

Graduate Employment Prediction with Bias

Teng Guo,¹ Feng Xia,^{1,2} Shihao Zhen,¹ Xiaomei Bai,³ Dongyu Zhang,^{1*} Zitao Liu,⁴ Jiliang Tang⁵

¹Key Laboratory for Ubiquitous Network and Service Software of Liaoning Province,
Dalian University of Technology, Dalian 116620, China

²School of Science, Engineering and Information Technology, Federation University Australia, Ballarat, VIC 3353, Australia

³Computing Center, Anshan Normal University, Anshan 114007, China

⁴TAL AI Lab, TAL Education Group, Beijing 100080, China

⁵Department of Computer Science and Engineering, Michigan State University, East Lansing, MI 48824, USA

Abstract

The failure of landing a job for college students could cause serious social consequences such as drunkenness and suicide. In addition to academic performance, unconscious biases can become one key obstacle for hunting jobs for graduating students. Thus, it is necessary to understand these unconscious biases so that we can help these students at an early stage with more personalized intervention. In this paper, we develop a framework, i.e., MAYA (Multi-mAyor emploYment stAtus) to predict students' employment status while considering biases. The framework consists of four major components. Firstly, we solve the heterogeneity of student courses by embedding academic performance into a unified space. Then, we apply a generative adversarial network (GAN) to overcome the class imbalance problem. Thirdly, we adopt Long Short-Term Memory (LSTM) with a novel dropout mechanism to comprehensively capture sequential information among semesters. Finally, we design a bias-based regularization to capture the job market biases. We conduct extensive experiments on a large-scale educational dataset and the results demonstrate the effectiveness of our prediction framework.

1 Introduction

Education, as the basic means of improving individual abilities, makes students competitive in recruitment. However, not every graduate can succeed in job hunting. The data from the statistical office of the European Union (EU) shows that the employment rate in the EU of 20-34 years old is 83.4% for tertiary education and 65.8% for upper secondary general education in 2018 (Eurostat 2019). The failure of job hunting could cause serious consequences like suicide (Drum et al. 2009; Westfeld et al. 2005). Therefore, detecting students with trouble in landing a job timely and providing personalized intervention and guidance at an early stage are greatly desired.

However, detecting these students faces tremendous challenges because recruitment might be impacted by various factors (Luo and Pardos 2018). Every recruiter aims to hire

the best employees. However, in addition to academic performance (Kong et al. 2018), recruitment decisions can be affected by unconscious biases, such as gender (Clauset, Arbesman, and Larremore 2015). These biases not only lead to imbalance in the hiring process, resulting in uniformity in the workplace instead of diversity, but also result in the inequality of employment (Giannakas, Fulton, and Awada 2017; Ford et al. 2018; Liang, Hong, and Gu 2018), especially for fresh graduates who have no work experience. Therefore, it is necessary to understand biases in recruitment, which can further be exploited for the prediction of graduates' employment. Nevertheless, previous related researches have been mainly based on questionnaires, which are time- and cost-consuming and hardly applicable to large-scale students.

Thanks to the advance of information technology, we are able to digitalize records of students in schools which product rich data about students. It enables the data-driven development (Wu et al. 2013; Liu et al. 2018) and provides us an opportunity to deepen our understandings on the employment of students. However, to achieve the goal, we face tremendous challenges. First, such data is much more complex than that based on questionnaires, thus advanced techniques are needed. Second, the number of graduates that cannot land a job is much smaller compared to those who can successfully obtain jobs, thus employment analysis and prediction are highly imbalanced. Third, there may exist biases in employment that can be varied by majors; while the majority of existing algorithms seldom consider possible biases.

In this paper, we are devoted to exploring the biases in different majors from demographics aspects, and predicting students with trouble in landing a job at an early stage. First, we analyze the employment biases for each major from 4 aspects including gender, nation, hometown, and enroll status. Second, based on possible employment biases, we propose a MAYA (Multi-mAyor emploYment stAtus) prediction framework, with four important components. In the first component, we solve the heterogeneity of students' courses through embedding academic performance into a space of unified dimension by autoencoder. Then GAN (Generative Adversarial Networks) is applied to generate data of the mi-

*Corresponding author: Dongyu Zhang (email: zhang-dongyu@dlut.edu.cn)

Copyright © 2020, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

nority class, to overcome the label imbalance problem. Next, considering the sequential information between semesters, the Long Short-Term Memory (LSTM) with a novel dropout mechanism is utilized. Finally, we design a model to capture the employment biases of different majors. Our contributions can be summarized as follows:

- We provide a comprehensive and systematic analysis on employment biases.
- We model the employment biases in different majors and incorporate them into our proposed prediction framework.
- We conduct comprehensive experiments on a large-scale educational dataset and the extensive results demonstrate the effectiveness of our prediction framework.

This paper is organized as follows. In Section 2, related work is reviewed. The problem formulation is presented in Section 3. In Section 4, we analyze the employment biases by majors. In Section 5, the MAYA prediction framework is introduced in detail. In Section 6, we analyze the results of our experiment. We present the discussion and conclusion of our work in Section 7.

2 Related Work

2.1 Employment of College Graduates

Whether a college student can find a job after graduation attracts tons of attention in recent decades. Liu et al (Liu, Silver, and Bemis 2018) develop a tool for career exploration based on the intuitiveness of node-link diagrams and the scalability of aggregation-based techniques to help the student understand the process of employment. Kong et al (Kong et al. 2018) carry out a series of experiments to explore the relationship between students’ academic performance and their graduation whereabouts. Uosaki et al (Uosaki et al. 2018) develop a career support system to help the international student find a job in Japan through their log record and eBook reading. Liu et al (Liu et al. 2017) design a job recommendation service framework for university students. According to a student profiling based re-ranking rule, users are recommended a list of potential jobs. Soumya et al (Soumya and Sugathan 2017) make a framework to identify student’s eligibility for a specific job by calculating the domain competencies and job competency score.

