

Singapore Management University

Institutional Knowledge at Singapore Management University

Research Collection School Of Computing and Information Systems

School of Computing and Information Systems

5-2022

Deep depression prediction on longitudinal data via joint anomaly ranking and classification

Guansong PANG

Singapore Management University, gspang@smu.edu.sg

Ngoc Thien Anh PHAM

Emma BAKER

Rebecca BENTLEY

Anton VAN DEN HENGEL

Follow this and additional works at: https://ink.library.smu.edu.sg/sis_research



Part of the [Databases and Information Systems Commons](#), and the [Longitudinal Data Analysis and Time Series Commons](#)

Citation

PANG, Guansong; PHAM, Ngoc Thien Anh; BAKER, Emma; BENTLEY, Rebecca; and VAN DEN HENGEL, Anton. Deep depression prediction on longitudinal data via joint anomaly ranking and classification. (2022). *Advances in Knowledge Discovery and Data Mining: 26th Pacific-Asia Conference PAKDD 2022, Chengdu, China, May 16-19: Proceedings*. 13281, 236-248.

Available at: https://ink.library.smu.edu.sg/sis_research/7544

This Conference Proceeding Article is brought to you for free and open access by the School of Computing and Information Systems at Institutional Knowledge at Singapore Management University. It has been accepted for inclusion in Research Collection School Of Computing and Information Systems by an authorized administrator of Institutional Knowledge at Singapore Management University. For more information, please email cherylids@smu.edu.sg.



Deep Depression Prediction on Longitudinal Data via Joint Anomaly Ranking and Classification

Guansong Pang¹(✉), Ngoc Thien Anh Pham², Emma Baker²,
Rebecca Bentley³, and Anton van den Hengel²

¹ Singapore Management University, Singapore, Singapore

gspang@smu.edu.sg

² University of Adelaide, Adelaide, Australia

{ngoc.t.pham,emma.baker,anton.vandenhengel}@adelaide.edu.au

³ University of Melbourne, Parkville, Australia

brj@unimelb.edu.au

Abstract. A wide variety of methods have been developed for identifying depression, but they focus primarily on measuring the degree to which individuals are suffering from depression currently. In this work we explore the possibility of predicting future depression using machine learning applied to longitudinal socio-demographic data. In doing so we show that data such as housing status, and the details of the family environment, can provide cues for predicting future psychiatric disorders. To this end, we introduce a novel deep multi-task recurrent neural network to learn time-dependent depression cues. The depression prediction task is jointly optimized with two auxiliary anomaly ranking tasks, including contrastive one-class feature ranking and deviation ranking. The auxiliary tasks address two key challenges of the problem: 1) *the high within class variance of depression samples*: they enable the learning of representations that are robust to highly variant in-class distribution of the depression samples; and 2) *the small labeled data volume*: they significantly enhance the sample efficiency of the prediction model, which reduces the reliance on large depression-labeled datasets that are difficult to collect in practice. Extensive empirical results on large-scale child depression data show that our model is sample-efficient and can accurately predict depression 2–4 years before the illness occurs, substantially outperforming eight representative comparators.

Keywords: Depression prediction · Anomaly detection · One-class classification · Deep learning

1 Introduction

Major Depressive Disorder (MDD), widely known as depression, is a mental disorder characterized by a severe and persistent feeling of sadness, loss of interest

in activities, or a sense of despair, causing significant impairment in daily life [9]. Globally over 300 million people of all ages are estimated to suffer from depression, and it is a major contributor to nearly 800 thousands suicide deaths per year [23]. There have been many studies [2,13,19,22] demonstrating effective automated diagnosis of depression using machine learning techniques. These studies focus on the detection of ongoing, current depression, using data that relates information about the current state of an individual. We focus here on the more challenging task of predicting depression in advance. This is achieved using longitudinal socio-demographic data. The motivation for this approach is that prediction far enough in advance of upcoming depression might enable early intervention before the condition arises.

The longitudinal socio-demographic data used contains a variety of non-medical information such as education level, socio-economic background, family environment, measured at intervals overtime. There are three major challenges in exploiting these cues for depression prediction. First, this is noisy high-dimensional temporal data that contains thousands of numeric and categorical features, among which only very selective set are relevant to depression prediction. Second, there is no single set of socio-demographic factors that cause depression, leading to high in-class variance for the depression samples. For example, for childhood depression, the root cause may be related to the mental health status of parents, dwelling conditions, or a medical condition. As a result, the depression cases are highly dissimilar and exhibit no single underlying mechanism, or common characteristics. A model that seeks a single common explanation for all cases cannot succeed. Lastly, in practice, only a small amount of labeled data is available, as it is difficult, if not impossible, to collect large depression samples. The available data thus often fails to cover the diverse types of depression cases.

