

See discussions, stats, and author profiles for this publication at: <https://www.researchgate.net/publication/319135442>

Collaborative Filtering using Denoising Auto-Encoders for Market Basket Data

Article · August 2017

CITATION

1

READS

67

2 authors, including:



[Andres G. Abad](#)

Escuela Superior Politécnica del Litoral (ESPOL)

23 PUBLICATIONS 163 CITATIONS

SEE PROFILE

Collaborative Filtering using Denoising Auto-Encoders for Market Basket Data

Andres G. Abad and Luis I. Reyes-Castro
Escuela Superior Politecnica del Litoral (ESPOL)
Guayaquil-Ecuador

Abstract

Recommender systems (RS) help users navigate large sets of items in the search for “interesting” ones. One approach to RS is Collaborative Filtering (CF), which is based on the idea that *similar* users are interested in *similar* items. Most model-based approaches to CF seek to train a machine-learning/data-mining model based on sparse data; the model is then used to provide recommendations. While most of the proposed approaches are effective for small-size situations, the combinatorial nature of the problem makes it impractical for medium-to-large instances. In this work we present a novel approach to CF that works by training a Denoising Auto-Encoder (DAE) on corrupted baskets, i.e., baskets from which one or more items have been removed. The DAE is then forced to learn to reconstruct the original basket given its corrupted input. Due to recent advancements in optimization and other technologies for training neural-network models (such as DAE), the proposed method results in a scalable and practical approach to CF. The contribution of this work is twofold: (1) to identify missing items in observed baskets and, thus, directly providing a CF model; and, (2) to construct a generative model of baskets which may be used, for instance, in simulation analysis or as part of a more complex analytical method.

Keywords

Recommender systems, collaborative filtering, denoising auto-encoders, market basket data, retail firms

1. Introduction

Recommender systems (RS) help users navigate large sets of items in the search for “interesting” ones. In retail firms, RS work by providing recommendations based on the analysis of sparse transactional data: market basket data. On-line retailers were one of the first sectors to adopt RS; its application was popularized by Amazon’s “*Customers who bought this item also bought*” feature [1]. In bricks-and-mortar stores, RS have been successfully applied to, for instance, designing up-selling strategies through customized discount coupons and targeted marketing campaigns. One approach to RS is Collaborative Filtering (CF), which is based on the idea that *similar* users are interested in *similar* items.

In the model-based approach to CF, data is used to train a machine-learning/data-mining model. Supervised-learning models proposed for CF include regression models, as in [2] where a regression model was proposed for predicting user’s ratings; and classification models, as in [3] where an inductive approach for classification was applied. In [4], a logistic regression model together with a PCA for dimensionality reduction were used. Other supervised-learning approaches proposed for CF include bayesian classifiers [5] and belief networks [6].

Unsupervised-learning models proposed for CF include clustering techniques, such as in [7] where *k*-means clustering was used. In [8] the CF problem was addressed as a sequence of decisions, where the optimal policy was learned using a Markov Decision Process (MDP); in [9] Latent Semantic Analysis was used for CF, reporting higher accuracy and constant time prediction as two of the main advantages of their method. Another unsupervised-learning proposed approach to CF is to estimate the frequency of occurrence of combinations of items in the search of interesting *association rules* [10]. A direct application of this type of analysis is the quantification of the complementary and supplementary relationship between items and their use for CF.

From a generative-model perspective, one may wish to learn the binary multivariate distribution of variables (representing the presence of items in each basket) to arrive to recommendations by performing probabilistic inference. In

[11] a graphical model, termed the dependency network, was used to estimate the conditional probability of each item given the others; probabilistic inference was made over the unconditional joint distribution using a Gibbs-sampling mechanism. In [12] a Restricted Boltzmann Machine (RBM) was used to learn the binary distribution of items, from which conditional-independence assumptions allowed for efficient estimation of the cross-category effect. In general, however, learning probability distributions with graphical models is intractable because of the normalization requirement.

While all these approaches are effective for small-size situations, the combinatorial nature of the problem makes it impractical for medium-to-large instances. In this work we present a novel approach to CF that works by training a Denoising Auto-Encoder (DAE) model on corrupted baskets, i.e., baskets from which one or more items have been removed. The DAE is then forced to learn to reconstruct the original basket given its corrupted input. The contribution of this work is twofold: (1) to identify missing items in observed baskets and, thus, directly providing a CF model; and, (2) to construct a generative model of baskets which may be used, for instance, in simulation analysis or as part of a more complex analytical method.