2.2 Dropout in Recurrent Neural Network

Dropout is a mechanism that stops a part of neurons to improve generalization performance (Srivastava et al. 2014). Zaremba et al (Zaremba, Sutskever, and Vinyals 2014) apply dropout technology on RNN with 0 memory loss through only applying the dropout operator on the non-recurrent connection. Moon et al (Moon et al. 2015) propose an effective solution to better preserve memory when applying dropout through generating a mask on input sequence and moving it at every time step. Gal et al (Gal and Ghahramani 2016) propose a Bayesian interpretation-based RNN dropout variant method. They generate a dropout mask according to the theory of the Bayesian posterior and keep the same mask for each time step in the sequence. Zhu et al (Zhu et al. 2016) propose a dropout method for multilayer LSTM. To keep

information stored in memory, they only allow the dropout mechanism to flow along with layers and prohibit it to flow along with the timeline. Billa Jayadev (Billa 2018) test the dropout mechanism on LSTM algorithm based on a speech recognition system and the results with two datasets are improved about 24.64% and 13.75%, respectively.

3 Problem Statement

In this section, we will introduce some notations and then formally define the problem in this work. In a university, let $\mathcal{M} = \{1, 2, \dots, M\}$ denote the set of majors and the set of students in every major is defined as $\mathcal{Q} = \{\mathcal{N}_1, \mathcal{N}_2, \dots, \mathcal{N}_M\}$. For student i in major m , we define the **academic vector** as $\mathbf{a}_i^m \in \mathbb{R}^n$ that will be introduced in the following section. The **feature vector** and **final employment status** are denoted as $\mathbf{d}_i^m \in \mathbb{R}^p$ and $y_i^m \in \{0, 1\}$. Let $\mathbf{D}^m = [\mathbf{d}_1^m, \mathbf{d}_2^m, \dots, \mathbf{d}_{|\mathcal{N}_M|}^m] \in \mathbb{R}^{|\mathcal{N}_M| \times p}$, $\mathbf{A}^m = [\mathbf{a}_1^m, \mathbf{a}_2^m, \dots, \mathbf{a}_{|\mathcal{N}_M|}^m] \in \mathbb{R}^{|\mathcal{N}_M| \times n}$ and $\mathbf{y}^m = [y_1^m, y_2^m, \dots, y_{|\mathcal{N}_M|}^m] \in \mathbb{R}^{|\mathcal{N}_M|}$ represent the feature matrix, the academic performance matrix and the employment status vector. The details of features used in this research are described in the following section.

Employment Status Prediction Problem: given the feature vector \mathbf{d}_i^m and the corresponding academic performance vector \mathbf{a}_i^m , then we predict the final employment status y_i^m .

4 Bias Analysis

4.1 Dataset

The dataset used in this experiment includes 2,133 students from a Chinese university. They all enrolled in 2013 and graduated in 2017. They are from 64 different majors involved in 13 colleges. This dataset consists of three types of information, which are described as follows:

Demographic Data Students are required to submit personal information at the time of admission, like hometown, gender, and nation. For the privacy concern, the students are already pseudonymous in the raw data. The demographic data includes 2,133 records.

Academic Performance Data Students’ academic performance data contains scores and credits of courses. There are in total 195,234 academic records.

Employment Data When finding a job, students need to sign tripartite agreements to guarantee their legal rights, thus universities own records of students’ employment status information about related companies and government agencies. This dataset consists of 2,133 records.

4.2 Bias in Employment

In this subsection, we analyze the bias in employment from two levels: major-level and college-level. We only show the results of the major-level and leave these of school-level in the Supplemental Material. We check the bias from four aspects: **gender** (Ford et al. 2018), **nation** (the minority or the majority) (Al-Ubaydli and List 2019), **administrative level**

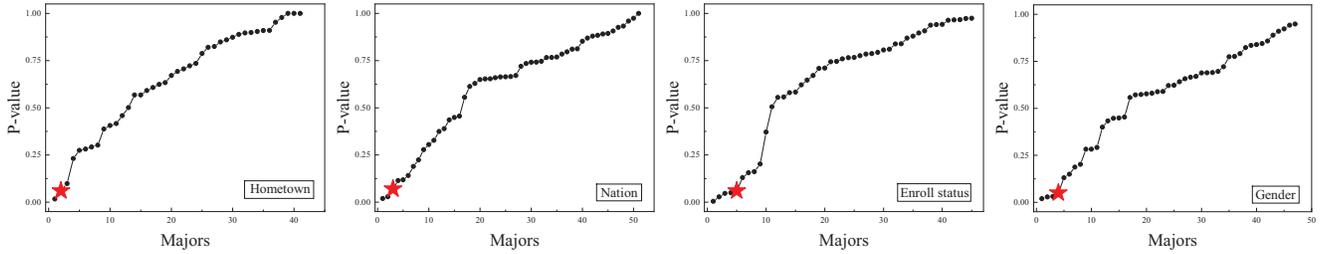


Figure 1: The distribution of p -value with respect to majors on employment status. Subfigures denote the results of chi-square test in terms of hometown, nation, enroll status and gender, respectively. Each black dot represents the p -value of a certain major. When p -value is less than the threshold (e.g., 0.05 shown as red stars), the hypothesis is acceptable, that is, bias exists.

