

# Preserving differential privacy in convolutional deep belief networks

NhatHai Phan<sup>1</sup> · Xintao Wu<sup>2</sup> · Dejing Dou<sup>3</sup>

Received: 18 November 2016 / Accepted: 24 June 2017 / Published online: 13 July 2017 © The Author(s) 2017

Abstract The remarkable development of deep learning in medicine and healthcare domain presents obvious privacy issues, when deep neural networks are built on users' personal and highly sensitive data, e.g., clinical records, user profiles, biomedical images, etc. However, only a few scientific studies on preserving privacy in deep learning have been conducted. In this paper, we focus on developing a private convolutional deep belief network (pCDBN), which essentially is a convolutional deep belief network (CDBN) under differential privacy. Our main idea of enforcing  $\epsilon$ -differential privacy is to leverage the functional mechanism to perturb the energy-based objective functions of traditional CDBNs, rather than their results. One key contribution of this work is that we propose the use of Chebyshev expansion to derive the approximate polynomial representation of objective functions. Our theoretical analysis shows that we can further derive the sensitivity and error bounds of the approximate polynomial representation. As a result, preserving differential privacy in CDBNs is feasible. We applied our model in a health social network, i.e., YesiWell data, and in a handwriting digit dataset, i.e., MNIST data, for human behavior prediction, human behavior classification, and handwriting digit recognition tasks. Theoretical analysis and rigorous experimental evaluations show that the pCDBN is highly effective. It significantly outperforms existing solutions.

Editors: Kurt Driessens, Dragi Kocev, Marko Robnik-Šikonja and Myra Spiliopoulou.

☑ NhatHai Phan phan@njit.edu

> Xintao Wu xintaowu@uark.edu

Dejing Dou dou@cs.uoregon.edu

- <sup>1</sup> New Jersey Institute of Technology, Newark, NJ, USA
- <sup>2</sup> University of Arkansas, Fayetteville, AR, USA
- <sup>3</sup> University of Oregon, Eugene, OR, USA

Keywords Deep learning  $\cdot$  Differential privacy  $\cdot$  Human behavior prediction  $\cdot$  Health informatics  $\cdot$  Image classification

## **1** Introduction

Today, amid rapid adoption of electronic health records and wearables, the global health care systems are systematically collecting longitudinal patient health information, e.g., diagnoses, medication, lab tests, procedures, demography, clinical notes, etc. The patient health information is generated by one or more encounters in any healthcare delivery systems (Jamoom et al. 2016). Healthcare data is now measured in exabytes, and it will reach the zettabyte and the yottabyte range in the near future (Fang et al. 2016). Although appropriate in a variety of situations, many traditional methods of analysis do not automatically capture complex and hidden features from large-scale and perhaps unlabeled data (Miotto et al. 2016). In practice, many health applications depend on including domain knowledge to construct relevant features, some of which are further based on supplemental data. This process is not straightforward and time consuming. That may result in missing opportunities to discover novel patterns and features.

This is where *deep learning*, which is one of the state-of-the-art machine learning techniques, comes in to take advantage of the potential that large-scale healthcare data holds, especially in the age of digital health. Deep neural networks can discover novel patterns and dependencies in both unlabeled and labeled data by applying state-of-the-art training algorithms, e.g., greedy-layer wise (Hinton et al. 2006), contrastive divergent algorithm (Hinton 2002), etc. That makes it easier to extract useful information when building classifiers and predictors (LeCun et al. 2015).

Deep learning has applications in a number of healthcare areas, e.g., phenotype extraction and health risk prediction (Cheng et al. 2016), prediction of the development of various diseases including schizophrenia, a variety of cancers, diabetes, heart failure, etc. (Choi et al. 2016; Li et al. 2015; Miotto et al. 2016; Roumia and Steinhubl 2014; Wu et al. 2010), prediction of risk of readmission (Wu et al. 2010), Alzheimer's diagnosis (Liu et al. 2014; Ortiz et al. 2016), risk prediction for chronic kidney disease progression (Perotte et al. 2015), physical activity prediction (Phan et al. 2015a, b, 2016a, c), feature learning from fMRI data (Plis et al. 2014), diagnosis code assignment (Gottlieb et al. 2013; Perotte et al. 2014), reconstruction of brain circuits (Helmstaedter et al. 2013), prediction of the activity of potential drug molecules (Ma et al. 2015), the effects of mutations in non-coding DNA on gene expressions (Leung et al. 2014; Xiong et al. 2015), and many more.

The development of deep learning in the domain of medicine and healthcare presents obvious privacy issues, when deep neural networks are built based on patients' personal and highly sensitive data, e.g., clinical records, user profiles, biomedical images, etc. To convince individuals to allow that their data be included in deep learning projects, principled and rigorous privacy guarantees must be provided. However, only a few deep learning techniques have yet been developed that incorporate privacy protections. In clinical trials, such lack of protection and efficacy may put patient data at high risk and expose healthcare providers to legal action based on HIPAA/HITECH law (U.S. Department of Health and Human Services 2016a, b). Motivated by this, we aim to develop an algorithm to preserve privacy in fundamental deep learning models in this paper.

Releasing sensitive results of statistical analyses and data mining while protecting privacy has been studied in the past few decades. One state-of-the-art privacy model is  $\epsilon$ -differential

privacy (Dwork et al. 2006). A differential privacy model ensures that the adversary cannot infer any information about any particular data record with high confidence (controlled by a privacy budget  $\epsilon$ ) from the released learning models. This strong standard for privacy guarantees is still valid, even if the adversary possesses all the remaining tuples of the sensitive data. The privacy budget  $\epsilon$  controls the amount by which the output distributions induced by two neighboring databases may differ. We say that two databases are neighboring if they differ in a single data record, that is, if one data record is present in one database and absent in the other. It is clear that the smaller values of  $\epsilon$  enforce a stronger privacy guarantee. This is because it is more difficult to infer any particular data record by distinguishing any two neighboring databases from the output distributions. Differential privacy research has been studied from the theoretical perspective, e.g., Chaudhuri and Monteleoni (2008a), Hay et al. (2010), Kifer and Machanavajjhala (2011) and Lee and Clifton (2012). Different types of mechanisms [e.g., the Laplace mechanism (Dwork et al. 2006), the smooth sensitivity (Nissim et al. 2007), the exponential mechanism (McSherry and Talwar 2007a), and the perturbation of objective function (Chaudhuri and Monteleoni 2008a)] have been studied to enforce differential privacy.

Combining differential privacy and deep learning, i.e., the two state-of-the-art techniques in privacy preserving and machine learning, is timely and crucial. This is a non-trivial task, and therefore only a few scientific studies have been conducted. In Shokri and Shmatikov (2015), the authors proposed a distributed training method, which directly injects noise into gradient descents of parameters, to preserve privacy in neural networks. The method is attractive for applications of deep learning on mobile devices. However, it may consume an unnecessarily large portion of the privacy budget to ensure model accuracy, as the number of training epochs and the number of shared parameters among multiple parties are often large. To improve this, based on the composition theorem (Dwork and Lei 2009; Abadi et al. 2016) proposed a privacy accountant, which keeps track of privacy spending and enforces applicable privacy policies. The approach is still dependent on the number of training epochs. With a small privacy budget  $\epsilon$ , only a small number of epochs can be used to train the model. In practice, that could potentially affect the model utility, when the number of training epochs needs to be large to guarantee the model accuracy.

Recently, Phan et al. (2016c) proposed deep private auto-encoders (dPAs), in which differential privacy is enforced by perturbing the objective functions of deep auto-encoders (Bengio 2009). It is worthy to note that the privacy budget consumed by dPAs is independent of the number of training epochs. A different method, named **CryptoNets**, was proposed in Dowlin et al. (2016) towards the application of neural networks to encrypted data. A data owner can send their encrypted data to a cloud service that hosts the network, and get encrypted predictions in return. This method is different from our context, since it does not aim at releasing learning models under privacy protections.

Existing differential privacy preserving algorithms in deep learning pose major concerns about their applicability. They are either designed for a specific deep learning model, i.e., deep auto-encoders (Phan et al. 2016c), or they are affected by the number of training epochs (Shokri and Shmatikov 2015; Abadi et al. 2016). Therefore, there is an urgent demand for the development of a privacy preserving framework, such that: (1) It is totally independent of the number of training epochs in consuming privacy budget; and (2) It has the potential to be applied in typical energy-based deep neural networks. Such frameworks will significantly promote the application of privacy preservation in deep learning.

Motivated by this, we aim at developing a *private convolutional deep belief network* (pCDBN), which essentially is a convolutional deep belief network (CDBN) (Lee et al. 2009) under differential privacy. CDBN is a typical and well-known deep learning model. It is an

energy-based model. Preserving differential privacy in CDBNs is non-trivial, since CDBNs are more complicated compared with other fundamental models, such as auto-encoders and Restricted Boltzmann Machines (RBM) (Smolensky 1986), in terms of structural designs and learning algorithms. In fact, there are multiple groups of hidden units in each of which parameters are shared in a CDBN. Inappropriate analysis might result in consuming too much of a privacy budget in training phases. The privacy consumption also must be independent of the number of training epochs to guarantee the potential to work with large datasets.

Our key idea is to apply Chebyshev Expansion (Rivlin 1990) to derive polynomial approximations of non-linear objective functions used in CDBNs, such that the design of differential privacy-preserving deep learning is feasible. Then, we inject noise into these polynomial forms, so that the  $\epsilon$ -differential privacy is satisfied in the training phases of each hidden layer by leveraging *functional mechanism* (Zhang et al. 2012). Third, hidden layers now become private hidden layers, which can be stacked on each other to produce a *private convolutional deep belief network* (pCDBN).

To demonstrate the effectiveness of our framework, we applied our model for binomial human behavior prediction and classification tasks in a health social network. A novel human behavior model based on the pCDBN is proposed to predict whether an overweight or obese individual will increase physical exercise in a real health social network. To illustrate the ability to work with large-scale datasets of our model, we also conducted additional experiments on the well-known handwriting digit dataset (MNIST data) (Lecun et al. 1998). We compare our model with the private stochastic gradient descent algorithm, denoted **pSGD**, from Abadi et al. (2016), and the deep private auto-encoders (**dPAs**) (Phan et al. 2016c). The pSGD and dPAs are the state-of-the-art algorithms in preserving differential privacy in deep learning. Theoretical analysis and rigorous experimental evaluations show that our model is highly effective. It significantly outperforms existing solutions.