In this work we introduce a novel multi-task learning approach to tackle these challenges, in which a depression classification task is jointly optimized with two auxiliary anomaly ranking tasks, contrastive one-class feature ranking and deviation ranking. Depression samples are treated as anomalous samples in our auxiliary tasks, because the cause-varying depression samples can be widely distributed, and as a result, are difficult to model as a single concrete class. Directly modeling these depression samples with classification models can easily overfit the given depression cases. The two anomaly ranking tasks are devised to enforce compact low-dimensional representations of normal samples and allow variations in the representations of depression samples, presenting effective inductive biases to regularize the classification models. This significantly improves the model’s generalization beyond the individual cases presented in the small volume of labeled data available.

In summary, this work makes two key contributions:

- We introduce a novel multi-task learning framework, which harnesses auxiliary anomaly detection tasks to empower the greater classification task. To the best of our knowledge, this is the first multi-task learning approach that

- jointly optimizes classification and anomaly ranking tasks, which is an important tool for application problems that are similar to depression prediction.
- We further instantiate the framework into a multi-task recurrent neural network model, termed MTNet, which optimizes the depression model with one-class constraints on its feature space and deviation constraints on its output layer. Extensive empirical results on large-scale child depression data show that MTNet can accurately predict depression two to four years before the illness occurs (*e.g.*, achieving a recall of 0.8), substantially outperforming eight competing methods. Additionally, MTNet can also outperform these competing methods even with largely reduced (50% less) training data.

2 Related Work

Longitudinal Studies. There have been many longitudinal studies of depression [6, 8], but they focus on association discovery that identifies factors or predictors associated with depression using traditional bivariate/multivariate statistic models. By contrast, our study is on learning prediction models to predict the occurrence/risk of future depression.

Automated Depression Diagnosis. Current studies focus on the detection of depression using classification models on vocal/visual data taken during clinical interviews, with vocal features like prosodic and cepstral features [2, 13, 24] and visual features like facial expression, gaze direction, and eye movement [13, 22, 26]. Depression detection based on social media data using linguistic and network features is also extensively studied [11, 19]. However, all these studies focus on the detection of ongoing depression, while we aim at predicting upcoming depression.

An exploratory task is introduced at the eRisk workshop [10] to facilitate the development and evaluation of models for early detection of signs of depression using social media data, resulting in a number of early depression detection models [1, 12, 18]. Like the aforementioned studies, they are also focused upon detecting *expression of depression cues/symptoms*, while we learn the *hidden root cause factors in socio-demographic features* which are more difficult to learn.

Multi-task Learning. Multi-task learning [16] has been successful in a range of applications. There have been many approaches introduced to jointly learn multiple related tasks (*e.g.*, classification, regression, and clustering) to improve the performance on small labeled data, such as feature learning, task relation learning, task clustering, low-rank and decomposition approaches [25]. Our work uses multi-task feature learning methods with hard parameter sharing as in [4, 20, 21], but our approach is the first work that uses anomaly ranking at the feature and output layers to regularize the classification.

Anomaly Detection. Anomaly detection techniques have been successfully applied to detect abnormal events/behaviors in many applications [14], but they are rarely used to regularize supervised learning models in multi-task learning as the two paradigms have rather different learning objectives. We show in this work that recent advanced anomaly detection models [15, 17] can be adapted to effectively regularize classification models and improve their sample efficiency.

3 Multi-task Recurrent Neural Networks

3.1 The Proposed Framework

Upcoming depression prediction aims to learn a binary depression classification mapping function $\phi : \mathcal{X} \rightarrow \mathcal{Y}$, where $\mathcal{X} = \{\mathbf{X}_1, \mathbf{X}_2, \dots, \mathbf{X}_N\}$ is a set of longitudinal data of N samples; each sample $\mathbf{X} \in \mathbb{R}^{w \times D}$ is a matrix input of an individual subject, which contains the socio-demographic features of the subject in the recent w waves (or time steps) of questionnaire data, *i.e.*, $\mathbf{X} = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_w\}$, where $\mathbf{x}_t \in \mathbb{R}^D$ is a feature vector derived from the t -th wave of questionnaire data; $\mathcal{Y} = \{0, 1\}$ is the output space, with ‘1’ indicating the subject being normal within all the recent w waves of questionnaire but having depression in the next waves of questionnaire, and with ‘0’ indicating the subject being normal in the recent and future waves of questionnaire.

