g., the deep auto-encoder. Nevertheless, the two-way classification points to important differences between DNNs and deep probabilistic models. The former is usually more efficient for training and testing, more flexible in its construction, less constrained (e.g., no normalization by the difficult partition function which can be replaced by sparsity), and is more suitable for end-to-end learning of complex systems (e.g., no approximate inference and learning). The latter, on the other hand, is easier to interpret and to embed domain knowledge, is easier to compose and to handle uncertainty, but is typically intractable in inference and learning for complex systems. This distinction, however, is retained also in the proposed three-way classification which is adopted throughout this paper.

Below we review representative work in each of the above three classes, where several basic definitions will be used as summarized in Table 1. Applications of these deep architectures are deferred to Section VII.

Table 1: Some basic deep learning terminologies

<ol style="list-style-type: none"> 1. Deep Learning: a class of machine learning techniques, where many layers of information processing stages in hierarchical architectures are exploited for unsupervised feature learning and for pattern analysis/classification. The essence of deep learning is to compute hierarchical features or representations of the observational data, where the higher-level features or factors are defined from lower-level ones. 2. Deep belief network (DBN): probabilistic generative models composed of multiple layers of stochastic, hidden variables. The top two layers have undirected, symmetric connections between them. The lower layers receive top-down, directed connections from the layer above. 3. Boltzmann machine (BM): a network of symmetrically connected, neuron-like units that
--

<p>make stochastic decisions about whether to be on or off.</p> <ol style="list-style-type: none"> 4. Restricted Boltzmann machine (RBM): a special BM consisting of a layer of visible units and a layer of hidden units with no visible-visible or hidden-hidden connections. 5. Deep Boltzmann machine (DBM): a special BM where the hidden units are organized in a deep layered manner, only adjacent layers are connected, and there are no visible-visible or hidden-hidden connections within the same layer. 6. Deep neural network (DNN): a multilayer network with many hidden layers, whose weights are fully connected and are often initialized (pre-trained) using stacked RBMs or DBN. (In the literature, DBN is sometimes used to mean DNN) 7. Deep auto-encoder: a DNN whose output target is the data input itself, often pre-trained with DBN or using distorted training data to regularize the learning. 8. Distributed representation: a representation of the observed data in such a way that they are modeled as being generated by the interactions of many hidden factors. A particular factor learned from configurations of other factors can often generalize well. Distributed representations form the basis of deep learning.
--

A. Generative architectures

Associated with this generative category, we often see “unsupervised feature learning”, since the labels for the data are not of concern. When applying generative architectures for pattern recognition (i.e., supervised learning), a key concept here is (unsupervised) pre-training. This concept arises from the need to learn deep networks but learning the lower levels of such networks is difficult due to the gradient dilution problem; see (Bengio, 2009). Therefore, it is desirable to learn each lower layer without relying on all the layers above and to learn all layers in a greedy, layer-by-layer manner from bottom up. This is the gist of “pre-training” before subsequent learning of all layers together.

Among the various subclasses of generative deep architecture, the energy-based deep models including autoencoders are the most common (e.g., LeCun et al., 2007; Ranzato et al., 2006; Ngiam et al., 2011a; Bengio, 2009). The original form of the deep autoencoder (Hinton and Salakhutdinov, 2006; Deng et al., 2010), which we will give more detail about in Section IV, is a typical example of the generative model category. Most other forms of deep autoencoders are also generative in nature, but with quite different properties and implementations. Examples are transforming auto-encoders (Hinton et al., 2010), predictive sparse coders and their stacked version, and de-noising autoencoders and their stacked versions (Vincent et al., 2010).

Specifically, in de-noising autoencoders, the input vectors are first corrupted; e.g., randomizing a percentage of the inputs and setting them to zeros. Then one designs the hidden encoding nodes to reconstruct the original, uncorrupted input data using criteria such as KL distance between the original inputs and the reconstructed inputs. Uncorrupted encoded representations are used as the inputs to the next level of the stacked de-noising autoencoder.

Another prominent type of generative model is deep Boltzmann machine or DBM (Salakhutdinov and Hinton, 2009, 2012; Srivastava and Salakhutdinov, 2012). A DBM contains many layers of hidden variables, and has no connections between the variables within the same layer. This is a special case of the general Boltzmann machine (BM), which is a network of symmetrically connected units that make stochastic decisions about whether to be on or off. While having very simple learning algorithm, the general BMs are very complex to study and very slow to compute in learning. In a DBM, each layer captures complicated, higher-order correlations between the activities of hidden features in the layer below. DBMs have the potential of learning internal representations that become increasingly complex, highly desirable for solving object and speech recognition problems. Further, the high-level representations can be built from a large supply of unlabeled sensory inputs and very limited labeled data can then be used to only slightly fine-tune the model for a specific task at hand.

When the number of hidden layers of DBM is reduced to one, we have Restricted Boltzmann Machine (RBM). Like DBM, there are no hidden-to-hidden and no visible-to-visible connections. The main virtue of RBM is that via composing many RBMs, many hidden layers can be learned efficiently using the feature activations of one RBM as the training data for the next. Such composition leads to Deep Belief Network (DBN), which we will describe in more detail, together with RBMs, in Section V.

The standard DBN has been extended to the factored higher-order Boltzmann machine in its bottom layer, with strong results for phone recognition obtained (Dahl et al., 2010). This model, called mean-covariance RBM or mcRBM, recognizes the limitation of the standard RBM in its ability to represent the covariance structure of the data. However, it is very difficult to train mcRBM and to use it at the higher levels of the deep architecture. Further, the strong results published are not easy to reproduce. In the architecture of (Dahl et al., 2010), the mcRBM parameters in the full DBN are not easy to be fine-tuned using the discriminative information as for the regular RBMs in the higher layers. However, recent work showed that when better features are used, e.g. cepstral speech features subject to linear discriminant analysis or to fMLLR transformation, then the mcRBM is not needed as covariance in the transformed data is already modeled (Mohamed et al., 2012b).

Another representative deep generative architecture is the sum-product network or SPN (Poon and Domingo, 2011; Gens and Domingo, 2012). An SPN is a directed acyclic graph with the data as leaves, and with sum and product operations as internal nodes in the deep architecture. The “sum” nodes give mixture models, and the “product” nodes build up the feature hierarchy. Properties of “completeness” and “consistency” constrain the SPN in a desirable way. The learning of SPN is carried out using the EM algorithm together with back-propagation. The learning procedure starts with a dense SPN. It then finds a SPN structure by learning its weights, where zero weights remove the connections. The main difficulty in learning is found to be the common one --- the learning signal (i.e., the gradient) quickly dilutes when it propagates to deep layers. Empirical solutions have been found to mitigate this difficulty reported in (Poon and Domingo, 2011), where it was pointed out that despite the many desirable generative properties in the SPN, it is difficult to fine tune its weights using the discriminative information, limiting its effectiveness in classification tasks. This difficulty has been overcome in the subsequent work reported in (Gens and Domingo, 2012), where an efficient backpropagation-style discriminative training algorithm for SPN was presented. It was pointed out that the standard gradient descent, computed by the derivative of the conditional likelihood, suffers from the same gradient diffusion problem well known for the regular deep networks. But when marginal inference is replaced by inferring the most probable state of the hidden variables, such a “hard” gradient descent can reliably estimate deep SPNs’ weights. Excellent results on (small-scale) image recognition tasks are reported.

Recurrent neural networks (RNNs) can be regarded as a class of deep generative architectures when they are used to model and generate sequential data (e.g., Sutskever et al., 2011). The “depth” of an RNN can be as large as the length of the input data sequence. RNNs are very powerful for modeling sequence data (e.g., speech or text), but until recently they had not been widely used partly because they are extremely difficult to train properly due to the well-known “vanishing gradient” problem. Recent advances in Hessian-free optimization (Martens, 2010) have partially overcome this difficulty using second-order information or stochastic curvature estimates. In the recent work of (Martens and Sutskever, 2011), RNNs that are trained with Hessian-free optimization are used as a generative deep architecture in the character-level language modeling tasks, where gated connections are introduced to allow the current input characters to predict the transition from one latent state vector to the next. Such generative RNN models are demonstrated to be well capable of generating sequential text characters. More recently, Bengio et al. (2013) and Sutskever (2013) have explored new optimization methods in training generative RNNs that modify stochastic gradient descent and show these modifications can outperform Hessian-free optimization methods. Mikolov et al. (2010) have reported

excellent results on using RNNs for language modeling, which we will review in Section VII.C.

As examples of a different type of generative deep models, there has been a long history in speech recognition research where human speech production mechanisms are exploited to construct dynamic and deep structure in probabilistic generative models; for a comprehensive review, see book (Deng, 2006). Specifically, the early work described in (Deng 1992, 1993; Deng et al., 1994; Ostendorf et al., 1996, Deng and Sameti, 1996) generalized and extended the conventional shallow and conditionally independent HMM structure by imposing dynamic constraints, in the form of polynomial trajectory, on the HMM parameters. A variant of this approach has been more recently developed using different learning techniques for time-varying HMM parameters and with the applications extended to speech recognition robustness (Yu and Deng, 2009; Yu et al., 2009). Similar trajectory HMMs also form the basis for parametric speech synthesis (Zen et al., 2011; Zen et al., 2012; Ling et al., 2013; Shannon et al., 2013). Subsequent work added a new hidden layer into the dynamic model so as to explicitly account for the target-directed, articulatory-like properties in human speech generation (Deng and Ramsay, 1997; Bridle et al., 1998; Deng, 1999; Picone et al., 1999; Deng, 2003; Minami et al., 2002; Deng and Huang, 2004). More efficient implementation of this deep architecture with hidden dynamics is achieved with non-recursive or FIR filters in more recent studies (Deng et al., 2006, 2006a; Deng and Yu, 2007). The above deep-structured generative models of speech can be shown as special cases of the more general dynamic Bayesian network model and even more general dynamic graphical models (Bilmes and Bartels, 2005; Bilmes, 2010). The graphical models can comprise many hidden layers to characterize the complex relationship between the variables in speech generation. Armed with powerful graphical modeling tool, the deep architecture of speech has more recently been successfully applied to solve the very difficult problem of single-channel, multi-talker speech recognition, where the mixed speech is the visible variable while the un-mixed speech becomes represented in a new hidden layer in the deep generative architecture (Rennie et al., 2010; Wohlmayr et al., 2011). Deep generative graphical models are indeed a powerful tool in many applications due to their capability of embedding domain knowledge. However, in addition to the weakness of using non-distributed representations for the classification categories, they also are often implemented with inappropriate approximations in inference, learning, prediction, and topology design, all arising from inherent intractability in these tasks for most real-world applications. This problem has been partly addressed in the recent work of (Stoyanov et al., 2011), which provides an interesting direction for making deep generative graphical models potentially more useful in practice in the future.

The standard statistical methods used for large-scale speech recognition and understanding combine (shallow) hidden Markov models for speech acoustics with higher layers of structure representing different levels of natural language hierarchy. This combined hierarchical model can be suitably regarded as a deep generative architecture, whose motivation and some technical detail may be found in Chapter 7 the recent book (Kurzweil, 2012) on “Hierarchical HMM” or HHMM. Related models with greater technical depth and mathematical treatment can be found in (Fine et al., 1998) for HHMM and (Oliver et al., 2004) for Layered HMM. These early deep models were formulated as directed graphical models, missing the key aspect of “distributed representation” embodied in the more recent deep generative architectures of DBN and DBM discussed earlier in this section.

Finally, dynamic or temporally recursive generative models for non-speech applications based on deep neural network architectures can be found in (Taylor et al., 2007) for human motion modeling, and in (Socher et al., 2011) for natural language and natural scene parsing. The latter model is particularly interesting because the learning algorithms are capable of automatically determining the optimal model structure. This contrasts with other deep architectures such as DBN where only the parameters are learned while the architectures need to be pre-defined. Specifically, as reported in (Socher et al., 2011), the recursive structure commonly found in natural scene images and in natural language sentences can be discovered using a max-margin structure prediction architecture. Not only the units contained in the images or sentences are identified but so is the way in which these units interact with each other to form the whole.

B. Discriminative architectures

Many of the discriminative techniques in signal and information processing apply to shallow architectures such as HMMs (e.g., Juang et al., 1997; Povey and Woodland, 2002; He et al., 2008; Jiang and Li, 2010; Xiao and Deng, 2010; Gibson and Hain, 2010) or conditional random fields or CRFs (e.g., Yang and Furui, 2009; Yu et al., 2010a; Hifny and Renals, 2009; Heintz et al., 2009; Zweig and Nguyen, 2009; Peng et al., 2009). Since a CRF is defined with the conditional probability on input data as well as on the output labels, it is intrinsically a shallow discriminative architecture. (Interesting equivalence between CRF and discriminatively trained Gaussian models and HMMs can be found in Heigold et al., 2011). More recently, deep-structured CRFs have been developed by stacking the output in each lower layer of the CRF, together with the original input data, onto its higher layer (Yu et al., 2010a). Various versions of deep-structured CRFs are usefully applied to phone recognition (Yu and Deng, 2010), spoken language identification (Yu et al., 2010), and natural language processing (Yu et al., 2010a). However, at least for the phone recognition task, the performance of deep-structured CRFs, which is purely discriminative (non-

generative), has not been able to match that of the hybrid approach involving DBN, which we will take on shortly.

The recent article of (Morgan, 2012) gives an excellent review on other major existing discriminative models in speech recognition based mainly on the traditional neural network or MLP architecture using back-propagation learning with random initialization. It argues for the importance of both the increased width of each layer of the neural networks and the increased depth. In particular, a class of deep neural network models forms the basis of the popular “tandem” approach, where a discriminatively learned neural network is developed in the context of computing discriminant emission probabilities for HMMs. For some representative recent work in this area, see (Pinto et al., 2011; Ketabdar and Bourlard, 2010). The tandem approach generates discriminative features for an HMM by using the activities from one or more hidden layers of a neural network with various ways of information combination, which can be regarded as a form of discriminative deep architectures (Morgan et al., 2005; Morgan 2012).

In the most recent work of (Deng et al, 2011; Deng et al., 2012a; Tur et al., 2012; Lena et al., 2012; Vinyals et al., 2012), a new deep learning architecture, sometimes called Deep Stacking Network (DSN), together with its tensor variant (Hutchinson et al, 2012, 2013) and its kernel version (Deng et al., 2012), are developed that all focus on discrimination with scalable, parallelizable learning relying on little or no generative component. We will describe this type of discriminative deep architecture in detail in Section V.

Recurrent neural networks (RNNs) have been successfully used as a generative model when the “output” is taken to be the predicted input data in the future, as discussed in the preceding sub-section. They can also be used as a discriminative model where the output is a label sequence associated with the input data sequence. Note that such discriminative RNNs were applied to speech a long time ago with limited success (e.g., Robinson, 1994). For training RNNs for discrimination, pre-segmented training data are typically required. Also, post-processing is needed to transform their outputs into label sequences. It is highly desirable to remove such requirements, especially the costly pre-segmentation of training data. Often a separate HMM is used to automatically segment the sequence during training, and to transform the RNN classification results into label sequences (Robinson, 1994). However, the use of HMM for these purposes does not take advantage of the full potential of RNNs.

An interesting method was proposed in (Graves et al., 2006, 2013; Graves, 2012) that enables the RNNs themselves to perform sequence classification, removing the need for pre-segmenting the training data and for post-processing the outputs. Underlying this method is the idea of interpreting RNN outputs as the conditional distributions over all possible

label sequences given the input sequences. Then, a differentiable objective function can be derived to optimize these conditional distributions over the correct label sequences, where no segmentation of data is required.