The rest of the paper is organized as follows. In Sect. 2, we introduce preliminaries and related works. We present our private convolutional deep belief network in Sect. 3. The experimental evaluation is in Sect. 4, and we conclude the paper in Sect. 5.

# 2 Preliminaries and related works

In this section, we briefly revisit the definition of differential privacy, functional mechanism (Zhang et al. 2012), convolutional deep belief networks (Lee et al. 2009), and the Chebyshev Expansion (Rivlin 1990). Let *D* be a database that contains *n* tuples  $t_1, t_2, \ldots, t_n$  and d+1 attributes  $X_1, X_2, \ldots, X_d, Y$ . For each tuple  $t_i = (x_{i1}, x_{i2}, \ldots, x_{id}, y_i)$ , we assume, without loss of generality,  $\sqrt{\sum_{j=1}^d x_{ij}^2} \le 1$  where  $x_{ij} \ge 0$ ,  $y_i$  follows a binomial distribution. Our objective is to construct a deep neural network  $\rho$  from *D* that (i) takes  $\mathbf{x}_i = (x_{i1}, x_{i2}, \ldots, x_{id})$  as input and (ii) outputs a prediction of  $y_i$  that is as accurate as possible.  $t_i$  and  $\mathbf{x}_i$  are used exchangeably to indicate the data tuple *i*. The model function  $\rho$  contains a model parameter vector *W*. To evaluate whether *W* leads to an accurate model, a cost function  $f_D(W)$  is often used to measure the difference between the original and predicted values of  $y_i$ . As the released model parameter *W* may disclose sensitive information of *D*, to protect the privacy, we require that the model training should be performed with an algorithm that satisfies  $\epsilon$ -differential privacy.

Differential privacy (Dwork et al. 2006) establishes a strong standard for privacy guarantees for algorithms, e.g., training algorithms of machine learning models, on aggregate databases. It is defined in the context of neighboring databases. We say that two databases are neighboring if they differ in a single data record. That is, if one data record is present in one database and absent in the other. The definition of differential privacy is as follows:

**Definition 1** [ $\epsilon$ -Different Privacy (Dwork et al. 2006)]. A randomized algorithm A fulfills  $\epsilon$ -differential privacy, iff for any two databases D and D' differing at most one tuple, and for all  $O \subseteq Range(A)$ , we have:

$$Pr[A(D) = O] \le e^{\epsilon} Pr[A(D') = O]$$
(1)

where the privacy budget  $\epsilon$  controls the amount by which the distributions induced by two neighboring datasets may differ. Smaller values of  $\epsilon$  enforce a stronger privacy guarantee of A.

A general method for computing an approximation to any function f (on D) while preserving  $\epsilon$ -differential privacy is the *Laplace mechanism* (Dwork et al. 2006), where the output of f is a vector of real numbers. In particular, the mechanism exploits the global sensitivity of f over any two neighboring databases (differing at most one record), which is denoted as  $GS_f(D)$ . Given  $GS_f(D)$ , the Laplace mechanism ensures  $\epsilon$ -differential privacy by injecting noise  $\eta$  into each value in the output of f(D) as

$$pdf(\eta) = \frac{\epsilon}{2GS_f(D)} exp\left(-|\eta| \cdot \frac{\epsilon}{GS_f(D)}\right)$$
(2)

where  $\eta$  is drawn i.i.d. from Laplace distribution with zero mean and scale  $GS_f(D)/\epsilon$ .

Research in differential privacy has been significantly studied, from both the theoretical perspective, e.g., Chaudhuri and Monteleoni (2008b), Kifer and Machanavajjhala (2011), and the application perspective, e.g., data collection (Erlingsson et al. 2014), data streams (Chan et al. 2012), stochastic gradient descents (Song et al. 2013), recommendation (McSherry and Mironov 2009), regression (Chaudhuri and Monteleoni 2008b), online learning (Jain et al. 2012), publishing contingency tables (Xiao et al. 2010), and spectral graph analysis (Wang et al. 2013). The mechanisms of achieving differential privacy mainly include the classic approach of adding Laplacian noise (Dwork et al. 2006), the exponential mechanism (McSherry and Talwar 2007b), and the functional perturbation approach (Chaudhuri and Monteleoni 2008b).

#### 2.1 Functional mechanism revisited

Functional mechanism (Zhang et al. 2012) is an extension of the Laplace mechanism. It achieves  $\epsilon$ -differential privacy by perturbing the objective function  $f_D(W)$  and then releasing the model parameter  $\overline{W}$  that minimizes the perturbed objective function  $\overline{f}_D(W)$  instead of the original one. The functional mechanism exploits the polynomial representation of  $f_D(W)$ . The model parameter W is a vector that contains d values  $W_1, \ldots, W_d$ . Let  $\phi(W)$  denote a product of  $W_1, \ldots, W_d$ , namely,  $\phi(W) = W_1^{c_1} \cdot W_2^{c_2} \cdots W_d^{c_d}$  for some  $c_1, \ldots, c_d \in \mathbb{N}$ . Let  $\Phi_j(j \in \mathbb{N})$  denote the set of all products of  $W_1, \ldots, W_d$  with degree j, i.e.,  $\Phi_j = \left\{ W_1^{c_1} \cdot W_2^{c_2} \cdots W_d^{c_d} \middle| \sum_{l=1}^d c_l = j \right\}$ . By the Stone-Weierstrass Theorem, any continuous and differentiable  $f(t_i, W) = \sum_{j=0}^J \sum_{\phi \in \Phi_j} \lambda_{\phi t_i} \phi(W)$  where  $\lambda_{\phi t_i} \in \mathbb{R}$  denotes the coefficient of  $\phi(W)$  in the polynomial. Note that  $t_i$  and  $\mathbf{x}_i$  are used exchangeably to indicate the data tuple i.

For instance, the polynomial expression of the loss function in the linear regression is as follows:  $f(\mathbf{x}_i, W) = (y_i - \mathbf{x}_i^T W)^2 = y_i^2 - \sum_{j=1}^d (2y_i x_{ij}) W_j + \sum_{1 \le j, l \le d} (x_{ij} x_{il}) W_j W_l$ .

P<sup>k</sup> (pooling layer)

H (hidden layer)

(visible layer)

**Fig. 1** Convolutional restricted Boltzmann machine (CRBM)

We can see that it only involves monomials in  $\Phi_0 = \{1\}, \Phi_1 = \{W_1, \ldots, W_d\}$ , and  $\Phi_2 = \{W_i W_j | i, j \in [1, d]\}$ . Each  $\phi(W)$  has its own coefficient, e.g., for  $W_j$ , its polynomial coefficient  $\lambda_{\phi_{i_i}} = -2y_i x_{ij}$ . Similarly,  $f_D(W)$  can also be expressed as a polynomial of  $W_1, \ldots, W_d$ .

$$f_D(W) = \sum_{j=0}^J \sum_{\phi \in \Phi_j} \sum_{t_i \in D} \lambda_{\phi t_i} \phi(W)$$
(3)

 $p_{\alpha}^{k}$ 

 $W^k$ 

С

**Lemma 1** (Zhang et al. 2012) Let D and D' be any two neighboring datasets. Let  $f_D(W)$  and  $f_{D'}(W)$  be the objective functions of regression analysis on D and D', respectively. The following inequality holds

$$\Delta = \sum_{j=1}^{J} \sum_{\phi \in \Phi_j}^{2} \left\| \sum_{t_i \in D} \lambda_{\phi t_i} - \sum_{t_i' \in D'} \lambda_{\phi t_i'} \right\|_1 \le 2 \max_t \sum_{j=1}^{J} \sum_{\phi \in \Phi_j} \|\lambda_{\phi t}\|_1$$

where  $t_i, t'_i$  or t is an arbitrary tuple.

To achieve  $\epsilon$ -differential privacy,  $f_D(W)$  is perturbed by injecting Laplace noise  $Lap(\frac{\Delta}{\epsilon})$  into its polynomial coefficients  $\lambda_{\phi}$ , and then the model parameter  $\overline{W}$  is derived to minimize the perturbed function  $\overline{f}_D(W)$ , where  $\Delta = 2 \max_t \sum_{j=1}^J \sum_{\phi \in \Phi_j} \|\lambda_{\phi t}\|_1$ , according to the Lemma 1.

#### 2.2 Convolutional deep belief networks

The basic convolutional restricted Boltzmann machine (CRBM) (Lee et al. 2009) consists of two layers: an input layer V and a hidden layer H (Fig. 1). The layer of hidden units consists of K groups, each of which is an  $N_H \times N_H$  array of binary units. There are  $N_H^2 K$ hidden units in total. Each group (in K groups) is associated with a  $N_W \times N_W$  filter, where  $N_W = N_V - N_H + 1$ . The filter weights are shared across all the hidden units within the group. In addition, each group of hidden units has a bias  $b_k$ , and all visible units share a single bias c. Training a CRBM is to minimize the following energy function E(v, h) as:

$$E(\mathbf{v},\mathbf{h}) = -\sum_{k=1}^{K} \sum_{i,j=1}^{N_H} \sum_{r,s=1}^{N_W} h_{ij}^k W_{rs}^k v_{i+r-1,j+s-1} - \sum_{k=1}^{K} b_k \sum_{i,j=1}^{N_H} h_{ij}^k - c \sum_{i,j=1}^{N_V} v_{ij}$$
(4)

🖄 Springer



Ν

Gibbs sampling can be applied using the following conditional distributions:

$$P(h_{ij}^k = 1 | \mathbf{V}) = \sigma\left((\tilde{W}^k * v)_{ij} + b_k\right)$$
(5)

$$P(v_{ij} = 1|\mathbf{h}) = \sigma\left(\left(\sum_{k} W^{k} * h^{k}\right)_{ij} + c\right)$$
(6)

where  $\sigma$  is the sigmoid function.