In this work we introduce a novel multi-task learning framework to tackle the problem. An overview of the approach is presented in Fig. 1. Depression classification is our primary task and is jointly optimized with two auxiliary tasks, including deviation score ranking and contrastive one-class feature learning. The auxiliary tasks treat depression samples as anomalies and enforce compact feature representations of normal samples and allow some variations in the representations of depression samples, serving as a regularizer of the classification model. This results in better generalized classification models than that in the single primary task. Formally, let $\tau : \mathcal{X} \rightarrow \mathbb{R}$ be an anomaly ranking function that assigns an anomaly score to each subject; $\psi : \mathcal{X} \rightarrow \mathcal{Q}$ be the one-class feature learning function, where $\mathcal{Q} \in \mathbb{R}^M$ with $M \ll D$ is a new feature space, then our overall objective function can be given as follows.

$$\arg \min_{\Theta_e, \Theta_a, \Theta_o} \sum_{i=1}^N \left[\ell_e(\phi(\mathbf{X}_i; \Theta_e), y_i) + \alpha \ell_a(\tau(\mathbf{X}_i; \Theta_a), y_i) + \beta \ell_o(\psi(\mathbf{X}_i; \Theta_o), y_i) \right], \quad (1)$$

where ℓ_e , ℓ_a and ℓ_o are respective loss functions for depression classification, deviation ranking and contrastive one-class metric learning, y_i is the class label of \mathbf{X}_i , $\Theta = \{\Theta_e, \Theta_a, \Theta_o\}$ is the set of parameters to be learned, α and β are hyperparameters to control the importance of the two auxiliary tasks.

We instantiate the framework into a model called MTNet that leverages a shared LSTM neural network [7] to learn critical temporal changes in longitudinal data for depression classification. Supervised anomaly deviation and one-class support vector data description loss functions are defined to improve the model’s generalization. A simple data augmentation method is also introduced to further enhance the generalizability. The modules of MTNet are presented as follows.

3.2 Primary Task: Depression Classification

Our classification model leverages an LSTM neural network layer to learn important temporal dependency in the longitudinal data. An LSTM layer consists of

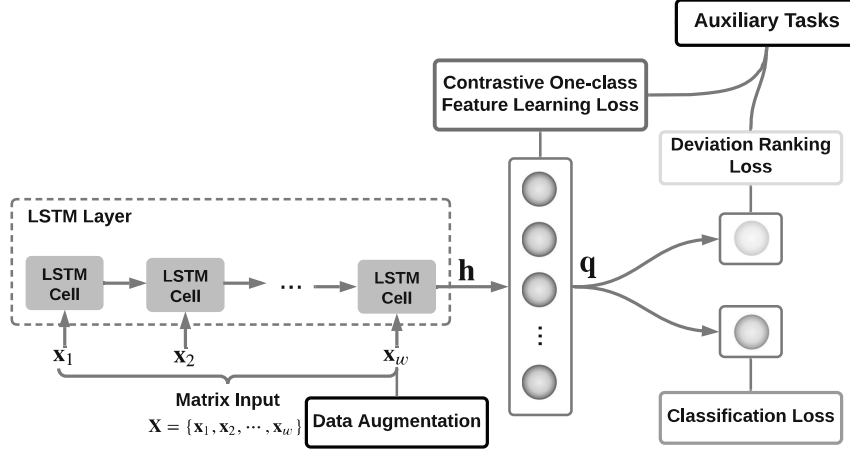


Fig. 1. The proposed multi-task learning framework.

multiple LSTM cells, with each LSTM cell learning temporal-dependent representations of the input data at a specific wave. The full LSTM layer uses the recurrent LSTM cells to encode important temporal changes across all different questionnaire waves into the output vector $\mathbf{h} \in \mathbb{R}^L$ in the last LSTM cell, *i.e.*, the LSTM layer is a mapping function η that performs $\mathbf{h} = \eta(\mathbf{X}; \Theta_l)$, where Θ_l contains all weight parameters in a standard LSTM (see *Supplementary Material*¹ for the full details of LSTM). To learn more expressive representations, a FC layer is used to project \mathbf{h} onto a lower-dimensional feature representation space:

$$\mathbf{q} = \psi(\mathbf{X}; \Theta_l, \mathbf{W}_s, \mathbf{b}_s) = g_s(\mathbf{W}_s \eta(\mathbf{X}; \Theta_l) + \mathbf{b}_s), \quad (2)$$

