



## NRC Publications Archive Archives des publications du CNRC

### **Parametric t-distributed stochastic exemplar-centered embeddin** Min, Martin Renqiang; Guo, Hongyu; Shen, Dinghan

This publication could be one of several versions: author's original, accepted manuscript or the publisher's version. /  
La version de cette publication peut être l'une des suivantes : la version prépublication de l'auteur, la version  
acceptée du manuscrit ou la version de l'éditeur.

#### **Publisher's version / Version de l'éditeur:**

*Computer Science, 2017-11-01*

#### **NRC Publications Record / Notice d'Archives des publications de CNRC:**

<https://nrc-publications.canada.ca/eng/view/object/?id=cd03cba4-a957-41fa-9c40-15d5466274ea>  
<https://publications-cnrc.canada.ca/fra/voir/objet/?id=cd03cba4-a957-41fa-9c40-15d5466274ea>

Access and use of this website and the material on it are subject to the Terms and Conditions set forth at  
<https://nrc-publications.canada.ca/eng/copyright>  
READ THESE TERMS AND CONDITIONS CAREFULLY BEFORE USING THIS WEBSITE.

L'accès à ce site Web et l'utilisation de son contenu sont assujettis aux conditions présentées dans le site  
<https://publications-cnrc.canada.ca/fra/droits>  
LISEZ CES CONDITIONS ATTENTIVEMENT AVANT D'UTILISER CE SITE WEB.

**Questions?** Contact the NRC Publications Archive team at  
PublicationsArchive-ArchivesPublications@nrc-cnrc.gc.ca. If you wish to email the authors directly, please see the  
first page of the publication for their contact information.

**Vous avez des questions?** Nous pouvons vous aider. Pour communiquer directement avec un auteur, consultez la  
première page de la revue dans laquelle son article a été publié afin de trouver ses coordonnées. Si vous n'arrivez  
pas à les repérer, communiquez avec nous à PublicationsArchive-ArchivesPublications@nrc-cnrc.gc.ca.



---

# Parametric t-Distributed Stochastic Exemplar-centered Embedding

---

**Martin Renqiang Min**  
 NEC Labs America  
 Princeton, NJ 08540  
 renqiang@nec-labs.com

**Hongyu Guo**  
 National Research Council Canada  
 Ottawa, ON K1A 0R6  
 hongyu.guo@nrc-cnrc.gc.ca

**Dinghan Shen**  
 Duke University  
 Durham, NC 27708  
 dinghan.shen@duke.edu

## Abstract

Parametric embedding methods such as parametric t-SNE (pt-SNE) have been widely adopted for data visualization and out-of-sample data embedding without further computationally expensive optimization or approximation. However, the performance of pt-SNE is highly sensitive to the hyper-parameter batch size due to conflicting optimization goals, and often produces dramatically different embeddings with different choices of user-defined perplexities. To effectively solve these issues, we present parametric t-distributed stochastic exemplar-centered embedding methods. Our strategy learns embedding parameters by comparing given data only with pre-computed exemplars, resulting in a cost function with linear computational and memory complexity, which is further reduced by noise contrastive samples. Moreover, we propose a shallow embedding network with high-order feature interactions for data visualization, which is much easier to tune but produces comparable performance in contrast to a deep neural network employed by pt-SNE. We empirically demonstrate, using several benchmark datasets, that our proposed methods significantly outperform pt-SNE in terms of robustness, visual effects, and quantitative evaluations.

## 1 Introduction

Unsupervised nonlinear dimensionality reduction methods, which embed high-dimensional data to a low-dimensional space, have been extensively deployed in many real-world applications for data visualization. Data visualization is

---

an important component of data exploration and data analytics, as it helps data analysts to develop intuitions and gain deeper understanding about the mechanisms underlying data generation. Comprehensive surveys about dimensionality reduction and data visualization methods can be found in van der Maaten et al. (2009) [26] and Burges (2010) [3]. Among these approaches, nonparametric neighbor embedding methods such as t-SNE [25] and Elastic Embedding [4] are widely adopted. They generate low-dimensional latent representations by preserving neighboring probabilities of high-dimensional data in a low-dimensional space, which involves pairwise data point comparisons and thus has quadratic computational complexity with respect to the size of a given data set. This prevents them from scaling to any dataset with a size beyond several thousand. Moreover, these methods are not designed for generating the embedding of out-of-sample data that are prevalent in modern big data analytics. To generate out-of-sample data embedding given an existing sample embedding, computationally expensive numerical optimization or Nyström approximation is often performed, which is undesirable in practice [2, 29, 5].

Parametric embedding methods, such as parametric t-SNE (pt-SNE) [22] employing a deep neural network (DNN), learn an explicit parametric mapping function from a high-dimensional data space to a low-dimensional embedding space, which can readily generate the embedding of out-of-sample data. The objective function of pt-SNE is the same as that of t-SNE with quadratic computational complexity. Fortunately, owing to the explicit mapping function defined by the DNN, optimization methods such as stochastic gradient descent or conjugate gradient descent based on mini-batches can be deployed when pt-SNE is applied to large-scale datasets.

However, on one hand, the objective function of pt-SNE is a sum of a quadratic number of terms over pairwise data points, which requires mini-batches with fairly large batch sizes to achieve reasonably good approximations to the original objective; On the other hand, optimizing the parameters of the DNN in pt-SNE also requires careful choices of batch sizes, which is often best served with small batch

sizes to avoid being stuck in a bad local minimum. These conflicting choices of batch sizes make the optimization of pt-SNE hard and render its performance sensitive to the chosen batch size. In addition, almost all previous stochastic neighbor embedding approaches including pt-SNE insist on accurately modeling pairwise neighboring probabilities of high-dimensional data regardless of global data density distributions, and often produce dramatically different embeddings with different choices of user-defined perplexities. Finally, although the mapping function of pt-SNE parameterized by a DNN is powerful, it is very hard to learn and requires complicated procedures such as tuning network architectures and tuning many hyper-parameters. For data embedding and visualization purposes, most users are reluctant to go through these complicated procedures.