Due to recent advancements in optimization and other technologies for training neural-network models, the proposed method results in a scalable and practical approach to CF. We train the model using the Adam algorithm, a recently proposed optimization method for effectively training neural networks. Model implementation and training was performed using TensorFlow library, version 0.11.0, on a linux machine endowed with an NVIDIA GK520 GPU with 3,072 cuda cores.

For illustrating the proposed methodology, we used a publicly-available data set \mathcal{S} consisting of 9,835 baskets on $p = 10$ item categories [13]. We represent each basket by a binary vector $\mathbf{x} = [x_1, \dots, x_p]$, with x_i equal to 1 if item i belongs in the basket, and 0 otherwise. A corrupted basket is represented by $\tilde{\mathbf{x}} = [\tilde{x}_1, \dots, \tilde{x}_p]$, with similar definition for the \tilde{x}_i 's. For training and evaluating the model we partition data set \mathcal{S} into training set \mathcal{T} (consisting of 6,885 baskets) and evaluation set \mathcal{E} (consisting of 2,950 baskets.)

2. Denoising Auto-Encoders for Basket Modeling

DAE are neural-network models that are trained to denoise a corrupted signal. In this section we describe its application to the analysis of market basket data in the context of CF.

2.1 Denoise Auto-Encoder Structure

DAE seek to reconstruct original basket \mathbf{x} from corrupted input basket $\tilde{\mathbf{x}}$ by *encoding* input $\tilde{\mathbf{x}}$ into $\mathbf{h} \in \mathbb{R}^N$ through mapping $\mathbf{h} = f(\mathbf{W}^i \tilde{\mathbf{x}} + \mathbf{b}^i)$, with $\mathbf{W}^i \in \mathbb{R}^{N \times p}$ and $\mathbf{b}^i \in \mathbb{R}^N$; and then *decoding* it back into the input space as $\mathbf{y} \in \mathbb{R}^p$ through $\mathbf{y} = g(\mathbf{W}^o \mathbf{h} + \mathbf{b}^o)$, with $\mathbf{W}^o \in \mathbb{R}^{p \times N}$ and $\mathbf{b}^o \in \mathbb{R}^p$ (see Figure 1a.) Vector \mathbf{h} is usually referred to as the hidden layer of the auto-encoder. For convenience, we bundle the model parameters into $\theta = \{\mathbf{W}^i, \mathbf{b}^i, \mathbf{W}^o, \mathbf{b}^o\}$.

The DAE model is, thus, composed of the encode and decode operations, combined as

$$\mathbf{y} = g(f(\tilde{\mathbf{x}})). \quad (1)$$

In this paper, the activation functions f and g are chosen to be the hyperbolic tangent and the sigmoid function, respectively, applied element-wise.




Figure 1: (a) Denoising Auto-encoder (DAE) model. (b) Miss-classification rate for trained DAEs with different number of hidden nodes. (c) Train and test errors for a DAE trained with 100 hidden nodes; batch size of 64 samples; 50,000 rounds; a learning rate of $1e-5$; and clipping thresholds of $\delta^{(l)} = 1$.

A key hyperparameter of the model is the dimension N of the hidden layer. In traditional (non-denoising) auto-encoders (AE), if N is equal or larger than the size of the input, the auto-encoder is able to trivially learn the identity function, providing no data compression. In these traditional AE's, to extract the salient features of interest from the input space, either N must be smaller than the size of the input, or we must use regularization on the model (e.g., weight decay) or on the learning algorithm (e.g., early-stop rules.) Because of the random noise added in DAE, even hidden layers with a dimension N larger than that the size of the input are able to extract relevant information from the input. This is because the DAE is not learning the input itself but, instead, it is learning the conditional density $\mathbb{P}(\mathbf{x}|\tilde{\mathbf{x}})$. Hyperparameter N can be chosen using cross-validation (see Figure 1b.)