The energy function given the dataset *D* is as follows:

$$E(D, W) = -\sum_{t \in D} \sum_{k=1}^{K} \sum_{i,j=1}^{N_H} \sum_{r,s=1}^{N_W} h_{ij}^{k,t} W_{rs}^k v_{i+r-1,j+s-1}^t - \sum_{t \in D} \sum_{k=1}^{K} b_k \sum_{i,j=1}^{N_H} h_{ij}^{k,t} - c \sum_{t \in D} \sum_{i,j=1}^{N_V} v_{ij}^t$$
(7)

The max-pooling layer plays the role of a signal filter. By stacking multiple CRBMs on top of each other, we can construct a *convolutional deep belief network* (CDBN) (Lee et al. 2009). Regarding the softmax layer, we use the cross-entropy error function for a binomial prediction task. Let  $Y_T$  be a set of labeled data points used to train the model, the cross-entropy error function is given by

$$C(Y_T, \theta) = -\sum_{i=1}^{|Y_T|} \left( y_i \log \hat{y}_i + (1 - y_i) \log(1 - \hat{y}_i) \right)$$
(8)

where 'T' in  $Y_T$  is used to denote "training" data.

We can use the layer-wise unsupervised training algorithm (Bengio et al. 2007) and *back-propagation* to train CDBNs.

#### 2.3 Chebyshev polynomials

In principle, many polynomial approximation techniques, e.g., Taylor Expansion, Bernoulli polynomial, Euler polynomial, Fourier series, Discrete Fourier transform, Legendre polynomial, Hermite polynomial, Gegenbauer polynomial, Laguerre polynomial, Jacobi polynomial, and even the state-of-the-art techniques in the twentieth century, including spectral methods and Finite Element methods (Harper 2012), can be applied to approximate non-linear energy functions used in CDBNs. However, figuring out an appropriate way to use each of them is non-trivial. First, estimating the lower and upper bounds of the approximation error incurred by applying a particular polynomial in deep neural networks is not straightforward; it is very challenging. It is significant to have a strong guarantee in terms of approximation errors incurred by the use of any approximation approach to ensure model utility in deep neural networks. In addition, the approximation error bounds must be independent of the number of data instances to guarantee the ability to be applied in large datasets without consuming excessive privacy budgets.

With these challenging issues, Chebyshev polynomial really stands out. The most important reason behind the usage of Chebyshev polynomial is that the upper and lower bounds of the error incurred by approximating activation functions and energy functions can be estimated and proved, as shown in the next section. Furthermore, these error bounds do not depend on the number of data instances, as we will present in Sect. 3.4. This is a substantial result when working with complex models, such as deep neural networks on large-scale datasets. In addition, Chebyshev polynomials are well-known, efficient, and widely used in many real-world applications (Mason and Handscomb 2002). Therefore, we propose to use Chebyshev polynomials in our work to preserve differential privacy in deep convolution belief networks.

The four kinds of Chebyshev polynomials can be generated from the two-term recursion formula:

$$T_{k+1}(x) = 2xT_k(x) - T_{k-1}(x), \quad T_0(x) = 1$$
(9)

with different choices of initial values  $T_1(x) = x$ , 2x, 2x - 1, 2x + 1.

According to the well-known result (Rivlin 1990), if a function f(x) is the Riemann integrable on interval [-1, 1], f(x) can be presented in a Chebyshev polynomial approximation as follows:

$$f(x) = \sum_{k=0}^{\infty} A_k T_k(x) = A' X(T(x))$$
(10)

where  $A_k = \frac{2}{\pi} \int_{-1}^{1} \frac{f(x)T_k(x)}{\sqrt{1-x^2}} dx$ ,  $k \in \mathbb{N}$ ,  $A' = \left[\frac{1}{2}A_0, \dots, A_k, \dots\right]$ ,  $T_k(x)$  is Chebyshev polynomial of degree k,  $X(T(x)) = [T_0(x) \dots T_k(x) \dots]$ .

The closed form expression for Chebyshev polynomials of any order is:

$$T_i(x) = \sum_{j=0}^{\lfloor i/2 \rfloor} (-1)^j \binom{i}{2j} x^{i-2j} (1-x^2)^j$$
(11)

where [i/2] is the integer part of  $\frac{i}{2}$ .

## 3 Private convolutional deep belief network

In this section, we formally present our framework (Algorithm 1) to develop a convolutional deep belief network under  $\epsilon$ -differential privacy. Intuitively, the algorithm used to develop dPAs can be applied to CDBNs. However, the main issue is that their approximation technique has been especially designed for cross-entropy error-based objective functions (Bengio 2009). There are many challenging issues in adapting their technique in CDBNs. The cross entropy error-based objective function (Eq. 7). As such: (1) It is difficult to derive its global sensitivity used in the functional mechanism, and (2) It is difficult to identify the approximation error bounds in CDBNs. To achieve private convolutional deep belief networks (pCDBNs), we figure out a new approach of using the Chebyshev Expansion (Rivlin 1990) to derive polynomial approximations of non-linear energy-based objective functions (Eq. 7), such that differential privacy can be preserved by leveraging the functional mechanism.

Our framework to construct the pCDBN includes four steps (Algorithm 1).

- First, we derive a polynomial approximation of energy-based function E(D, W) (Eq. 7), using the Chebyshev Expansion. The polynomial approximation is denoted as  $\widehat{E}(D, W)$ .
- Second, the functional mechanism is used to perturb the approximation function  $\widehat{E}(D, W)$ ; the perturbed function is denoted as  $\overline{E}(D, W)$ . We introduce a new result of *sensitivity* computation for CDBNs. Next, we train the model to obtain the optimal perturbed parameters  $\overline{W}$  by using gradient descent. That results in private hidden layers, which are used to produce max-pooling layers. Note that we do not need to enforce differential privacy in max-pooling layers. This is because max-pooling layers play roles as signal filters only.

#### **Algorithm 1: Private Convolutional Deep Belief Network**

1: Derive a polynomial approximation of the energy function E(D, W) (Eq. 7), denoted as  $\widehat{E}(D, W)$ 2: The function  $\widehat{E}(D, W)$  is perturbed by using *functional mechanism* (FM) (Zhang et al. 2012), the perturbed function is denoted as  $\overline{E}(D, W)$ 

3: Stack the private hidden and pooling layers

4: By using the technique in Phan et al. (2016c), we derive and perturb the polynomial approximation of the softmax layer  $\widehat{C}(\theta)$  (Eq. 17), the perturbed function is denoted as  $\overline{C}(\theta)$ , Return  $\overline{\theta} = \arg \min_{\theta} \overline{C}(\theta)$ 

- Third, we stack multiple pairs of a private hidden layer and a max-pooling layer (H, P) on top of each other to construct the private convolutional deep belief network (pCDBN).
- Finally, we apply the technique presented in Phan et al. (2016c) to enforce differential privacy in the softmax layer for prediction and classification tasks.

Let us first derive the polynomial approximation form of E(D, W) by applying Chebyshev Expansion, as follows.

#### 3.1 Polynomial approximation of the energy function

There are two challenges in the energy function E(D, W) that prevent us from applying it for private data reconstruction analysis: (1) Gibbs sampling is used to estimate the value of every  $h_{ij}^k$ ; and (2) The probability of every  $h_{ij}^k$  equal to 1 is a sigmoid function which is not a polynomial function with parameters  $W^k$ . Therefore, it is difficult to derive the sensitivity and error bounds of the approximation polynomial representation of the energy function E(D, W). Perturbing Gibbs sampling is challenging. Meanwhile, injecting noise in the results of Gibbs sampling will significantly affect the properties of hidden variables, i.e., values of hidden variables might be out of their original bounds, i.e., [0, 1].

To address this, we propose to preserve differential privacy in the model before applying Gibbs sampling. The generality is still guaranteed since Gibbs sampling is applied for all hidden units. In addition, we need to derive an effective polynomial approximation of the energy function, so that differential privacy preserving is feasible. First, we propose to consider the probability  $P(h_{ij}^k = 1 | \mathbf{v}) = \sigma((W^k * v)_{ij} + b_k)$  instead of  $h_{ij}^k$  in the energy function E(D, W). The main goal of minimizing the energy function, i.e., "the better the reconstruction of v is, the better the parameters W are," remains the same. Therefore, the generality of our proposed approach is still guaranteed. The energy function can be rewritten as follows:

$$\widetilde{E}(D, W) = \sum_{t \in D} \left[ -\sum_{k=1}^{K} \sum_{i,j=1}^{N_H} \sum_{r,s=1}^{N_W} \sigma \left( (W^k * v^t)_{ij} + b_k \right) \times W^k_{rs} v^t_{i+r-1,j+s-1} - \sum_{k=1}^{K} b_k \sum_{i,j=1}^{N_H} \sigma \left( (W^k * v^t)_{ij} + b_k \right) - c \sum_{i,j=1}^{N_V} v^t_{ij} \right]$$
(12)

As the sigmoid function  $\sigma(\cdot)$  in neural networks satisfies the Reimann integrable condition (Vlcek 2012), it can be approximated by the Chebyshev series. We propose to derive a Chebyshev polynomial approximation function for the  $\sigma((W^k * v^t)_{ij} + b_k)$  that results in a polynomial approximation function for our energy function  $\widetilde{E}(\cdot)$ . To make our sigmoid function satisfy the Riemann integrable condition on [-1, 1], we rewrite it as:  $\sigma(\frac{(W^k * v^t)_{ij} + b_k}{Z_{ij}^k})$  where  $Z_{ij}^k$  is a local response normalization (LRN) term which can be computed as follows:

 $Z_{ij}^{k} = \max\left(\left|(W^{k} * v^{t})_{ij} + b_{k}\right|, \left[q + \alpha \sum_{m=\max(0,k-l/2)}^{\min(K-1,k+l/2)} \left((W^{m} * v^{t})_{ij} + b_{m}\right)^{2}\right]^{\beta}\right), \text{ where the constants } q, l, \alpha, \text{ and } \beta \text{ are hyper-parameters, } K \text{ is the total number of feature maps. As in Krizhevsky et al. (2012), we used } q = 2, l = 5, \alpha = 10^{-4}, \text{ and } \beta = 0.75 \text{ in our experiments.}$ 

From Eq. (10), the Chebyshev polynomial approximation of our sigmoid function is as follows:  $\sum_{k=1}^{\infty} \sum_{j=1}^{\infty} \sum_{k=1}^{\infty} \sum_{j=1}^{\infty} \sum_{k=1}^{\infty} \sum_{j=1}^{\infty} \sum_{j=1$ 

$$\sigma\left(\frac{(W^k * v^t)_{ij} + b_k}{Z_{ij}^k}\right) = \sum_{l=0}^{\infty} A_l T_l\left(\frac{(W^k * v^t)_{ij} + b_k}{Z_{ij}^k}\right)$$
(13)

where  $A_l$  and  $T_l$  can be computed using Eqs. (10) and (11).