where $\mathbf{W}_s \in \mathbb{R}^{M \times L}$ and $\mathbf{b}_s \in \mathbb{R}^M$ are the learnable parameters, g_s is an activation function, and $\mathbf{q} \in \mathcal{Q}$ is the final feature representation of \mathbf{X} . We then train a classifier on the \mathcal{Q} representation space with a standard binary cross-entropy loss function:

$$\ell_e(\mathbf{X}, y) = -(y \log(p) + (1 - y) \log(1 - p)), \quad (3)$$

where y is the class label of \mathbf{X} and $p = \phi(\mathbf{X}; \Theta_e) = g_e(\mathbf{W}_e \mathbf{q} + b_e)$, where g_e is a sigmoid activation, $\mathbf{W}_e \in \mathbb{R}^{1 \times M}$ and $b_e \in \mathbb{R}$. $\Theta_e = \{\Theta_l, \mathbf{W}_s, \mathbf{b}_s, \mathbf{W}_e, b_e\}$ contains all the network parameters that can be learned in an end-to-end manner.

3.3 Auxiliary Tasks

Two auxiliary tasks, contrastive one-class-based feature ranking and deviation ranking, are incorporated to introduce an inductive bias (*i.e.*, to learn compact normal representations and allow large variations in abnormal representations) to learn more generalized representations of depression samples.

Deviation Ranking. A partially-supervised anomaly ranking task is introduced to enforce the model to assign significantly larger anomaly scores for depression

¹ Supplementary material is available at <https://tinyurl.com/MTNetPAKDD22>.

samples than that of non-depression samples. Inspired by [15], a prior-driven anomaly ranking loss function, called deviation loss, is used to fulfill this goal. Particularly, a Gaussian prior $\mathcal{N}(\mu, \sigma^2)$ is imposed on the anomaly scores of all samples, which posits that the anomaly scores of non-depression samples are centered around a Gaussian mean value μ while the anomaly scores of depression samples have at least $a * \sigma$ deviations from μ . Formally, we add another network output head with one linear unit to learn an anomaly score for each sample:

$$\tau(\mathbf{X}; \Theta_a) = \mathbf{W}_a \psi(\mathbf{X}; \Theta_r) + b_a, \quad (4)$$

where $\Theta_a = \{\Theta_r, \mathbf{W}_a, b_a\}$ are the parameters to be learned. We then define the deviation using the well-known Z-Score:

$$dev(\mathbf{X}) = \frac{\tau(\mathbf{X}; \Theta_a) - \mu}{\sigma}, \quad (5)$$

The deviation function is then used to define our anomaly ranking loss function:

$$\ell_a(\mathbf{X}, y) = (1 - y)|dev(\mathbf{X})| + y \max(0, a - dev(\mathbf{X})). \quad (6)$$

By minimizing ℓ_a , our model pushes the anomaly scores of normal samples as close as possible to μ while enforcing at least $a * \sigma$ between μ and the anomaly scores of depression samples in the upper tail of the Gaussian distribution. Following [15], the prior $\mathcal{N}(0, 1)$ is used with $a = 5$ to guarantee significant deviations of depression samples from normal samples.

Contrastive One-class Feature Learning. Unlike the anomaly ranking task that introduces the inductive bias using an output layer independent from the classification output, the one-class feature learning complements the deviation score learning and exerts directly on the feature layer. Particularly, we introduce a contrastive one-class feature learning method, in which we devise a supervised variant of support vector data description (SVDD) [17] by contrasting the one-class center of the normal samples and the depression samples.

$$\ell_o(\mathbf{X}, y) = (1 - y)\|\psi(\mathbf{X}; \Theta_r) - \mathbf{n}\|_2 + y \max(0, m - \|\psi(\mathbf{X}; \Theta_r) - \mathbf{n}\|_2), \quad (7)$$

where $\Theta_o = \{\Theta_r\}$, $\mathbf{n} \in \mathbb{R}^M$ is the one-class center vector of normal samples and m is a hyperparameter to control the contrast margin. The first term is the original SVDD objective. The second term is added to enforce a large margin between non-depression and depression samples in the ψ -induced representation space, while minimizing the \mathbf{n} -centered hypersphere's volume. We found empirically that MTNet can perform well with varying settings of \mathbf{n} , *e.g.*, $\mathbf{n} \sim \mathcal{N}(0, 1)$ or $\mathbf{n} \sim \mathcal{U}(0, 1)$. We use $\mathbf{n} \sim \mathcal{N}(0, 1)$ by default, *i.e.*, generating \mathbf{n} by randomly drawing a vector from a standard Gaussian distribution. $m = 1$ is used to enforce a sufficiently large distance margin in the feature representation space.