Another type of discriminative deep architecture is convolutional neural network (CNN), with each module consisting of a convolutional layer and a pooling layer. These modules are often stacked up with one on top of another, or with a DNN on top of it, to form a deep model. The convolutional layer shares many weights, and the pooling layer subsamples the output of the convolutional layer and reduces the data rate from the layer below. The weight sharing in the convolutional layer, together with appropriately chosen pooling schemes, endows the CNN with some “invariance” properties (e.g., translation invariance). It has been argued that such limited “invariance” or equi-variance is not adequate for complex pattern recognition tasks and more principled ways of handling a wider range of invariance are needed (Hinton et al., 2011). Nevertheless, the CNN has been found highly effective and been commonly used in computer vision and image recognition (LeCun et al., 1998; Ciresan et al., 2012; Le et al., 2012; Dean et al., 2012; Krizhevsky et al., 2012). More recently, with appropriate changes from the CNN designed for image analysis to that taking into account speech-specific properties, the CNN is also found effective for speech recognition (Abdel-Hamid et al., 2012, 2013; Sainath et al., 2013; Deng et al., 2013). We will discuss such applications in more detail in Section VII.

It is useful to point out that time-delay neural networks (TDNN, Lang et al., 1990) developed for early speech recognition are a special case of the CNN when weight sharing is limited to one of the two dimensions, i.e., time dimension. It was not until recently that researchers have discovered that time is the wrong dimension to impose “invariance” and frequency dimension is more effective in sharing weights and pooling outputs (Abdel-Hamid et al., 2012, 2013; Deng et al., 2013). Analysis and the underlying reasons are described in (Deng et al., 2013), together with a new strategy for designing the CNN’s pooling layer demonstrated to be more effective than nearly all previous CNNs in phone recognition.

It is also useful to point out that the model of hierarchical temporal memory (HTM, Hawkins and Blakeslee, 2004; Hawkins et al., 2010; George, 2008) is another variant and extension of the CNN. The extension includes the following aspects: 1) Time or temporal dimension is introduced to serve as the “supervision” information for discrimination (even for static images); 2) Both bottom-up and top-down information flow are used, instead of just bottom-up in the CNN; and 3) A Bayesian probabilistic formalism is used for fusing information and for decision making.

Finally, the learning architecture developed for bottom-up, detection-based speech recognition proposed in (Lee, 2004)

and developed further since 2004, notably in (Yu et al, 2012a; Siniscalchi et al., 2013, 2013a) using the DBN-DNN technique, can also be categorized in the discriminative deep architecture category. There is no intent and mechanism in this architecture to characterize the joint probability of data and recognition targets of speech attributes and of the higher-level phone and words. The most current implementation of this approach is based on multiple layers of neural networks using back-propagation learning (Yu et al, 2012). One intermediate neural network layer in the implementation of this detection-based framework explicitly represents the speech attributes, which are simplified entities from the “atomic” units of speech developed in the early work of (Deng and Sun, 1994). The simplification lies in the removal of the temporally overlapping properties of the speech attributes or articulatory-like features. Embedding such more realistic properties in the future work is expected to improve the accuracy of speech recognition further.

C. Hybrid generative-discriminative architectures

The term “hybrid” for this third category refers to the deep architecture that either comprises or makes use of both generative and discriminative model components. In many existing hybrid architectures published in the literature (e.g., Hinton and Salakhutdinov, 2006; Mohamed et al., 2010; Dahl et al., 2012; Sainath et al., 2012), the generative component is exploited to help with discrimination, which is the final goal of the hybrid architecture. How and why generative modeling can help with discrimination can be examined from two viewpoints:

- 1) The optimization viewpoint where generative models can provide excellent initialization points in highly nonlinear parameter estimation problems (The commonly used term of “pre-training” in deep learning has been introduced for this reason); and/or
- 2) The regularization perspective where generative models can effectively control the complexity of the overall model.

The study reported in (Erhan et al., 2010) provided an insightful analysis and experimental evidence supporting both of the viewpoints above.

When the generative deep architecture of DBN discussed in Subsection III.A is subject to further discriminative training, commonly called “fine-tuning” in the literature, we obtain an equivalent architecture of deep neural network (DNN, which is sometimes also called DBN or deep MLP in the literature). In the DNN or the hybrid DBN with fine tuning, the weights of the network are “pre-trained” from stacked RBMs or DBN instead of the usual random initialization. See (Mohamed et al., 2012) for a detailed explanation of the equivalence relationship and the use of the often confusing terminology. We will review details of the DNN in the context of

RBM/DBN pre-training as well as its interface with the most commonly used shallow generative architecture of HMM (DNN-HMM) in Section IV.

Another example of the hybrid deep architecture is developed in (Mohamed et al., 2010), where again the generative DBN is used to initialize the DNN weights but the fine tuning is carried out not using frame-level discriminative information (e.g., cross-entropy error criterion) but sequence-level one. This is a combination of the static DNN with the shallow discriminative architecture of CRF. Here, the overall architecture of DNN-CRF is learned using the discriminative criterion of the conditional probability of full label sequences given the input sequence data. It can be shown that such DNN-CRF is equivalent to a hybrid deep architecture of DNN and HMM whose parameters are learned jointly using the full-sequence maximum mutual information (MMI) between the entire label sequence and the input vector sequence. A closely related full-sequence training method is carried out with success for a shallow neural network (Kingbury, 2009) and for a deep one (Kingbury et al., 2012).

Here, it is useful to point out a connection between the above hybrid discriminative training and a highly popular MPE (minimum phone error) training technique for the HMM (Povey and Woodland, 2002). In the iterative MPE training procedure using extended Baum-Welch, the initial HMM parameters cannot be arbitrary. One commonly used initial parameter set is that previously trained generatively using Baum-Welch algorithm for maximum likelihood. Further, an interpolation term taking the values of generatively trained HMM parameters is needed in the extended Baum-Welch updating formula, which may be analogous to “fine tuning” in the DNN training discussed earlier. Such I-smoothing (Povey and Woodland, 2002) has a similar spirit to DBN pre-training in the “hybrid” DNN learning.

Along the line of using discriminative criteria to train parameters in generative models as in the above HMM training example, we here briefly discuss the same method applied to learning other generative architectures. In (Larochelle and Bengio, 2008), the generative model of RBM is learned using the discriminative criterion of posterior class/label probabilities when the label vector is concatenated with the input data vector to form the overall visible layer in the RBM. In this way, RBM can be considered as a stand-alone solution to classification problems and the authors derived a discriminative learning algorithm for RBM as a shallow generative model. In the more recent work of (Ranzato et al., 2011), the deep generative model of DBN with the gated MRF at the lowest level is learned for feature extraction and then for recognition of difficult image classes including occlusions. The generative ability of the DBN model facilitates the discovery of what information is captured and what is lost at each level of representation in the deep model, as demonstrated in (Ranzato et al., 2011). A related work on using the discriminative criterion of empirical

risk to train deep graphical models can be found in (Stoyanov et al., 2011).

A further example of the hybrid deep architecture is the use of the generative model of DBN to pre-train deep convolutional neural networks (deep DNN) (Lee et al., 2009, 2010, 2011; Abdel-Hamid et al., 2013). Like the fully-connected DNN discussed earlier, the DBN pre-training is also shown to improve discrimination of the deep CNN over random initialization.

The final example given here of the hybrid deep architecture is based on the idea and work of (Ney, 1999; He and Deng, 2011), where one task of discrimination (speech recognition) produces the output (text) that serves as the input to the second task of discrimination (machine translation). The overall system, giving the functionality of speech translation -- translating speech in one language into text in another language --- is a two-stage deep architecture consisting of both generative and discriminative elements. Both models of speech recognition (e.g., HMM) and of machine translation (e.g., phrasal mapping and non-monotonic alignment) are generative in nature. But their parameters are all learned for discrimination. The framework described in (He and Deng, 2011) enables end-to-end performance optimization in the overall deep architecture using the unified learning framework initially published in (He et al., 2008). This hybrid deep learning approach can be applied to not only speech translation but also all speech-centric and possibly other information processing tasks such as speech information retrieval, speech understanding, cross-lingual speech/text understanding and retrieval, etc. (e.g., Yamin et al., 2008; Tur et al., 2012; He and Deng, 2012, 2013; Deng et al., 2012; Deng et al., 2013a; He et al., 2013).

After briefly surveying a wide range of work in each of the three classes of deep architectures above, in the following three sections, I will elaborate on three prominent models of deep learning, one from each of the three classes. While ideally they should represent the most influential architectures giving state of the art performance, I have chosen the three that I am most familiar with as being responsible for their developments and that may serve the tutorial purpose well with the simplicity of the architectural and mathematical descriptions. The three architectures described in the following three sections may not be interpreted as the most representative and influential work in each of the three classes. For example, in the category of generative architectures, the highly complex deep architecture and generative training methods developed by and described in Le et al. (2002), which is beyond the scope of this tutorial, performs extremely well in image recognition. Likewise, in the category of discriminative architectures, the even more complex architecture and learning described in Kingsbury et al. (2012) gives the state of the art performance in large-scale speech recognition.

IV. GENERATIVE ARCHITECTURE: DEEP AUTOENCODER

A. Introduction

Deep autoencoder is a special type of DNN whose output is the data input itself, and is used for learning efficient encoding or dimensionality reduction for a set of data. More specifically, it is a nonlinear feature extraction method involving no class labels; hence generative. An autoencoder uses three or more layers in the neural network:

- An input layer of data to be efficiently coded (e.g., pixels in image or spectra in speech);
- One or more considerably smaller hidden layers, which will form the encoding.
- An output layer, where each neuron has the same meaning as in the input layer.

When the number of hidden layers is greater than one, the autoencoder is considered to be deep.

An auto-encoder is often trained using one of the many backpropagation variants (e.g., conjugate gradient method, steepest descent, etc.) Though often reasonably effective, there are fundamental problems with using back-propagation to train networks with many hidden layers. Once the errors get back-propagated to the first few layers, they become minuscule, and quite ineffective. This causes the network to almost always learn to reconstruct the average of all the training data. Though more advanced backpropagation methods (e.g., the conjugate gradient method) help with this to some degree, it still results in very slow learning and poor solutions. This problem is remedied by using initial weights that approximate the final solution. The process to find these initial weights is often called pretraining.

A successful pretraining technique developed in (Hinton et al., 2006) for training deep auto-encoders involves treating each neighboring set of two layers like an RBM for pre-training to approximate a good solution and then using a backpropagation technique to fine-tune so as to minimize the “coding” error. This training technique is applied to construct a deep autoencoder to map images to short binary code for fast, content-based image retrieval. It is also applied to coding documents (called semantic hashing), and to coding spectrogram-like speech features which we review below.

B. Use of deep autoencoder to extract speech features

Here we review the more recent work of (Deng et al., 2010) in developing a similar type of autoencoder for extracting bottleneck speech instead of image features. Discovery of efficient binary codes related to such features can also be used in speech information retrieval. Importantly, the potential benefits of using discrete representations of speech constructed by this type of deep autoencoder can be derived from an almost unlimited supply of unlabeled data in future-generation speech recognition and retrieval systems.

A deep generative model of patches of spectrograms that contain 256 frequency bins and 1, 3, 9, or 13 frames is illustrated in Fig. 1. An undirected graphical model called a Gaussian-binary restricted Boltzmann machine (RBM) is built that has one visible layer of linear variables with Gaussian noise and one hidden layer of 500 to 3000 binary latent variables. After learning the Gaussian-Binary RBM, the activation probabilities of its hidden units are treated as the data for training another binary-binary RBM. These two RBM's can then be composed to form a deep belief net (DBN) in which it is easy to infer the states of the second layer of binary hidden units from the input in a single forward pass. The DBN used in this work is illustrated on the left side of Fig. 1, where the two RBMs are shown in separate boxes. (See more detailed discussions on RBM and DBN in the next section.)

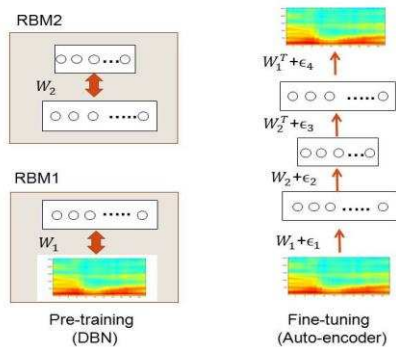


Fig. 1: The architecture of the deep autoencoder used in (Deng et al., 2010) for extracting “bottle-neck” speech features from high-resolution spectrograms.

The deep autoencoder with three hidden layers is formed by “unrolling” the DBN using its weight matrices. The lower layers of this deep autoencoder use the matrices to encode the input and the upper layers use the matrices in reverse order to decode the input. This deep autoencoder is then fine-tuned using back-propagation of error-derivatives to make its output as similar as possible to its input, as shown on the right side of Fig. 1. After learning is complete, any variable-length spectrogram can be encoded and reconstructed as follows. First, N consecutive overlapping frames of 256-point log power spectra are each normalized to zero-mean and unit-variance to provide the input to the deep autoencoder. The first hidden layer then uses the logistic function to compute real-valued activations. These real values are fed to the next, coding layer to compute “codes”. The real-valued activations of hidden units in the coding layer are quantized to be either zero or one with 0.5 as the threshold. These binary codes are then used to reconstruct the original spectrogram, where individual fixed-frame patches are reconstructed first using the two upper layers of network weights. Finally, overlap-and-add technique is used to reconstruct the full-length speech spectrogram from the outputs produced by applying the deep autoencoder to every possible window of N consecutive

frames. We show some illustrative encoding and reconstruction examples below.

C. Illustrative examples

At the top of Fig. 2 is the original speech, followed by the reconstructed speech utterances with forced binary values (zero or one) at the 312 unit code layer for encoding window lengths of $N=1, 3, 9,$ and $13,$ respectively. The lower coding errors for $N=9$ and $N=13$ are clearly seen.

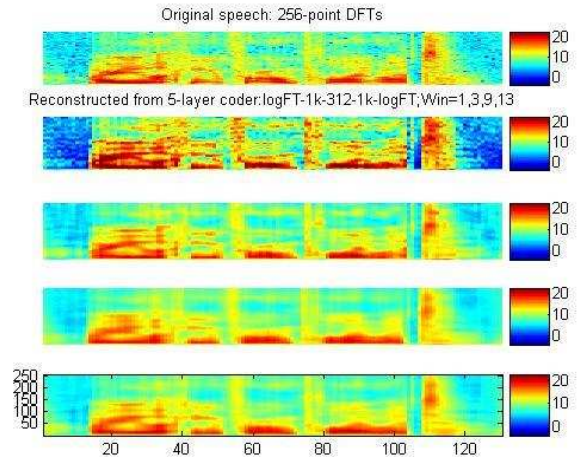


Fig. 2. Top to Bottom: Original spectrogram; reconstructions using input window sizes of $N= 1, 3, 9,$ and 13 while forcing the coding units to be zero or one (i.e., a binary code). The y-axis values indicate FFT bin numbers (i.e. 256-point FFT is used for constructing all spectrograms).

Encoding accuracy of the deep autoencoder is qualitatively examined to compare with the more traditional codes via vector quantization (VQ). Fig. 3 shows various aspects of the encoding accuracy. At the top is the original speech utterance’s spectrogram. The next two spectrograms are the blurry reconstruction from the 312-bit VQ and the much more faithful reconstruction from the 312-bit deep autoencoder. Coding errors from both coders, plotted as a function of time, are shown below the spectrograms, demonstrating that the auto-encoder (red curve) is producing lower errors than the VQ coder (blue curve) throughout the entire span of the utterance. The final two spectrograms show the detailed coding error distributions over both time and frequency bins.