Now, there is still a challenge that prevents us from applying the functional mechanism to preserve differential privacy in applying Eq. (13): The equation involves an infinite summation. To address this problem, we remove all orders greater than *L*. Based on the Chebyshev series, the polynomial approximation of the energy function  $\tilde{E}(\cdot)$  can be written as:

$$\widehat{E}(D, W) = \sum_{t \in D} \left[ -\sum_{k=1}^{K} \sum_{i,j=1}^{N_H} \sum_{r,s=1}^{N_W} \left( \sum_{l=0}^{L} A_l T_l \left( \frac{(W^k * v^t)_{ij} + b_k}{Z_{ij}^k} \right) \right) \times W_{rs}^k v_{i+r-1,j+s-1}^t - \sum_{k=1}^{K} b_k \sum_{i,j=1}^{N_H} \sum_{l=0}^{L} A_l T_l \left( \frac{(W^k * v^t)_{ij} + b_k}{Z_{ij}^k} \right) - c \sum_{i,j=1}^{N_V} v_{ij}^t \right]$$
(14)

 $\widehat{E}(\cdot) \text{ is a polynomial approximation function of the original energy function } E(\cdot) \text{ in } Eq. (7).$  Furthermore, the term  $\sum_{l=0}^{L} A_l T_l \left( \frac{(W^k * v^l)_{ij} + b_k}{Z_{ij}^k} \right)$  can be rewritten as:  $\sum_{l=0}^{L} \alpha_l \left( \frac{(W^k * v^l)_{ij} + b_k}{Z_{ij}^k} \right)^l$ , where  $\alpha$  are the Chebyshev polynomial coefficients. For instance, given L = 7, we have  $\sum_{l=0}^{L=7} A_l T_l(X) = \frac{1}{2^5} (-5X^7 + 21X^5 - 35X^3 + 35X + 16)$ , where  $X = \frac{(W^k * v^l)_{ij} + b_k}{Z_{ij}^k}$ .

### 3.2 Perturbation of objective functions

We employ the functional mechanism (Zhang et al. 2012) to perturb the objective function  $\widehat{E}(\cdot)$  by injecting Laplace noise into its polynomial coefficients. The hidden layer contains K groups of hidden units. Each group is trained with a local region of input neurons, which will not be merged with each other in the learning process. Therefore, it is not necessary to aggregate sensitivities of the training algorithm in K groups to the sensitivity of the function  $\widehat{E}(\cdot)$ . Instead, the sensitivity of the function  $\widehat{E}(\cdot)$  can be considered the maximal sensitivity given any single group. As a result, the sensitivity of the function  $\widehat{E}(\cdot)$  can be computed in the following lemma.

**Lemma 2** Let D and D' be any two neighboring datasets. Let  $\widehat{E}(D, W)$  and  $\widehat{E}(D', W)$  be the objective functions of regression analysis on D and D', respectively.  $\alpha$  are Chebyshev

polynomial coefficients. The following inequality holds:

$$\Delta \leq 2 \max_{t,k} \sum_{i,j=1}^{N_H} \sum_{l=0}^{L} |\alpha_l| \left[ \left( \frac{\sum_{r,s=1}^{N_W} v_{ij,rs}^{t,k} + 1}{Z_{ij}^k} \right)^l + \sum_{r,s=1}^{N_W} \left( \frac{\sum_{r',s'=1}^{N_W} v_{ij,rs'}^{t,k} + 1}{Z_{ij}^k} \right)^l |v_{ij,rs}^{t,k}| \right] + \sum_{i,j=1}^{N_V} |v_{ij}^t|$$
(15)

*Proof* By replacing  $W_{rs}^t$  (i.e.,  $\forall r, s$ ),  $b_k$ , and c in  $\widehat{E}(D, W)$  with 1, we have the function with only polynomial coefficients of  $\widehat{E}(D, W)$ , denoted  $\lambda_{\phi D}$ . We have that

$$\lambda_{\phi D} = \sum_{t \in D} \lambda_{\phi t}$$

where

$$\lambda_{\phi t} = -\sum_{k=1}^{K} \sum_{i,j=1}^{N_{H}} \sum_{r,s=1}^{N_{W}} \left( \sum_{l=0}^{L} \alpha_{l} \left( \frac{\sum_{r',s'=1}^{N_{W}} v_{ij,r's'}^{t} + 1}{Z_{ij}^{k}} \right)^{l} \right) v_{ij,rs}^{t}$$
$$-\sum_{k=1}^{K} b_{k} \sum_{i,j=1}^{N_{H}} \sum_{l=0}^{L} \alpha_{l} \left( \frac{\sum_{r,s=1}^{N_{W}} v_{ij,rs}^{t} + 1}{Z_{ij}^{k}} \right)^{l} - \sum_{i,j=1}^{N_{V}} v_{ij}^{t}$$

The  $\widehat{E}(\cdot)$ 's sensitivity can be computed as follows:

$$\Delta = \left\| \sum_{t_i \in D} \lambda_{\phi t_i} - \sum_{t_i' \in D'} \lambda_{\phi t_i'} \right\|_1 \le 2 \max_t \|\lambda_{\phi t}\|_1$$

$$\le 2 \max_t \sum_{k=1}^K \sum_{i,j=1}^{N_H} \sum_{l=0}^L |\alpha_l| \left[ \left( \frac{\sum_{r,s=1}^{N_W} v_{ij,rs}^{t,k} + 1}{Z_{ij}^k} \right)^l + \sum_{r,s=1}^{N_W} \left( \frac{\sum_{r',s'=1}^{N_W} v_{ij,rs'}^{t,k} + 1}{Z_{ij}^k} \right)^l |v_{ij,rs}^{t,k}| \right] + \sum_{i,j=1}^{N_V} |v_{ij}^t|$$
(16)

The current sensitivity is an aggregation of sensitivities from all K groups of hidden units. However, each of them is trained with a local region of input neurons, which will not be merged with the others in the learning process. Therefore, the sensitivity of the function  $\widehat{E}(\cdot)$ can be considered the maximal sensitivity given any single group of hidden units in a hidden layer. From Eq. (16), the final sensitivity of the function  $\widehat{E}(\cdot)$  is as follows:

$$\Delta \leq 2 \max_{t,k} \sum_{i,j=1}^{N_H} \sum_{l=0}^{L} |\alpha_l| \left[ \left( \frac{\sum_{r,s=1}^{N_W} v_{ij,rs}^{t,k} + 1}{Z_{ij}^k} \right)^l + \sum_{r,s=1}^{N_W} \left( \frac{\sum_{r',s'=1}^{N_W} v_{ij,rs'}^{t,k} + 1}{Z_{ij}^k} \right)^l |v_{ij,rs}^{t,k}| \right] + \sum_{i,j=1}^{N_V} |v_{ij}^t|$$

Consequently, the Eq. (15) holds.

Deringer

We use gradient descent to train the perturbed model  $\overline{E}(\cdot)$ . That results in *private hidden layers*. To construct a private convolutional deep belief network (pCDBN), we stack multiple private hidden layers and max-pooling layers on top of each other. The pooling layers only play the roles of signal filters of the private hidden layers. Therefore, there is no need to enforce privacy in max-pooling layers.

#### 3.3 Perturbation of softmax layer

On top of the pCDBN, we add an output layer, which includes a single binomial variable to predict *Y*. The output variable  $\hat{y}$  is fully linked to the hidden variables of the highest hidden (pooling) layer, denoted  $p_{(o)}$ , by weighted connections  $W_{(o)}$ , where *o* is the number of hidden (pooling) layers in the CDBNs. We use the logistic function as an activation function of  $\hat{y}$ , i.e.,  $\hat{y} = \sigma(W_{(o)}p_{(o)})$ . *Cross-entropy error*, which is a typical objective function in deep learning (Bengio 2009), is used as a loss function. The cross-entropy error function has been widely used and applied in real-world applications (Bengio 2009). Therefore, it is critical to preserve differential privacy under the use of the cross-entropy error function. However, other loss functions, e.g., square errors, can be applied in the softmax layer as well. Let  $Y_T$  be a set of labeled data points used to train the model, the cross-entropy error function is given by:

$$C(Y_T, \theta) = -\sum_{i=1}^{|Y_T|} \left( y_i \log(1 + e^{-W_{(o)}p_{i(o)}}) + (1 - y_i) \log(1 + e^{W_{(o)}p_{i(o)}}) \right)$$
(17)

By applying the technique in Phan et al. (2016c), based on Taylor Expansion (Arfken 1985), we can derive the polynomial approximation of the cross-entropy error function as follows:

$$\widehat{C}(Y_T,\theta) = \sum_{i=1}^{|Y_T|} \sum_{l=1}^{2} \sum_{R=0}^{2} \frac{f_l^{(R)}(0)}{R!} \left( W_{(o)} p_{i(o)} \right)^R$$
(18)

where

$$g_1(t_i, W_{(o)}) = W_{(o)} p_{i(o)} , \quad g_2(t_i, W_{(o)}) = W_{(o)} p_{i(o)}$$
  
$$f_1(z) = y_i \log(1 + e^{-z}) , \quad f_2(z) = (1 - y_i) \log(1 + e^{z})$$

To preserve the differential privacy, the softmax layer is perturbed by using the *functional* mechanism (Phan et al. 2016c; Zhang et al. 2012). The sensitivity of the softmax layer,  $\Delta_C$ , is estimated as  $\Delta_C = |p_{(0)}| + \frac{1}{4}|p_{(0)}|^2$  (Phan et al. 2016c).