3.4 Data Augmentation

A simple data augmentation method is introduced to augment depression samples and further enhance the model's generalizability. Specifically, a pair of

depression samples are randomly selected, and then a small percentage of randomly selected values in the last wave data of one sample are replaced with the corresponding values in another sample to create a new depression sample. The augmented sample can well retain the original depression-relevant information while at the same time enriching the depression samples. By using this method, we increase the number of depression samples in the training data by a factor of 10. In our experiment, we randomly replaced 5% of the feature values by default.

3.5 The Algorithm of MTNet

The algorithmic procedure of our model MTNet is presented in Algorithm 1. After random initialization of the network parameters in Step 1, stochastic gradient descent is used to optimize the model in Steps 2–8. In Step 4, as the number of depression samples is typically far smaller than that of non-depression samples, we generate sample batches with balanced class distribution to achieve more effective optimization. This shares the same spirit as oversampling in imbalanced learning [5]. Step 5 calculates the batch-wise loss for the three tasks. Step 6 performs gradient descent steps to learn the parameters Θ . Note that Θ_r are shared parameters in $\{\Theta_e, \Theta_a, \Theta_o\}$, and thus, the feature representations in MTNet are jointly optimized by all three tasks. At the inference stage, only the classification function ϕ is used to produce the class label.

Algorithm 1. MTNet

Input: $\mathcal{X} \in \mathbb{R}^{w \times D}$ - training samples, and binary class labels \mathbf{y}

Output: $\phi : \mathcal{X} \rightarrow \mathcal{Y}$ - a depression classification network

1: Randomly initialize $\Theta = \{\Theta_e, \Theta_a, \Theta_o\}$

2: **for** $j = 1$ to $\#epochs$ **do**

3: **for** $k = 1$ to $\#batches$ **do**

4: $\mathcal{B} \leftarrow$ Randomly sample the same number of depression and non-depression samples

5: Calculate the loss using $\frac{1}{|\mathcal{B}|} \sum_{\mathbf{X}_i \in \mathcal{B}} \left[\ell_e(\phi(\mathbf{X}_i; \Theta_e), y_i) + \alpha \ell_a(\tau(\mathbf{X}_i; \Theta_a), y_i) + \beta \ell_o(\psi(\mathbf{X}_i; \Theta_o), y_i) \right]$

6: Perform a gradient descent step w.r.t. the parameters in Θ

7: **end for**

8: **end for**

9: **return** ϕ

4 Experiments

4.1 Datasets

Our model is evaluated on a large child depression data dataset based on the Longitudinal Study of Australian Children (LSAC) data [3]. LSAC consists of

multiple bi-annual waves of questionnaire-based interview data of 10,090 children across Australia. Initially, children aged from infant to 5 years and their families are interviewed between August 2003 and February 2004. This routine is repeated every two years afterwards. At the time of writing seven waves of data are available. LSAC provides a dataset of 4,983 children aged 4 to 5 years in the first wave of questionnaire. These children are all healthy until 287 children are confirmed to have depression at the 6/7-th wave of interview in the years 2013 to 2015. Depression is measured using parental self-report data. Particularly, in the waves 6 and 7, the primary caregiver is asked “does study child have any of these ongoing (depression) conditions²?”; the child is confirmed to have depression if the answer is ‘yes’. After a simple feature screening to remove uninformative features (*e.g.*, features with very large percentage of missing values), 210 social-demographic features related to individual growth and development (*e.g.*, age, gender, living location, schooling performance), and family environment (*e.g.*, social, educational, economic, employment, household income, housing conditions) are used (see *Supplementary Material* for the list of the interview questions and a sample of the questionnaire). In the selected features, missing values are filled with the mean/mode value in each feature; categorical features are then converted into numeric features by using one-hot encoding. The resulting dataset contains 762 features in each wave of data. Thus, the dataset used has 4,983 samples, with each sample represented by a 7×762 matrix. We further perform a stratified random split of the dataset into three subsets, including 60% data as a training set, and respective 20% data for validation and testing sets.

4.2 Experimental Setup

We evaluate the performance of predicting the possible occurrence of depression in the near future. To this end, the model is trained and tested using only the first five waves of data when all children are reported to be mentally healthy (*i.e.*, absence of depressive symptoms). The task is to predict whether a child will have depression at the upcoming wave 6 or 7. MTNet is compared with eight temporal and non-temporal methods.