To address the aforementioned problems, in this paper, we present unsupervised parametric t-distributed stochastic exemplar-centered embedding methods. Instead of modeling pairwise neighboring probabilities, our strategy learns embedding parameters by comparing high-dimensional data only with precomputed representative high-dimensional exemplars, resulting in an objective function with linear computational and memory complexity with respect to the size of a given dataset. The exemplars are identified by a small number of iterations of k-means updates, taking into account both local data density distributions and global clustering patterns of high-dimensional data. These nice properties make the parametric exemplar-centered embedding insensitive to batch size and perplexity, and scalable to large-scale datasets. We further use noise contrastive samples to avoid comparing data points with all exemplars, which further reduces computational/memory complexity and increases scalability. Moreover, we propose a shallow embedding network with high-order feature interactions for data visualization, which is much easier to tune but produces comparable performance in contrast to a deep neural network employed by pt-SNE. Experimental results on several benchmark datasets show that, our proposed parametric exemplar-centered embedding methods for data visualization significantly outperform pt-SNE in terms of robustness, visual effects, and quantitative evaluations. We call our proposed deep t-distributed stochastic exemplar-centered embedding method (dt-SEE) and high-order t-distributed exemplar-centered embedding method (hot-SEE).

Our contributions in this paper are summarized as follows: (1) We propose a scalable unsupervised parametric data embedding strategy with an objective function of linear computational complexity, avoiding pairwise training data comparisons in existing methods; (2) With the help of exemplars, our methods eliminate the instability and sensitivity issues caused by batch sizes and perplexities haunting other unsupervised embedding approaches including pt-SNE; (3) Our proposed approach hot-SEE learns a

simple shallow high-order parametric embedding function, beating state-of-the-art unsupervised deep parametric embedding method pt-SNE on several benchmark datasets in terms of both qualitative and quantitative evaluations.

## 2 Related Work

Dimensionality reduction and data visualization have been extensively studied in the last twenty years [12, 3]. SNE [10], its variant t-SNE [12], and Elastic Embedding [6] are among the most successful approaches. To efficiently generate the embedding of out-of-sample data, SNE and t-SNE were, respectively, extended to take a parametric embedding form of a shallow neural network [14] and a deep neural network [22]. As is discussed in the introduction, the objective functions of neighbor embedding methods have  $O(n^2)$  computational complexity for  $n$  data points, which limits their applicability only to small datasets. Recently, with the growing importance of big data analytics, several research efforts have been devoted to enhancing the scalability of nonparametric neighbor embedding methods [23, 24, 30, 28]. These methods mainly borrowed ideas from efficient approximations developed for N-body force calculations based on Barnes-Hut trees [23] or fast multipole methods [8]. Iterative methods with auxiliary variables and second-order methods have been developed to optimize the objective functions of neighbor embedding approaches [27, 28, 31, 7]. Particularly, the alternating optimization method with auxiliary variables was shown to achieve faster convergence than mini-batch based conjugate gradient method for optimizing the objective function of pt-SNE. All these scalability handling and optimization research efforts are orthogonal to our development in this paper, because all these methods are designed for the embedding approaches modeling the neighboring relationship between pairwise data points. Therefore, they still have the sensitivity and instability issues, and we can readily borrow these speedup methods to further accelerate our approaches modeling the relationship between data points and exemplars.

Our proposed method hot-SEE learns a shallow parametric embedding function by considering high-order feature interactions. High-order feature interactions have been studied for learning Boltzmann Machines, autoencoders, structured outputs, feature selection, and biological sequence classification [13, 17, 15, 20, 21, 9, 19, 11, 16]. To the best of our knowledge, our work here is the first successful one to model input high-order feature interactions for unsupervised data embedding and visualization.

Our work in this paper is also related to a recent supervised data embedding method called en-HOPE [16]. Unlike en-HOPE, our proposed methods here are unsupervised and have a completely different objective function with different motivations.

### 3 Methods

In this section, we introduce the objective of pt-SNE at first. Then we describe the parametric embedding functions of our methods based on a deep neural network as in pt-SNE and a shallow neural network with high-order feature interactions. Finally, we present our proposed parametric stochastic exemplar-centered embedding methods dt-SEE and hot-SEE with linear computational cost.

#### 3.1 Parametric t-SNE using a Deep Neural Network and a Shallow High-order Neural Network

Given a set of data points  $\mathcal{D} = \{\mathbf{x}^{(i)} : i = 1, \dots, n\}$ , where  $\mathbf{x}^{(i)} \in \mathbb{R}^H$  is the input feature vector. pt-SNE learns a deep neural network as a nonlinear feature transformation from the high-dimensional input feature space to a low-dimensional latent embedding space  $\{f(\mathbf{x}^{(i)}) : i = 1, \dots, n\}$ , where  $f(\mathbf{x}^{(i)}) \in \mathbb{R}^h$ , and  $h < H$ . For data visualization, we set  $h = 2$ .