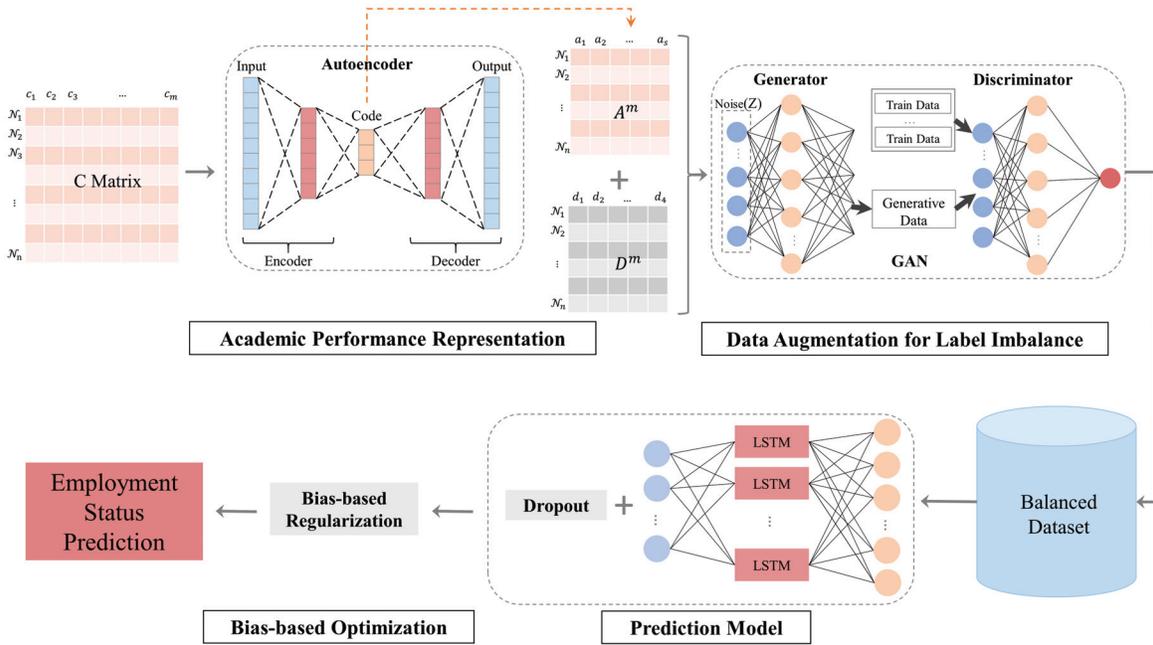


Figure 2: The illustration of MAYA.

of hometown (city or county) (Liang, Hong, and Gu 2018), and **enroll status** (whether passing the college entrance examination at one time). Chi-square test is used here to examine the impact of these features on 64 different majors involved in our dataset. Bias analysis of employment status is shown in Figure 1. From the figure, the majors with bias in recruitment are shown as follows:

- **Gender:** English, Applied Psychology, Electronic Information Science and Technology.
- **Administrative Level of Hometown:** Physical Education.
- **Enroll Status:** Preschool education, English, Electronic Information Science and Technology, Food Science and Engineering.
- **Nation:** Information and Computing Science, Computer Science and Technology.

These observations suggest that employment bias does exist in some majors and the existence of bias does affect graduates' employment. Note that we also analyze the bias in employment choice and results are provided in the Supplemental Material.

5 The Proposed MAYA Prediction Framework

In this section, we provide a detailed description of the proposed framework, MAYA. Figure 2 shows an illustration of the MAYA framework. The framework has four components including representation learning of academic performance, data augmentation for label imbalance, prediction model, and bias-based optimization. Next we detail each component.

5.1 Academic Performance Representation

When taking academic performance as features, the heterogeneity of curriculum is always a challenge, due to the difference in students who take courses for each semester. In this work, we propose a C matrix and based on C , we use an auto-encoder to get the embedding representation to tackle the heterogeneity.

C Matrix To solve the problem caused by the heterogeneity, we create the matrix $C_s \in \mathbb{R}^{n_s \times m_s}$ where n_s and m_s represent the number of students and the number of courses, respectively. In our dataset, $s = 1, 2, \dots, 6$ since students have valid grades of 6 semesters except for the two-semester graduation project and social practice. c_{ij} is the grade of student i on course j . If a student does not attend a particular course, the corresponding element remains 0. The size of this matrix is different for each semester as shown in the following matrix. For example, there are 300 students and 500 courses in the first semester, then the size of C_1 matrix is 300×500 .

$$\begin{pmatrix} c_{11} & c_{12} & \cdots & c_{1n} \\ c_{21} & c_{22} & \cdots & c_{2n} \\ \vdots & \vdots & \ddots & \vdots \\ c_{m1} & c_{m2} & \cdots & c_{mn} \end{pmatrix}$$