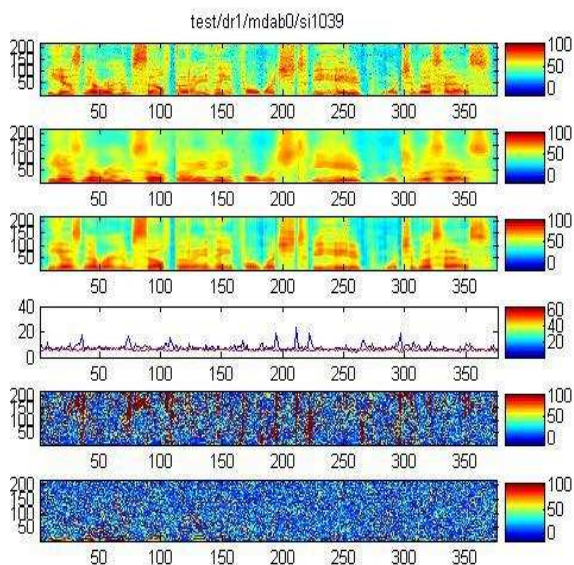


Fig 3. Top to bottom: Original spectrogram from the test set; reconstruction from the 312-bit VQ coder; reconstruction from the 312-bit auto-encoder; coding errors as a function of time for the VQ coder (blue) and auto-encoder (red); spectrogram of the VQ coder residual; spectrogram of the deep autoencoder’s residual.

D. Transforming Auto-encoder

The deep auto-encoder described above can extract a compact code for a feature vector due to its many layers and the non-linearity. But the extracted code would change unpredictably when the input feature vector is transformed. It is desirable to be able to have the code change predictably that reflects the underlying transformation invariant to the perceived content. This is the goal of transforming auto-encoder proposed in for image recognition (Hinton et al., 2011).

The building block of the transforming auto-encoder is a “capsule”, which is an independent sub-network that extracts a single parameterized feature representing a single entity, be it visual or audio. A transforming auto-encoder receives both an input vector and a target output vector which is related to the input vector by a simple global transformation; e.g., the translation of a whole image or frequency shift due to vocal tract length differences for speech. An explicit representation of the global transformation is known also. The bottleneck or coding layer of the transforming autoencoder consists of the outputs of several capsules.

During the training phase, the different capsules learn to extract different entities in order to minimize the error between the final output and the target.

In addition to the deep autoencoder architectures described above in detail in this section, there are many other types of generative architectures in the literature, all characterized by the use of data alone (i.e., free of classification labels) to automatically derive higher-level features. Although such more complex architectures have produced state of the art

results (e.g., Le et al., 2012), their complexity does not permit detailed treatment in this tutorial paper; rather I have included a brief survey of a broader range of the generative deep architectures in Section III.A.

V. HYBRID ARCHITECTURE: DEEP NEURAL NETWORK PRE-TRAINED WITH DEEP BELIEF NETWORK

A. Basics

In this section, we present the most widely studied hybrid deep architecture of deep neural networks (DNN), consisting of both pre-training (using generative deep belief network or DBN) and fine-tuning stages in its parameter learning. Part of this review is based on the recent publication of (Hinton et al., 2012; Yu and Deng, 2011) and (Dahl et al., 2012).

As the generative component of the DBN, it is a probabilistic model composed of multiple layers of stochastic, latent variables. The unobserved variables can have binary values and are often called hidden units or feature detectors. The top two layers have undirected, symmetric connections between them and form an *associative memory*. The lower layers receive top-down, directed connections from the layer above. The states of the units in the lowest layer, or the visible units, represent an input data vector.

There is an efficient, layer-by-layer procedure for learning the top-down, generative weights that determine how the variables in one layer depend on the variables in the layer above. After learning, the values of the latent variables in every layer can be inferred by a single, bottom-up pass that starts with an observed data vector in the bottom layer and uses the generative weights in the reverse direction.

DBNs are learned one layer at a time by treating the values of the latent variables in one layer, when they are being inferred from data, as the data for training the next layer. This efficient, greedy learning can be followed by, or combined with, other learning procedures that fine-tune all of the weights to improve the generative or discriminative performance of the full network. This latter learning procedure constitutes the discriminative component of the DBN as the hybrid architecture.

Discriminative fine-tuning can be performed by adding a final layer of variables that represent the desired outputs and back-propagating error derivatives. When networks with many hidden layers are applied to highly-structured input data, such as speech and images, back-propagation works much better if the feature detectors in the hidden layers are initialized by learning a DBN to model the structure in the input data as originally proposed in (Hinton and Salakhutdinov, 2006).

A DBN can be viewed as a composition of simple learning modules via stacking them. This simple learning module is

called restricted Boltzmann machines (RBMs) that we introduce next.

B. Restricted Boltzmann Machine

An RBM is a special type of Markov random field that has one layer of (typically Bernoulli) stochastic hidden units and one layer of (typically Bernoulli or Gaussian) stochastic visible or observable units. RBMs can be represented as bipartite graphs, where all visible units are connected to all hidden units, and there are no visible-visible or hidden-hidden connections.

In an RBM, the joint distribution $p(\mathbf{v}, \mathbf{h}; \theta)$ over the visible units \mathbf{v} and hidden units \mathbf{h} , given the model parameters θ , is defined in terms of an energy function $E(\mathbf{v}, \mathbf{h}; \theta)$ of

$$p(\mathbf{v}, \mathbf{h}; \theta) = \frac{\exp(-E(\mathbf{v}, \mathbf{h}; \theta))}{Z},$$

where $Z = \sum_{\mathbf{v}} \sum_{\mathbf{h}} \exp(-E(\mathbf{v}, \mathbf{h}; \theta))$ is a normalization factor or partition function, and the marginal probability that the model assigns to a visible vector \mathbf{v} is

$$p(\mathbf{v}; \theta) = \frac{\sum_{\mathbf{h}} \exp(-E(\mathbf{v}, \mathbf{h}; \theta))}{Z}.$$

For a Bernoulli (visible)-Bernoulli (hidden) RBM, the energy function is defined as

$$E(\mathbf{v}, \mathbf{h}; \theta) = - \sum_{i=1}^I \sum_{j=1}^J w_{ij} v_i h_j - \sum_{i=1}^I b_i v_i - \sum_{j=1}^J a_j h_j,$$

where w_{ij} represents the symmetric interaction term between visible unit v_i and hidden unit h_j , b_i and a_j the bias terms, and I and J are the numbers of visible and hidden units. The conditional probabilities can be efficiently calculated as

$$p(h_j = 1 | \mathbf{v}; \theta) = \sigma \left(\sum_{i=1}^I w_{ij} v_i + a_j \right),$$

$$p(v_i = 1 | \mathbf{h}; \theta) = \sigma \left(\sum_{j=1}^J w_{ij} h_j + b_i \right),$$

where $\sigma(x) = 1/(1 + \exp(x))$.

Similarly, for a Gaussian (visible)-Bernoulli (hidden) RBM, the energy is

$$E(\mathbf{v}, \mathbf{h}; \theta) = - \sum_{i=1}^I \sum_{j=1}^J w_{ij} v_i h_j - \frac{1}{2} \sum_{i=1}^I (v_i - b_i)^2 - \sum_{j=1}^J a_j h_j,$$

The corresponding conditional probabilities become

$$p(h_j = 1 | \mathbf{v}; \theta) = \sigma \left(\sum_{i=1}^I w_{ij} v_i + a_j \right),$$

$$p(v_i | \mathbf{h}; \theta) = \mathcal{N} \left(\sum_{j=1}^J w_{ij} h_j + b_i, 1 \right),$$

where v_i takes real values and follows a Gaussian distribution with mean $\sum_{j=1}^J w_{ij} h_j + b_i$ and variance one. Gaussian-Bernoulli RBMs can be used to convert real-valued stochastic variables to binary stochastic variables, which can then be further processed using the Bernoulli-Bernoulli RBMs.

The above discussion used two most common conditional distributions for the visible data in the RBM --- Gaussian (for continuous-valued data) and binomial (for binary data). More general types of distributions in the RBM can also be used. See (Welling et al., 2005) for the use of general exponential-family distributions for this purpose.

Taking the gradient of the log likelihood $\log p(\mathbf{v}; \theta)$ we can derive the update rule for the RBM weights as:

$$\Delta w_{ij} = E_{data}(v_i h_j) - E_{model}(v_i h_j),$$

where $E_{data}(v_i h_j)$ is the expectation observed in the training set and $E_{model}(v_i h_j)$ is that same expectation under the distribution defined by the model. Unfortunately, $E_{model}(v_i h_j)$ is intractable to compute so the contrastive divergence (CD) approximation to the gradient is used where $E_{model}(v_i h_j)$ is replaced by running the Gibbs sampler initialized at the data for one full step. The steps in approximating $E_{model}(v_i h_j)$ is as follows:

- Initialize \mathbf{v}_0 at data
- Sample $\mathbf{h}_0 \sim p(\mathbf{h} | \mathbf{v}_0)$
- Sample $\mathbf{v}_1 \sim p(\mathbf{v} | \mathbf{h}_0)$
- Sample $\mathbf{h}_1 \sim p(\mathbf{h} | \mathbf{v}_1)$

Then $(\mathbf{v}_1, \mathbf{h}_1)$ is a sample from the model, as a very rough estimate of $E_{model}(v_i h_j) = (v_\infty, h_\infty)$, which is a true sample from the model. Use of $(\mathbf{v}_1, \mathbf{h}_1)$ to approximate $E_{model}(v_i h_j)$ gives rise to the algorithm of CD-1. The sampling process can be pictorially depicted as below in Fig. 4 below.

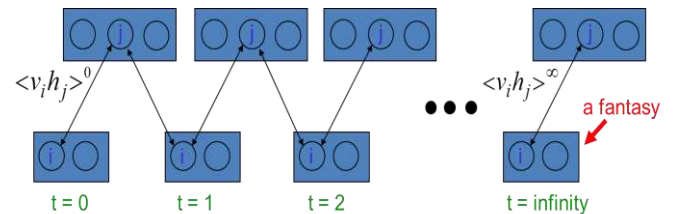


Fig. 4. A pictorial view of sampling from a RBM during the “negative” learning phase of the RBM (courtesy of G. Hinton)

Careful training of RBMs is essential to the success of applying RBM and related deep learning techniques to solve practical problems. See the Technical Report (Hinton 2010) for a very useful practical guide for training RBMs.

The RBM discussed above is a generative model, which characterizes the input data distribution using hidden variables and there is no label information involved. However, when the label information is available, it can be used together with the data to form the joint “data” set. Then the same CD learning can be applied to optimize the approximate “generative” objective function related to data likelihood. Further, and more interestingly, a “discriminative” objective function can be defined in terms of conditional likelihood of labels. This discriminative RBM can be used to “fine tune” RBM for classification tasks (Larochelle and Bengio, 2008).

Note the SESM architecture by Ranzato et al. (2007a) surveyed in Section III is quite similar to the RBM described above. While they both have a symmetric encoder and decoder, and a logistic non-linearity on the top of the encoder, the main difference is that RBM is trained using (approximate) maximum likelihood, but SESM is trained by simply minimizing the average energy plus an additional code sparsity term. SESM relies on the sparsity term to prevent flat energy surfaces, while RBM relies on an explicit contrastive term in the loss, an approximation of the log partition function. Another difference is in the coding strategy in that the code units are “noisy” and binary in RBM, while they are quasi-binary and sparse in SESM.

C. Stacking up RBMs to form a DBN/DNN architecture

Stacking a number of the RBMs learned layer by layer from bottom up gives rise to a DBN, an example of which is shown in Fig. 5. The stacking procedure is as follows. After learning a Gaussian-Bernoulli RBM (for applications with continuous features such as speech) or Bernoulli-Bernoulli RBM (for applications with nominal or binary features such as black-white image or coded text), we treat the activation probabilities of its hidden units as the data for training the Bernoulli-Bernoulli RBM one layer up. The activation probabilities of the second-layer Bernoulli-Bernoulli RBM are then used as the visible data input for the third-layer Bernoulli-Bernoulli RBM, and so on. Some theoretical justification of this efficient layer-by-layer greedy learning strategy is given in (Hinton et al., 2006), where it is shown that the *stacking* procedure above improves a variational lower bound on the likelihood of the training data under the composite model. That is, the greedy procedure above achieves approximate maximum likelihood learning. Note that this learning procedure is unsupervised and requires no class label.

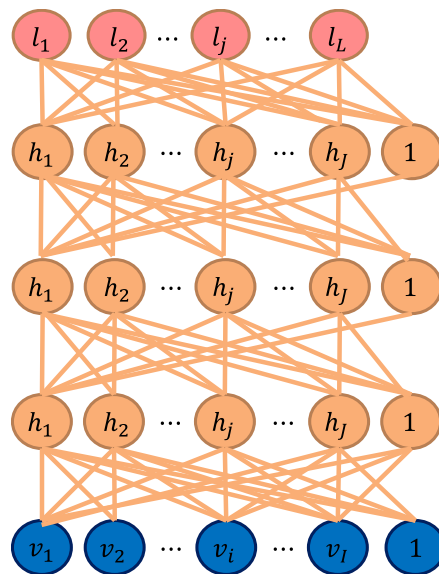


Fig. 5. Illustration of a DBN/DNN architecture.

When applied to classification tasks, the generative pre-training can be followed by or combined with other, typically discriminative, learning procedures that fine-tune all of the weights jointly to improve the performance of the network. This discriminative fine-tuning is performed by adding a final layer of variables that represent the desired outputs or labels provided in the training data. Then, the back-propagation algorithm can be used to adjust or fine-tune the DBN weights and use the final set of weights in the same way as for the standard feed-forward neural network. What goes to the top, label layer of this DNN depends on the application. For speech recognition applications, the top layer, denoted by “ $l_1, l_2, \dots, l_j, \dots, l_L$,” in Fig. 5, can represent either syllables, phones, sub-phones, phone states, or other speech units used in the HMM-based speech recognition system.

The generative pre-training described above has produced excellent phone and speech recognition results on a wide variety of tasks, which will be surveyed in Section VII. Further research has also shown the effectiveness of other pre-training strategies. As an example, greedy layer-by-layer training may be carried out with an additional discriminative term to the generative cost function at each level. And without generative pre-training, purely discriminative training of DNNs from random initial weights using the traditional stochastic gradient descent method has been shown to work very well when the scales of the initial weights are set carefully and the mini-batch sizes, which trade off noisy gradients with convergence speed, used in stochastic gradient descent are adapted prudently (e.g., with an increasing size over training epochs). Also, randomization order in creating mini-batches needs to be judiciously determined. Importantly, it was found effective to learn a DNN by starting with a shallow neural net with a single hidden layer. Once this has been trained discriminatively (using early stops to avoid overfitting), a second hidden layer is inserted between the first

hidden layer and the labeled softmax output units and the expanded deeper network is again trained discriminatively. This can be continued until the desired number of hidden layers is reached, after which a full backpropagation “fine tuning” is applied. This discriminative “pre-training” procedure is found to work well in practice (e.g., Seide et al., 2011).

This type of discriminative “pre-training” procedure is closely related to the learning algorithm developed for the deep architectures called deep convex/stacking network, to be described in Section VI, where interleaving linear and nonlinear layers are used in building up the deep architectures in a modular manner, and the original input vectors are concatenated with the output vectors of each module consisting of a shallow neural net. Discriminative “pre-training” is used for positioning a subset of weights in each module in a reasonable space using parallelizable convex optimization, followed by a batch-mode “fine tuning” procedure which is also parallelizable due to the closed-form constraint between two subsets of weights in each module.

It is useful to point out that historically the use of the generative model of DBN to facilitate the training of DNNs played an important role in igniting the interest of deep learning for speech feature coding and for speech recognition (Deng et al., 2010; Dahl et al., 2011, 2012; Hinton et al., 2012). After this effectiveness was demonstrated, however, further research showed many other alternative, simpler ways. With a large amount of training data, we can now effectively learn a DNN by starting with a shallow neural network (i.e., with one hidden layer). After this net has been trained discriminatively, a new hidden layer is inserted between the previous hidden layer and the softmax output layer. Then the full network can be again discriminatively trained. This process can be continued until the desired number of hidden layers is reached. And finally, full backpropagation fine-tuning is carried out to complete the DNN training. This procedure is sometimes referred to as “discriminative pre-training” as reviewed and discussed in some recent papers including (Hinton et al., 2012).

Further, purely discriminative training of the full DNN from random initial weights is now known to work much better than had been thought in early days, provided that the scales of the initial weights are set carefully, a large amount of labeled training data is available, and mini-batch sizes over training epochs are set appropriately. Nevertheless, generative pretraining still improves test performance, sometimes by a significant amount especially for small tasks. Layer-by-layer generative pretraining was originally done using RBMs, but various types of autoencoder with one hidden layer can also be used.