- **Non-temporal Methods.** Three popular classification methods, including logistic regression (LR), support vector machines (SVM) and multi-layer perceptron (MLP) neural networks, are used as baseline methods that are not designed to capture temporal dependence. They are two main ways to apply these methods to the longitudinal data. One way is to build the classification model using the most recent single (*i.e.*, 5-th) wave data only. The second way is to use the data from all the five waves, in which for each subject we concatenate the feature vectors derived from all the waves into one lengthy unified feature vector; the classifiers are then built upon this concatenated data. This way helps capture some temporal-dependent changes. All three

² ‘Ongoing conditions’ means that the conditions “*exist for some period of time (weeks, months or years) or re-occur regularly. They do not have to be diagnosed by a doctor*”.

methods are evaluated in both ways, with LR/SVM/MLP-s denoting the classifier using the single wave data and LR/SVM/MLP-m denoting the use of the concatenated multi-wave data.

- **Temporal Methods.** The LSTM-based deep classifier, is used as a competing temporal method. The standard binary cross-entropy loss function is used to train the model. We also compare MTNet with the state-of-the-art anomaly detection model DevNet [15] that is adapted to temporal data with LSTM network.

4.3 Implementation Details

MTNet is implemented with one LSTM layer with 200 units, followed by a fully-connected (FC) layer with 20 units and a classification output layer. The sigmoid function is used in g_r in the LSTM layer by default; the widely-used ReLU activation function is used in g_s in the FC layer. A dropout layer with a dropout rate of 0.5 is applied to the LSTM and FC layers. The competing methods LSTM and DevNet use exactly the same network architecture as MTNet. MLP uses a similar network structure with two hidden layers of respectively 200 and 20 units, with each layer having a dropout rate of 0.5. MTNet, DevNet, LSTM, and MLP are implemented using Keras³ and optimized using RMSprop with a batch size of 256 and 20 batches per epoch. They are trained with 30 epochs as their performance can converge early. $\alpha = 0.5$ and $\beta = 2.0$ are used in MTNet by default. LR and SVM are taken from the open-source scikit-learn package⁴. Due to a large percentage of irrelevant features presented in the data, our extensive results showed that applying the l_1 -norm regularizer to MLP, LR and SVM obtains significantly better performance than the l_2 -norm regularizer. Thus, the l_1 -norm regularizer is applied to these three classifiers to bring sparsity to the model. The regularization hyperparameter is probed with $\{0.001, 0.01, 0.1, 1\}$, with the best performance reported. The oversampling method in MTNet is used in all competing methods to alleviate the class-imbalanced problem.

4.4 Performance Evaluation Measures

Three widely-used evaluation measures are used, including the Area Under Receiver Operating Characteristic Curve (AUC-ROC), Area Under Precision-Recall Curve (AUC-PR), and F₁-score (F-score for brevity). AUC-ROC summarizes the ROC curve of true positives against false positives, while AUC-PR summarizes the curve of precision against recall. AUC-ROC is popular due to its good interpretability. AUC-PR is more indicative than AUC-ROC in evaluating performance on imbalanced data. F-score is the harmonic mean of precision and recall. We also report the precision and recall results to gain more insights into the performance. The reported results are averaged over five independent runs.

³ <https://keras.io/>.

⁴ <https://scikit-learn.org/>.

4.5 Empirical Results

Effectiveness on Real-World Data. The results of depression prediction are shown in Table 1. Our model MTNet is the best performer in AUC-ROC, AUC-PR and F-score. MTNet substantially outperforms all of its competing methods by 2%–38% in AUC-PR and 3%–20% in F-score. Impressively, MTNet obtains a recall of 0.8, achieving at least 7.8% improvement over its contenders. Given a sufficiently high precision of 0.7, the high recall rate in MTNet would enable accurate intervention and pre-treatment of up to 80% upcoming depression cases at a very early stage (2–4 years before the depression occurs), effectively preventing and reducing the depression cases.

Table 1. Performance results (mean \pm std) of depression prediction.