pt-SNE assumes,  $p_{j|i}$ , the probability of each data point  $i$  chooses every other data point  $j$  as its nearest neighbor in the high-dimensional space follows a Gaussian distribution. The joint probabilities measuring the pairwise similarities between data points  $\mathbf{x}^{(i)}$  and  $\mathbf{x}^{(j)}$  are defined by symmetrizing two conditional probabilities,  $p_{j|i}$  and  $p_{i|j}$ , as follows,

$$p_{j|i} = \frac{\exp(-\|\mathbf{x}^{(i)} - \mathbf{x}^{(j)}\|^2 / 2\sigma_i^2)}{\sum_{k \neq i} \exp(-\|\mathbf{x}^{(i)} - \mathbf{x}^{(k)}\|^2 / 2\sigma_i^2)}, \quad (1)$$

$$p_{i|i} = 0, \quad (2)$$

$$p_{ij} = \frac{p_{j|i} + p_{i|j}}{2n}, \quad (3)$$

where the variance of the Gaussian distribution,  $\sigma_i$ , is set such that the perplexity of the conditional distribution  $P_i$  equals a user-specified perplexity  $u$  that can be interpreted as the expected number of nearest neighbors of data point  $i$ . With the same  $u$  set for all data points,  $\sigma_i$ 's tend to be smaller in regions of higher data densities than the ones in regions of lower data densities. The optimal value of  $\sigma_i$  for each data point  $i$  can be easily found by a simple binary search [10]. Although the user-specified perplexity  $u$  makes the variance  $\sigma_i$  for each data point  $i$  adaptive, the embedding performance is still very sensitive to this hyperparameter, which will be discussed later. In the low-dimensional space, pt-SNE assumes, the neighboring probability between pairwise data points  $i$  and  $j$ ,  $q_{ij}$ , follows a heavy-tailed student t-distribution. The student t-distribution is able to, on one hand, measure the similarities between pairwise low-dimensional points, on the other hand, allow dissimilar objects to be modeled far apart in

the embedding space, avoiding crowding problems.

$$q_{ij} = \frac{(1 + \|f(\mathbf{x}^{(i)}) - f(\mathbf{x}^{(j)})\|^2)^{-1}}{\sum_{kl:k \neq l} (1 + \|f(\mathbf{x}^{(k)}) - f(\mathbf{x}^{(l)})\|^2)^{-1}}, \quad (4)$$

$$q_{ii} = 0. \quad (5)$$

To learn the parameters of the deep embedding function  $f(\cdot)$ , pt-SNE minimizes the following Kullback-Leibler divergence between the joint distributions  $P$  and  $Q$  using Conjugate Gradient descent,

$$\ell = KL(P||Q) = \sum_{ij:i \neq j} p_{ij} \log \frac{p_{ij}}{q_{ij}}. \quad (6)$$

The above objective function has  $O(n^2)$  terms defined over pairwise data points, which is computationally prohibitive and prevents pt-SNE from scaling to a fairly big dataset. To overcome such scalability issue, heuristic mini-batch approximation is often used in practice. However, as will be shown in our experiments, pt-SNE is unstable and highly sensitive to the chosen batch size to achieve reasonable performance. This is due to the dilemma of the quadratic cost function approximation and DNN optimization through mini-batches: approaching the true objective requires large batch sizes but finding a good local minimum benefits from small batch sizes.

Although pt-SNE based on a deep neural network has a powerful nonlinear feature transformation, parameter learning is hard and requires complicated procedures such as tuning network architectures and tuning many hyperparameters. Most users who are only interested in data embedding and visualization are reluctant to go through these complicated procedures. Here we propose to use high-order feature interactions, which often capture structural knowledge of input data, to learn a shallow parametric embedding model instead of a deep model. The shallow model is much easier to train and does not have many hyperparameters. In the following, the shallow high-order parametric embedding function will be presented. We expand each input feature vector  $\mathbf{x}$  to have an additional component of 1 for absorbing bias terms, that is,  $\mathbf{x}' = [\mathbf{x}; 1]$ , where  $\mathbf{x}' \in \mathbb{R}^{H+1}$ . The  $O$ -order feature interaction is the product of all possible  $O$  features  $\{x_{i_1} \times \dots \times x_{i_t} \times \dots \times x_{i_O}\}$  where,  $t \in \{1, \dots, O\}$ , and  $\{i_1, \dots, i_t, \dots, i_O\} \in \{1, \dots, H\}$ . Ideally, we want to use each  $O$ -order feature interaction as a coordinate and then learn a linear transformation to map all these high-order feature interactions to a low-dimensional embedding space. However, it's very expensive to enumerate all possible  $O$ -order feature interactions. For example, if  $H = 1000$ ,  $O = 3$ , we must deal with a  $10^9$ -dimensional vector of high-order features. We approximate a Sigmoid-transformed high-order feature mapping  $\mathbf{y} = f(\mathbf{x})$  by constrained tensor factorization as follows,

$$y_s = \sum_{k=1}^m V_{sk} \sigma \left( \sum_{f=1}^F W_{fk} (\mathbf{C}_f^T \mathbf{x}')^O + b_k \right), \quad (7)$$

where  $b_k$  is a bias term,  $\mathbf{C} \in \mathbb{R}^{(H+1) \times F}$  is a factorization matrix,  $\mathbf{C}_f$  is the  $f$ -th column of  $\mathbf{C}$ ,  $\mathbf{W} \in \mathbb{R}^{F \times m}$  and  $\mathbf{V} \in \mathbb{R}^{h \times m}$  are projection matrices,  $y_s$  is the  $s$ -th component of  $\mathbf{y}$ ,  $F$  is the number of factors,  $m$  is the number of high-order hidden units, and  $\sigma(x) = \frac{1}{1+e^{-x}}$ . Because the last component of  $\mathbf{x}'$  is 1 for absorbing bias terms, the full polynomial expansion of  $(\mathbf{C}_f^T \mathbf{x}')^O$  essentially captures all orders of input feature interactions up to order  $O$ . Empirically, we find that  $O = 2$  works best for all datasets we have and set  $O = 2$  for all our experiments. The hyperparameters  $F$  and  $m$  are set by users. Combining Equation 6, Equation 1, Equation 4 and the feature transformation function in Equation 7 leads to a method called high-order t-SNE (hot-SNE). As pt-SNE, the objective function of hot-SNE involves comparing pairwise data points and thus has quadratic computational complexity with respect to the sample size. The parameters of hot-SNE are learned by Conjugate Gradient descent as in pt-SNE.