D. Interfacing DNN with HMM

A DBN/DNN discussed above is a static classifier with input vectors having a fixed dimensionality. However, many

practical pattern recognition and information processing problems, including speech recognition, machine translation, natural language understanding, video processing and bio-information processing, require sequence recognition. In sequence recognition, sometimes called classification with structured input/output, the dimensionality of both inputs and outputs are variable.

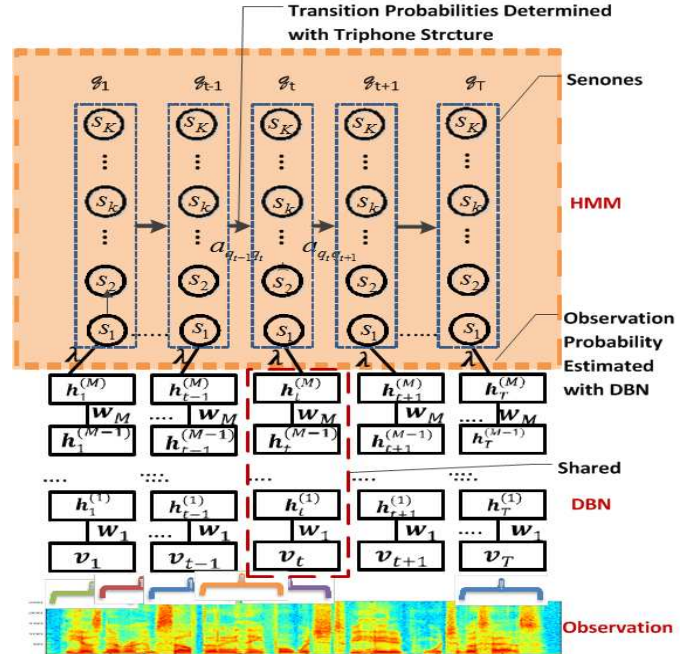


Fig. 6. Interface between DBN-DNN and HMM to form a DNN-HMM. This architecture has been successfully used in speech recognition experiments reported in (Dahl et al., 2012).

The HMM, based on dynamic programming operations, is a convenient tool to help port the strength of a static classifier to handle dynamic or sequential patterns. Thus, it is natural to combine DBN/DNN and HMM to bridge the gap between static and sequence pattern recognition. An architecture that shows the interface between a DNN and HMM is provided in Fig. 6. This architecture has been successfully used in speech recognition experiments as reported in (Dahl et al., 2012).

It is important to note that the unique elasticity of temporal dynamic of speech as elaborated in (Deng, 2006) would require temporally-correlated models better than HMM for the ultimate success of speech recognition. Integrating such dynamic models having realistic co-articulatory properties with the DNN and possibly other deep learning models to form the coherent dynamic deep architecture is a challenging new research.

VI. DISCRIMINATIVE ARCHITECTURE: DEEP STACKING NETWORK AND ITS TENSOR VERSION

A. Introduction

While the DNN just reviewed has been shown to be extremely powerful in connection with performing recognition and classification tasks including speech recognition and image classification, training a DBN has proven to be more difficult computationally. In particular, conventional techniques for training DNN at the fine tuning phase involve the utilization of a stochastic gradient descent learning algorithm, which is extremely difficult to parallelize across machines. This makes learning at large scale practically impossible. For example, it has been possible to use one single, very powerful GPU machine to train DNN-based speech recognizers with dozens to a few hundreds of hours of speech training data with remarkable results. It is very difficult, however, to scale up this success with thousands or more hours of training data.

Here we describe a new deep learning architecture, Deep Stacking Network (DSN), which attacks the learning scalability problem. This section is based in part on the recent publications of (Deng and Yu, 2011; Deng et al., 2012; Hutchinson et al., 2012, 2013) with expanded discussions.

The central idea of DSN design relates to the concept of stacking, as proposed originally in (Wolpert, 1992), where simple modules of functions or classifiers are composed first and then they are “stacked” on top of each other in order to learn complex functions or classifiers. Various ways of implementing stacking operations have been developed in the past, typically making use of supervised information in the simple modules. The new features for the stacked classifier at a higher level of the stacking architecture often come from concatenation of the classifier output of a lower module and the raw input features. In (Cohen and de Carvalho, 2005), the simple module used for stacking was a conditional random field (CRF). This type of deep architecture was further developed with hidden states added for successful natural language and speech recognition applications where segmentation information is unknown in the training data (Yu et al., 2010a). Convolutional neural networks, as in (Jarrett, 2009), can also be considered as a stacking architecture but the supervision information is typically not used until in the final stacking module.

The DSN architecture was originally presented in (Deng and Yu, 2011), which also used the name Deep Convex Network or DCN to emphasize the convex nature of the principal learning algorithm used for learning the network, and discussed in this section makes use of supervision information for stacking each of the basic modules, which takes the simplified form of multilayer perceptron. In the basic module, the output units are linear and the hidden units are sigmoidal nonlinear. The linearity in the output units permits highly efficient, parallelizable, and closed-form estimation (a result of convex optimization) for the output network weights given the hidden units’ activities. Due to the closed-form constraints between the input and output weights, the input weights can

also be elegantly estimated in an efficient, parallelizable, batch-mode manner.

The name “convex” used in (Deng and Yu, 2011) accentuates the role of convex optimization in learning the output network weights given the hidden units’ activities in each basic module. It also points to the importance of the closed-form constraints, derived from the convexity, between the input and output weights. Such constraints make the learning the remaining network parameters (i.e., the input network weights) much easier than otherwise, enabling batch-mode learning of DSN that can be distributed over CPU clusters. And in more recent publications, DSN was used when the key operation of stacking is emphasized.

B. An architectural overview of deep stacking network

A DSN, shown in Fig. 7, includes a variable number of layered modules, wherein each module is a specialized neural network consisting of a single hidden layer and two trainable sets of weights. In Fig. 7, only four such modules are illustrated, where each module is shown with a separate color. (In practice, up to a few hundreds of modules have been efficiently trained and used in image and speech classification experiments.)

The lowest module in the DSN comprises a first linear layer with a set of linear input units, a non-linear layer with a set of non-linear hidden units, and a second linear layer with a set of linear output units. For instance, if the DSN is utilized in connection with recognizing an image, the input units can correspond to a number of pixels (or extracted features) in the image, and can be assigned values based at least in part upon intensity values, RGB values, or the like corresponding to the respective pixels. If the DSN is utilized in connection with speech recognition, the set of input units may correspond to samples of speech waveform, or the extracted features from speech waveforms, such as power spectra or cepstral coefficients.

The hidden layer of the *lowest module* of a DSN comprises a set of non-linear units that are mapped to the input units by way of a first, lower-layer weight matrix, which we denote by \mathbf{W} . For instance, the weight matrix may comprise a plurality of randomly generated values between zero and one, or the weights of an RBM trained separately. The non-linear units may be sigmoidal units that are configured to perform non-linear operations on weighted outputs from the input units (weighted in accordance with the first weight matrix \mathbf{W}).

The second, *linear* layer in any module of a DSN includes a set of output units that are representative of the targets of classification. For instance, if the DSN is configured to perform digit recognition, then the plurality of output units may be representative of the values 1, 2, 3, and so forth up to 10 with a 0-1 coding scheme. If the DSN is configured to perform speech recognition, then the output units may be

representative of phones, HMM states of phones, or context-dependent HMM states of phones. The non-linear units in each module of the DSN may be mapped to a set of the linear output units by way of a second, upper-layer weight matrix, which we denote by U . This second weight matrix can be learned by way of a batch learning process, such that learning can be undertaken in parallel. Convex optimization can be employed in connection with learning U . For instance, U can be learned based at least in part upon the first weight matrix W , values of the coded classification targets, and values of the input units.

As indicated above, the DSN includes a set of serially connected, overlapping, and layered modules, wherein each module includes the aforementioned three layers -- a first linear layer that includes a set of linear input units whose number equals the dimensionality of the input features, a hidden layer that comprises a set of non-linear units whose number is a tunable hyper-parameter, and a second linear layer that comprises a plurality of linear output units whose number equals that of the target classification classes (e.g., the total number of context-dependent phones clustered by a decision tree used in). The modules are referred to herein as being layered because the output units of a lower module are a subset of the input units of an adjacent higher module in the DSN. More specifically, in a second module that is directly above the lowest module in the DSN, the input units can include the output units of the lower module(s). The input units can additionally include the raw training data – in other words, the output units of the lowest module can be appended to the input units in the second module, such that the input units of the second module also include the output units of the lowest module.

The pattern discussed above of including output units in a lower module as a portion of the input units in an adjacent higher module in the DSN and thereafter learning a weight matrix that describes connection weights between hidden units and linear output units via convex optimization can continue for many modules. A resultant learned DSN may then be deployed in connection with an automatic classification task such as frame-level speech phone or state classification. Connecting DSN’s output to an HMM or any dynamic programming device enables continuous speech recognition and other forms of sequential pattern recognition.

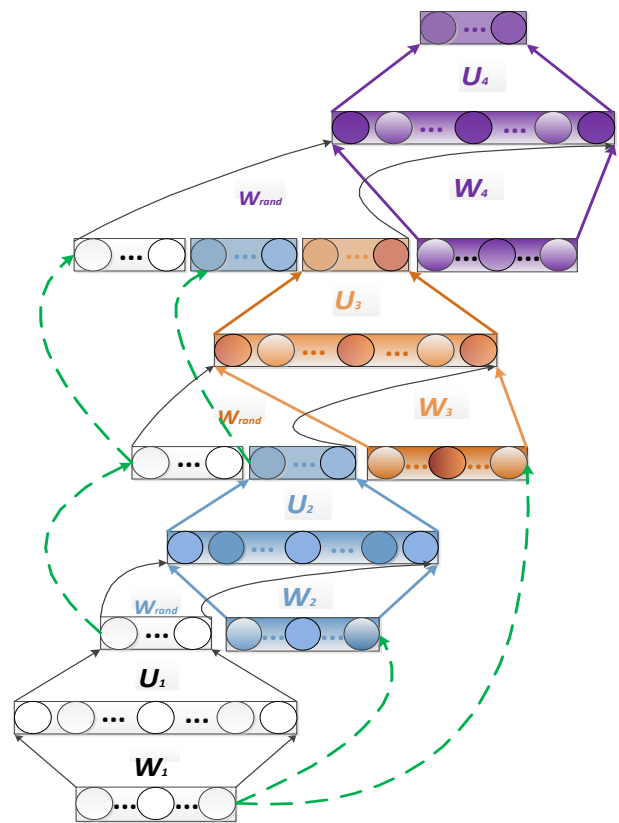


Fig. 7: A DSN architecture. Only four modules are illustrated, each with a distinct color. Dashed lines denote copying layers.

C. Tensorized deep stacking network

The above DSN architecture has recently been generalized to its tensorized version, which we call TDSN (Hutchinson et al., 2012, 2013). It has the same scalability as DSN in terms of parallelizability in learning, but it generalizes DSN by providing higher-order feature interactions missing in DSN.

The architecture of TDSN is similar to that of DSN in the way that stacking operation is carried out. That is, modules of the TDSN are stacking up in a similar way to form a deep architecture. The differences of TDSN and DSN lie mainly in how each module is constructed. In DSN, we have one set of hidden units forming a hidden layer, as denoted at the left panel of Fig. 8. In contrast, each module of a TDSN contains two independent hidden layers, denoted as “Hidden 1” and “Hidden 2” in the middle and right panels of Fig. 8. As a result of this difference, the upper-layer weights, denoted by “ U ” in Fig. 8, changes from a matrix (a two dimensional array) in DSN to a tensor (a three dimensional array) in TDSN, shown as a cube labeled by “ U ” in the middle panel.

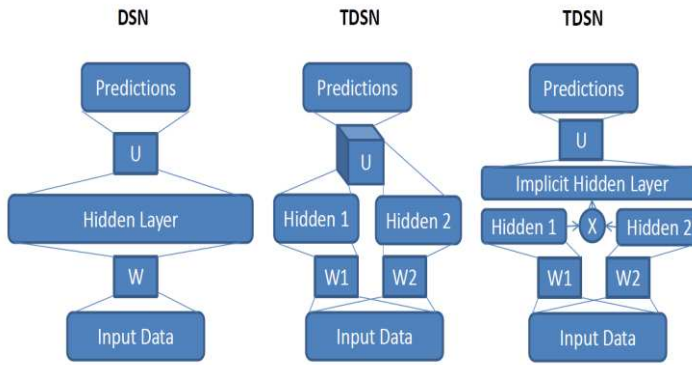


Fig. 8: Comparisons of one single module of a DSN (left) and that of a tensorized-DSN (TDSN). Two equivalent forms of a TDSN module are shown to the right.

The tensor U has a three-way connection, one to the prediction layer and the remaining to the two separate hidden layers. An equivalent form of this TDSN module is shown in the right panel of Fig. 8, where the implicit hidden layer is formed by expanding the two separate hidden layers into their outer product. The resulting large vector contains all possible pair-wise products for the two sets of hidden-layer vectors. This turns tensor U into a matrix again whose dimensions are 1) size of the prediction layer; and 2) product of the two hidden layers' sizes. Such equivalence enables the same convex optimization for learning U developed for DSN to be applied to learning tensor U . Importantly, higher-order hidden feature interactions are enabled in TDSN via the outer product construction for the large, implicit hidden layer.

Stacking TDSN modules to form a deep architecture pursues in a similar way to DSN by concatenating various vectors. Two examples are shown in Figs. 9 and 10. Note stacking by concatenating hidden layers with input (Fig. 10) would be difficult for DSN since its hidden layer tends to be too large for practical purposes.

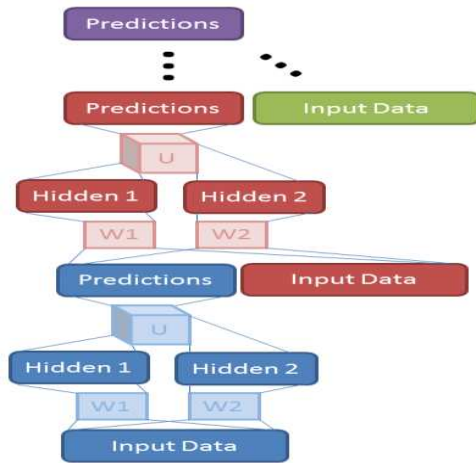


Fig. 9: Stacking of TDSN modules by concatenating prediction vector with input vector.

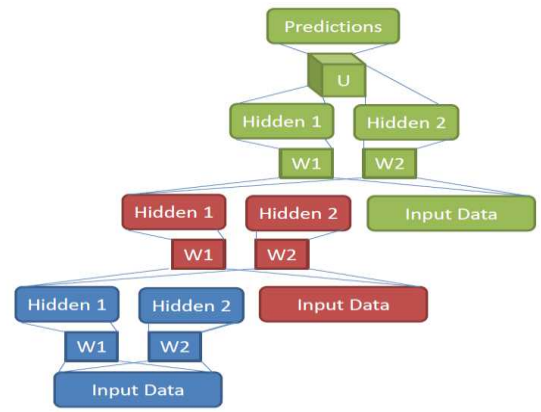


Fig. 10: Stacking of TDSN modules by concatenating two hidden-layers' vectors with the input vector.

VII. APPLICATIONS OF DEEP LEARNING TO SIGNAL AND INFORMATION PROCESSING

In the expanded technical scope of signal processing, the *signal* is endowed with not only the traditional types such as audio, speech, image and video, but also text, language, and document that convey high-level, semantic information for human consumption. In addition, the scope of *processing* has been extended from the conventional coding, enhancement, analysis, and recognition to include more human-centric tasks of interpretation, understanding, retrieval, mining, and user interface (Deng, 2008). Signal processing researchers have been working on one or more of the signal processing areas defined by the matrix constructed with the two axes of *signal* and *processing* discussed here. The deep learning techniques discussed in this article have recently been applied to a large number of traditional and extended signal processing areas, with the most recent interesting application of predicting protein structure (Lena et al., 2012, which we will not cover here). We now provide a brief survey of this body of work in four main categories pertaining closely to signal and information processing.