#### 3.4 Approximation error bounds

The following lemma shows how much error our approximation approaches incur. The average error of the approximations is always bounded, as presented in the following lemma:

**Lemma 3** Approximation Error bounds. Let  $S_L(E) = ||E(D, W) - \widehat{E}(D, W)||$ ,  $U_L(E) = ||E(D, W) - E^*(D, W)||$  where E(D, W) is the target function,  $\widehat{E}(D, W)$  is the approximation function learned by our model, and  $E^*(D, W)$  is the best uniform approximation. The lower and upper bounds of the sum square error are as follows:

$$\left(4 + \frac{4}{\pi^2} \log L\right) N_H^2 K \times U_L(E) > S_L(E) \ge U_L(E) \ge \frac{\pi}{4} N_H^2 K |A_{L+1}|$$
(19)

*Proof* As the well-known results in Rivlin (1990), given a target sigmoid function  $\sigma$ , a polynomial approximation function  $\hat{\sigma}$  learned by the model, and the best uniform approximation of  $\sigma$ ,  $S_L(\sigma) = \|\sigma - \hat{\sigma}\|$ ,  $U_L(\sigma) = \|\sigma - \sigma^*\|$ , we have that:

$$S_L(\sigma) \ge U_L(\sigma) \ge \frac{\pi}{4} |A_{L+1}| \tag{20}$$

Since there are  $N_H^2 K$  hidden units in our pCDBN model, we have  $S_L(E) \ge U_L(E) \ge \frac{\pi}{4}N_H^2 K|A_{L+1}|$ . Similarly, in Rivlin (1990), we also have

$$U_L(\sigma) \le S_L(\sigma) < \left(4 + \frac{4}{\pi^2} \log L\right) U_L(\sigma)$$
(21)

Since there are  $N_H^2 K$  hidden units in our pCDBN model, we have

$$\left(4 + \frac{4}{\pi^2} \log L\right) N_H^2 K \times U_L(E) > S_L(E) \ge U_L(E)$$
(22)

Therefore, the Eq. (19) holds.

The approximation error depends on the structure of the energy function E(D, W), i.e., the number of hidden neurons  $N_H^2 K$  and  $|A_{L+1}|$ , and the number of attributes of the dataset. Lemma 3 can be used to determine when it should stop learning the approximation model. For each group of  $N_H^2$  hidden units, the upper bound of the sum square error is only  $\frac{\pi}{4}N_H^2|A_{L+1}|$ , i.e.,  $|A_{L+1}|$  is tiny when L is large enough.

Importantly, Lemmas 2 and 3 show that the sensitivity  $\Delta$  and the approximation error bounds of the energy-based function are entirely independent of the number of data instances. This sufficiently guarantees that our differential privacy preserving framework can be applied in large datasets without consuming excessive privacy budgets. This is a substantial result when working with complex models, such as deep neural networks on large-scale datasets. It is worth noting that non-linear activation functions, which are continuously differentiable [Stone-Weierstrass Theorem (Rudin 1976)] and satisfy the Riemann-integrable condition, can be approximated by using Chebyshev Expansion. Therefore, our framework can be applied given such activation functions as, e.g., tanh, arctan, sigmoid, softsign, sinusoid, sinc, Gaussian, etc. (Wikipedia 2016). In the experiment section, we will show that our approach leads to accurate results.

Note that the proofs of Lemmas 2 and 3 do not depend on the assumption of the data features being non-negative, and that the target follows by a binomial distribution. The proofs are generally applicable for inputs and the target, which are not restricted by any constraint. As shown in the next section, our approach efficiently works with a multi-class classification task on the MNIST dataset (Lecun et al. 1998). The cross-entropy error function is applied in the softmax layer.

# 4 Experiments

To validate our approach, we have conducted an extensive experiment on well-known and large-scale datasets, including a health social network, YesiWell data (Phan et al. 2016c), and a handwriting digit dataset, MNIST (Lecun et al. 1998). Our task of validation focuses on four key issues: (1) The effectiveness and robustness of our pCDBN model; (2) The effects of our model and hyper-parameter selections, including the use of Chebyshev polynomial, the impact of the polynomial degree *L*, and the effect of probabilities  $P(h_{ij}^k = 1|v)$  in approximating the energy function; (3) The ability to work on large-scale datasets of our model; and

(4) The benefits of being independent of the number of training epochs in consuming privacy budget.

We carry out the validation through three approaches. One is by conducting the human behavior prediction with various settings of data cardinality, privacy budget  $\epsilon$ , noisy vs. noiseless models, and original versus approximated models. By this we rigorously examine the effectiveness of our model compared with the state-of-the-art algorithms, i.e., Phan et al. (2016c) and Abadi et al. (2016). The second approach is to discover gold standards in our model configuration by examining various settings of hyper-parameters. The third approach is to access the benefits of being independent of the number of training epochs in terms of consuming privacy budget of our pCDBN model. In fact, we present the prediction accuracies of our pCDBN and existing algorithms as a function of the number of training epochs.

# 4.1 Human behavior modeling

In this experiment, we have developed a *private convolutional deep belief network* (pCDBN) for human behavior prediction and classification tasks in the YesiWell health social network (Phan et al. 2016c).

*Health social network data* To be able to compare our model with the state-of-the-art deep private auto-encoders for human behavior prediction (dPAH), we use the same dataset used in Phan et al. (2016c). Data were collected from Oct 2010 to Aug 2011 as a collaboration between PeaceHealth Laboratories, SK Telecom Americas, and the University of Oregon to record daily physical activities, social activities (text messages, competitions, etc.), biomarkers, and biometric measures (cholesterol, BMI, etc.) for a group of 254 overweight and obese individuals. Physical activities, including information about the number of walking and running steps, were reported via a mobile device carried by each user. All users enrolled in an online social network, allowing them to friend and communicate with each other. Users' biomarkers and biometric measures were recorded via daily/weekly/monthly medical tests performed at home individually or at our laboratories.

In total, we consider three groups of attributes:

- Behaviors: #competitions joined, #exercising days, #goals set, #goals achieved, \sum (distances), avg(speeds);
- #Inbox Messages: Encouragement, Fitness, Followup, Competition, Games, Personal, Study protocol, Progress report, Technique, Social network, Meetups, Goal, Wellness meter, Feedback, Heckling, Explanation, Invitation, Notice, Technical fitness, Physical;
- Biomarkers and Biometric Measures: Wellness Score, BMI, BMI slope, Wellness Score slope.

*pCDBN for human behavior modeling* Our starting observation is that a human behavior is the outcome of behavior determinants such as *self-motivation, social influences*, and *environmental events*. This observation is rooted in *human agency in social cognitive theory* (Bandura 1989). In our model, these human behavior determinants are combined together to model human behaviors. Given a tuple  $t_i, x_{i1}, \ldots, x_{id}$  are the personal attributes and  $y_i$  is a binomial parameter that indicates whether a user increases or decreases his/her exercises. To describe the pCDBN model, we adjust the notations  $x_{i1}$  and  $y_i$  a little bit to denote the temporal dimension, and our social network information. Specifically,  $x_u^t = \{x_{1u}^t, \ldots, x_{du}^t\}$  is used to denote the *d* attributes of user *u* at time point *t*. Meanwhile,  $y_u^t$  is used to denote the status of the user *u* at time point *t*.  $y_u^t = 1$  denotes *u* increases exercises at time *t*; otherwise  $y_u^t = 0$ .

In fact, the current behavior at time-stamp t of a user u is conditional on his/her behavior in the past N time-stamps, i.e.,  $t - N, \ldots, t - 1$ . To model this effect [i.e., also considered as a form of self-motivation in social cognitive theory (Bandura 1989)], we first aggregate his personal attributes in the last N timestamps into a  $d \times N$  matrix, which will be considered the visible input V. Then, to model self-motivation and social influence, we add an aggregation of his/her attributes and the effects from his/her friends at the current timestamp t into the dynamic biases, i.e.,  $\hat{b}_k^t$  and  $\hat{c}^t$ , of the hidden and visible units (Eqs. 23–26). The hidden and visible variables at time t are

$$h_{i,j,t}^{k} = \sigma \left( (\tilde{W}^{k} * v^{t})_{ij} + \hat{b}_{k}^{t} \right)$$
(23)

$$v_{i,j,t} = \sigma\left(\left(\sum_{k} W^{k} * h^{k}\right)_{ij}^{t} + \hat{c}^{t}\right)$$
(24)

where

$$\hat{b}_{k}^{t} = b_{k} + \sum_{e=1}^{d} B_{e}^{k} x_{eu}^{t} + \frac{\eta_{k}}{|F_{u}|} \sum_{v \in F_{u}} \psi_{t}(v, u)$$
(25)

$$\hat{c}^{t} = c + \sum_{e=1}^{u} A_{e} x_{eu}^{t} + \frac{\eta}{|F_{u}|} \sum_{v \in F_{u}} \psi_{t}(v, u)$$
(26)

where  $\hat{b}_k^t$  and  $\hat{c}^t$  are dynamic biases,  $B_e^k$  is a matrix of weights which connects  $x_u^t$  with hidden variables in the group k.  $\psi_t(v, u)$  is the probability that v influences u on physical activity at time t.  $\psi_t(v, u)$  is derived from the TaCPP model (Phan et al. 2016b).  $F_u$  is a set of friends of u in the social network.  $\eta$  and  $\eta_k$  are parameters which present the ability to observe the explicit social influences from neighboring users.

The model includes two hidden layers. We trained 10 first layer bases, each  $4 \times 12$  variables v, and 10 second layer bases, each  $2 \times 6$ . The pooling ratio was 2 for both layers. In our work, contrastive divergent algorithm (Hinton 2002) was used to optimize the energy function, and back-propagation was used to optimize the cross-entropy error function in the softmax layer. The implementations of our models using Tensorflow<sup>1</sup> and Python were made publicly available on GitHub.<sup>2</sup> The results and algorithms can be reproduced on either a single workstation or a Hadoop cluster. To examine the effectiveness of our pCDBN, we established two experiments, i.e., prediction and classification, as follows.

#### 4.1.1 Human behavior prediction

*Experimental setting* Our pCDBN model is used to predict the statuses of all the users in the next time point t + 1 given  $\mathcal{M}$  past time points  $t - \mathcal{M} + 1, ..., t$ . The model has been trained on *daily* and *weekly* datasets. Both datasets contain 300 time points, 30 attributes, 254 users, 2766 messages, 1383 friend connections, 11,458 competitions, etc. For each dataset, we have, in total, 254 users ×300 timestamps = 76, 200 data points.

The number of previous time intervals N is set to 4. N is used as a time window to generate training samples. For instance, given 10 days of data (M = 10), a time window of 4 days N = 4, and d data features, e.g., BMI, #steps, etc., a single input V will be a

<sup>&</sup>lt;sup>1</sup> https://www.tensorflow.org.

<sup>&</sup>lt;sup>2</sup> https://github.com/haiphanNJIT/PrivateDeepLearning/releases/tag/V1.0.