### 3.2 Parametric t-Distributed Stochastic Exemplar-centered Embedding

To address the instability, sensitivity, and unscalability issues of pt-SNE, we present deep t-distributed stochastic exemplar-centered embedding (dt-SEE) and high-order t-distributed stochastic exemplar-centered embedding (hot-SEE) building upon pt-SNE and hot-SNE for parametric data embedding described earlier. The resulting objective function has linear computational complexity with respect to the size of training set. The underlying intuition is that, instead of comparing pairwise training data points, we compare training data only with a small number of representative exemplars in the training set for neighborhood probability computations. To this end, we simply precompute the exemplars by running a fixed number of iterations of k-means with scalable k-means++ seeding on the training set, which has at most linear computational complexity with respect to the size of training set [1].

Formally, given the same dataset  $\mathcal{D}$  with formal descriptions as introduced in Section 3.1, we perform a fixed number of iterations of k-means updates on the training data to identify  $z$  exemplars from the whole dataset, where  $z$  is a user-specified free parameter and  $z \ll n$  (please note that k-means often converges within a dozen iterations and shows linear computational cost in practice). Before performing k-means updates, the exemplars are carefully seeded by scalable k-means++, which will make our methods robust under abnormal conditions, although our experiments show that random seeding works equally well. We denote these exemplars by  $\{\mathbf{e}^{(j)} : j = 1, \dots, z\}$ . The high-dimensional neighboring probabilities is calculated

through a Gaussian distribution,

$$p_{j|i} = \frac{\exp(-\|\mathbf{x}^{(i)} - \mathbf{x}^{(j)}\|^2 / 2\sigma_i^2)}{\sum_{k \neq i} \exp(-\|\mathbf{x}^{(i)} - \mathbf{x}^{(k)}\|^2 / 2\sigma_i^2)}, \quad (8)$$

$$p_{j|i} = \frac{p_{j|i}}{n}, \quad (9)$$

where  $i = 1, \dots, n, j = 1, \dots, z$ , and the variance of the Gaussian distribution,  $\sigma_i$ , is set such that the perplexity of the conditional distribution  $P_i$  equals a user-specified perplexity  $u$  that can be interpreted as the expected number of nearest exemplars, not neighboring data points, of data instance  $i$ . Since the high-dimensional exemplars capture both local data density distributions and global clustering patterns, different choices of perplexities over exemplars will not change the embedding too much, resulting in much more robust visualization performance than that of other embedding methods insisting on modeling local pairwise neighboring probabilities.

Similarly, the low-dimensional neighboring probabilities is calculated through a t-distribution,

$$q_{j|i} = \frac{(1 + d_{ij})^{-1}}{\sum_{i=1}^n \sum_{k=1}^z (1 + d_{ik})^{-1}}, \quad (10)$$

$$d_{ij} = \|\mathbf{f}(\mathbf{x}^{(i)}) - \mathbf{f}(\mathbf{e}^{(j)})\|^2, \quad (11)$$

where  $\mathbf{f}(\cdot)$  denotes a deep neural network for dt-SEE or the high-order embedding function as described in Equation 7 for hot-SEE.

Then we minimize the following objective function to learn the embedding parameters  $\Theta$  of dt-SEE and hot-SEE while keeping the exemplars  $\{\mathbf{e}^{(j)}\}$  fixed,

$$\min \ell(\Theta, \{\mathbf{e}^{(j)}\}) = \sum_{i=1}^n \sum_{j=1}^z p_{j|i} \log \frac{p_{j|i}}{q_{j|i}} \quad (12)$$

where  $i$  indexes training data points,  $j$  indexes exemplars,  $\Theta$  denotes the high-order embedding parameters  $\{\{b_k\}_{k=1}^m, \mathbf{C}, \mathbf{W}, \mathbf{V}\}$  in Equation 7.

Note that unlike the probability distribution in Equation 4,  $q_{j|i}$  here is computed only using the pairwise distances between training data points and exemplars. This small modification has significant benefits. Because  $z \ll n$ , compared to the quadratic computational complexity with respect to  $n$  of Equation 6, the objective function in Equation 12 has a linear computational complexity with respect to  $n$ .

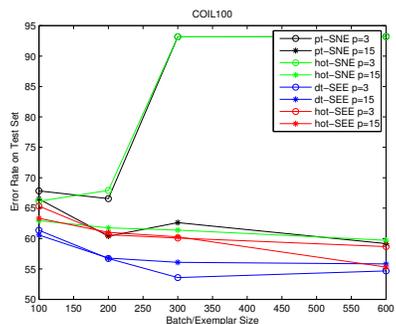
### 3.3 Further Reduction on Computational Complexity and Memory Complexity by Noise Contrastive Estimation

We can even further reduce the computational complexity and memory complexity of dt-SEE and hot-SEE using noise contrastive estimation (NCE). Instead of computing neighboring probabilities between each data point  $i$  and all

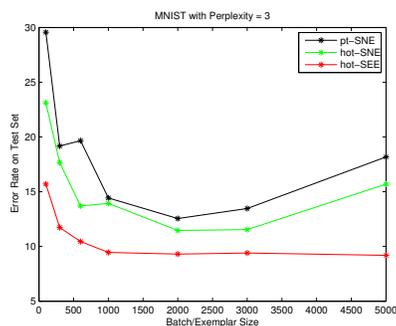
Methods	Error Rates
pt-SNE (batchsize = 1000)	32.48
pt-SNE (batchsize = 2000)	32.04
hot-SNE (batchsize = 1000)	31.29
hot-SNE (batchsize = 2000)	31.82
dt-SEE (batchsize = 1000)	29.42
dt-SEE (batchsize = 2000)	28.30
hot-SEE (batchsize = 1000)	29.06
hot-SEE (batchsize = 2000)	<b>28.18</b>

Table 1: Error rates (%) by 1NN on the 2-dimensional representations produced by different methods with perplexity = 3 on the Fashion dataset.