A. Speech and audio

The traditional neural network or MLP has been in use for speech recognition for many years. When used alone, its performance is typically lower than the state-of-the-art HMM systems with observation probabilities approximated with Gaussian mixture models (GMMs). Recently, the deep learning technique was successfully applied to phone recognition (Mohamed et al., 2009, 2010, 2012; Sivaram and Hermansky, 2012) and large vocabulary speech recognition tasks (Yu et al., 2012; Seide et al., 2011, 2011a; Dahl et al., 2011, 2012; Kubo et al., 2012) by integrating the powerful discriminative training ability of the DNNs with the sequential modeling ability of the HMMs.

More specifically, in the work of (Mohamed et al., 2009), a five-layer DNN (called DBN in the paper) was used to replace

the Gaussian mixture component of the GMM-HMM and the monophone state was used as the modeling unit. Although monophones are generally accepted as a weaker phonetic representation than triphones, the DBN-HMM approach with monophones was shown to achieve higher phone recognition accuracy than the state-of-the-art triphone GMM-HMM systems.

The technique of (Mohamed et al., 2009) was improved in the later work reported in (Mohamed et al., 2010) by using the CRF instead of the HMM to model the sequential speech data and by applying the maximum mutual information (MMI) training technique successfully developed in speech recognition to the resultant DBN-CRF training. The sequential discriminative learning technique developed jointly optimizes the DBN weights, transition weights, and phone language model and achieved higher accuracy than the DBN-HMM phone recognizer with the frame-discriminative training criterion implicit in the DBN's fine tuning procedure implemented in (Mohamed et al., 2009).

Here we elaborate on the method of (Mohamed et al., 2010). The early phonetic recognition experiments made use of the standard frame-based objective function in static pattern classification, cross-entropy, to optimize the DNN weights. The transition parameters and language model scores were obtained from an HMM-like approach and were trained independently of the DNN weights. However, it has been known during the long history of HMM research that sequence classification criteria can be very helpful in improving speech and phone recognition accuracy. This is because the sequence classification criteria are more directly correlated with the performance measure (e.g., the overall word or phone error rate) than cross entropy.

The benefit of using such sequence classification criteria was shown on shallow neural networks in (Kingsbury 2009; Prabhavalkar and Fosler-Lussier, 2010). In more recent work of (Mohamed et al., 2010), one popular type of sequence classification criterion, maximum mutual information or MMI, was successfully applied to learn DNN weights for the TIMIT phone recognition task. Use of cross entropy to train DNN for phone sequence recognition does not explicitly take into account the fact that the neighboring frames have smaller distances between the assigned probability distributions over phone class labels. To overcome this deficiency, one can optimize the conditional probability of the whole sequence of labels, given the whole visible feature utterance or equivalent the hidden feature sequence extracted by DNN. To optimize the log conditional probability on the training data, we take the gradient over the activation parameters, transition parameters and lower-layer weights, and then pursue back-propagation of the error defined at the sentence level.

In implementing the above learning algorithm for MMI-DNN, the DNN weights can be initialized using the frame-based training with cross entropy as the objective. The transition

parameters can be initialized from the combination of the HMM transition matrices and the "phone language" model scores, and can be further optimized by tuning the transition features while fixing the DNN weights before the joint optimization. Using the joint optimization with careful scheduling, it was shown that MMI-DNN can outperform the DNN trained with cross entropy by approximately 5% relative.

In (Dahl et al., 2011, 2012), the DNN-HMM was extended from the monophone phonetic representation to the triphone or context-dependent counterpart and from phone recognition to large vocabulary speech recognition. Experiments on the Bing mobile voice search dataset collected under the real usage scenario demonstrate that the triphone DNN-HMM significantly outperforms the state-of-the-art HMM system. Three factors contribute to the success: the use of triphones as the DNN modeling units, the use of the best available triphone GMM-HMM to generate the alignment with each state in the tri-phones, and the tuning of the transition probabilities. Experiments also indicate that the decoding time of a five-layer DNN-HMM is almost the same as that of the state-of-the-art triphone GMM-HMM.

In other large vocabulary speech recognition tasks including those defined on Google voice search and YouTube, and in the Switchboard and Broadcast News databases, DNN-HMM also produces strong results (Seide et al., 2011; Sainath et al., 2011; Jaitly et al., 2012; Hinton et al., 2012).

A similar strength of DNN (called DBN in the paper) was reported in (Vinyals and Ravuri, 2011) in a somewhat different setup called Tandem approach and for mis-matched noisy speech. It was reported that DNNs outperform the MLPs with a single hidden layer under the clean condition, but the gains slowly diminish as the noise level is increased. Furthermore, using MFCCs in conjunction with the posteriors computed from DNNs outperforms using single DNNs alone in low to moderate noise conditions.

In (Deng et al., 2010), a type of deep auto-encoder developed originally for image feature coding was explored and modified to solve the speech feature coding problem, which we reviewed in Section IV. The goal is to extract "bottle-neck" speech features by compressing the high-resolution speech spectrogram data to a pre-defined number of bits with minimal reproduction error. DBN pre-training is found to be crucial for high coding efficiency. When the DBN pre-training is used, the deep auto-encoder is shown to significantly outperform a traditional vector quantization technique. If the weights in the deep auto-encoder are randomly initialized, the performance is observed to be substantially degraded. The architecture of the deep auto-encoder used in that work was shown in Fig. 1 earlier.

The more recent work of (Deng and Yu, 2011; Yu and Deng, 2011; Deng, Yu, and Platt, 2012), which we reviewed in Section VI, makes use of the discriminative DSN architecture

to perform frame-level phone classification as well as phone recognition. Higher accuracy than DBN/DNN is reported, especially after a discriminative “fine-tuning” technique developed in (Yu and Deng, 2011) is exploited. The effectiveness of DSN is particularly enabled by convex optimization for discriminative “pre-training”, and is recently found to beat strong baseline results also in a speech understanding task (Tur et al., 2012).

The basic DSN architecture has recently been extended to its tensor version, TDSN, as reviewed in Section VI, to incorporate correlation structure in the hidden layer activities with promising results on phone recognition and computational properties reported in (Hutchinson et al., 2012). While the work reported in these papers has not fully developed parallel implementation of the basic learning algorithm in the DSN and TDSN architectures, research is currently underway to enable high scalability of learning DSN and TDSN via parallelization, with initial results shown in (Deng et al., 2012).

In the work reported in (Lee et al., 2009) and the follow-up work, the convolutional structure is imposed on the RBM while building up a DBN. Convolution is made in time by sharing weights between hidden units in an attempt to detect the same “invariant” feature over different times. Then a max-pooling operation is performed where the maximal activations over small temporal neighborhoods of hidden units are obtained, proving some local temporal invariance. The resulting convolutional DBN is applied to audio and speech data for a number of tasks including music artist and genre classification, speaker identification, speaker gender classification, and phone classification, with promising results presented.

As discussed in Section III.B, the concept of convolution in time in the above work was originated in TDNN as a shallow neural net (Lang et al., 1990) developed in early speech recognition. Only recently and when deep architectures (e.g. deep CNN) are used, it has been found that frequency-dimension weight sharing is more effective for high-performance phone recognition than time domain as in the previous TDNN (Abdel-Hamid et al., 2012, 2013; Deng et al., 2013). These studies also show that designing the pooling in deep CNN to properly trade-off between invariance to vocal tract length and discrimination between speech sounds (together with a regularization technique of “dropout” (Hinton et al., 2012a) leads to even better phone recognition performance. This set of work also points to the direction of trading-off between trajectory discrimination and invariance expressed in the whole dynamic pattern of speech defined in mixed time and frequency domains using convolution and pooling. Moreover, the most recent work of (Sainath et al., 2013) shows that CNNs are also useful for large vocabulary continuous speech recognition and further demonstrates that multiple convolutional layers provide even more

improvement when the convolutional layers use a large number of convolution kernels or feature maps.

The more recent work reported in (Jaitly and Hinton, 2011) makes use of speech sound waves as the raw input feature to an RBM with a convolutional structure as the classifier. With the use of rectifier linear units in the hidden layer (Glorot et al., 2011), it is possible to automatically normalize the amplitude variation in the waveform signal, thus overcoming the difficulty encountered in the earlier attempt of using the same raw feature in the HMM-based approach (Sheikhzadeh and Deng, 1994).

In addition to RBM, DBN, and DSN, other deep models have also been developed and reported in the literature for speech processing and related applications. For example, the deep-structured CRF, which stacks many layers of CRFs, have been successfully used in the task of language identification (Yu et al., 2010), phone recognition (Yu and Deng, 2010), sequential labeling in natural language processing (Yu et al., 2010a), and confidence calibration in speech recognition (Yu et al., 2010b). Further, while RNN has early success in phone recognition (Robinson, 1994), it was not easy to duplicate due to the intricacy in training, let alone to scale up for larger speech recognition tasks. Learning algorithms for RNN have been dramatically improved since then, and much better results have been obtained recently using RNN (Graves, et al., 2006; 2013; Maas et al., 2012). RNN has also been recently applied to music processing applications (Bengio et al., 2013), where the use of rectified linear hidden units instead of logistic or tanh nonlinearities is explored in RNN. Rectified linear units compute $y = \max(x, 0)$, and lead to sparser gradients, less diffusion of credit and blame in the RNN, and faster training.

B. Image, video, and multimodality

The original DBN and deep auto-encoder were developed and demonstrated with success on the simple image recognition and dimensionality reduction (coding) tasks (MNIST) in (Hinton and Salakhutdinov, 2006). It is interesting to note that the gain of coding efficiency using the DBN-based auto-encoder on the image data over the conventional method of principal component analysis as demonstrated in (Hinton and Salakhutdinov, 2006) is very similar to the gain reported in (Deng et al., 2010) on the speech data over the traditional technique of vector quantization.

In (Nair and Hinton, 2009), a modified DBN is developed where the top-layer model uses a third-order Boltzmann machine. This type of DBN is applied to the NORB database – a 3-dimensional object recognition task. An error rate close to the best published result on this task is reported. In particular, it is shown that the DBN substantially outperforms shallow models such as SVMs.

In (Tang and Eliasmith, 2010), two strategies to improve the robustness of the DBN are developed. First, sparse

connections in the first layer of the DBN are used as a way to regularize the model. Second, a probabilistic de-noising algorithm is developed. Both techniques are shown to be effective in improving robustness against occlusion and random noise in a noisy image recognition task.

DBNs have also been successfully applied to create compact but meaningful representations of images (Taralba et al., 2008) for retrieval purposes. On this large collection image retrieval task, deep learning approaches also produced strong results.

Deep architectures with convolution structure have been found highly effective and been commonly used in computer vision and image recognition (Bengio and LeCun, 1995; LeCun et al., 1998; Jarrett et al., 2009; Kavukcuoglu et al., 2010; Ciresan et al., 2012; Le et al., 2012; Dean et al., 2012; Krizhevsky et al., 2012). The most notable advance was recently achieved in the 2012 ImageNet LSVRC contest, where 1000 different image classes are the targets with 1.2 million high-resolution images in the training set. On the test set consisting of 150,000 images, the deep CNN approach described in (Krizhevsky et al., 2012) achieved the error rates considerably lower than the previous state-of-the-art. Very large deep CNNs are used, consisting of 60 million weights, and 650,000 neurons, and five convolutional layers together with max-pooling layers. Additional three fully-connected layers as in the DNN described previously are used on top of the deep CNN layers. Although all the above structures were developed separately in earlier work, their best combination accounted for part of the success. Additional factors contributing to the final success are: 1) a powerful regularization technique called “dropout” (see details in Hinton et al., 2012); and 2) use of non-saturating neurons or rectified linear units (ReLU) that compute $y = \max(x, 0)$, significantly speeding up the training process especially with a very efficient GPU implementation.

The use of a temporally conditional DBN for video sequence and human motion synthesis is reported in (Taylor et al., 2007). The conditional DBN makes the DBN weights associated with a fixed time window conditioned on the data from previous time steps. The computational tool offered in this type of temporal DBN and the related recurrent networks may provide the opportunity to improve the DBN-HMMs towards efficient integration of temporal-centric human speech production mechanisms into DBN-based speech production model.

A very interesting study appeared in (Ngiam et al., 2011), where the authors propose and evaluate a novel application of deep networks to learn features over both audio and video modalities. A similar deep auto-encoder architecture that we described in Section IV and in (Deng et al., 2010) is used but it can be considered as a generalization from a single modality to two modalities. Cross modality feature learning has been demonstrated --- better features for video can be learned if both audio and video information sources are available at

feature learning time. The authors further show how to learn a shared audio and video representation, and evaluate it on a fixed task, where the classifier is trained with audio-only data but tested with video-only data and vice-versa. The work concludes that deep learning architectures are generally effective in learning multimodal features from unlabeled data and in improving single modality features through cross modality learning. One exception is the cross-modality setting using the CUAVE dataset. The results presented in (Ngiam et al., 2011) show that there is an improvement by learning video features with both video and audio compared to learning features with only video data. However, the same paper also shows that a model of (Papandreou, 2009) in which a sophisticated signal processing technique for extracting visual features, together with the uncertainty-compensation method developed originally from robust speech recognition (Deng et al., 2005), gives the best classification accuracy in the cross-modeling learning task, beating the features derived from the generative deep architecture designed for this task.

While the deep generative architecture for multimodal learning described in (Ngiam et al., 2011) is based on non-probabilistic auto-encoder neural nets, a probabilistic version based on deep Boltzmann machine (DBM) has appeared more recently for the same multimodal application. In (Srivastava and Salakhutdinov, 2012), a DBM is used to extract a unified representation integrating separate modalities, useful for both classification and information retrieval tasks. Rather than using the “bottleneck” layers in the deep auto-encoder to represent multimodal inputs, here a probability density is defined on the joint space of multimodal inputs, and states of suitably defined latent variables are used for the representation. The advantage of this probabilistic formulation, lacking in the deep auto-encoder, is that the missing modality’s information can be filled in naturally by sampling from its conditional distribution. For the bi-modal data consisting of image and text, the multimodal DBM is shown to outperform deep multimodal auto-encoder as well as multimodal DBN in classification and information retrieval tasks.

C. Language modeling

Research in language, document, and text processing has seen increasing popularity recently in the signal processing community, and has been designated as one of the main focus areas by the society’s audio, speech, and language processing technical committee. There has been a long history (e.g., Bengio et al., 2000; Zamora et al., 2009) of using (shallow) neural networks in language modeling (LM) – an important component in speech recognition, machine translation, text information retrieval, and in natural language processing. Recently, deep neural networks have been attracting more and more attention in statistical language modeling.

An LM is a function that captures the salient statistical characteristics of the distribution of sequences of words in a natural language. It allows one to make probabilistic

predictions of the next word given preceding ones. A neural network LM is one that exploits the neural network ability to learn distributed representations to reduce the impact of the curse of dimensionality.

A distributed representation of a symbol is a vector of features which characterize the meaning of the symbol. With a neural network LM, one relies on the learning algorithm to discover meaningful, continuous-valued features. The basic idea is to learn to associate each word in the dictionary with a continuous-valued vector representation, where each word corresponds to a point in a feature space. One can imagine that each dimension of that space corresponds to a semantic or grammatical characteristic of words. The hope is that functionally similar words get to be closer to each other in that space, at least along some directions. A sequence of words can thus be transformed into a sequence of these learned feature vectors. The neural network learns to map that sequence of feature vectors to the probability distribution over the next word in the sequence.

The distributed representation approach to LM has the advantage that it allows the model to generalize well to sequences that are not in the set of training word sequences, but that are similar in terms of their features, i.e., their distributed representation. Because neural networks tend to map nearby inputs to nearby outputs, the predictions corresponding to word sequences with similar features are mapped to similar predictions.