	AUC-ROC	AUC-PR	F-score	Precision	Recall
LR-s	0.648 \pm 0.020	0.595 \pm 0.018	0.632 \pm 0.021	0.659 \pm 0.032	0.611 \pm 0.037
LR-m	0.648 \pm 0.019	0.596 \pm 0.016	0.620 \pm 0.023	0.670 \pm 0.027	0.579 \pm 0.038
SVM-s	0.666 \pm 0.012	0.610 \pm 0.012	0.646 \pm 0.005	0.682 \pm 0.023	0.614 \pm 0.011
SVM-m	0.684 \pm 0.013	0.631 \pm 0.010	0.646 \pm 0.026	0.729 \pm 0.013	0.582 \pm 0.045
MLP-s	0.771 \pm 0.007	0.780 \pm 0.010	0.706 \pm 0.023	0.679 \pm 0.032	0.742 \pm 0.071
MLP-m	0.814 \pm 0.009	0.808 \pm 0.019	0.718 \pm 0.022	0.730 \pm 0.049	0.718 \pm 0.082
LSTM	0.779 \pm 0.012	0.775 \pm 0.010	0.662 \pm 0.034	0.751 \pm 0.031	0.593 \pm 0.039
DevNet	0.785 \pm 0.013	0.786 \pm 0.017	0.688 \pm 0.035	0.704 \pm 0.044	0.684 \pm 0.094
MTNet	0.818 \pm 0.009	0.823 \pm 0.008	0.743 \pm 0.037	0.697 \pm 0.006	0.800 \pm 0.080

Sample Efficiency. This section examines the model’s generalizability from the sample efficiency aspect, *i.e.*, how is the performance if less labeled training data is available? Specifically, each model is trained on a new training set that is a random subset of the original training data, and then it is evaluated using the same test data as that used in Table 1. We focus on comparing to the better contenders LR/SVM/MLP-m and omit the less effective ones – LR/SVM/MLP-s. The results are shown in Fig. 2. Remarkably, MTNet is substantially more sample-efficient than its competing methods; it performs much better than, or comparably well to, the best performance of its competing methods even when it uses 50% less training data. This superiority of MTNet benefits from the integrated anomaly ranking tasks in its multi-task objective function, which enable better generalization to diverse depression cases. This is manifested by the large performance gap between MTNet and LSTM, since the only difference between them is the two auxiliary tasks integrated into MTNet.

Parameter Sensitivity. This section evaluates the sensitivity of MTNet w.r.t. its two hyperparameters, α and β , which respectively adjust the importance of the anomaly ranking loss and one-class metric loss. The test results are presented in Fig. 3. The results show that MTNet generally performs stably with both α

and β in a wide range of setting choices. Relatively small α and large β are needed for MTNet to achieve the best performance. This indicates that MTNet is dependent more on the one-class feature learning than the deviation ranking.




Fig. 2. Sample efficiency




Fig. 3. Parameter sensitivity

Ablation Study. This section evaluates the importance of each module in MTNet. LSTM is used as a baseline to evaluate the effect of incorporating one or more of the following modules, including the deviation ranking loss ℓ_a , the one-class metric learning loss ℓ_o , and the data augmentation (DA). The results are reported in Table 2. It is clear that i) the multi-task learning performs substantially better than the individual tasks, ii) all three modules in MTNet make important contribution to its overall performance, and iii) the deviation ranking loss ℓ_a and the one-class metric learning loss ℓ_o are complementary to each other.

Table 2. Ablation study results. DA is short for data augmentation.

Method	AUC-ROC	AUC-PR	F-score	Precision	Recall
LSTM	0.779	0.775	0.662	0.751	0.593
LSTM+ ℓ_a	0.785	0.786	0.688	0.704	0.684
LSTM+ ℓ_o	0.804	0.821	0.702	0.673	0.737
LSTM+ ℓ_a + ℓ_o	0.817	0.834	0.721	0.709	0.737
LSTM+ ℓ_a + ℓ_o +DA	0.818	0.823	0.743	0.697	0.800

5 Conclusions and Future Work

In this work we propose a novel multi-task learning framework and its instantiation MTNet for depression prediction. MTNet effectively leverages the two

auxiliary anomaly ranking tasks to improve the depression prediction model's representations and generalizability. Remarkably, our empirical results show that MTNet is able to accurately predict 80% depression cases 2–4 years before the depression actually occurs in children. The improved generalizability of MTNet is supported by the sample efficiency experiment, in which MTNet requires significantly less labeled depression samples to perform comparably well to, or substantially better than, the competing methods. It should be noted the model has around 70% precision only, indicating 30% false positive predictions. Thus, caution must be taken when using the the model in practice. In future work, to improve the model's accountability and its collaboration with clinical psychologists, we plan to incorporate an interpretation module into our model to provide insightful explanation for each of its depression prediction result.