$z$  exemplars, we can simply only compute the probabilities between data point  $i$  and its  $z_e$  nearest exemplars for both  $P$  and  $Q$ . For high-dimensional probability distribution  $P_i$ , we simply set the probabilities between  $i$  and other exemplars 0; for low-dimensional probability distribution  $Q_i$ , we randomly sample  $z_n$  non-neighboring exemplars outside of these  $z_e$  neighboring exemplars, and use the sum of these  $z_n$  non-neighboring probabilities multiplied by a constant  $K_e$  and the  $z_e$  neighboring probabilities to approximate the normalization terms involving data point  $i$  in Equation 10. Since this strategy based on noise contrastive estimation eliminates the need of computing neighboring probabilities between data points and all exemplars, it further reduces computational and memory complexity.



(a) COIL100

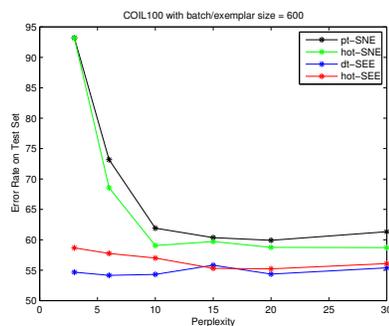


(b) MNIST

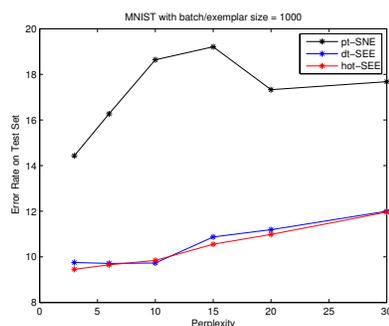
Figure 1: batch size sensitivity test on COIL100 and MNIST

## 4 Experiments

In this section, we evaluate the effectiveness of dt-SEE and hot-SEE by comparing them against state-of-the-art unsupervised parametric embedding method pt-SNE based upon three datasets, *i.e.*, COIL100, MNIST, and Fashion. The COIL100 data<sup>1</sup> contains 7200 images with 100 classes, where 3600 samples for training and 3600 for test. The MNIST dataset<sup>2</sup> consists of 60,000 training and 10,000 test gray-level 784-dimensional images. The Fashion dataset<sup>3</sup> has the same number of classes, training and test data points as that of MNIST, but is designed to classify 10 fashion products, such as boot, coat, and bag, where each contains a set of pictures taken by professional photographers from different aspects of the product, such as looks from front, back, with model, and in an outfit.



(a) COIL100



(b) MNIST

Figure 2: perplexity sensitivity test on COIL100 and MNIST

To make computational procedures and tuning procedures for data visualization simpler, none of these models was pre-trained using any unsupervised learning strategy, although hot-SNE, hot-SEE, dt-SEE, and pt-SNE could all be pre-trained by autoencoders or variants of Restricted Boltzmann Machines [18, 11].

For hot-SNE and hot-SEE, we set  $F = 800$  and  $m = 400$

<sup>1</sup><http://www1.cs.columbia.edu/CAVE/software/softlib/coil-100.php>

<sup>2</sup><http://yann.lecun.com/exdb/mnist/>

<sup>3</sup><https://github.com/zalandoresearch/fashion-mnist>

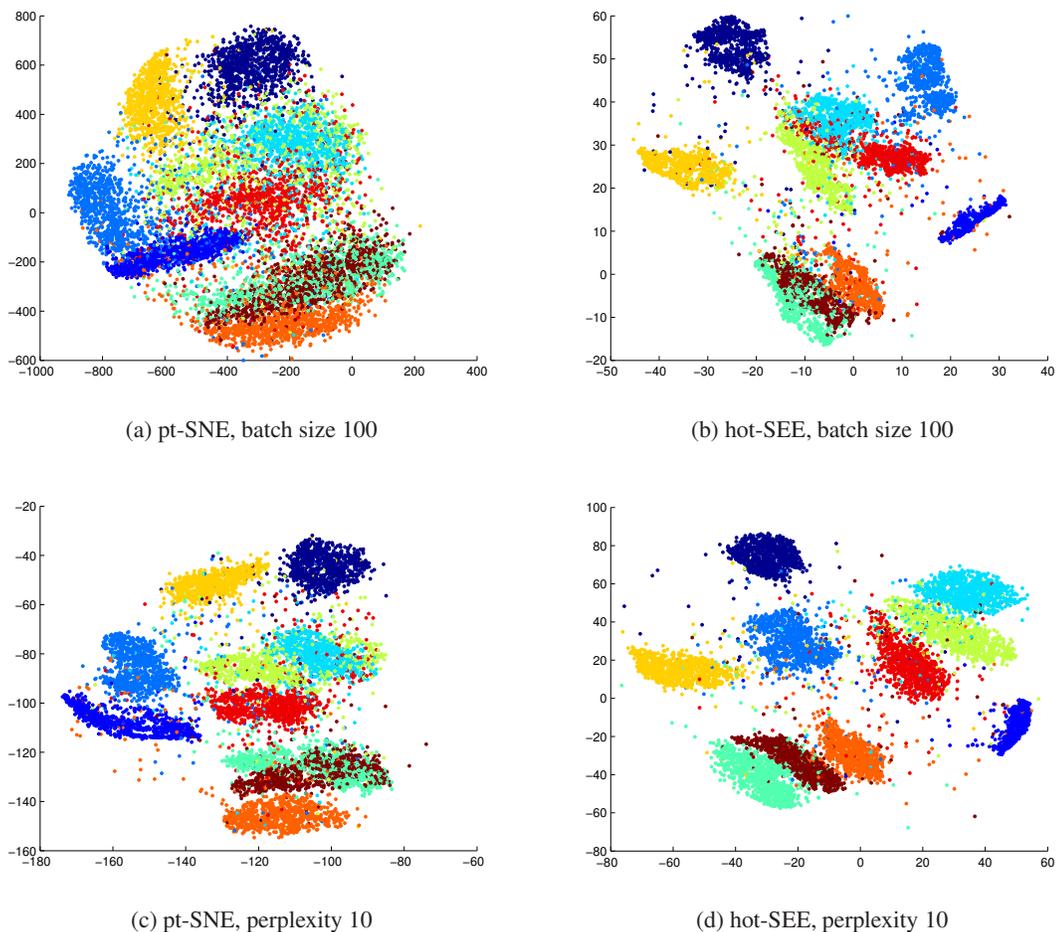


Figure 3: Comparing pt-SNE to hot-SEE with a small batch size = 100 (perplexity = 3) or a reasonable perplexity = 10 (batch size = 1000) to illustrate pt-SNE’s unstable visual performance.