 $d \times N (= d \times 4)$  matrix. A single input V is considered as a data sample to model human behavior in our prediction model. If we move the window N on 10 days of data, i.e.,  $\mathcal{M}$ , we will have  $\mathcal{M} - N + 1$  training samples for each individual, i.e., 10 - 4 + 1 = 7 in this example. So, we have, in total,  $254(\mathcal{M} - N + 1) = 254 \times 7 = 1$ , 778 training samples for every 10 days of data  $\mathcal{M}$  to predict whether an individual will increase physical activity in the next day t + 1.

The Chebyshev polynomial approximation degree L and learning rates are set to 7 and  $10^{-3}$ . To avoid over-fitting, we apply the L1-regularization and the dropout technique (Srivastava et al. 2014), i.e., the dropout probability is set to 0.5. Regarding  $\mathcal{K}$ -fold cross-validation or bootstrapping, it is either unnecessary or impractical to apply them in deep learning, and particularly in our study (Bengio 2017; Reed et al. 2014). This is because: (1) It is too expensive and time consuming to train  $\mathcal{K}$  deep neural networks, each of which usually has a large number of parameters, e.g., hundreds of thousands of parameters (Bengio 2017); and (2) Bootstrapping is only used to train neural networks when class labels may be missing, objects in the image may not be localized, and in general, the labeling may be subjective, noisy, and incomplete (Reed et al. 2014). This is out of the scope of our focus. Our models were trained on a graphic card NVIDIA GTX TITAN X, 12 GB RAM with 3072 CUDA cores.

*Competitive models* We compare our pCDBN with two types of state-of-the-art models, as follows:

- (a) Deep learning models for human behavior prediction, including: (1) The original convolutional deep neural network (*CDBN*) for human behavior prediction without enforcing differential privacy; (2) The truncated version of the CDBNs, in which the energy function is approximated without injecting noise to preserve differential privacy, denoted **TCDBN**; and (3) The conditional Restricted Boltzmann Machine, denoted **SctRBM** (Li et al. 2014). None of these models enforces  $\epsilon$ -differential privacy.
- (b) Deep Private Auto-Encoder (dPAH) (Phan et al. 2016c), which is the state-of-the-art deep learning model under differential privacy for human behavior prediction. The dPAH model outperforms general methods for regression analysis under ε-differential privacy, i.e., functional mechanism (Zhang et al. 2012), DPME (Lei 2011), and filter-priority (Cormode 2011). Therefore, we only compare our model with the dPAH.

• Accuracy versus dataset cardinality Fig. 2 shows the prediction accuracy of each algorithm as a function of the dataset cardinality. We vary the size of  $\mathcal{M}$ , which also can be considered as the sampling rate of the dataset.  $\epsilon$  is 1.0 in this experiment. In both datasets, there is a gap between the prediction accuracy of pCDBN and the original convolutional deep



Fig. 2 Prediction accuracy versus dataset cardinality



Fig. 3 Prediction accuracy versus privacy budget  $\epsilon$ 

belief network (CDBN). However, the gap dramatically gets smaller with the increase of the dataset cardinality ( $\mathcal{M}$ ). In addition, our pCDBN outperforms the state-of-the-art dPAH in most of the cases, and the results are statistically significant ( $p = 5.3828e^{-05}$ , performed by paired *t* test). It also is significantly better than the SctRBM when the sampling rate goes just a bit higher, i.e., > 0.2 or > 0.3 ( $p = 8.8350e^{-04}$ , performed by paired *t* test). Either 0.2 or 0.3 is a small sampling rate; thus, this is a remarkable result.

• Accuracy versus privacy budget Fig. 3 illustrates the prediction accuracy of each model as a function of the privacy budget  $\epsilon$ .  $\mathcal{M}$  is set to  $12 \approx 0.32\%$ . The prediction accuracies of privacy non-enforcing models remain unchanged for all  $\epsilon$ . Since a smaller  $\epsilon$  requires a larger amount of noise to be injected, privacy enforcing models incur higher inaccurate prediction results when  $\epsilon$  decreases. pCDBN outperforms dPAH in all cases, and the results are statistically significant ( $p = 2.7266e^{-12}$ , performed by paired t test). In addition, it is relatively robust against the change of  $\epsilon$ . In fact, the pCDBN model is competitive even with privacy non-enforcing models, i.e., SctRBM.

• Probabilities  $P(h_{ij}^k = 1|v)$  and Gibbs sampling To approximate the energy function E(D, W), we propose to use the probabilities  $P(h_{ij}^k = 1|v)$  instead of the values of  $h_{ij}^k$ , which are estimated by applying Gibbs Sampling on the  $P(h_{ij}^k = 1|v)$ . To illustrate the effect of our approach, we conducted both theoretical analysis and experimental evaluations as follows. Let's use  $h_{ij}^k$  to estimate the sensitivity  $\Delta$  of the energy function E(D, W) (Eq. 7) by following Lemma 1 as follows:

$$\Delta = 2 \max_{t,k} \sum_{i,j=1}^{N_H} \sum_{r,s=1}^{N_W} \left| h_{ij}^{k,t} v_{i+r-1,j+s-1}^t \right| + \sum_{i,j=1}^{N_H} \left| h_{ij}^{k,t} \right| + \sum_{i,j=1}^{N_V} \left| v_{ij}^t \right|$$
(27)

There are several issues in the Eq. (27) that prevent us from applying it. First,  $h_{ij}^k$  cannot be considered an observed variable, since its value can only be estimated by applying Gibbs sampling from observed variables v and parameters W. In other words, the value of  $\Delta$ is significantly dependent on Gibbs sampling given  $P(h_{ij}^k = 1|v)$ . Therefore,  $\Delta$  can be uncertain in every sampling step. That may lead to a violation of the guarantee of privacy protection under a differential privacy mechanism. To address this issue, one may set all the hidden variables  $h_{ij}^k$  to 1. That leads to the use of a maximal value of the sensitivity  $\Delta$  as follows:

$$\Delta = 2 \max_{t,k} \sum_{i,j=1}^{N_H} \sum_{r,s=1}^{N_W} \left| v_{i+r-1,j+s-1}^t \right| + N_H^2 + \sum_{i,j=1}^{N_V} \left| v_{ij}^t \right|$$
(28)

Deringer



Fig. 4 Results of classification accuracy for different noise levels, different approximation approaches, and the number of training epochs

The maximal value of  $\Delta$  (Eq. 28) is huge and is not an optimal bound. In other words, the model efficiency will be affected, since too much noise will be unnecessarily injected into the model.

To tackle this challenge, our solution is to consider the probabilities  $P(h_{ij}^k = 1|v)$  instead of  $h_{ij}^k$ . As a result, the sensitivity  $\Delta$  in Lemma 2 is only dependent on observed variables vinstead of Gibbs samplings. That leads to a smaller amount of noise injected into the model. To demonstrate the effect of this approach, our model is compared with its truncated version, in which the energy function is approximated without injecting noise to preserve differential privacy, denoted **TCDBN**. Experimental results illustrated in Figs. 2 and 3 show that the impact of our approach on the original model **CDBN** is marginal in terms of prediction accuracy. On average, the prediction accuracy is only less than 1% lower compared with the original model. This is a practical result.

#### 4.1.2 Human behavior classification

In this experiment, we aim to examine: (1) The robustness of our approach when it is trained with a large number of epochs at different noise levels; and (2) The effectiveness of different approximation approaches, including Chebyshev, Taylor, and Piecewise approximations. Our experiment setting is as follows:

We consider every pair (u, t) is a data point. Given t is a week, we have, in total, 9652 data points (254 users × 38 weeks). We randomly select 10% data points as a testing set, and the remaining data points are used as a training set. At each training step, the model is trained with 111 randomly selected data points, i.e., batch size = 111. To avoid the imbalance in the data, each training batch consists of a balanced number of data samples from different data classes. With this technique, data points in the under-represented class can be incidentally sampled more than the others (Brownlee 2015). The model is used to classify the statuses of all the users given their features. In this experiment, we compare our model with state-of-the-art polynomial approximation approaches in digital implementations, including truncated Taylor Expansion:  $\sigma(x) = \tanh x \approx x - \frac{x^3}{3} + \frac{2x^5}{15}$  (Lee and Jeng 1998; Vlcek 2012) (**pCDBN\_TE**), and linear piecewise approximation:  $\sigma(x) \approx c_1 x + c_2$  (Armato et al. 2009) (**pCDBN\_PW**). Other baseline models, i.e., dPAH and SctRBM, cannot be directly applied to this task; so, we do not include them in this experiment.

• Figure 4 shows classification accuracies for different levels of privacy budget  $\epsilon$ . Each plot illustrates the evolution of the testing accuracy of each algorithm and its power fit curve as a function of the number of epochs. After 600 epochs, our pCDBN can achieve 88% with  $\epsilon = 0.1, 92\%$  with  $\epsilon = 2$ , and 94% with  $\epsilon = 8$ . In addition, our model outperforms baseline approaches, i.e., pCDBN\_TE and pCDBN\_PW, and the results are statistically significant ( $p = 4.4293e^{-07}$ , performed by paired t test). One of the important observations we acquire from this result is that: The Chebyshev polynomial approximation is more effective than the competitive approaches in preserving differential privacy in convolutional deep belief networks. One of the reasons is that Chebyshev polynomial approximation incurs fewer errors than the other two approaches (Harper 2012; Vlcek 2012). Similar to Layer-wise Relevance Propagation (Bach et al. 2015), the approximation errors will propagate across neural layers. Therefore, the smaller the error, the more accurate the models will be.

Note that our observations (i.e., data points) in the YesiWell data are not strictly independent. Therefore, the simple use of paired t test may not give rigorous conclusions. However, the very small p values under the paired t test can still indicate the significant improvement of our approach over baselines.

#### 4.2 Handwriting digit recognition

To further demonstrate the ability to work on large-scale datasets, we conducted additional experiments on the well-known MNIST dataset (Lecun et al. 1998). The MNIST database of handwritten digits consists of 60,000 training examples, and a test set of 10,000 examples (Lecun et al. 1998). Each example is a  $28 \times 28$  size gray-level image. The MNIST dataset is completely balanced, with 6000 images for each category, with 10 categories in total.

We compare our model with the private stochastic gradient descent algorithm, denoted **pSGD**, from Abadi et al. (2016). The pSGD is the state-of-the-art algorithm in preserving differential privacy in deep learning. pSGD is an advanced version of Shokri and Shmatikov (2015); therefore, there is no need to include the work proposed by Shokri and Shmatikov (2015) in our experiments. The two approaches, i.e., our proposed algorithm and the pSGD, are built on the same structure of a convolutional deep belief network. As in prior work (Abadi et al. 2016), two convolution layers, one with 32 features and one with 64 features, and each hidden neuron which connects with a  $5 \times 5$  unit patch are applied. On top of the convolution layers, there are a fully-connected layer with 1024 units, and a softmax of 10 classes (corresponding to the 10 digits) with cross-entropy loss.