Representation Learning C matrix is quite sparse, thus we use the autoencoder to get the embedding representation that is the academic performance matrix A . The hidden layers of the autoencoder are divided into two parts: the encoder part and the decoder part. The layers consistently encode and decode the input data. The input of the i th layer is considered as the output of $(i - 1)$ th layer. The hidden layers can automatically capture the characteristics of input data and keep them unchanged. To capture the temporality among semesters, we use autoencoder to embed the matrix of each semester, respectively. In each hidden layer, we adopt the following nonlinear transformation function:

$$\begin{aligned} \mathbf{h}_{(2)} &= f(\mathbf{W}_{(2)}\mathbf{h}_{(1)} + \mathbf{b}_{(2)}) \\ \mathbf{h}_{(3)} &= f(\mathbf{W}_{(3)}\mathbf{h}_{(2)} + \mathbf{b}_{(3)}) \\ \mathbf{h}_{(i)} &= f(\mathbf{W}_{(i)}\mathbf{h}_{(i-1)} + \mathbf{b}_{(i)}), i = 1, 2, \dots, k \end{aligned} \quad (1)$$

where f is the activation function and $\mathbf{W}_{(i)}$, $\mathbf{b}_{(i)}$ are the transformation matrix and the bias vector. We use C as the input and minimize the reconstruction error between the output and the original input. Then, we take the output of the encoder as the academic performance matrix A .

5.2 Data Augmentation for Label Imbalance

In general, the number of students who fail to land a job is smaller, leading to a label imbalance problem. Thus, we employ generative adversarial networks (GAN) (Goodfellow et al. 2014) to augment data in order to improve the generalization performance. GAN consist of two components: a generator G and a discriminator D that compete in a two-player mini-max game on $V(D, G)$:

$$\begin{aligned} \min_G \max_D V(D, G) &= \mathbb{E}_{x \sim p_{data}} [\log D(x)] + \\ &\mathbb{E}_{x \sim p_z(z)} [\log(1 - D(G(z)))] \end{aligned} \quad (2)$$

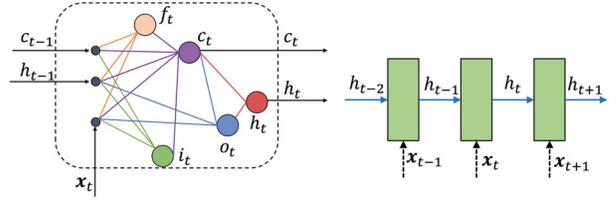


Figure 3: The diagram of LSTM.

The generator G shown in Figure 2 takes a random vector from a uniform distribution as input. It outputs a vector including all features of the corresponding class (i.e., the students failed in job hunting). Then, the generated data and the real data are entered into discriminator D for classification. Through repeated training, the D can not identify the generated data from the real data. Then we use G to generate data of students failed in job hunting until the two categories are balanced. In other words, we aim to implicitly learn the distribution of data of students failed in job hunting, to further generate new samples.

5.3 Prediction Model

We utilize LSTM to capture the sequentiality between semesters for prediction. LSTM is an RNN architecture that uses a vector of cells $\mathbf{c}_t \in \mathbb{R}^n$ with several element-wise multiplication gates to manage information. Generally, dropout aims to combine many 'thinned' networks to improve the performance of prediction. Some neurons are randomly dropped during the training stage in order to force the remaining sub-network to compensate. Then all the neurons are used for predictions during the testing stage.

In this work, we design a temporal dropout structure to improve the generalization performance. Previous researches show that it's not expected to erase information from a unit who remembers events that occurred many timestamps back in the past (Pham et al. 2014; Zhu et al. 2016). However, using complex models on a relatively simple dataset can easily cause overfitting. Moreover, the time span is not very long in our problem. Thus we allow the information of dropout in LSTM to flow along the time dimension. We utilize the classical LSTM framework shown as follows:

$$\begin{aligned} \mathbf{i}_t &= \sigma(\mathbf{W}_{xi}\mathbf{x}_t + \mathbf{W}_{hi}\mathbf{h}_{t-1} + \mathbf{W}_{ci}\mathbf{c}_{t-1} + \mathbf{b}_i) \\ \mathbf{f}_t &= \sigma(\mathbf{W}_{xf}\mathbf{x}_t + \mathbf{W}_{hf}\mathbf{h}_{t-1} + \mathbf{W}_{cf}\mathbf{c}_{t-1} + \mathbf{b}_f) \\ \mathbf{c}_t &= (\mathbf{f}_t \odot \mathbf{c}_{t-1} + \mathbf{i}_t \odot \tanh(\mathbf{W}_{xc} + \mathbf{W}_{hc}\mathbf{h}_{t-1} + \mathbf{b}_c)) \\ \mathbf{o}_t &= \sigma(\mathbf{W}_{xo}\mathbf{x}_t + \mathbf{W}_{ho}\mathbf{h}_{t-1} + \mathbf{W}_{co}\mathbf{c}_{t-1} + \mathbf{b}_o) \\ \mathbf{h}_t &= \mathbf{o}_t \odot \tanh(\mathbf{c}_t) \end{aligned} \quad (3)$$

where $\sigma(x)$ is the sigmoid function defined as $\sigma(x) = \frac{1}{1+e^{-x}}$. $\mathbf{W}_{\alpha\beta}$ denotes the weight matrix between α and β (e.g., \mathbf{W}_{xi} is the weight matrix from input \mathbf{x}_t to the input gate \mathbf{i}_t), \mathbf{b}_α is the bias term of $\alpha \in \{i, f, c, o\}$.