References

1. Burdisso, S.G., Errecalde, M., Montes-y Gómez, M.: τ -ss3: a text classifier with dynamic n-grams for early risk detection over text streams. *Pattern Recogn. Lett.* **138**, 130–137 (2020)
2. Gong, Y., Poellabauer, C.: Topic modeling based multi-modal depression detection. In: AVEC, pp. 69–76 (2017)
3. Gray, M., Sanson, A., et al.: Growing up in Australia: the longitudinal study of australian children. *Family Matters* **72**, 4 (2005)
4. Hassani, K., Haley, M.: Unsupervised multi-task feature learning on point clouds. In: ICCV, pp. 8160–8171 (2019)
5. He, H., Garcia, E.A.: Learning from imbalanced data. *IEEE Trans. Knowl. Data Eng.* **21**(9), 1263–1284 (2009)
6. Henry, M., et al.: A screening algorithm for early detection of major depressive disorder in head and neck cancer patients post-treatment: longitudinal study. *Psycho-oncology* **27**(6), 1622–1628 (2018)
7. Hochreiter, S., Schmidhuber, J.: Long short-term memory. *Neural Comput.* **9**(8), 1735–1780 (1997)
8. Korsten, L.H., et al.: Factors associated with depression over time in head and neck cancer patients: a systematic review. *Psycho-oncology* **28**(6), 1159–1183 (2019)
9. Lamers, F., Milaneschi, Y., Smit, J.H., Schoevers, R.A., Wittenberg, G., Penninx, B.W.: Longitudinal association between depression and inflammatory markers: results from the Netherlands study of depression and anxiety. *Biol. Psychiatry* **85**(10), 829–837 (2019)
10. Losada, D.E., Crestani, F., Parapar, J.: Overview of eRisk: early risk prediction on the internet. In: Bellot, P., et al. (eds.) CLEF 2018. LNCS, vol. 11018, pp. 343–361. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-98932-7_30
11. Mann, P., Paes, A., Matsushima, E.H.: See and read: detecting depression symptoms in higher education students using multimodal social media data. In: ICWSM, vol. 14, pp. 440–451 (2020)
12. Masood, R.: Adapting models for the case of early risk prediction on the internet. In: Azzopardi, L., Stein, B., Fuhr, N., Mayr, P., Hauff, C., Hiemstra, D. (eds.) ECIR 2019. LNCS, vol. 11438, pp. 353–358. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-15719-7_48

13. Nasir, M., Jati, A., Shivakumar, P.G., Nallan Chakravarthula, S., Georgiou, P.: Multimodal and multiresolution depression detection from speech and facial landmark features. In: AVEC, pp. 43–50 (2016)
14. Pang, G., Shen, C., Cao, L., Hengel, A.V.D.: Deep learning for anomaly detection: a review. *ACM Comput. Surv.* **54**(2), 1–38 (2021)
15. Pang, G., Shen, C., van den Hengel, A.: Deep anomaly detection with deviation networks. In: KDD, pp. 353–362 (2019)
16. Ruder, S.: An overview of multi-task learning in deep neural networks. arXiv preprint arXiv:1706.05098 (2017)
17. Ruff, L., et al.: Deep one-class classification. In: International Conference on Machine Learning, pp. 4393–4402. PMLR (2018)
18. Sadeque, F., Xu, D., Bethard, S.: Measuring the latency of depression detection in social media. In: WSDM, pp. 495–503 (2018)
19. Shen, T., et al.: Cross-domain depression detection via harvesting social media. In: IJCAI, pp. 1611–1617 (2018)
20. Standley, T., Zamir, A., Chen, D., Guibas, L., Malik, J., Savarese, S.: Which tasks should be learned together in multi-task learning? In: ICML, pp. 9120–9132. PMLR (2020)
21. Strezoski, G., Noord, N.V., Worring, M.: Many task learning with task routing. In: ICCV, pp. 1375–1384 (2019)
22. Uddin, M.A., Joolee, J.B., Lee, Y.K.: Depression level prediction using deep spatiotemporal features and multilayer bi-ITSM. *IEEE Trans. Affect. Comput.* (2020)
23. WHO. Depression and Other Common Mental Disorders: Global Health Estimates. Technical Report, World Health Organization (2017)
24. Yang, Y., Fairbairn, C., Cohn, J.F.: Detecting depression severity from vocal prosody. *IEEE Trans. Affect. Comput.* **4**(2), 142–150 (2012)
25. Zhang, Y., Yang, Q.: A survey on multi-task learning. arXiv preprint arXiv:1707.08114 (2017)
26. Zhou, X., Jin, K., Shang, Y., Guo, G.: Visually interpretable representation learning for depression recognition from facial images. *IEEE Trans. Affect. Comput.* (2018)