The above ideas of neural network LM have been implemented in various studies, some involving deep architecture. In (Mnih and Hinton, 2007), temporally factored RBM was used for language modeling. Unlike the traditional N-gram model the factored RBM uses distributed representations not only for context words but also for the words being predicted. This approach is generalized to deeper structures as reported in (Mnih and Hinton, 2008).

More recent work on neural network LM with deep architectures can be found in (Le et al., 2010, 2011a; Mikolov et al., 2010; Mikolov et al., 2011; Mikolov, 2012). In particular, the work described in (Mikolov et al., 2011) and (Mikolov, 2012) makes use RNNs to build large scale language models. It achieves stability and fast convergence in training, helped by capping the growing gradient in training RNNs. It also develops adaptation schemes for the RNN-based LM by sorting the training data with respect to their relevance and by training the model during processing of the test data. Empirical comparisons with other LM state-of-the-art show much better performance of RNN especially in the perplexity measure. A separate work on applying RNN as an LM on the unit of characters instead of words can be found in (Sutskever et al., 2011). Very interesting properties such as predicting long-term dependency (e.g. making open and closing quotes in a paragraph) are demonstrated. But its usefulness in practical applications has not been clear because

word is such a powerful representation for natural language and changing word to character in LM limits most practical application scenarios.

Further, the use of hierarchical Bayesian priors in building up deep and recursive structure in LM appeared in (Huang and Renals, 2010). Specifically, Pitman-Yor process is exploited as the Bayesian prior, from which a deep (four layers) probabilistic generative model is built. It offers a principled approach to LM smoothing by incorporating the power-law distribution for natural language. As discussed in Section III, this type of prior knowledge embedding is more readily achievable in the probabilistic modeling setup than in the neural network one.

D. Natural language processing

In the well-known and sometimes debatable work on natural language processing, Collobert and Weston (2008) developed and employed a convolutional DBN as the common model to simultaneously solve a number of classic problems including part-of-speech tagging, chunking, named entity tagging, semantic role identification, and similar word identification. More recent work reported in (Collobert, 2010) further developed a fast purely discriminative approach for parsing based on the deep recurrent convolutional architecture called Graph Transformer Network. Collobert et al., (2011) provides a comprehensive review on this line of work, specifically on ways of applying a unified neural network architectures and related deep learning algorithms to solve natural language processing problems from “scratch”. The theme of this line of work is to avoid task-specific, “man-made” feature engineering while providing versatility and unified features constructed automatically from deep learning applicable to all natural language processing tasks. The system described in (Collobert et al., 2011) automatically learns internal representation from vast amounts of mostly unlabeled training data.

One most important aspects of the work described in (Collobert and Weston, 2008) and (Collobert et al., 2011) is the transformation of raw word representations in terms of sparse vectors with a very high dimension (vocabulary size or its square or even its cubic) into low-dimensional, real-valued vectors for processing by subsequent neural network layers. This is known as “word embedding”, widely used in natural language processing and language modeling nowadays. Unsupervised learning is used where “context” of the word is used as the learning signal in neural networks. An excellent tutorial was recently given (Socher et al, 2012) that explains how the neural network is trained to perform word embedding originally proposed in (Collobert and Weston, 2008). More recent work proposes new ways of doing word embedding that better capture the semantics of words by incorporating both local and global document context and better account for homonymy and polysemy by learning multiple embeddings per word (Huang et al., 2012). Also, there is evidence that the

use of RNN can also provide empirically good performance in word embedding (Mikolov, 2012).

Another study on applying multi-task learning techniques with DBN is provided in (Deselaers et al., 2009) to attack a machine transliteration problem. This type of deep architectures and learning may be generalized to the more difficult machine translation problem.

In the work of (Sarikaya et al., 2011), DNNs (called DBNs in the paper) are used to perform a natural language call-routing task. The DNNs use unsupervised learning to discover multiple layers of features that are then used in a feed-forward neural network and fine-tuned to optimize discrimination. Unsupervised feature discovery is found to make DBNs far less prone to overfitting than the neural networks initialized with random weights. It also makes it easier to train neural networks with many hidden layers. DBNs are found to produce better classification results than several other widely used learning techniques, e.g., Maximum Entropy and Boosting based classifiers.

An interesting recent work on applying deep learning to natural language processing appears in (Socher et al., 2011), where a recursive neural network is used to build a deep architecture. The network is shown to be capable of successful merging of natural language words based on the learned semantic transformations of their original features. This deep learning approach provides an excellent performance on natural language parsing. The same approach is also demonstrated by the same authors to be successful in parsing natural scene images. In related studies, a similar recursive deep architecture is used for paraphrase detection (Socher et al., 2011a), and for predicting sentiment distributions from text (Socher et al., 2011b).

E. Information Retrieval

Here we discuss very interesting applications of the DBN and the related deep auto-encoder to document indexing and information retrieval as published in (Salakhutdinov and Hinton, 2007; Hinton and Salakhutdinov, 2010). It is shown that the hidden variables in the final layer of a DBN not only are easy to infer but also give a much better representation of each document, based on the word-count features, than the widely used latent semantic analysis and the traditional TF-IDF approach for information retrieval. Using the compact code produced by deep networks, documents are mapped to memory addresses in such a way that semantically similar text documents are located at nearby address to facilitate rapid document retrieval. And the mapping from a word-count vector to its compact code is highly efficient, requiring only a matrix multiplication and a subsequent sigmoid function evaluation for each hidden layer in the encoder part of the network.

A deep generative model of DBN is exploited for the above purpose, with details provided in (Hinton and Salakhutdinov,

2010). Briefly, the lowest layer of the DBN represents the word-count vector of a document and the top layer represents a learned binary code for that document. The top two layers of the DBN form an undirected associative memory and the remaining layers form a Bayesian (also called belief) network with directed, top-down connections. This DBN, composed of a set of stacked RBMs as we reviewed in Section V, produces a feed-forward “encoder” network that converts word-count vectors to compact codes. By composing the RBMs in the opposite order, a “decoder” network is constructed that maps compact code vectors into reconstructed word-count vectors. Combining the encoder and decoder, one obtains a deep auto-encoder (subject to further fine-tuning as discussed in Section IV) for document coding and subsequent retrieval.

After the deep model is trained, the retrieval process starts with mapping each query document into a 128-bit binary code by performing a forward pass through the model with thresholding. Then, the similarity, with Hamming distance, between the query binary code and all other documents’ 128-bit binary codes are computed efficiently.

The same idea above for coding text documents for information retrieval has been explored for audio document retrieval and speech feature coding problems with some initial exploration reported in (Deng et al., 2010) discussed in Section IV.

VIII. SUMMARY AND DISCUSSIONS

This paper presents a brief history of deep learning, and develops a categorization scheme to analyze the existing deep architectures in the literature into generative, discriminative, and hybrid classes. The deep auto-encoder, DSN, and DBN-DNN architectures, one in each of the three classes, are discussed and analyzed in detail, as they appear to be popular and promising approaches with author’s personal research experience. Applications of deep learning in five broad areas of information processing are then reviewed.

The literature on deep learning is vast, mostly coming from the machine learning community. The signal processing community embraced deep learning only within the past four years or so and the momentum is growing fast. This overview paper is written mainly from the signal processing perspective. Beyond just surveying existing deep learning work, a classificatory scheme based on the architecture and the nature of learning algorithms is developed and in-depth analysis with concrete examples conducted. This will hopefully provide insight for readers to better understand the capability of the various deep learning systems discussed in the paper, the connection among different but similar deep learning methods, and how to design proper deep learning algorithms under different circumstances.

Throughout this review, the important message is conveyed that building/learning deep architectures and hierarchies of features is highly desirable. We have discussed the difficulty

of learning parameters in all layers at once due to pervasive local optimum and gradient dilution. The generative, pre-training method in the hybrid architecture of DBN-DNN, which we reviewed in detail in Section V in this paper, appears to have offered a useful, albeit empirical, solution to poor local optima in optimization and to regularization for the deep model containing massive parameters.

Deep learning is an emerging technology. Despite the empirical promising results reported so far, much needs to be developed. Importantly, it has not been the experience of deep learning researchers that a single deep learning technique can be successful for all classification tasks. For example, while the popular learning strategy of generative pre-training followed by discriminative fine-tuning seems to work well empirically for many tasks, it failed to work for some other tasks that have been explored (e.g., language identification). For these tasks, the features extracted at the generative pre-training phase seem to describe the underlining speech variations well but do not contain sufficient information to distinguish between different languages. A learning strategy that can extract discriminative yet also invariant features is expected to provide better solutions. Extracting discriminative features may also greatly reduce the model size needed in the many current deep learning systems. Also, domain knowledge such as what kind of invariance is useful for a specific task in hand (e.g., vision, speech, or language) and what kind of regularization in terms of parameter constraints is key to the success of applying deep learning methods.

Recent published work shows that there is vast room to improve the current optimization techniques for learning deep architectures (Martens, 2010; Le et al., 2011; Martens and Sutskever, 2011; Dean et al., 2012; Sutskever, 2013). To what extent pre-training is essential to learning the full set of parameters in deep architectures has been currently under investigation, especially when very large amounts of labeled training data are available which reduces or even obliterates the need for model regularization. Some preliminary results have been discussed in this paper and in (Hinton et al., 2012).

Effective and scalable parallel algorithms are critical for training deep models with very large data, as in many common information processing applications such as speech recognition and machine translation. The popular mini-batch stochastic gradient technique is known to be difficult to be parallelized over computers. The common practice nowadays is to use graphical processing units (GPUs) to speed up the learning process, although recent advance in developing asynchronous stochastic gradient learning has shown promises by using large-scale CPU clusters (e.g. Le et al., 2011; Dean et al., 2012). In this interesting computing architecture, many different replicas of the DNN to compute gradients on different subsets of the training data in parallel. These gradients are communicated to a central parameter server that updates the shared weights. Even though each replica typically computes gradients using parameter values

not immediately updated, stochastic gradient descent is robust to the slight errors this has introduced. To make deep learning techniques scalable to very large training data, theoretically sound parallel learning algorithms or novel architectures need to be further developed (e.g., Bottou and LeCun, 2004; Dean et al., 2012; Hutchinson et al., 2013; Sutskever, 2013; Bengio et al., 2013).

One major barrier to the application of DNNs and related deep models is that it currently requires considerable skill and experience to choose sensible values for hyper-parameters such as the learning rate schedule, the strength of the regularizer, the number of layers and the number of units per layer, etc. Sensible values for one hyper-parameter may depend on the values chosen for other hyper-parameters and hyper-parameter tuning in DNNs is especially expensive. Some interesting methods for solving the problem have been developed recently, including random sampling (Bergstra et al., 2012) and Bayesian optimization procedure (Snoek et al., 2012). Further research is needed in this important area.

Finally, solid theoretical foundations of deep learning need to be established in a myriad of aspects. As an example, the success of deep learning in unsupervised learning has not been demonstrated as much as for supervised learning; yet the essence and major motivation of deep learning lie right in unsupervised learning for automatically discovering data representation. What are the appropriate objectives for learning effective representations? How may the deep learning architectures and algorithms use distributed representations to effectively disentangle the hidden explanatory factors of variation in the data? How can computational neuroscience models about hierarchical brain structure and learning style help improve engineering deep learning architectures and algorithms? All these important questions will need intensive research in order to further push the frontier of deep learning.

REFERENCES

- Abdel-Hamid, O., Mohamed, A., Jiang, H., and G. Penn, "Applying convolutional neural networks concepts to hybrid NN-HMM model for speech recognition," ICASSP, 2012.
- Abdel-Hamid, O., Deng, L., and Yu. D. "Exploring convolutional neural network structures and optimization for speech recognition," Interspeech, 2013..
- Arel, I., Rose, C., and Karnowski, T. "Deep Machine Learning - A New Frontier in Artificial Intelligence," IEEE Computational Intelligence Mag., Nov., 2010.
- Baker, J., et al. "Research developments and directions in speech recognition and understanding," IEEE Sig. Proc. Mag., vol. 26, no. 3, May 2009, pp. 75-80.
- Baker, J., et al. "Updated MINS report on speech recognition and understanding," IEEE Sig. Proc. Mag., vol. 26, no. 4, July 2009a.
- Bengio, Y., Boulanger, N., and Pascanu. R. "Advances in optimizing recurrent networks," Proc. ICASSP, 2013.
- Bengio Y., Courville, A., and Vincent, P. "Representation learning: A review and new perspectives," IEEE Trans. PAMI, 2013a.