Inspired by (Zhu et al. 2016), we design an LSTM variant that allows the information of dropout in LSTM to flow

along the time dimension through designing the mask vector \mathbf{m} to drop the gates. The structure is shown in Figure 3, defined by the following equations:

$$\begin{aligned}
\mathbf{i}_t &= \sigma(\mathbf{W}_{xi}\mathbf{x}_t + \mathbf{W}_{hi}\mathbf{h}_{t-1} + \mathbf{W}_{ci}\mathbf{c}_{t-1} + \mathbf{b}_i) \odot \mathbf{m}_i \\
\mathbf{f}_t &= \sigma(\mathbf{W}_{xf}\mathbf{x}_t + \mathbf{W}_{hf}\mathbf{h}_{t-1} + \mathbf{W}_{cf}\mathbf{c}_{t-1} + \mathbf{b}_f) \odot \mathbf{m}_f \\
\mathbf{c}_t &= (\mathbf{f}_t \odot \mathbf{c}_{t-1} + \mathbf{i}_t \odot \tanh(\mathbf{W}_{xc} + \mathbf{W}_{hc}\mathbf{h}_{t-1} + \mathbf{b}_c)) \\
&\odot \mathbf{m}_c \\
\mathbf{o}_t &= \sigma(\mathbf{W}_{xo}\mathbf{x}_t + \mathbf{W}_{ho}\mathbf{h}_{t-1} + \mathbf{W}_{co}\mathbf{c}_t + \mathbf{b}_o) \odot \mathbf{m}_o \\
\mathbf{h}_t &= \mathbf{o}_t \odot \tanh(\mathbf{c}_t) \odot \mathbf{m}_h
\end{aligned} \tag{4}$$

where \odot represents element-wise product and \mathbf{m}_f , \mathbf{m}_c , \mathbf{m}_o and \mathbf{m}_h are dropout binary mask vectors, with an element value of 0 indicating that dropout happens, for input gates, forget gates, cells, output gates and output gates, respectively.

In our case, since we use single-layer LSTM, we only need to consider error back-propagation in the same network layer and the errors to output responses \mathbf{h}_t are:

$$\epsilon_h^t = \epsilon_{h+1}^t \odot \mathbf{m}_h \tag{5}$$

where ϵ_{h+1} represents the back-propagation error vectors from the next time in the same network layer.

Based on the Eq. 4, we get the errors from \mathbf{h}_t to \mathbf{o}_t which represents the errors from next time with dropout involved:

$$\epsilon_o^t = (\epsilon_h^t \odot \frac{\partial \mathbf{h}_t}{\partial \mathbf{o}_t}) \odot \mathbf{m}_o = \epsilon_h^t \odot \tanh(\mathbf{c}_t) \odot \mathbf{m}_o \tag{6}$$

Using the same approach, we can get the back-propagation errors of other gates and the details are provided in the Supplemental Material.

5.4 Bias-based Optimization

Modeling Employment Bias As mentioned above, bias is varied by majors in employment. This finding motivates us to eliminate the influence of bias in various majors. Therefore, we propose a smoothed regularization for the steady change of weight, defined as follow:

$$\Omega_M = \frac{1}{2} \sum_{m=1}^M \sum_{n>m}^M \|\mathbf{W} \cdot (\mathbf{u}_m - \mathbf{u}_n)\|_F^2 \tag{7}$$

where \mathbf{W} is the weight matrix of LSTM mentioned above. $\|\cdot\|_F$ denotes the Frobenius norm. $\mathbf{u}_m \in \mathbb{R}^p$ indicates the importance of the p tested aspects in prediction for students in major m . It is represented by the p -value of chi-square test calculated in the Section 4 through a transformation function for emphasizing the importance of biases. In other words, the lower the p value, the greater the weight of the bias. The transformation function is defined as follows:

$$f(u) = \frac{e^{1-u} - e^{1+u}}{e^{1-u} + e^{1+u}} \tag{8}$$

Optimization Based on the discussion above, we formulate the whole loss function of our MAYA prediction framework as follows:

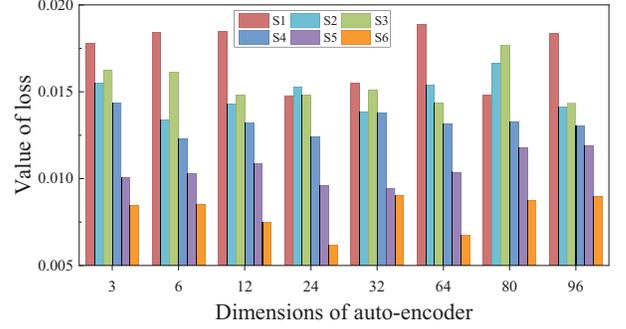


Figure 4: The results of representation learning.

$$\mathcal{L} = \frac{1}{2} \sum_{m=1}^M (\mathbf{W} \mathbf{x}_i^m - \mathbf{y}_i^m)^2 + \Omega_M \tag{9}$$

Its corresponding gradient is shown as follows:

$$\begin{aligned}
\nabla \mathcal{L}(\mathbf{W}) &= \sum_{m=1}^M \mathbf{W} \mathbf{x}_i^m (\mathbf{x}_i^m)^\top - \mathbf{y}_i^m (\mathbf{x}_i^m)^\top + \\
&\sum_{m=1}^M \sum_{n>m}^N \mathbf{W} (\mathbf{u}_m - \mathbf{u}_n) (\mathbf{u}_m - \mathbf{u}_n)^\top
\end{aligned} \tag{10}$$

6 Experiment

In this section, we would present the experimental results in detail to demonstrate the effectiveness of our proposed MAYA. We first introduce the experimental settings, then present comparison results and finally investigate important parameters of the proposed framework.