for all the datasets used. For pt-SNE and dt-SEE, we set the deep neural network architecture to input dimensionality  $H-500-500-2000-2$  for all datasets, following the architecture design in van der Maaten (2009) [22]. For hot-SEE and dt-SEE, when the exemplar size is smaller than 1000, we set batch size to 100; otherwise, we set it 1000. With the above architecture design, the shallow high-order neural network used in hot-SNE and hot-SEE is as fast as 2.5 times of the deep neural network used in pt-SNE and dt-SEE for embedding 10,000 MNIST test data.

For all the experiments, the predictive accuracies were obtained by the 1NN approach on top of the 2-dimensional representations generated by different methods. The error rate was calculated by the number of misclassified test data points divided by the total number of test data points.

#### 4.1 Performance Comparisons with Different Batch Sizes and Perplexities on COIL100 and MNIST

Our first experiment aims at examining the robustness of different testing methods with respect to the batch size and the perplexity used. Figures 1 and 2 depict our results on the COIL100 and MNIST datasets when varying the batch size and perplexity, respectively, used by the testing methods.

The top subfigure of Figure 1 suggests that, for the COIL100 data, the pt-SNE was very sensitive to the selection of the batch size; efforts were needed to find a right batch size in order to obtain good performance. On the other hand, the use of different batch sizes had very minor impact on the predictive performance of both the dt-SEE and hot-SEE strategies. Similarly, for the MNIST data, as shown in the bottom subfigure of Figure 1, in order to obtain good predictive performance, the pt-SNE needed to have a batch size not too big and not too small. On the contrary, the hot-SEE methods was insensitive to the size

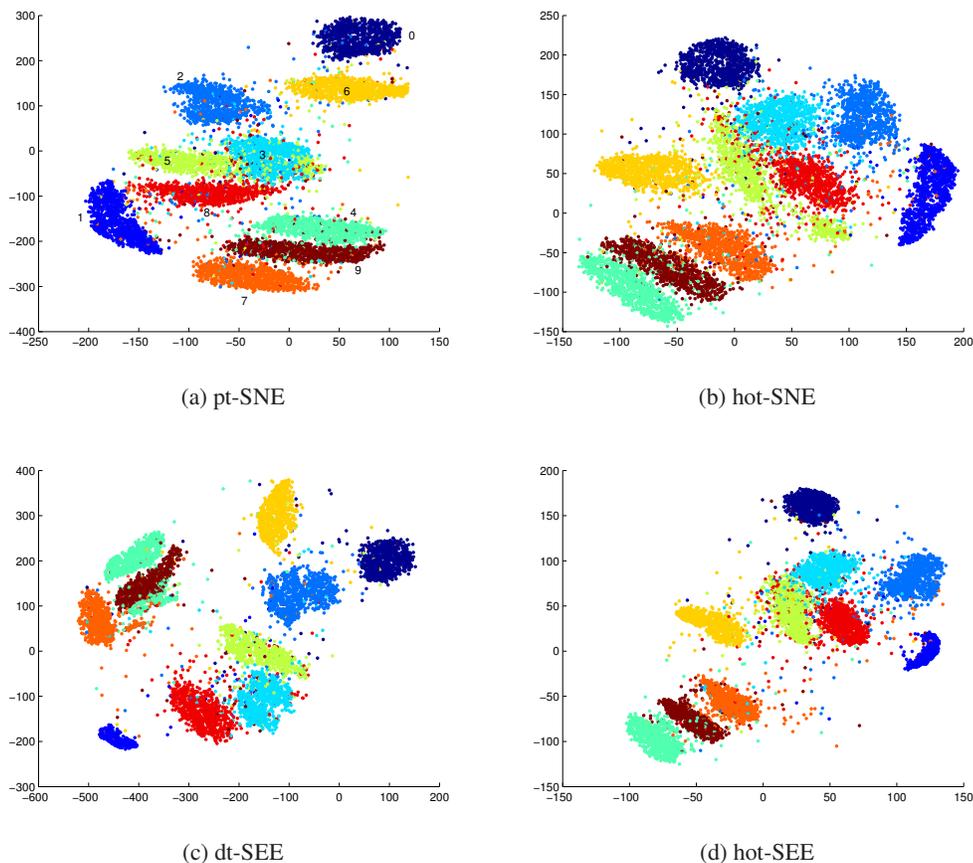


Figure 4: MNIST embedding figures for pt-SNE, hot-SNE, dt-SEE, and hot-SEE

of batch larger than 300.

Based on the results in Figure 1, we selected the best batch sizes for both the COIL100 and MNIST data sets, with 600 and 1000, respectively, but we varied the values of the perplexities used. In the top subfigure of Figure 2, one can observe that, the performance of the pt-SNE and hot-SNE could dramatically change due to the use of different perplexities, but that was not the case for both the dt-SEE and hot-SEE. Similarly, for the MNIST data, as depicted in the bottom subfigure of Figure 2, in order to obtain good predictive performance, one would need to carefully tune for the right perplexity. On the contrary, both the dt-SEE and hot-SEE methods performed quite robust with respect to different selected perplexities.

## 4.2 Experimental Results on the Fashion dataset

We also further evaluated the predictive performance of the testing methods using the Fashion data set. We used batch sizes of 1000 and 2000, along with perplexity of 3 in all the experiments since both pt-SNE and hot-SNE favored these settings as suggested in Figures 1 and 2. The achieved accuracies are shown in Table 1.