- Bengio Y. "Learning deep architectures for AI," in Foundations and Trends in Machine Learning, Vol. 2, No. 1, 2009, pp. 1-127.
- Bengio, Y., Ducharme, R., Vincent, P. and Jauvin, C. "A neural probabilistic language model," Proc. NIPS, 2000, pp. 933-938.
- Bengio, Y., De Mori, R., Flammia, G. and Kompe, F. "Global optimization of a neural network—Hidden Markov model hybrid," in Proc. Proc. Eurospeech, 1991.
- Bergstra J. and Bengio, Y. "Random search for hyper-parameter optimization," J. Machine Learning Research," Vol. 3, pp. 281-305, 2012.
- Bottou L. and LeCun. Y. "Large scale online learning," Proc. NIPS, 2004.
- Bilmes, J. "Dynamic graphical models," IEEE Signal Processing Mag., vol. 33, pp. 29–42, 2010.
- Bilmes, J. and Bartels, C. "Graphical model architectures for speech recognition," IEEE Signal Processing Mag., vol. 22, pp. 89–100, 2005.
- Bouvard H. and Morgan, N., Connectionist Speech Recognition: A Hybrid Approach, Norwell, MA: Kluwer, 1993.
- Bouvier, J. "Hierarchical Learning: Theory with Applications in Speech and Vision," Ph.D. thesis, MIT, 2009.
- Bridle, J., L. Deng, J. Picone, H. Richards, J. Ma, T. Kamm, M. Schuster, S. Pike, and R. Reagan, "An investigation of segmental hidden dynamic models of speech coarticulation for automatic speech recognition," Final Report for 1998 Workshop on Language Engineering, CLSP, Johns Hopkins, 1998.
- Cho Y. and Saul L. "Kernel methods for deep learning," NIPS, pp. 342–350, 2009.
- Ciresan D., Giusti, A., Gambardella, L., and Schmidhuber, J. "Deep neural networks segment neuronal membranes in electron microscopy images," Proc. NIPS, 2012.
- Cohen W. and R. V. de Carvalho. Stacked sequential learning. In Proc. IJCAI, pp. 671–676, 2005.
- Collobert R. "Deep learning for efficient discriminative parsing," Proc. NIPS Workshop on Deep Learning and Unsupervised Feature Learning, 2010.
- Collobert R. and Weston J. "A unified architecture for natural language processing: Deep neural networks with multitask learning," Proc. ICML, 2008.
- Collobert, R., Weston, J., Bottou, L., Karlen, M., Kavukcuoglu, K., and Kuksa, P. "Natural language processing (almost) from scratch," J. Machine Learning Research, Vo. 12, pp. 2493-2537, 2011.
- Dahl, G., Yu, D., Deng, L., and Acero, A. "Context-dependent DBN-HMMs in large vocabulary continuous speech recognition," Proc. ICASSP, 2011.
- Dahl, G., Yu, D., Deng, L., and Acero, A. "Context-dependent DBN-HMMs in large vocabulary continuous speech recognition," IEEE Trans. Audio, Speech, & Language Proc. Vol. 20 (1), January 2012.
- Dahl, G., Ranzato, M., Mohamed, A. and Hinton, G. "Phone recognition with the mean-covariance restricted Boltzmann machine," Proc. NIPS, vol. 23, 2010, 469-477.
- Dean, J., Corrado, G., R. Monga, K. Chen, M. Devin, Q. Le, M. Mao, M. Ranzato, A. Senior, P. Tucker, K. Yang, A. Ng " Large Scale Distributed Deep Networks". Proc. NIPS, 2012.
- Deng, L. and Li, X. "Machine learning paradigms in speech recognition: An overview," IEEE Trans. Audio, Speech, & Language, July 2013.
- Deng, L., Abdel-Hamid, O., and Yu, D. "A deep convolutional neural network using heterogeneous pooling for trading acoustic invariance with phonetic confusion," Proc. ICASSP, 2013.
- Deng, L., He, X, and Gao, J. "Deep stacking networks for information retrieval," Proc. ICASSP, 2013a.
- Deng, L., Tur, G, He, X, and Hakkani-Tur, D. "Use of Kernel Deep Convex Networks and End-To-End Learning for Spoken Language Understanding," Proc. IEEE Workshop on Spoken Language Technologies, December 2012.
- Deng, L., Yu, D., and Platt, J. "Scalable stacking and learning for building deep architectures," Proc. ICASSP, 2012a.
- Deng, L., Hutchinson, B., and Yu, D. "Parallel training of deep stacking networks," Proc. Interspeech, 2012b.
- Deng, L. "An Overview of Deep-Structured Learning for Information Processing, in Proceedings of Asian-Pacific Signal & Information Processing Annual Summit and Conference (APSIPA-ASC), October 2011.
- Deng, L. and Yu, D. "Deep Convex Network: A scalable architecture for deep learning," Proc. Interspeech, 2011.
- Deng, L., Seltzer, M., Yu, D., Acero, A., Mohamed, A., and Hinton, G. "Binary coding of speech spectrograms using a deep auto-encoder," Proc. Interspeech, 2010.
- Deng, L. "Expanding the scope of signal processing," IEEE Signal Processing Magazine, vol. 25, no. 3, May 2008.
- Deng, L. and Yu, D. "Use of differential cepstra as acoustic features in hidden trajectory modeling for phonetic recognition, Proc. ICASSP, April 2007.
- Deng, L. DYNAMIC SPEECH MODELS --- Theory, Algorithm, and Application, Morgan & Claypool, December 2006.
- Deng, L., Yu, D. and Acero, A. "Structured speech modeling," IEEE Trans. on Audio, Speech and *Language Processing*, vol. 14, no. 5, pp. 1492-1504, September 2006.
- Deng, L., Yu, D. and Acero, A. "A bidirectional target filtering model of speech coarticulation: Two-stage implementation for phonetic recognition," IEEE Transactions on Audio and Speech Processing, vol. 14, no. 1, pp. 256-265, January 2006a.
- Deng, L., Wu, J., Droppo, J., and Acero, A. "Dynamic Compensation of HMM Variances Using the Feature Enhancement Uncertainty Computed From a Parametric Model of Speech Distortion," IEEE Transactions on Speech and Audio Processing, vol. 13, no. 3, pp. 412–421, 2005.
- Deng, L. and Huang, X.D. "Challenges in Adopting Speech Recognition, Communications of the ACM, vol. 47, no. 1, pp. 11-13, January 2004.
- Deng, L. and O'Shaughnessy, D. SPEECH PROCESSING --- A Dynamic and Optimization-Oriented Approach, Marcel Dekker, 2003.
- Deng, L. "Switching dynamic system models for speech articulation and acoustics," in Mathematical Foundations of Speech and Language Processing, pp. 115–134. Springer-Verlag, New York, 2003.
- Deng, L. "Computational Models for Speech Production," in Computational Models of Speech Pattern Processing, pp. 199-213, Springer Verlag, 1999.
- Deng, L., Ramsay, G., and Sun, D. "Production models as a structural basis for automatic speech recognition," Speech Communication, vol. 33, no. 2-3, pp. 93–111, Aug 1997.
- Deng, L. and Sameti, H. "Transitional speech units and their representation by regressive Markov states: Applications to speech recognition," IEEE Transactions on speech and audio processing, vol. 4, no. 4, pp. 301–306, July 1996.

- Deng, L., Aksmanovic, M., Sun, D., and Wu, J. "Speech recognition using hidden Markov models with polynomial regression functions as nonstationary states," *IEEE Transactions on Speech and Audio Processing*, vol. 2, no. 4, pp. 507-520, 1994.
- Deng L. and Sun, D. "A statistical approach to automatic speech recognition using the atomic speech units constructed from overlapping articulatory features," *Journal of the Acoustical Society of America*, vol. 85, no. 5, pp. 2702-2719, 1994.
- Deng, L. "A stochastic model of speech incorporating hierarchical nonstationarity," *IEEE Transactions on Speech and Audio Processing*, vol. 1, no. 4, pp. 471-475, 1993.
- Deng, L. "A generalized hidden Markov model with state-conditioned trend functions of time for the speech signal," *Signal Processing*, vol. 27, no. 1, pp. 65-78, 1992.
- Deselaers, T., Hasan, S., Bender, O. and Ney, H. "A deep learning approach to machine transliteration," *Proc. 4th Workshop on Statistical Machine Translation*, pp. 233-241, Athens, Greece, March 2009.
- Erhan, D., Bengio, Y., Courville, A., Manzagol, P., Vencent, P., and Bengio, S. "Why does unsupervised pre-training help deep learning?" *J. Machine Learning Research*, 2010, pp. 201-208.
- Fine, S., Singer, Y. and Tishby, N. "The Hierarchical Hidden Markov Model: Analysis and Applications," *Machine Learning*, vol. 32, p. 41-62, 1998.
- Gens R. and Domingo, P. "Discriminative learning of sum-product networks," *NIPS*, 2012.
- George, D. "How the Brain Might Work: A Hierarchical and Temporal Model for Learning and Recognition," Ph.D. thesis, Stanford University, 2008.
- Gibson, M. and Hain, T. "Error approximation and minimum phone error acoustic model estimation," *IEEE Trans. Audio, Speech, and Language Proc.*, vol. 18, no. 6, August 2010, pp. 1269-1279.
- Glorot, X., Bordes, A., and Bengio, Y. "Deep sparse rectifier neural networks," *Proc. AISTAT*, April 2011.
- Glorot, X. and Bengio, Y. "Understanding the difficulty of training deep feedforward neural networks" *Proc. AISTAT*, 2010.
- Graves, A., Fernandez, S., Gomez, F., and Schmidhuber, J. "Connectionist temporal classification: Labeling unsegmented sequence data with recurrent neural networks," *Proc. ICML*, 2006.
- Graves, A. "Sequence Transduction with Recurrent Neural Networks," *Representation Learning Workshop, ICML 2012*.
- Graves, A., Mahamed, A., and Hinton, G. "Speech recognition with deep recurrent neural networks," *Proc. ICASSP*, 2013.
- Hawkins, J. and Blakeslee, S. *On Intelligence: How a New Understanding of the Brain will lead to the Creation of Truly Intelligent Machines*, Times Books, New York, 2004.
- Hawkins, G., Ahmad, S. and Dubinsky, D. "Hierarchical Temporal Memory including HTM Cortical Learning Algorithms," *Numenta Tech. Report*, December 10, 2010.
- He, X., Deng, L., Tur, G., and Hakkani-Tur, D. "Multi-style adaptive training for robust cross-lingual spoken language understanding," *Proc. ICASSP*, 2013.
- He, X., Deng, L., Chou, W. "Discriminative learning in sequential pattern recognition --- A unifying review for optimization-oriented speech recognition," *IEEE Sig. Proc. Mag.*, vol. 25, 2008, pp. 14-36.
- He, X. and Deng, L. "Speech recognition, machine translation, and speech translation --- A unifying discriminative framework," *IEEE Sig. Proc. Magazine*, Vol. 28, November, 2011.
- He, X. and Deng, L. "Optimization in speech-centric information processing: Criteria and techniques," *Proc. ICASSP*, 2012.
- He, X. and Deng, L. "Speech-centric information processing: An optimization-oriented approach," *Proc. of the IEEE*, 2013.
- Heigold, G., Ney, H., Lehnen, P., Gass, T., Schluter, R. "Equivalence of generative and log-liner models," *IEEE Trans. Audio, Speech, and Language Proc.*, vol. 19, no. 5, February 2011, pp. 1138-1148.
- Heintz, I., Fosler-Lussier, E., and Brew, C. "Discriminative Input Stream Combination for Conditional Random Field Phone Recognition," *IEEE Trans. Audio, Speech, and Language Proc.*, vol. 17, no. 8, Nov. 2009, pp. 1533-1546.
- Hifny, Y. and Renals, S. "Speech recognition using augmented conditional random fields," *IEEE Trans. Audio, Speech, and Language Proc.*, vol. 17, no. 2, February 2009, pp. 354-365.
- Hinton, G., Deng, L., Yu, D., Dahl, G., Mohamed, A., Jaitly, N., Senior, A., Vanhoucke, V., Nguyen, P., Sainath, T., and Kingsbury, B., "Deep Neural Networks for Acoustic Modeling in Speech Recognition," *IEEE Signal Processing Magazine*, vol. 29, no. 6, pp. 82-97, November 2012.
- Hinton, G., Srivastava, N., Krizhevsky, A., Sutskever, I., and Salakhutdinov, R. "Improving neural networks by preventing co-adaptation of feature detectors," *arXiv: 1207.0580v1*, 2012a.
- Hinton, G., Krizhevsky, A., and Wang, S. "Transforming auto-encoders," *Proc. Intern. Conf. Artificial Neural Networks*, 2011.
- Hinton, G. "A better way to learn features," *Communications of the ACM*, Vol. 54, No. 10, October, 2011, pp. 94.
- Hinton, G. and Salakhutdinov, R. "Discovering binary codes for documents by learning deep generative models," *Topics in Cognitive Science*, pp. 1-18, 2010.
- Hinton, G. "A practical guide to training restricted Boltzmann machines," *UTML Tech Report 2010-003*, Univ. Toronto, August 2010.
- Hinton, G., Osindero, S., and Teh, Y. "A fast learning algorithm for deep belief nets," *Neural Computation*, vol. 18, pp. 1527-1554, 2006.
- Hinton, G. and Salakhutdinov, R. "Reducing the dimensionality of data with neural networks," *Science*, vol. 313. no. 5786, pp. 504 - 507, July 2006.
- Huang, S. and Renals, S. "Hierarchical Bayesian language models for conversational speech recognition," *IEEE Trans. Audio, Speech, and Language Proc.*, vol. 18, no. 8, November 2010, pp. 1941-1954.
- Huang, E., Socher, R., Manning, C, and Ng, A. "Improving Word Representations via Global Context and Multiple Word Prototypes," *Proc. ACL*, 2012.
- Hutchinson, B., Deng, L., and Yu, D. "A deep architecture with bilinear modeling of hidden representations: Applications to phonetic recognition," *Proc. ICASSP*, 2012.
- Hutchinson, B., Deng, L., and Yu, D. "Tensor deep stacking networks," *IEEE Trans. Pattern Analysis and Machine Intelligence*, 2013.
- Jaitly, N. and Hinton, G. "Learning a better representation of speech sound waves using restricted Boltzmann machines," *Proc. ICASSP*, 2011.
- Jaitly, N., Nguyen, P., and Vanhoucke, V. "Application of pre-trained deep neural networks to large vocabulary speech recognition," *Proc. Interspeech*, 2012.
- Jarrett, K., Kavukcuoglu, K. and LeCun, Y. "What is the best multistage architecture for object recognition?" *Proc. Intl. Conf. Computer Vision*, pp. 2146-2153, 2009.

- Jiang, H. and Li, X. "Parameter estimation of statistical models using convex optimization: An advanced method of discriminative training for speech and language processing," *IEEE Signal Processing Magazine*, vol. 27, no. 3, pp. 115–127, 2010.
- Juang, B.-H., Chou, W., and Lee, C.-H. "Minimum classification error rate methods for speech recognition," *IEEE Trans. On Speech and Audio Processing*, vol. 5, pp. 257–265, 1997.
- Kavukcuoglu, K., Sermanet, P., Boureau, Y., Gregor, K., Mathieu M., and LeCun, Y. "Learning Convolutional Feature Hierarchies for Visual Recognition," *Proc. NIPS*, 2010.
- Ketabdar, H. and Bourlard, H. "Enhanced phone posteriors for improving speech recognition systems," *IEEE Trans. Audio, Speech, and Language Proc.*, vol. 18, no. 6, August 2010, pp. 1094-1106.
- Kingsbury, B. "Lattice-based optimization of sequence classification criteria for neural-network acoustic modeling," *Proc. ICASSP*, 2009.
- Kingsbury, B., Sainath, T., and Soltau, H. "Scalable minimum Bayes risk training of deep neural network acoustic models using distributed Hessian-free optimization," *Proc. Interspeech*, 2012.
- Krizhevsky, A., Sutskever, I. and Hinton, G. "ImageNet classification with deep convolutional neural Networks," *Proc. NIPS* 2012.
- Kubo, Y., Hori, T., and Nakamura, A. "Integrating deep neural networks into structural classification approach based on weighted finite-state transducers," *Proc. Interspeech*, 2012.
- Kurzweil R. *How to Create a Mind*. Viking Books, Dec., 2012.
- Lang, K., Waibel, A., and Hinton, G. "A time-delay neural network architecture for isolated word recognition," *Neural Networks*, Vol. 3(1), pp. 23-43, 1990.
- Larochelle, H. and Bengio, Y. "Classification using discriminative restricted Boltzmann machines," *Proc. ICML*, 2008.
- Le, H., Allauzen, A., Wisniewski, G., and Yvon, F. "Training continuous space language models: Some practical issues," in *Proc. of EMNLP*, 2010, pp. 778–788.
- Le, H., Oparin, I., Allauzen, A., Gauvain, J., and Yvon, F. "Structured output layer neural network language model," *Proc. ICASSP*, 2011a.
- Le, Q., Ngiam, J., Coates, A., Lahiri, A., Prochnow, B., and Ng, A. "On optimization methods for deep learning," *Proc. ICML*, 2011.
- Le, Q., Ranzato, M., Monga, R., Devin, M., Corrado, G., Chen, K., Dean, J., Ng, A. "Building High-Level Features Using Large Scale Unsupervised Learning," *Proc. ICML* 2012.
- LeCun, Y., Bottou, L., Bengio, Y., and Haffner, P. "Gradient-based learning applied to document recognition," *Proceedings of the IEEE*, Vol. 86, pp. 2278-2324, 1998.
- LeCun, Y. and Bengio, Y. "Convolutional networks for images, speech, and time series," in *The Handbook of Brain Theory and Neural Networks* (M. A. Arbib, ed.), pp. 255- 258, Cambridge, Massachusetts: MIT Press, 1995.
- LeCun Y., Chopra S., Ranzato, M., and Huang, F. "Energy-based models in document recognition and computer vision," *Proc. Intern. Conf. Document Analysis and Recognition, (ICDAR)*, 2007.
- Lee, C.-H. "From knowledge-ignorant to knowledge-rich modeling: A new speech research paradigm for next-generation automatic speech recognition," *Proc. ICSLP*, 2004, p. 109-111.
- Lee, H., Grosse, R., Ranganath, R., and Ng, A. "Unsupervised learning of hierarchical representations with convolutional deep belief networks," *Communications of the ACM*, Vol. 54, No. 10, October, 2011, pp. 95-103.
- Lee, H., Grosse, R., Ranganath, R., and Ng, A. "Convolutional Deep Belief Networks for Scalable Unsupervised Learning of Hierarchical Representations," *Proc. ICML*, 2009.
- Lee, H., Largman, Y., Pham, P., Ng, A. "Unsupervised feature learning for audio classification using convolutional deep belief networks," *Proc. NIPS*, 2010.
- Lena, P., Nagata, K., and Baldi, P. "Deep spatiotemporal architectures and learning for protein structure prediction," *Proc. NIPS*, 2012.
- Ling, Z., Richmond, K., and Yamagishi, J. "Articulatory control of HMM-based parametric speech synthesis using feature-space-switched multiple regression," *IEEE Trans. Audio, Speech, and Language Proc.*, Vol. 21, Jan, 2013.
- Maas, A., Le, Q., O'Neil, R., Vinyals, O., Nguyen, P., and Ng, Y. "Recurrent Neural Networks for Noise Reduction in Robust ASR," *Proc. Interspeech*, 2012
- Markoff, J. "Scientists See Promise in Deep-Learning Programs," *New York Times*, Nov 24, 2012.
- Martens J. "Deep learning with Hessian-free optimization," *Proc. ICML*, 2010.
- Martens J. and Sutskever, I. "Learning recurrent neural networks with Hessian-free optimization," *Proc. ICML*, 2011.
- Mikolov, T. "Statistical Language Models based on Neural Networks," PhD thesis, Brno University of Technology, 2012.
- Mikolov, T., Deoras, A., Povey, D., Burget, L., and Cernocky, J. "Strategies for training large scale neural network language models," *Proc. IEEE ASRU*, 2011.
- Mikolov, T., Karafiat, M., Burget, L., Cernocky, J., and Khudanpur, S. "Recurrent neural network based language model," *Proc. ICASSP*, 2010, 1045–1048.
- Minami, Y., McDermott, E. Nakamura, A. and Katagiri, S. "A recognition method with parametric trajectory synthesized using direct relations between static and dynamic feature vector time series," *Proc. ICASSP*, pp. 957-960, 2002.
- Mnih A. and Hinton G. "Three new graphical models for statistical language modeling," *Proc. ICML*, 2007, pp. 641-648.
- Mnih A. and Hinton G. "A scalable hierarchical distributed language model" *Proc. NIPS*, 2008, pp. 1081-1088.
- Mohamed A., Yu, D., and Deng, L. "Investigation of full-sequence training of deep belief networks for speech recognition," *Proc. Interspeech*, Sept. 2010.
- Mohamed, A., Dahl, G., and Hinton, G. "Deep belief networks for phone recognition," in *Proc. NIPS Workshop Deep Learning for Speech Recognition and Related Applications*, 2009.
- Mohamed, A., Dahl, G. and Hinton, G. "Acoustic modeling using deep belief networks", *IEEE Trans. Audio, Speech, & Language Proc.* Vol. 20 (1), January 2012.
- Mohamed, A., Hinton, G., and Penn, G., "Understanding how deep belief networks perform acoustic modelling," *Proc. ICASSP*, 2012a.
- Mohamed, A., Hinton, G., Penn, G. "Understanding how deep belief networks perform acoustic modeling," *Proc. ICASSP*, 2012b.
- Morgan, N. "Deep and Wide: Multiple Layers in Automatic Speech Recognition," *IEEE Trans. Audio, Speech, & Language Proc.* Vol. 20 (1), January 2012.
- Morgan, N., Q. Zhu, A. Stolcke, K. Sonmez, S. Sivasdas, T. Shinozaki, M. Ostendorf, P. Jain, H. Hermansky, D. Ellis, G. Doddington, B. Chen, O. Cretin, H. Bourlard, , and M. Athineos, "Pushing the