2.2 Corruption Mechanism

In DAE models there is a stochastic corruption mechanism, defined by conditional distribution $C(\tilde{\mathbf{x}}|\mathbf{x})$, generating corrupted baskets $\tilde{\mathbf{x}}$. To define $C(\tilde{\mathbf{x}}|\mathbf{x})$ let us first define the empirical naïve distribution $\pi(\mathbf{x}) = \prod_{i=1}^p \pi_i$, where π_i is the probability that item i is contained in a particular basket (π_i is usually referred to as the support of item(set) i in the affinity-analysis literature.) Marginal probabilities π_i are estimated as the observed frequency of the presence of item i in training set \mathcal{T} . To obtain *meaningful* corrupted inputs, we seek to corrupt items proportionally to their occurrence in the training set. Specifically, we define conditional-marginal distribution $C(\tilde{x}_i|\mathbf{x})$ as

$$C(\tilde{x}_i|\mathbf{x}) = \begin{cases} 0 & \text{if } x_i = 0 \text{ or } u_i \leq \pi_i; \\ 1 & \text{else,} \end{cases} \quad (2)$$

where $u_i \sim \text{Uniform}(0, 1)$. We further reject generated corrupted baskets $\tilde{\mathbf{x}}$ with elements $\tilde{x}_i = 0$ for every i .

2.3 Loss function

Loss function L quantifies the degree to which we are dissatisfied with a model output. While we are ultimately interested in maximizing the rate at which we correctly identify missing items, such a counting-like function is, nevertheless, not an appropriate loss function for learning algorithms due to its non-differentiability. As a surrogate function, we will use the cross-entropy between baskets $\mathbf{x} = [x_1, \dots, x_p]$ and $\bar{\mathbf{x}} = [\bar{x}_1, \dots, \bar{x}_p]$ given by

$$L(\mathbf{x}, \bar{\mathbf{x}}) = - \sum_{i=1}^p [x_i \log(\bar{x}_i) + (1 - x_i) \log(1 - \bar{x}_i)]. \quad (3)$$

Note that minimizing loss function L is equivalent to maximizing the log-likelihood of Bernoulli random variables x_i 's, each with parameter \bar{x}_i , as

$$\log p(\mathbf{x}|\bar{\mathbf{x}}) = \log \prod_{i=1}^p \bar{x}_i^{x_i} (1 - \bar{x}_i)^{1-x_i}. \quad (4)$$

This result shows that the DAE model is, in fact, trained to learn conditional density $\mathbb{P}(\mathbf{x}|\tilde{\mathbf{x}})$.

The network is trained by searching the parameter space for the values that minimize the loss function L , i.e., by solving optimization problem

$$\min_{\theta} \sum_{\mathbf{x} \in \mathcal{T}, \tilde{\mathbf{x}}|\mathbf{x} \sim C} L(\mathbf{x}, g(f(\tilde{\mathbf{x}}))). \quad (5)$$

2.4 Training the model

Optimization Problem (5) is usually solved using the back-propagation (BP) algorithm, a gradient-descent method that takes advantage of patterns appearing in the composition of the chain rule of the gradient of the connections between layers. With large data sets and high-dimensional parameter space, BP is usually performed using the stochastic gradient descent (SGD) method, in which at each learning step only a randomly-chosen subset of the original training set is utilized. Specifically, in this work we used the Adam algorithm [14], an implementation of SGD with first- and second-order parameter's momentum on the learning rule (for further details the reader is referred to [15].)

During early stages of the training process, we may need to avoid *exploiting* gradients, which arbitrarily enlarge the learning steps, jeopardizing convergence. To avoid this situation (and based on the discussion in [16]) we clip the gradient as follows

$$\nabla f(\mathbf{x}^{(t)}) = \begin{cases} \nabla f(\mathbf{x}^{(t)}) & \text{if } \|\nabla f(\mathbf{x}^{(t)})\| \leq \delta^{(t)}; \\ \left(\delta^{(t)} / \|\nabla f(\mathbf{x}^{(t)})\| \right) \cdot \nabla f(\mathbf{x}^{(t)}) & \text{else,} \end{cases} \quad (6)$$

where $\|\cdot\|$ is chosen to be the euclidean norm. We further require that $\delta^{(t)} \rightarrow 0$ as $t \rightarrow \infty$.