Results in Table 1 further confirmed the superior performance of our methods. Both the dt-SEE and hot-SEE significantly outperformed the pt-SNE and hot-SNE.

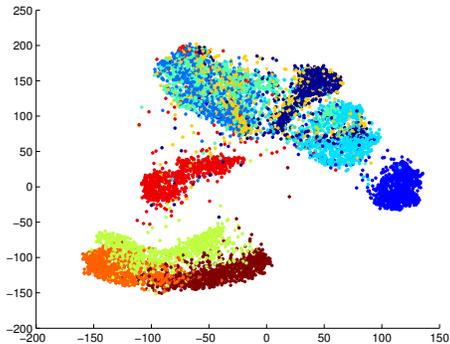
## 4.3 Two-dimensional Visualization of Embeddings

This section provides the visual results of the embeddings formed by the pt-SNE and hot-SEE methods.

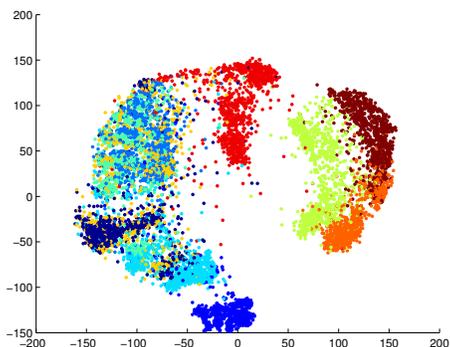
The top and bottom subfigures in Figure 3 depicts the 2D embeddings on the MNIST data set created by pt-SNE and hot-SEE, with batch size of 100 (perplexity = 3) and perplexity of 10 (batch size = 1000), respectively. From these visual figures, one may conclude that the hot-SEE was more stable compared to its competitor pt-SNE.

In Figure 4, we also provided the visual results of the MNIST embeddings created by pt-SNE, hot-SNE, dt-SEE, and hot-SEE, with batch size of 2000. These results imply that the dt-SEE and hot-SEE produced the best visualization: the data points in each cluster were close to each other but with large separation between different clusters, compared to that of the pt-SNE and hot-SNE methods.

Also, in Figures 5, we depicted our visual 2D embedding



(a) pt-SNE



(b) hot-SEE

Figure 5: Fashion embedding figures for pt-SNE and hot-SEE

results on the Fashion data set. These figures further confirmed the better clustering quality generated by the hot-SEE method, compared to that of the pt-SNE strategy.

#### 4.4 Noise Contrastive Estimation

In this section, we evaluated the performance of the noise contrastive estimation (NCE) strategy applied to our method hot-SEE with perplexity 3 and 2000 exemplars. We set  $z_e = z_n = 100$  and  $K_e = 18$ . Table 2 show the error rates (%) obtained by 1NN on the two-dimensional representations produced by hot-SEE with or without NCS, respectively, on the MNIST and Fashion datasets.

Results in Table 2 suggest that the NCE was able to further reduce the computational and memory complexity of our method without sacrificing the predictive performance. As shown in the table, the accuracy difference of the hot-SEE method with and without NCE was less than 0.4% for both the MNIST and Fashion data sets.

MNIST		Fashion	
standard	w/ NCE	standard	w/ NCE
9.30	9.69	28.18	28.19

Table 2: Error rates (%) obtained by 1NN on the two-dimensional representations produced by hot-SEE (perplexity = 3 and 2000 exemplars) with or without further computational complexity reduction based on Noise Contrastive Estimation (NCE), respectively, on the MNIST and Fashion datasets.

#### 4.5 Careful Exemplar Seeding vs. Random Initialization

We also further evaluate the performance of our methods in terms of different exemplar initializations used. We compared the performance of using careful seeding based on scalable K-means++ and randomly initialized exemplars. We presented the results in Table 3. From Table 3, one can observe that our methods were insensitive to the exemplar seeding approach used. That is, very similar predictive performances (less than 0.4%) were obtained by our methods on all the three testing data sets, i.e., COIL100, MNIST, and Fashion.

## 5 Conclusion and Future Work

In this paper, we present unsupervised parametric t-distributed stochastic exemplar-centered data embedding and visualization approaches, leveraging a deep neural network or a shallow neural network with high-order feature interactions. Owing to the benefit of a small number of precomputed high-dimensional exemplars, our approaches avoid pairwise training data comparisons and only have linear computational cost for the given data. In addition, the high-dimensional exemplars reflect local data density distributions and global clustering patterns. With these nice properties, the resulting embedding approaches solved the important problem of embedding performance being sensitive to hyperparameters such as batch sizes and perplexities, which have haunted other neighbor embedding methods for a long time. Experimental results on several benchmark datasets demonstrate that our proposed methods significantly outperform state-of-the-art unsupervised deep parametric embedding method pt-SNE in terms of robustness, visual effects, and quantitative evaluations.

In the future, we plan to extend our exemplar-centered embedding framework to unsupervised nonparametric data embedding and visualization, and we plan to incorporate recent neighbor-embedding speedup developments based on efficient N-body force approximations into our exemplar-centered embedding framework.

COIL100		MNIST		Fashion	
careful seeding	random seeding	careful seeding	random seeding	careful seeding	random seeding
58.67	58.44	9.30	9.19	28.18	28.53

Table 3: Error rates (%) obtained by 1NN on the 2-dimensional representations produced by hot-SEE (perplexity = 3) with careful seeding or random seeding on the COIL100 (with 600 exemplars), MNIST (with 2000 exemplars), and Fashion (with 2000 exemplars) datasets.