- envelope - aside [speech recognition]," *IEEE Signal Processing Magazine*, vol. 22, no. 5, pp. 81–88, Sep 2005.
- Nair, V. and Hinton, G. "3-d object recognition with deep belief nets," *Proc. NIPS*, 2009.
- Ney, H. "Speech translation: Coupling of recognition and translation," *Proc. ICASSP*, 1999.
- Ngiam, J., Khosla, A., Kim, M., Nam, J., Lee, H., and Ng, A. "Multimodal deep learning," *Proc. ICML*, 2011.
- Ngiam, J., Chen, Z., Koh, P., and Ng, A. "Learning deep energy models," *Proc. ICML*, 2011.
- Oliver, N., Garg, A., and Horvitz, E. "Layered Representations for Learning and Inferring Office Activity from Multiple Sensory Channels," *Computer Vision and Image Understanding*, vol. 96, pp. 163-180, 2004.
- Ostendorf, V. Digalakis, and O. Kimball, "From HMM's to segment models: A unified view of stochastic modeling for speech recognition," *IEEE Trans. Speech and Audio Proc.*, vol. 4, no. 5, September 1996.
- Papandreou, G., Katsamanis, A., Pitsikalis, V., and Maragos, P. "Adaptive multimodal fusion by uncertainty compensation with application to audiovisual speech recognition," *IEEE Trans. Audio, Speech, and Lang. Processing*, Vol. 17(3), pp. 423-435, 2009.
- Peng, J., Bo, L., and Xu, J. "Conditional neural fields," *Proc. NIPS*, 2009.
- Picone, P., S. Pike, R. Regan, T. Kamm, J. bridle, L. Deng, Z. Ma, H. Richards, and M. Schuster, "Initial evaluation of hidden dynamic models on conversational speech," *Proc. ICASSP*, 1999.
- Pinto, J., Garimella, S., Magimai-Doss, M., Hermansky, H., and Bourlard, H. "Analysis of MLP-based hierarchical phone posterior probability estimators," *IEEE Trans. Audio, Speech, and Language Proc.*, vol. 19, no. 2, Feb. 2011.
- Poggio, T. "How the Brain Might Work: The Role of Information and Learning in Understanding and Replicating Intelligence," In: *Information: Science and Technology for the New Century*, Editors: G. Jacovitt, A. Pettorossi, R. Consolo and V. Senni, Lateran University Press, pp. 45-61, 2007.
- Poon, H. and Domingos, P. "Sum-product networks: A new deep architecture," *Proc. Twenty-Seventh Conference on Uncertainty in Artificial Intelligence*, 2011. Barcelona, Spain.
- Povey, D. and Woodland, P. "Minimum phone error and i-smoothing for improved discriminative training," *Proc. ICASSP*, 2002, pp. 105–108.
- Prabhavalkar, R. and Fosler-Lussier, E. "Backpropagation training for multilayer conditional random field based phone recognition", *Proc. ICASSP 2010*, pp. 5534-5537.
- Ranzato, M., Poultney, C., Chopra, S. and LeCun, Y. "Efficient learning of sparse representations with an energy-based model," *Proc. NIPS*, 2006.
- Ranzato, M., Boureau, Y., and LeCun, Y. "Sparse Feature Learning for Deep Belief Networks," *Proc. NIPS*, 2007a.
- Ranzato, M., Susskind, J., Mnih, V., and Hinton, G. "On deep generative models with applications to recognition," *Proc. CVPR*, 2011.
- Rennie, S., Hershey, H., and Olsen, P. "Single-channel multitalker speech recognition — Graphical modeling approaches," *IEEE Signal Processing Mag.*, vol. 33, pp. 66–80, 2010.
- Rifai, S., Vincent, P., X. Muller, X. Glorot, and Y. Bengio, "Contractive autoencoders: Explicit invariance during feature extraction," *Proc. ICML*, 2011, pp. 833-840.
- Robinson, A. "An application of recurrent nets to phone probability estimation," *IEEE Trans. Neural Networks*, Vol. 5, pp. 298-305, 1994.
- Sainath, T., Mohamed, A., Kingsbury, B., and Ramabhadran, B. "Convolutional neural networks for LVCSR," *Proc. ICASSP*, 2013.
- Sainath, T., Kingsbury, B., Ramabhadran, B., Novak, P., and Mohamed, A. "Making deep belief networks effective for large vocabulary continuous speech recognition," *Proc. IEEE ASRU*, 2011.
- Sainath, T., Kingsbury, B., and Ramabhadran, B. "Improving training time of deep belief networks through hybrid pre-training and larger batch sizes," *Proc. NIPS Workshop on Log-linear Models*, Dec. 2012
- Salakhutdinov R. and Hinton, G. "Semantic hashing," *Proc. SIGIR Workshop on Information Retrieval and Applications of Graphical Models*, 2007.
- Salakhutdinov R. and Hinton, G. "Deep Boltzmann machines," *Proc. AISTATS*, 2009.
- Salakhutdinov R. and Hinton, G. "A better way to pretrain deep Boltzmann machines," *Proc. NIPS*, 2012.
- Sarikaya, R., Hinton, G., Ramabhadran, B. "Deep belief nets for natural language call-routing," *Proc. ICASSP*, pp. 5680-5683, 2011.
- Seide, F., Li, G., and Yu, D. "Conversational Speech Transcription Using Context-Dependent Deep Neural Networks," *Interspeech 2011*, pp. 437-440.
- Seide, F., Li, G., Chen, X., and Yu, D. "Feature engineering in context-dependent deep neural networks for conversational speech transcription," *Proc. ASRU 2011a*, pp. 24-29.
- Shannon, M., Zen, H., and Byrne W. "Autoregressive models for statistical parametric speech synthesis," *IEEE Trans. Audio, Speech, Language Proc.*, Vol. 21, No. 3, 2013, pp. 587-597.
- Sheikhzadeh, H. and Deng, L. "Waveform-based speech recognition using hidden filter models: Parameter selection and sensitivity to power normalization," *IEEE Trans. on Speech and Audio Processing*, Vol. 2, pp. 80-91, 1994.
- Siniscalchi, M., Yu, D., Deng, L., and Lee, C.-H. "Exploiting deep neural networks for detection-based speech recognition," *Neurocomputing*, 2013.
- Siniscalchi, M., T. Svendsen, and Lee, C.-H. "A bottom-up modular search approach to large vocabulary continuous speech recognition," *IEEE Trans. Audio, Speech, Language Proc.*, Vol. 21, 2013a.
- Sivaram G. and Hermansky, H. "Sparse Multilayer Perceptron for Phoneme Recognition," *IEEE Trans. Audio, Speech, & Language Proc.* Vol. 20 (1), January 2012.
- Snoek, J., Larochelle, H., and Adams, R. "Practical Bayesian Optimization of Machine Learning Algorithms," *Proc. NIPS*, 2012.
- Socher, R., Lin, C., Ng, A., and Manning, C. "Learning continuous phrase representations and syntactic parsing with recursive neural networks," *Proc. ICML*, 2011.
- Socher, R., Pennington, J., Huang, E., Ng, A., and Manning, C. "Semi-Supervised Recursive Autoencoders for Predicting Sentiment Distributions," *Proc. EMNLP*, 2011a.
- Socher, R., Pennington, J., Huang, E., Ng, A., and Manning, C. "Dynamic Pooling and Unfolding Recursive Autoencoders for Paraphrase Detection," *Proc. NIPS 2011b*.
- Socher, R., Bengio, Y., and Manning, C. "Deep learning for NLP," *Tutorial at ACL, 2012*, <http://www.socher.org/index.php/DeepLearningTutorial/DeeplearningTutorial>.

- Stoyanov, V., Ropson, A. and Eisner, J. "Empirical Risk Minimization of Graphical Model Parameters Given Approximate Inference, Decoding, and Model Structure," Proc. AISTAT, 2011.
- Srivastava, N. and Salakhutdinov R. "Multimodal learning with deep Boltzmann machines," Proc. NIPS, 2012.
- Sutskever. I. "Training Recurrent Neural Networks," Ph.D. Thesis, University of Toronto, 2013.
- Sutskever, I., Martens J., and Hinton, G. "Generating text with recurrent neural networks", Proc. ICML, 2011.
- Taylor, G., Hinton, G. E., and Roweis, S. "Modeling human motion using binary latent variables." Proc. NIPS, 2007.
- Tang, Y. and Elasmith, C. "Deep networks for robust visual recognition," Proc. ICML, 2010.
- Taralba, A, Fergus R, and Weiss, Y. "Small codes and large image databases for recognition," Proc. CVPR, 2008.
- Tur, G., Deng, L., Hakkani-Tür, D., and X. He. "Towards deep understanding: Deep convex networks for semantic utterance classification," Proc. ICASSP, 2012.
- Vincent, P., Larochelle, H., Lajoie, I., Bengio, Y., and Manzagol, P. "Stacked denoising autoencoders: Learning useful representations in a deep network with a local denoising criterion," J. Machine Learning Research, Vol. 11, 2010, pp. 3371-3408.
- Vinyals, O., Jia, Y., Deng, L., and Darrell, T. "Learning with recursive perceptual representations," Proc. NIPS, 2012.
- Vinyals O., and Ravuri, S. "Comparing multilayer perceptron to deep belief network tandem features for robust ASR," Proc. ICASSP, 2011.
- Welling, M., Rosen-Zvi, M., and Hinton, G. "Exponential family harmoniums with an application to information retrieval," Proc. NIPS, Vol. 20, 2005.
- Wohlmayr, M., Stark, M., Pernkopf, F. "A probabilistic interaction model for multipitch tracking with factorial hidden Markov model," IEEE Trans. Audio, Speech, and Language Proc., vol. 19, no. 4, May. 2011.
- Wolpert, D. "Stacked generalization," Neural Networks, 5(2), pp 241-259, 1992.
- Xiao, L. and Deng, L. "A geometric perspective of large-margin training of Gaussian models," IEEE Signal Processing Magazine, vol. 27, no. 6, pp. 118-123, IEEE, November 2010.
- Yamin, S., Deng, L., Wang, Y., and Acero, A. "An integrative and discriminative technique for spoken utterance classification," IEEE Trans. Audio, Speech, and Language Proc., 2008.
- Yang, D., Furui, S. "Combining a two-step CRF model and a joint source channel model for machine transliteration," Proc. ACL, Uppsala, Sweden, 2010, pp. 275-280.
- Yu, D., Deng, L., and Seide, F. "The deep tensor neural network with applications to large vocabulary speech recognition," IEEE Trans. Audio, Speech, Lang. Proc., 2013.
- Yu, D. and Deng, L. "Efficient and effective algorithms for training single-hidden-layer neural networks," Pattern Recognition Letters, 2012.
- Yu, D., Seide, F., Li, G., Deng, L. "Exploiting sparseness in deep neural networks for large vocabulary speech recognition," Proc. ICASSP 2012.
- Yu, D., Siniscalchi, S., Deng, L., and Lee, C. "Boosting attribute and phone estimation accuracies with deep neural networks for detection-based speech recognition", Proc. ICASSP 2012a.
- Yu, D. and Deng, L. "Deep learning and its applications to signal and information processing," IEEE Signal Processing Magazine, January 2011, pp. 145-154.
- Yu, D. and Deng, L. "Accelerated parallelizable neural networks learning algorithms for speech recognition," Proc. Interspeech 2011a.
- Yu, D., Wang, S., Karam, Z., Deng, L. "Language recognition using deep-structured conditional random fields," Proc. ICASSP, 2010, pp. 5030-5033.
- Yu, D., Wang, S., Deng, L., "Sequential labeling using deep-structured conditional random fields", J. of Selected Topics in Signal Processing, 2010a.
- Yu, D., Li, J.-Y., and Deng, L. "Calibration of confidence measures in speech recognition," IEEE Trans. Audio, Speech and Language, 2010b.
- Yu, D. and Deng, L. "Deep-structured hidden conditional random fields for phonetic recognition," Proc. Interspeech, Sept. 2010.
- Yu, D, Deng, L., Gong, Y. and Acero, A. "A novel framework and training algorithm for variable-parameter hidden Markov models," IEEE Transactions on Audio, Speech and Language Processing, vol. 17, no. 7, September 2009, pp. 1348-1360.
- Yu, D. and Deng, L. "Solving nonlinear estimation problems using Splines," IEEE Signal Processing Magazine, vol. 26, no. 4, pp. 86-90, July 2009.
- Zamora-Martínez, F., Castro-Bleda, M., España-Boquera, S. "Fast evaluation of connectionist language models," Intern. Conf. Artificial Neural Networks, 2009, pp. 144-151.
- Zen, H., Nankaku, Y., and Tokuda, K. "Continuous stochastic feature mapping based on trajectory HMMs," IEEE Trans. Audio, Speech, and Language Proc., vol. 19, no. 2, Feb. 2011, pp. 417-430.
- Zen, H. Gales, M. J. F. Nankaku, Y. Tokuda, K. "Product of experts for statistical parametric speech synthesis," IEEE Trans. Audio, Speech, and Language Proc., vol. 20, no. 3, March, 2012, pp. 794-805.
- Zweig, G. and Nguyen, P. "A segmental CRF approach to large vocabulary continuous speech recognition," Proc. ASRU, 2009.