A seemingly more efficient alternative to direct training DAEs is to use *walkback training*, a variant approach proposed in [17] in which the training set is augmented by including samples with corrupted basket $\tilde{\mathbf{x}}$ obtained based on the currently learned version of $\mathbb{P}_\theta(\mathbf{x}|\tilde{\mathbf{x}})$; the additional samples are obtained from a generative mechanism as the one presented in Section 4. The walkback training procedure is similar in spirit to Contrastive Divergence with k MCMC steps, proposed for training RBMs in [18].

2.5 Evaluating the model

While the model is trained on minimizing loss function L , its performance is, however, evaluated using a misclassification rate scheme on individual items. For this, the output \mathbf{y} of the model is discretized into $\hat{\mathbf{y}}$ by the rule $\hat{y}_i = \mathbb{I}(y_i > \eta)$, for a threshold $\eta \in [0, 1]$, where \mathbb{I} is the indicator function. Threshold η is another hyperparameter of the model that can be tuned using cross-validation.

3. Collaborative filtering

Motivating the application of the DAE model described in Section 2 as a probabilistic CF is straightforward. In the retail context, original basket \mathbf{x} corresponds to the set of items that the customer consciously and unconsciously wants, i.e.: items the customer consciously wants; the ones he forgot or did not know he wants; and the ones he picked but are somehow “unusual” among *similar* customers. Corruption process $C(\tilde{\mathbf{x}}|\mathbf{x})$ models the limitations of our human capabilities to either remember or identify every item we do want. Given observed basket $\tilde{\mathbf{x}}$, the DAE can compute probabilities for each possible basket \mathbf{x} based on learned conditional distribution $\mathbb{P}_\theta(\mathbf{x}|\tilde{\mathbf{x}})$, thus providing a set of probabilistic recommendations.

Using our trained DAE (see Figure 1c and its caption) we evaluate its performance for recommendations using the ROC curve shown in Figure 2a. As it can be seen in the figure, points mainly concentrate around a False Positive Rate (FPR) of 0.12 and a True Positive Rate (TPR) of 0.77. Note that for a CF application we should mainly consider the FPR performance.

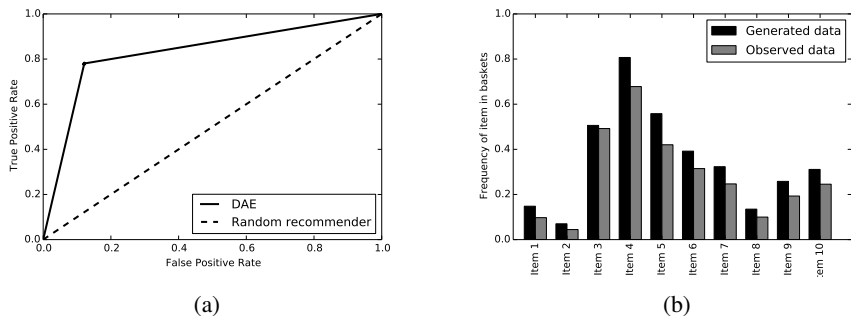


Figure 2: (a) ROC curve for the DAE model. (b) Frequency of items in the training data set and from a generated data set obtained using Algorithm 1.

Table 1 shows the confusion matrix for our DAE model. The matrix was obtained by considering the DAE-based CF performance on six thousand randomly-chosen baskets. The corresponding FPR is 0.122 and the TPR is 0.784.

Table 1: Confusion matrix for a DAE-based CF.

	Predicted Missing	Predicted Present	Total
Observed Missing	16,702	4,732	21,434
Observed Present	4,601	33,965	38,566
Total	21,303	38,697	60,000

4. Market-Basket Generative Model

In this section we describe an iterative procedure to sample from distribution of baskets $\mathbb{P}(\mathbf{x})$ using a trained DAE. The procedure works by first sampling $\tilde{\mathbf{x}}^{(t)}$ from $C(\tilde{\mathbf{x}}|\mathbf{x}^{(t)})$; then sampling $\mathbf{x}^{(t+1)}$ from $\mathbb{P}_\theta(\mathbf{x}|\tilde{\mathbf{x}}^{(t)})$; and repeat alternating.