## References

- [1] B. Bahmani, B. Moseley, A. Vattani, R. Kumar, and S. Vassilvitskii. Scalable k-means++. *Proceedings of the VLDB Endowment*, 5(7):622–633, 2012.
- [2] Y. Bengio, J.-f. Paiement, P. Vincent, O. Delalleau, N. L. Roux, and M. Ouimet. Out-of-sample extensions for lle, isomap, mds, eigenmaps, and spectral clustering. In *Advances in neural information processing systems*, pages 177–184, 2004.
- [3] C. J. Burges. Dimension reduction: A guided tour, January 2010.
- [4] M. A. Carreira-Perpiñán. The elastic embedding algorithm for dimensionality reduction. In *Proceedings of the 27th International Conference on International Conference on Machine Learning, ICML’10*, pages 167–174, USA, 2010. Omnipress.
- [5] M. A. Carreira-Perpiñán and M. Vladymyrov. A fast, universal algorithm to learn parametric nonlinear embeddings. In *Proceedings of the 28th International Conference on Neural Information Processing Systems, NIPS’15*, pages 253–261, Cambridge, MA, USA, 2015. MIT Press.
- [6] M. A. Carreira-Perpiñán. The elastic embedding algorithm for dimensionality reduction. In *ICML*, volume 10, pages 167–174, 2010.
- [7] M. A. Carreira-Perpiñán and M. Vladymyrov. A fast, universal algorithm to learn parametric nonlinear embeddings. In *Advances in Neural Information Processing Systems*, pages 253–261, 2015.
- [8] L. Greengard and V. Rokhlin. A fast algorithm for particle simulations. *Journal of computational physics*, 73(2):325–348, 1987.
- [9] H. Guo, X. Zhu, and M. R. Min. A deep learning model for structured outputs with high-order interaction. *CoRR*, abs/1504.08022, 2015.
- [10] G. Hinton and S. Roweis. Stochastic neighbor embedding. *Advances in neural information processing systems*, 15:833–840, 2003.
- [11] P. P. Kuksa, M. R. Min, R. Dugar, and M. Gerstein. High-order neural networks and kernel methods for peptide-MHC binding prediction. *Bioinformatics*, 31(22):3600–3607, 2015.
- [12] L. v. d. Maaten and G. Hinton. Visualizing data using t-sne. *Journal of Machine Learning Research*, 9(Nov):2579–2605, 2008.
- [13] R. Memisevic. Gradient-based learning of higher-order image features. In *ICCV*, pages 1591–1598, 2011.
- [14] M. R. Min. A non-linear dimensionality reduction method for improving nearest neighbour classification. In *Master Thesis. Department of Computer Science, University of Toronto*, 2005.
- [15] M. R. Min, S. Chowdhury, Y. Qi, A. Stewart, and R. Ostroff. An integrated approach to blood-based cancer diagnosis and biomarker discovery. In *Pacific Symposium on Biocomputing (PSB)*, pages 87–98, 2014.
- [16] M. R. Min, H. Guo, and D. Song. Exemplar-centered supervised shallow parametric data embedding. In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI 2017, Melbourne, Australia, August 19-25, 2017*, pages 2479–2485, 2017.
- [17] M. R. Min, X. Ning, C. Cheng, and M. Gerstein. Interpretable sparse high-order boltzmann machines. In *Proceedings of the Seventeenth International Conference on Artificial Intelligence and Statistics*, pages 614–622, 2014.
- [18] M. R. Min, L. van der Maaten, Z. Yuan, A. J. Bonner, and Z. Zhang. Deep supervised t-distributed embedding. In *Proceedings of the 27th International Conference on Machine Learning*, pages 791–798, 2010.
- [19] S. Purushotham, M. R. Min, C.-C. J. Kuo, and R. Ostroff. Factorized sparse learning models with interpretable high order feature interactions. In *KDD, New York, USA*, 2014.
- [20] M. Ranzato and G. E. Hinton. Modeling pixel means and covariances using factorized third-order boltzmann machines. In *CVPR*, 2010.

- [21] M. Ranzato, A. Krizhevsky, and G. E. Hinton. Factored 3-way restricted boltzmann machines for modeling natural images. In *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics, Chia Laguna Resort, Sardinia, Italy, May 13-15, 2010*, pages 621–628, 2010.
- [22] L. van der Maaten. Learning a parametric embedding by preserving local structure. In *Proceedings of the 12th International Conference on Artificial Intelligence and Statistics*, pages 384–391, 2009.
- [23] L. Van Der Maaten. Barnes-hut-sne. *arXiv preprint arXiv:1301.3342*, 2013.
- [24] L. Van Der Maaten. Accelerating t-sne using tree-based algorithms. *Journal of machine learning research*, 15(1):3221–3245, 2014.
- [25] L. van der Maaten and G. Hinton. Visualizing data using t-sne. *The Journal of Machine Learning Research*, 9:2579–2605, 2008.
- [26] L. van der Maaten, E. O. Postma, and H. J. van den Herik. Dimensionality reduction: A comparative review, 2008.
- [27] M. Vladymyrov and M. Carreira-Perpinan. Partial-hessian strategies for fast learning of nonlinear embeddings. *arXiv preprint arXiv:1206.4646*, 2012.
- [28] M. Vladymyrov and M. Carreira-Perpinan. Linear-time training of nonlinear low-dimensional embeddings. In *Artificial Intelligence and Statistics*, pages 968–977, 2014.
- [29] M. Vladymyrov and M. Á. Carreira-Perpiñán. Linear-time training of nonlinear low-dimensional embeddings. In *Proceedings of the Seventeenth International Conference on Artificial Intelligence and Statistics, AISTATS 2014, Reykjavik, Iceland, April 22-25, 2014*, pages 968–977, 2014.
- [30] Z. Yang, J. Peltonen, and S. Kaski. Scalable optimization of neighbor embedding for visualization. In *International Conference on Machine Learning*, pages 127–135, 2013.
- [31] Z. Yang, J. Peltonen, and S. Kaski. Majorization-minimization for manifold embedding. In *Artificial Intelligence and Statistics*, pages 1088–1097, 2015.