6.1 Experimental Settings

To deal with the heterogeneity of courses enrolled by students each semester, we design a \mathcal{C} matrix to denote students' academic performance. The four-year university life involves 8 semesters. Campus recruitment takes place densely at the beginning of the last year. Hence, only the academic performance of the previous three years (or six semesters $S1$ to $S6$) would affect students' employment. Autoencoder is applied to embed academic performance data to overcome the heterogeneity of course selection. We test the different dimensions including 3, 6, 12, 24, 32, 64, 80, 96 and the performance is shown in Figure 4. The value of loss function fluctuates slightly. That is, even vectors with low dimensions can still effectively represent the academic performance of each student. Thus, we choose 3 as the dimension of representation for computational efficiency.

6.2 Prediction Results

We predict employment status with features including academic performance, gender, nation, enroll status, hometown and their major. To verify the effectiveness of our MAYA framework, we design prediction experiments including two settings: comparison with LSTM-based MAYA's variants and comparison with representative baselines.

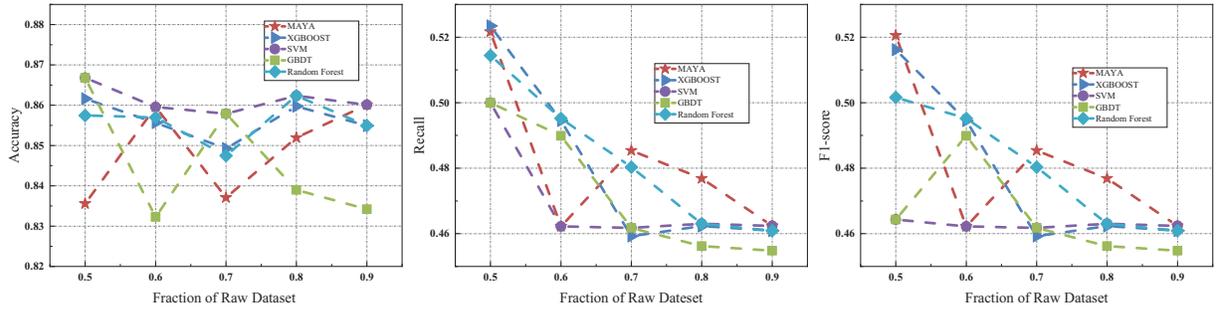


Figure 5: Prediction performance on raw training dataset a .

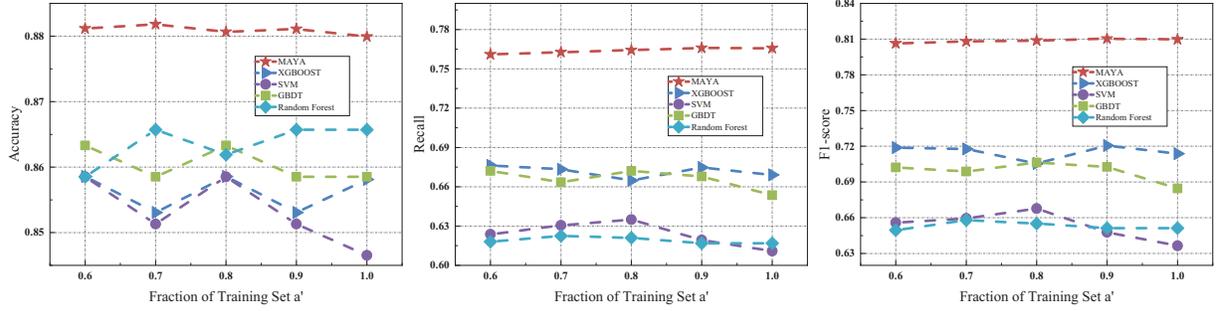


Figure 6: Prediction performance on balanced training dataset a' .

Variants	Accuracy	Recall	F1-score
LSTM+Raw Data	0.862	0.500	0.463
LSTM+GAN	0.869	0.670	0.717
LSTM+Dropout+GAN	0.876	0.712	0.761
LSTM+Dropout+GAN+New Loss	0.880	0.766	0.810

Table 1: Prediction performance of MAYA variants.

Comparison with LSTM-based MAYA’s Variants Table 1 displays the prediction performance of MAYA and its variants. We design a four-step experiment to test the performance with metrics, i.e., accuracy, recall and F1-score, to understand the results collectively.

In the first step, we use raw data to fit the LSTM algorithm. The imbalance label issue exists in raw data, leading to the occurrence of unexpected results on precision and recall. First, the algorithm could achieve a low loss value by ignoring the minority class and predicting all the samples into the majority class. It results in a recall with 0.5. Second, the precision is relatively low because no samples are predicted as a minority. In other words, for the minority class, the correct ratio of prediction results is 0.

In the second step, GAN is used to solve the label imbalance and the data generation process is shown as follows:

- First, raw data is divided into two categories: training set a and testing set b by stratified sampling.
- Second, we use GAN on the training set a to generate samples of the minority class. Then in the new training

set a' , the number of students in the two classes is equal.

Then, we use the training set a' to fit the model and test it on the original testing set b . The performance shown in Table 1 verifies its effectiveness.

In the third step, a new dropout mechanism of LSTM is employed to alleviate the overfitting problem caused by the relatively small experimental dataset. We add the dropout mechanism based on the experiment of the second step. In the final step, we add the bias-based regularization into the optimization loss based on the last step and the results suggest its importance.

Comparison with Baseline Methods In addition to the comparison with deep learning-based variants, we compare the MAYA framework with several popular algorithms shown as follows:

- **SVM** (Scholkopf and Smola 2001): SVM is a classic algorithm and is widely used in the field of data mining.
- **Random Forest** (Breiman 2001): is a classic ensemble algorithm that achieves good performance in various applications.
- **GBDT** (Friedman 2001): GBDT is an additive regression model consisting of regression trees.
- **XGBoost** (Chen and Guestrin 2016): XGBoost is a boosting-tree-based method and is widely used in various data mining scenarios with good performance.