An initial point $\mathbf{x}^{(0)}$ may be picked arbitrarily. This procedure produces a homogeneous markov chain with transition probabilities $\mathbf{T}(\mathbf{x}^{(t+1)}|\mathbf{x}^{(t)}) = \int \mathbb{P}_\theta(\mathbf{x}^{(t+1)}|\tilde{\mathbf{x}})C(\tilde{\mathbf{x}}|\mathbf{x}^{(t)})d\tilde{\mathbf{x}}$. It was shown in [17] that if $\mathbb{P}_\theta(\mathbf{x}|\tilde{\mathbf{x}})$ is a consistent estimator of $\mathbb{P}(\mathbf{x}|\tilde{\mathbf{x}})$, then the asymptotic marginal distribution of \mathbf{T} (if it exists) is a consistent estimator of $\mathbb{P}(\mathbf{x})$. Algorithm 1 provides a methodology for generating random baskets with a convergent distribution $\mathbb{P}(\mathbf{x})$ using the afore-mentioned result.

Algorithm 1: Algorithm for Generating Random Baskets

Input: Dataset $\mathcal{X} = \{\mathbf{x}_1, \dots, \mathbf{x}_N\}$ consisting of N baskets

Output: Randomly-generated basket \mathbf{x} from distribution $\mathbb{P}(\mathbf{x})$

- 1 Sample $\mathbf{x}^{(t+1)} \sim \pi(\mathbf{x})$
 - 2 **repeat**
 - 3 $\mathbf{x}^{(t)} = \mathbf{x}^{(t+1)}$
 - 4 Sample $\tilde{\mathbf{x}}^{(t)}|\mathbf{x}^{(t)} \sim C(\tilde{\mathbf{x}}|\mathbf{x}^{(t)})$
 - 5 Sample $\mathbf{x}^{(t+1)}|\tilde{\mathbf{x}}^{(t)} \sim \mathbb{P}_\theta(\mathbf{x}|\tilde{\mathbf{x}}^{(t)})$
 - 6 **until** convergence;
 - 7 **return** Basket $\mathbf{x}^{(t+1)} \sim \mathbb{P}(\mathbf{x})$
-

Sampling $\mathbf{x}^{(t+1)} \sim \pi(\mathbf{x})$ (line 1 of Algorithm 1) can be readily done by randomly sampling a basket from data set \mathcal{S} . On the other hand, sampling $\mathbf{x}^{(t+1)}|\tilde{\mathbf{x}}^{(t)} \sim \mathbb{P}_\theta(\mathbf{x}|\tilde{\mathbf{x}}^{(t)})$ (line 5 of Algorithm 1) can be achieved by sampling $\tilde{x}_i^{(t+1)}$ from a Bernoulli distribution with parameter $y_i^{(t)} = [g(f(\mathbf{x}^{(t)}))]_i$, for $i = 1, \dots, p$. Figure 2b shows the item's frequency from the original data set and from a generated data set obtained using Algorithm 1.

While the procedure described here resembles that of a Gibbs-sampling chain, it should be noted that it is not a proper one since there is no guarantee that $C(\tilde{\mathbf{x}}|\mathbf{x})$ and $\mathbb{P}_\theta(\mathbf{x}|\tilde{\mathbf{x}})$ coherently describe a unique joint distribution. The procedure is, in fact, similar to the one proposed in [11] for training dependency networks.

5. Conclusions

We presented a CF model obtained by training a DAE seeking to recover corrupted baskets. We showed that after careful tuning of the parameters, good results can be obtained and the DAE model used as a recommendation mechanism. Based on the DAE trained, a generative model was also presented, which can be used to generate synthetic baskets. This may be used for direct simulation analysis or as part of a more complex analytical method.

Acknowledgements

The authors are thankful to the Board of Directors of Tiendas Industriales Asociadas Sociedad Anónima (TIA S.A.), a leading supermarket retail chain in Ecuador, for authorizing their company to provide historical sales data for hundreds of their products and allowing invaluable discussions with their personnel, which help to shape the present research effort.

References

- [1] G. Linden, B. Smith, and J. York. Amazon.com recommendations: item-to-item collaborative filtering. *IEEE Internet Computing*, 7(1):76–80, January 2003.
- [2] Slobodan Vucetic and Zoran Obradovic. Collaborative Filtering Using a Regression-Based Approach. *Knowledge and Information Systems*, 7(1):1–22, January 2005.
- [3] Chumki Basu, Haym Hirsh, and William Cohen. Recommendation as Classification: Using Social and Content-Based Information in Recommendation. In *In Proceedings of the Fifteenth National Conference on Artificial Intelligence*, pages 714–720. AAAI Press, 1998.
- [4] Jong-Seok Lee, Chi-Hyuck Jun, Jaewook Lee, and Sooyoung Kim. Classification-based collaborative filtering using market basket data. *Expert Systems with Applications*, 29(3):700–704, October 2005.