• Figure 5a illustrates the prediction accuracies of each algorithm as a function of the privacy budget  $\epsilon$ . We can see that our model pCDBN outperforms the pSGD in terms of prediction accuracies with small values of the privacy budget  $\epsilon$ , i.e.,  $\epsilon \leq 1.0$ . This is a substantial result, since smaller values of  $\epsilon$  enforce a stronger privacy guarantee of the model. With higher values of  $\epsilon (>1.0)$ , i.e., small injected noise, the two models converge to similar prediction accuracies.

• Figure 5b demonstrates the benefit of being independent of the number of training epochs in consuming the privacy budget of our mechanism. In this experiment,  $\epsilon$  is set



Fig. 5 Accuracy for different noise levels on the MNIST dataset



to 0.5, i.e., large injected noise. The pSGD achieves higher prediction accuracies after using a small number of training epochs, i.e., 88.75% after 18 epochs, compared with the pCDBN. More epochs cannot be used to train the pSGD, since it will violate the privacy protection guarantee. Meanwhile, our model, the pCDBN, can be trained with an unlimited number of epochs. After a certain number of training epochs, i.e., 162 epochs, the pCDBN outperforms the pSGD in terms of prediction accuracy, with 91.71% compared with 88.75%.

Our experimental results clearly show the ability to work with large-scale datasets using our mechanism. In addition, it is significant to be independent of the number of training epochs in consuming privacy budget  $\epsilon$ . Our mechanism is the first of its kind offering this distinctive ability.

• The impact of polynomial degree L Figure 6 shows the prediction accuracies of our model by using different values of L on the MNIST dataset (Lecun et al. 1998). After a certain number of training epochs, it is clear that the impact of L is not significant when L is larger than or equal to 3. In fact, the models with  $L \ge 3$  converge to similar prediction accuracies after 162 training epochs. The difference is notable with small numbers of training epochs. With L larger than 7, the prediction accuracies are very much the same. Therefore we did not show them in Fig. 6. Our observation can be used as a gold standard in selecting L when approximating energy functions based on Chebyshev polynomials.

• Computational performance Given the MNIST dataset, it takes an average of 761 seconds to train our model, after 162 epochs, by using a GPU (NVIDIA GTX TITAN X, 12 GB RAM with 3072 CUDA cores). Meanwhile, training the pSGD is faster than our model, since only a small number of training epochs is needed to train the pSGD. On average, training the pSGD takes 86s, after 18 training epochs. For the YesiWell dataset, training our pCDBN

model takes an average of 2910s, after 600 epochs, compared with 2141s of the dPAH model.

# 5 Conclusions and discussions

In this paper, we propose a novel framework for developing convolutional deep belief networks under differential privacy. Our approach conducts both sensitivity analysis and noise insertion on the energy-based objective functions. Distinctive characteristics offered by our model include: (1) It is totally independent of the number of training epochs in consuming privacy budget; (2) It has the potential to be applied in typical energy-based deep neural networks; (3) Non-linear activation functions, which are continuously differentiable [Stone-Weierstrass Theorem (Rudin 1976)] and satisfy the Riemann-integrable condition, e.g., tanh, arctan, sigmoid, softsign, sinusoid, sinc, Gaussian, etc. (Wikipedia 2016), can be applied; and (4) It has the ability to work with large-scale datasets. With these fundamental abilities, our framework could significantly improve the applicability of differential privacy preservation in deep learning. To illustrate the effectiveness of our framework, we propose a novel model based on our private convolutional deep belief network (pCDBN), for human behavior modeling. Experimental evaluations on a health social network, YesiWell data, and a handwriting digit dataset, MNIST data, validate our theoretical results and the effectiveness of our approach.

In future work, it is worthwhile to study how we might be able to extract private information from deep neural networks. We will also examine potential approaches to preserve differential privacy in more complex deep learning models, such as Long Short-Term Memory (LSTM) (Hochreiter and Schmidhuber 1997). Another open direction is how to adapt our framework to multiparty computational settings, in which multiple parties can jointly train a deep learning model under differential privacy. Innovative multiparty computational protocols for deep learning under differential privacy must have the ability to work with large-scale datasets.

In principle, our mechanism can be applied on rectified linear units (ReLUs) (Glorot et al. 2011) and on parametric rectified linear units (PReLUs) (He et al. 2015). The main difference is that we do not need to approximate the energy function. This is because the energy function is a polynomial function when applying ReLU units. However, we need to add a local response normalization (LRN) layer (Krizhevsky et al. 2012) to bound the values of hidden neurons. This is a common step when dealing with ReLU units. The implementation of this layer and ReLU units under differential privacy is an exciting opportunity for other researchers in future work.

Another challenging problem is identifying the exact risk of re-identification/reconstruction of the data under differential privacy. In Lee and Clifton (2012), the authors proposed differential identifiability to link individual identifiability to  $\epsilon$  differential privacy. However, this is still a non-trivial question. A fancy solution is to design innovative approaches to reconstruct original models from noisy deep neural networks. Then, one could use the original models to infer sensitive information in the training data. However, how to reconstruct the original models from differentially private deep neural networks is an open question. Of course, it is very challenging and will require a significant effort of the whole community to answer.

Acknowledgements This work is supported by the NIH Grant R01GM103309 to the SMASH project. Wu is also supported by NSF Grant 1502273 and 1523115. Dou is also supported by NSF Grant 1118050. We thank Xiao Xiao and Rebeca Sacks for their contributions.

# References

- Abadi, M., Chu, A., Goodfellow, I., McMahan, H. B., Mironov, I., Talwar, K., & Zhang, L. (2016). Deep learning with differential privacy. arXiv:1607.00133.
- Arfken, G. (1985). Mathematical methods for physicists (3rd ed.). Cambridge: Academic Press.
- Armato, A., Fanucci, L., Pioggia, G., & Rossi, D. D. (2009). Low-error approximation of artificial neuron sigmoid function and its derivative. *Electronics Letters*, 45(21), 1082–1084.
- Bach, S., Binder, A., Montavon, G., Klauschen, F., Müller, K. R., & Samek, W. (2015). On pixel-wise explanations for non-linear classifier decisions by layer-wise relevance propagation. *PLoS ONE*, 10(7), e0130,140.
- Bandura, A. (1989). Human agency in social cognitive theory. The American Psychologist, 44(9), 1175.
- Bengio, Y. (2009). Learning deep architectures for AI. Foundation and Trends in Machine Learning, 2(1), 1–127. doi:10.1561/220000006.
- Bengio, Y. (2017). Is cross-validation heavily used in deep learning or is it too expensive to be used? Quora. https://wwwquoracom/Is-cross-validation-heavily-used-in-Deep-Learning-or-is-it-tooexpensive-to-be-used.
- Bengio, Y., Lamblin, P., Popovici, D., Larochelle, H., Montreal, U. D., & Quebec, M. (2007). Greedy layer-wise training of deep networks. In NIPS.
- Brownlee, J. (2015). 8 tactics to combat imbalanced classes in your machine learning dataset. http:// machinelearningmastery.com/tactics-to-combat-imbalanced-classes-in-your-machine-learningdataset/.
- Chan, T. H. H., Li, M., Shi, E., & Xu, W. (2012). Differentially private continual monitoring of heavy hitters from distributed streams. In *PETS'12* (pp. 140–159).
- Chaudhuri, K., & Monteleoni, C. (2008a). Privacy-preserving logistic regression. In NIPS (pp. 289-296).
- Chaudhuri, K., & Monteleoni, C. (2008b). Privacy-preserving logistic regression. In NIPS'08 (pp. 289-296).
- Cheng, Y., Wang, F., Zhang, P., & Hu, J. (2016). Risk prediction with electronic health records: A deep learning approach. In SDM'16.
- Choi, E., Schuetz, A., Stewart, W. F., & Sun, J. (2016). Using recurrent neural network models for early detection of heart failure onset. *Journal of the American Medical Informatics Association*,. doi:10.1093/ jamia/ocw112.
- Cormode, G. (2011). Personal privacy vs population privacy: Learning to attack anonymization. In KDD'11 (pp. 1253–1261).
- Dowlin, N., Gilad-Bachrach, R., Laine, K., Lauter, K., Naehrig, M., & Wernsing, J. (2016). Cryptonets: Applying neural networks to encrypted data with high throughput and accuracy. In *Proceedings of the* 33rd international conference on machine learning, PMLR, proceedings of machine learning research (Vol. 48, pp. 201–210).
- Dwork, C., & Lei, J. (2009). Differential privacy and robust statistics. In STOC'09 (pp. 371-380).
- Dwork, C., McSherry, F., Nissim, K., & Smith, A. (2006). Calibrating noise to sensitivity in private data analysis. *Theory of Cryptography*, 3876, 265–284.
- Erlingsson, U., Pihur, V., & Korolova, A. (2014). Rappor: Randomized aggregatable privacy-preserving ordinal response. In CCS'14 (pp. 1054–1067).
- Fang, R., Pouyanfar, S., Yang, Y., Chen, S. C., & Iyengar, S. S. (2016). Computational health informatics in the big data age: A survey. ACM Computing Surveys, 49(1), 12:1–12:36. doi:10.1145/2932707.
- Glorot, X., Bordes, A., & Bengio, Y. (2011). Deep sparse rectifier neural networks. In Aistats (Vol. 15, p. 275).
- Gottlieb, A., Stein, G. Y., Ruppin, E., Altman, R. B., & Sharan, R. (2013). A method for inferring medical diagnoses from patient similarities. *BMC Medicine*, 11(1), 194. doi:10.1186/1741-7015-11-194.
- Harper, T. (2012). A comparative study of function approximators involving neural networks. Thesis, Master of Science, University of Otago. http://hdl.handle.net/10523/2397.
- Hay, M., Rastogi, V., Miklau, G., & Suciu, D. (2010). Boosting the accuracy of differentially private histograms through consistency. *Proceedings of the VLDB Endowment*, 3(1), 1021–1032.
- He, K., Zhang, X., Ren, S., & Sun, J. (2015). Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. CoRR abs/1502.01852. http://arxiv.org/abs/1502.01852.
- Helmstaedter, M., Briggman, K. L., Turaga, S. C., Jain, V., Seung, H. S., & Denk, W. (2013). Connectomic reconstruction of the inner plexiform layer in the mouse retina. *Nature*, 500(7461), 168–174.
- Hinton, G. (2002). Training products of experts by minimizing contrastive divergence. *Neural Computation*, 14(8), 1771–1800.
- Hinton, G. E., Osindero, S., & Teh, Y. W. (2006). A fast learning algorithm for deep belief nets. *Neural Computation*, 18(7), 1527–1554. doi:10.1162/neco.2006.18.7.1527.
- Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. *Neural Computation*, 9(8), 1735–1780. doi:10.1162/neco.1997.9.8.1735.