Note that, we test the performance of these algorithms from two aspects. On one hand, we fit algorithms based on the raw training set a and test them on testing set b . The results are shown in Figure 5. It’s shown that the prediction is

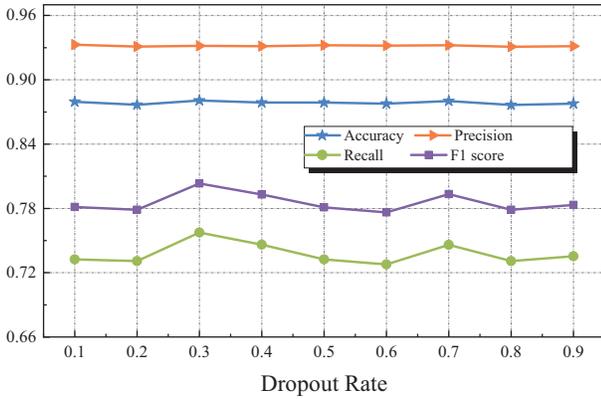


Figure 7: Performance of MAYA with different dropout proportions.

Semester	Accuracy	Precision	Recall	F1
1	0.73469	0.60848	0.62668	0.61484
2	0.82287	0.71878	0.65398	0.67470
3	0.86666	0.92909	0.654762	0.69820
4	0.86758	0.92944	0.658824	0.70311
5	0.87631	0.93264	0.69898	0.74856
6	0.88073	0.93172	0.757463	0.80326

Table 2: Performance of MAYA with different numbers of semesters.

not accurate and the fluctuation is quite large due to the label imbalance of raw data. To overcome this problem, we fit algorithms based on the balanced training set a' and test them on b . The results are shown in Figure 6. The performance is improved significantly.

6.3 Parameter Sensitivity

Dropout Proportions As mentioned above, a dropout mechanism is utilized to improve the generalization performance and we test the sensitivity of MAYA framework on dropout proportions (Figure 7). The change of dropout proportions generates a slight impact and 0.3 is the best.

Input Features It's of great significance to distinguish early the students who might encounter difficulties in employment. Then teachers can intervene at an early stage. Therefore, we conduct a test on the number of semesters involved in academic performance data. As shown in Table 2, prediction performance grows slowly since the fourth semester. In other words, we can predict the students with trouble in landing a job with high accuracy at the end of the second year. We also design an experiment to test the effectiveness of demographic features and leave the results in Supplemental Material.

Learning Rate Learning rate that controls the update speed of the model is an important parameter in the MAYA framework. In Figure 8, we analyze the performance of various learning rates and find that the model can achieve the best prediction performance when the learning rate is set to 0.01.

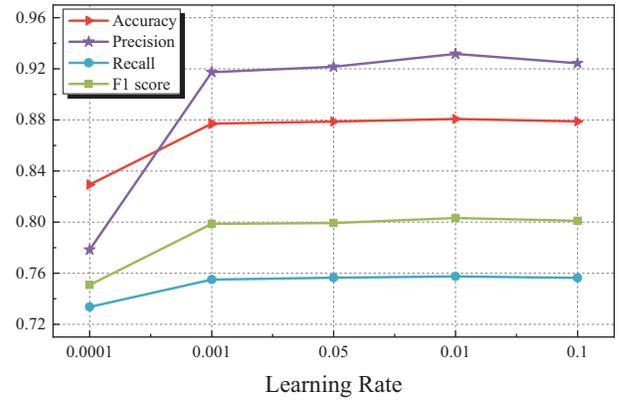


Figure 8: Performance of MAYA with different learning rates.

Optimization Function	Accuracy	Recall	F1-score
Eq.11	0.873	0.689	0.738
Eq.9	0.880	0.766	0.810

Table 3: Performance of Different Optimization.

Bias-based Regularization We design an experiment to test the effectiveness of bias based regularization. We use Eq.11 and Eq.9 as loss function separately and the prediction performance is shown in Table 3. The bias-based regularization can improve the performance remarkably.

$$\mathcal{L}_1 = \frac{1}{2} \sum_{m=1}^M (\mathbf{W} \mathbf{x}_i^m - \mathbf{y}_i^m)^2 + \|\mathbf{W}\|_F^2 \quad (11)$$

7 Conclusion and Discussion

In this paper, we analyze a large-scale educational data for predicting graduates' employment status. For the reason that bias cannot be ignored in employment, we analyze the employment bias of different majors first of all and verify the existence of employment bias. Then based on such bias, MAYA, a prediction framework, is proposed in this paper to predict graduates' employment status. We incorporate the autoencoder to ease the data-sparse issue and deal with the imbalance label data using GAN. LSTM is combined with dropout and a bias-based regularization to overcome the over-fitting problem and capture the impact of biases. Our extensive experiments based on an education dataset demonstrate the proposed framework can improve the prediction performance significantly and MAYA outperforms other baselines like LSTM and XGBoost significantly.

There are multiple directions for future work. Firstly, we plan to expand our dataset and explore this issue from more aspects. Secondly, we would conduct data acquisition from various companies and further study this issue from the company's perspective. Last but not least, we also intend to integrate MAYA framework into the modern educational management system and apply it to detect the employment status of graduating students.