- [5] Koji Miyahara and Michael J. Pazzani. Collaborative Filtering with the Simple Bayesian Classifier. In Riichiro Mizoguchi and John Slaney, editors, *PRICAI 2000 Topics in Artificial Intelligence*, Lecture Notes in Computer Science, pages 679–689. Springer Berlin Heidelberg, August 2000. DOI: 10.1007/3-540-44533-1_68.
- [6] X. Su and T. M. Khoshgoftaar. Collaborative Filtering for Multi-class Data Using Belief Nets Algorithms. In *2006 18th IEEE International Conference on Tools with Artificial Intelligence (ICTAI'06)*, pages 497–504, November 2006.
- [7] Songjie Gong. An Efficient Collaborative Recommendation Algorithm Based on Item Clustering. In Qi Luo, editor, *Advances in Wireless Networks and Information Systems*, number 72 in Lecture Notes in Electrical Engineering, pages 381–387. Springer Berlin Heidelberg, 2010. DOI: 10.1007/978-3-642-14350-2_48.
- [8] Guy Shani, David Heckerman, and Ronen I. Brafman. An MDP-Based Recommender System. *Journal of Machine Learning Research*, 6(Sep):1265–1295, 2005.
- [9] Thomas Hofmann. Latent Semantic Models for Collaborative Filtering. *ACM Transactions on Information Systems*, 22(1):89–115, January 2004.
- [10] Rakesh Agrawal, Tomasz Imieliński, and Arun Swami. Mining Association Rules Between Sets of Items in Large Databases. In *Proceedings of the 1993 ACM SIGMOD International Conference on Management of Data*, SIGMOD '93, pages 207–216, New York, NY, USA, 1993. ACM.
- [11] David Heckerman, David Maxwell Chickering, Christopher Meek, Robert Rounthwaite, and Carl Kadie. Dependency Networks for Inference, Collaborative Filtering, and Data Visualization. *Journal of Machine Learning Research*, 1(Oct):49–75, 2000.
- [12] Harald Hruschka. Analyzing market baskets by restricted Boltzmann machines. *OR Spectrum*, 36(1):209–228, January 2014.
- [13] Michael Hahsler, Kurt Hornik, and Thomas Reutterer. Implications of Probabilistic Data Modeling for Mining Association Rules. In Professor Dr Myra Spiliopoulou, Professor Dr Rudolf Kruse, Dr Christian Borgelt, Jun-Professor Dr Andreas Nürnberger, and Professor Dr Wolfgang Gaul, editors, *From Data and Information Analysis to Knowledge Engineering*, Studies in Classification, Data Analysis, and Knowledge Organization, pages 598–605. Springer Berlin Heidelberg, 2006. DOI: 10.1007/3-540-31314-1_73.
- [14] Diederik Kingma and Jimmy Ba. Adam: A Method for Stochastic Optimization. *arXiv:1412.6980 [cs]*, December 2014. arXiv: 1412.6980.
- [15] Yoshua Bengio. Practical Recommendations for Gradient-Based Training of Deep Architectures. In Grégoire Montavon, Geneviève B. Orr, and Klaus-Robert Müller, editors, *Neural Networks: Tricks of the Trade*, number 7700 in Lecture Notes in Computer Science, pages 437–478. Springer Berlin Heidelberg, 2012. DOI: 10.1007/978-3-642-35289-8_26.
- [16] Razvan Pascanu, Tomas Mikolov, and Yoshua Bengio. On the difficulty of training recurrent neural networks. pages 1310–1318, 2013.
- [17] Yoshua Bengio, Li Yao, Guillaume Alain, and Pascal Vincent. Generalized Denoising Auto-Encoders as Generative Models. In C. J. C. Burges, L. Bottou, M. Welling, Z. Ghahramani, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 26*, pages 899–907. Curran Associates, Inc., 2013.
- [18] Geoffrey E. Hinton, Simon Osindero, and Yee-Whye Teh. A Fast Learning Algorithm for Deep Belief Nets. *Neural Computation*, 18(7):1527–1554, July 2006.