- Jain, P., Kothari, P., & Thakurta, A. (2012). Differentially private online learning. In COLT'12 (pp. 24.1–24.34).
- Jamoom, E. W., Yang, N., & Hing, E. (2016). Adoption of certified electronic health record systems and electronic information sharing in physician offices: United states, 2013 and 2014. NCHS Data Brief, 236, 1–8.
- Kifer, D., & Machanavajjhala, A. (2011). No free lunch in data privacy. In SIGMOD'11 (pp. 193-204).
- Krizhevsky, A., Sutskever, I., & Hinton, G. E. (2012). Imagenet classification with deep convolutional neural networks. In Advances in neural information processing systems (pp. 1097–1105).
- LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature*, 521(7553), 436–444. doi:10.1038/ nature14539.
- Lecun, Y., Bottou, L., Bengio, Y., & Haffner, P. (1998). Gradient-based learning applied to document recognition. Proceedings of the IEEE, 86(11), 2278–2324. doi:10.1109/5.726791.
- Lee, J., & Clifton, C. (2012). Differential identifiability. In *The 18th ACM SIGKDD international conference* on knowledge discovery and data mining, KDD '12, Beijing, China, 12–16 August 2012 (pp. 1041–1049).
- Lee, T., & Jeng, J. (1998). The chebyshev-polynomials-based unified model neural networks for function approximation. *IEEE Transactions on Systems, Man, and Cybernetics, Part B*, 28(6), 925–935.
- Lee, H., Grosse, R., Ranganath, R., & Ng, A. Y. (2009). Convolutional deep belief networks for scalable unsupervised learning of hierarchical representations. In *ICML'09* (pp. 609–616).
- Lei, J. (2011). Differentially private m-estimators. In NIPS (pp. 361-369).
- Leung, M. K. K., Xiong, H. Y., Lee, L. J., & Frey, B. J. (2014). Deep learning of the tissue-regulated splicing code. *Bioinformatics*, 30(12), i121–i129. doi:10.1093/bioinformatics/btu277.
- Li, H., Li, X., Ramanathan, M., & Zhang, A. (2015). Prediction and informative risk factor selection of bone diseases. *IEEE/ACM Transactions on Computational Biology and Bioinformatics*, 12(1), 79–91. doi:10. 1109/TCBB.2014.2330579.
- Li, X., Du, N., Li, H., Li, K., Gao, J., & Zhang, A. (2014). A deep learning approach to link prediction in dynamic networks. In SIAM'14 (pp. 289–297).
- Liu, S., Liu, S., Cai, W., Pujol, S., Kikinis, R., & Feng, D. (2014). Early diagnosis of Alzheimer's disease with deep learning. In *IEEE 11th international symposium on biomedical imaging*, *ISBI 2014*, Beijing, China (pp. 1015–1018). doi:10.1109/ISBI.2014.6868045.
- Ma, J., Sheridan, R. P., Liaw, A., Dahl, G. E., & Svetnik, V. (2015). Deep neural nets as a method for quantitative structure-activity relationships. *Journal of Chemical Information and Modeling*, 55(2), 263–274. doi:10. 1021/ci500747n.
- Mason, J., & Handscomb, D. (2002). Chebyshev polynomials. Boca Raton: CRC Press. https://books.google. com/books?id=8FHf0P3to0UC.
- McSherry, F., & Mironov, I. (2009). Differentially private recommender systems. In KDD'09, ACM.
- McSherry, F., & Talwar, K. (2007a). Mechanism design via differential privacy. In 48th annual IEEE symposium on foundations of computer science (FOCS 2007), 20-23 October 2007, Providence, RI, USA, Proceedings (pp. 94–103).
- McSherry, F., & Talwar, K. (2007b). Mechanism design via differential privacy. In FOCS '07 (pp. 94-103).
- Miotto, R., Li, L., Kidd, B. A., & Dudley, J. T. (2016). Deep patient: An unsupervised representation to predict the future of patients from the electronic health records. *Scientific Reports*, 6, 26094. doi:10. 1038/srep26094.
- Nissim, K., Raskhodnikova, S., & Smith, A. (2007). Smooth sensitivity and sampling in private data analysis. In Proceedings of the thirty-ninth annual ACM symposium on theory of computing (pp. 75–84), ACM.
- Ortiz, A., Munilla, J., Grriz, J. M., & Ramrez, J. (2016). Ensembles of deep learning architectures for the early diagnosis of the alzheimers disease. *International Journal of Neural Systems*, 26(07), 1650,025. doi:10. 1142/S0129065716500258.
- Perotte, A., Pivovarov, R., Natarajan, K., Weiskopf, N., Wood, F., & Elhadad, N. (2014). Diagnosis code assignment: models and evaluation metrics. *Journal of the American Medical Informatics Association*, 21(2), 231–237. doi:10.1136/amiajnl-2013-002159.
- Perotte, A., Ranganath, R., Hirsch, J. S., Blei, D., & Elhadad, N. (2015). Risk prediction for chronic kidney disease progression using heterogeneous electronic health record data and time series analysis. *Journal* of the American Medical Informatics Association, 22(4), 872–880. doi:10.1093/jamia/ocv024.
- Phan, N., Dou, D., Piniewski, B., & Kil, D. (2015a). Social restricted boltzmann machine: Human behavior prediction in health social networks. In ASONAM'15 (pp. 424–431).
- Phan, N., Dou, D., Wang, H., Kil, D., & Piniewski, B. (2015b). Ontology-based deep learning for human behavior prediction in health social networks. In *Proceedings of the 6th ACM conference on bioinformatics*, *computational biology and health informatics* (pp. 433–442). doi:10.1145/2808719.2808764.
- Phan, N., Dou, D., Piniewski, B., & Kil, D. (2016a). A deep learning approach for human behavior prediction with explanations in health social networks: social restricted boltzmann machine (SRBM+). Social Network Analysis and Mining, 6(1), 79:1–79:14. doi:10.1007/s13278-016-0379-0.

- Phan, N., Ebrahimi, J., Kil, D., Piniewski, B., & Dou, D. (2016b). Topic-aware physical activity propagation in a health social network. *IEEE Intelligent Systems*, 31(1), 5–14.
- Phan, N., Wang, Y., Wu, X., & Dou, D. (2016c). Differential privacy preservation for deep auto-encoders: An application of human behavior prediction. In AAAI'16 (pp. 1309–1316).
- Plis, S. M., Hjelm, D. R., Salakhutdinov, R., Allen, E. A., Bockholt, H. J., Long, J. D., et al. (2014). Deep learning for neuroimaging: A validation study. *Frontiers in Neuroscience*, 8, 229. doi:10.3389/fnins. 2014.00229.
- Reed, S. E., Lee, H., Anguelov, D., Szegedy, C., Erhan, D., & Rabinovich, A. (2014). Training deep neural networks on noisy labels with bootstrapping. CoRR abs/1412.6596.
- Rivlin, T. J. (1990). Chebyshev polynomials form approximation theory to algebra and number theory (2nd ed.). New York: Wiley.
- Roumia, M., & Steinhubl, S. (2014). Improving cardiovascular outcomes using electronic health records. *Current Cardiology Reports*, 16(2), 451. doi:10.1007/s11886-013-0451-6.
- Rudin, W. (1976). Principles of mathematical analysis. New York: McGraw-Hill.
- Shokri, R., & Shmatikov, V. (2015). Privacy-preserving deep learning. In CCS'15 (pp. 1310–1321).
- Smolensky, P. (1986). Information processing in dynamical systems: Foundations of harmony theory. In Parallel distributed processing: Explorations in the microstructure of cognition (Vol. 1, pp. 194–281).
- Song, S., Chaudhuri, K., & Sarwate, A. D. (2013). Stochastic gradient descent with differentially private updates. In *GlobalSIP* (pp. 245–248).
- Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., & Salakhutdinov, R. (2014). Dropout: A simple way to prevent neural networks from overfitting. *Journal of Machine Learning Research*, 15, 1929–1958. http:// jmlr.org/papers/v15/srivastava14a.html.
- U.S. Department of Health and Human Services. (2016a). Health information technology for economic and clinical health (hitech) act. https://www.hhs.gov/hipaa/for-professionals/special-topics/HITECH-actenforcement-interim-final-rule/.
- U.S. Department of Health and Human Services. (2016b). Health insurance portability and accountability act of 1996. http://www.hhs.gov/hipaa/.
- Vlcek, M. (2012). Chebyshev polynomial approximation for activation sigmoid function. *Neural Network World*, 4, 387–393.
- Wang, Y., Wu, X., & Wu, L. (2013). Differential privacy preserving spectral graph analysis. In PAKDD (2) (pp. 329–340).
- Wikipedia. (2016). Activation function. https://en.wikipedia.org/wiki/Activation\_function.
- Wu, J., Roy, J., & Stewart, W. F. (2010). Prediction modeling using EHR data: Challenges, strategies, and a comparison of machine learning approaches. *Medical Care*, 48(6 Suppl), S106–S113. doi:10.1097/mlr. 0b013e3181de9e17.
- Xiao, X., Wang, G., & Gehrke, J. (2010). Differential privacy via wavelet transforms. In *ICDE'10* (pp. 225–236).
- Xiong, H. Y., Alipanahi, B., Lee, L. J., Bretschneider, H., Merico, D., Yuen, R. K. C., et al. (2015). The human splicing code reveals new insights into the genetic determinants of disease. *Science*, 347(6218), 1254806. doi:10.1126/science.1254806.
- Zhang, J., Zhang, Z., Xiao, X., Yang, Y., & Winslett, M. (2012). Functional mechanism: Regression analysis under differential privacy. PVLDB, 5(11), 1364–1375.