References

- Al-Ubaydli, O., and List, J. A. 2019. How natural field experiments have enhanced our understanding of unemployment. *Nature Human Behaviour* 3(1):33–39.
- Billa, J. 2018. Dropout approaches for lstm based speech recognition systems. In *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 5879–5883. IEEE.
- Breiman, L. 2001. Random forests. *Machine learning* 45(1):5–32.
- Chen, T., and Guestrin, C. 2016. Xgboost: A scalable tree boosting system. In *Proceedings of the 22nd acm sigkdd international conference on knowledge discovery and data mining*, 785–794. ACM.
- Clauset, A.; Arbesman, S.; and Larremore, D. B. 2015. Systematic inequality and hierarchy in faculty hiring networks. *Science advances* 1(1):e1400005.
- Drum, D. J.; Brownson, C.; Burton Denmark, A.; and Smith, S. E. 2009. New data on the nature of suicidal crises in college students: Shifting the paradigm. *Professional Psychology: Research and Practice* 40(3):213.
- Eurostat. 2019. Employment rates of recent graduates. https://ec.europa.eu/eurostat/statistics-explained/index.php/Employment_rates_of_recent_graduates#Employment_rates_of_recent_graduates.
- Ford, H. L.; Brick, C.; Blaufuss, K.; and Dekens, P. S. 2018. Gender inequity in speaking opportunities at the american geophysical union fall meeting. *Nature communications* 9(1):1358.
- Friedman, J. H. 2001. Greedy function approximation: a gradient boosting machine. *Annals of statistics* 29(5):1189–1232.
- Gal, Y., and Ghahramani, Z. 2016. A theoretically grounded application of dropout in recurrent neural networks. In *Advances in neural information processing systems*, 1019–1027. Curran Associates Inc.
- Giannakas, K.; Fulton, M.; and Awada, T. 2017. Hiring leaders: Inference and disagreement about the best person for the job. *Palgrave Communications* 3(1):17.
- Goodfellow, I.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; and Bengio, Y. 2014. Generative adversarial nets. In *Advances in neural information processing systems*, 2672–2680.
- Kong, J.; Ren, M.; Lu, T.; and Wang, C. 2018. Analysis of college students employment, unemployment and enrollment with self-organizing maps. In *International Conference on E-Learning and Games*, 318–321. Springer.
- Liang, C.; Hong, Y.; and Gu, B. 2018. Home bias in hiring: Evidence from an online labor market. In *PACIS*, 49.
- Liu, R.; Rong, W.; Ouyang, Y.; and Xiong, Z. 2017. A hierarchical similarity based job recommendation service framework for university students. *Frontiers of Computer Science* 11(5):912–922.
- Liu, J.; Kong, X.; Xia, F.; Bai, X.; Wang, L.; Qing, Q.; and Lee, I. 2018. Artificial intelligence in the 21st century. *IEEE Access* 6:34403–34421.
- Liu, L.; Silver, D.; and Bemis, K. 2018. Application-driven design: Help students understand employment and see the big picture. *IEEE computer graphics and applications* 38(3):90–105.
- Luo, Y., and Pardos, Z. A. 2018. Diagnosing university student subject proficiency and predicting degree completion in vector space. In *Thirty-Second AAAI Conference on Artificial Intelligence*, 7920–7927. AAAI Press.
- Moon, T.; Choi, H.; Lee, H.; and Song, I. 2015. Rnndrop: A novel dropout for rnns in asr. In *2015 IEEE Workshop on Automatic Speech Recognition and Understanding (ASRU)*, 65–70. IEEE.
- Pham, V.; Bluche, T.; Kermorvant, C.; and Louradour, J. 2014. Dropout improves recurrent neural networks for handwriting recognition. In *2014 14th International Conference on Frontiers in Handwriting Recognition*, 285–290. IEEE.
- Scholkopf, B., and Smola, A. J. 2001. *Learning with kernels: support vector machines, regularization, optimization, and beyond*. MIT press.
- Soumya, M., and Sugathan, T. 2017. Improve student placement using job competency modeling and personalized feedback. In *2017 International Conference on Advances in Computing, Communications and Informatics (ICACCI)*, 1751–1755. IEEE.
- Srivastava, N.; Hinton, G.; Krizhevsky, A.; Sutskever, I.; and Salakhutdinov, R. 2014. Dropout: a simple way to prevent neural networks from overfitting. *The journal of machine learning research* 15(1):1929–1958.
- Uosaki, N.; Mouri, K.; Yin, C.; and Ogata, H. 2018. Seamless support for international students’ job hunting in japan using learning log system and ebook. In *2018 7th International Congress on Advanced Applied Informatics (IIAI-AAI)*, 374–377. IEEE.
- Westefeld, J. S.; Homaifar, B.; Spotts, J.; Furr, S.; Range, L.; and Werth, J. L. 2005. Perceptions concerning college student suicide: Data from four universities. *Suicide and Life-Threatening Behavior* 35(6):640–645.
- Wu, X.; Zhu, X.; Wu, G.-Q.; and Ding, W. 2013. Data mining with big data. *IEEE transactions on knowledge and data engineering* 26(1):97–107.
- Zaremba, W.; Sutskever, I.; and Vinyals, O. 2014. Recurrent neural network regularization. *arXiv preprint arXiv:1409.2329*.
- Zhu, W.; Lan, C.; Xing, J.; Zeng, W.; Li, Y.; Shen, L.; and Xie, X. 2016. Co-occurrence feature learning for skeleton based action recognition using regularized deep lstm networks. In *Thirtieth AAAI Conference on Artificial Intelligence*, 3697–3703. AAAI Press.