

Being Automated or Not? Risk Identification of Occupations with Graph Neural Networks

Dawei Xu¹, Haoran Yang¹, Marian-Andrei Rizoiu², and Guandong Xu²

¹ University of Technology Sydney, Sydney, Australia
{dawei.xu, haoran.yang-2}@student.uts.edu.au

² University of Technology Sydney, Sydney, Australia
{marian-andrei.rizoiu, guandong.xu}@uts.edu.au

Abstract. The rapid advances in automation technologies, such as artificial intelligence (AI) and robotics, pose an increasing risk of automation for occupations, with a likely significant impact on the labour market. Recent social-economic studies suggest that nearly 50% of occupations are at high risk of being automated in the next decade. However, the lack of granular data and empirically informed models have limited the accuracy of these studies and made it challenging to predict which jobs will be automated. In this paper, we study the automation risk of occupations by performing a classification task between automated and non-automated occupations. The available information is 910 occupations' task statements, skills and interactions categorised by Standard Occupational Classification (SOC). To fully utilize this information, we propose a graph-based semi-supervised classification method named **Automated Occupation Classification based on Graph Convolutional Networks (AOC-GCN)** to identify the automated risk for occupations. This model integrates a heterogeneous graph to capture occupations' local and global contexts. The results show that our proposed method outperforms the baseline models by considering the information of both internal features of occupations and their external interactions. This study could help policymakers identify potential automated occupations and support individuals' decision-making before entering the job market.

Keywords: Automated occupation identification · Graph convolutional network · Semi-supervised classification.

1 Introduction

In light of rapid development in the fields of artificial intelligence and robotics, the number of jobs for certain occupations is decreasing alarmingly. Also, massive applications of automation technologies, especially in the workplace, are raising fears of large-scale technological unemployment and a renewed appeal for policymakers to handle the consequences of technological change. From a historical perspective, every disappearance of occupations comes with increasing unemployment [2,16] which brings individual existence crisis, social instability,

and techno-phobia that may hinder the future development of science and technologies. Therefore, occupations' automated risk identification is critical to both policymakers and individuals.

Generally, an occupation at risk of being automated has attributes containing routine and rule-based information, which are prone to be replaced by automation technologies [5,6]. Predicting the automated risk of occupations can be performed by capturing the routine and rule-based features in their attributes, more specifically, their tasks and skills information [4]. Therefore, this prediction task can be conducted as a binary classification between automated and non-automated occupations by mining their internal attributes and external interactions.




Fig. 1: BLS Projected Declining Occupations in 2030

However, this task becomes more challenging due to the following three characteristics of the data:

- Limited Data Source.** Predicting automated risk of occupations requires tasks and skills data that are at their most granular level to reflect micro differences among occupations. Standard labor data from U.S. Bureau of Labor Statistics (BLS) or Australian Bureau of Statistics (ABS) focus on demographic, economic or other aggregate statistics that lack resolutions into the defining features to distinguish different occupations. Extracting granular tasks and skills data from job recruitment platforms is also not feasible since they are dynamic, computing expensive and not all jobs are posted online. The numbers of job posts are also various across occupations which may lead to imbalanced data [8]. Furthermore, the lack of ground truth data is also a barrier hindering the prediction of automated risk.
- Non-numerical Data Type.** Most of the tasks and skills data are in text form, which are single words, grouped words or sentences. These text data cannot be directly processed by statistical models. Current studies [2][6] use importance or level value attached to each task or skill as its numerical representations; however, this kind of representation is aggregated and coarse

which may obfuscate the differential impact of various technologies and skill requirements then further affects models’ performance in distinguishing occupations. Therefore, choosing a method that can not only encode the text data into a numerical format but maintain the features hidden in the context are necessary.

• **Complex Data Structure.** Since occupations are likely to interact by attributes, e.g., they may share similar skillsets, the prediction of automated risk cannot be limited to occupations’ internal attributes but also the connections with other occupations [4]. Moreover, various attributes may lead to different data dimension sizes, and the relationships among occupations can be multiple paths [17]. Therefore, linear data structures are not capable of capturing all the information, and a more complicated non-linear structure and computational model should be designed to recognise the patterns in the information.

Moreover, existing methods designed for the task have several limitations. For instance, current studies conduct experiments on BLS wages and education data [35] or the importance level, a numerical value of task statements [2][3], and they utilise linear data structure to represent the various attributes. According to the defects mentioned above, the accuracy of their results is unavoidable compromised. A better solution is to use Natural Language Processing (NLP) technologies to generate numerical embeddings for tasks and skills text data and use a non-linear data structure such as a graph to capture not only the features of occupations but also the interactions between occupations. Therefore, we propose a semi-supervised classification method named Automated Occupation Classification based on Graph Convolutional Networks (AOC-GCN) to identify the automated risk for occupations. We will be using tasks and skills datasets extracted from the O*NET database, which is a publicly available occupational information network containing hundreds of job definitions, to understand the labour market.

In summary, the contributions of our work are listed below:

- We use Natural Language Processing (NLP) technologies – Word2Vec and Doc2Vec to generate embeddings for tasks and skills data, which is, to the best of our knowledge, the first work in this domain using NLP to generate numerical representations for text data.
- We propose a GCN-based heterogeneous graph classification method that aggregates both node features and network-wide behaviours, which is also the first work in this domain.
- Our proposed method significantly outperforms baseline models, and the results are verified through comparison with government statistical data.

The rest of the paper is organized as follows. Section 2 lists the related work. We elaborate on the proposed AOC-GCN method in Section 3. Section 4 compares and analyzes the experimental results of our method. In Section 5, we discuss the results and compare them with government statistical data. In Section 6, we discuss limitations and future work, and the paper is concluded in Section 7.

2 Related Work

2.1 Occupations Automated Risk identification.

Most existing studies about automated occupation identification are based on expert-based assessments, or basic statistical models focusing on probability space. Frey et al. [2] assessed the automated probability of occupations using a Gaussian process classification and concluded that 47% of current US employment is at high risk of being automated. Similar studies have been conducted at the impact of automation on occupations in other countries and reached cautionary conclusions: automated occupations rate will be 35% in Finland [10], 59% in Germany [9], and 45 to 60 % across Europe [11]. On the contrary, Arntz et al. [3] followed the experiments in [2] but used job-level data and a different coefficient estimates method, and concluded that a less alarming 9% of employment is at risk. Vermeulen et al. [6] and Arntz et al. [12] also conducted job-level automatability scorings and get similar findings that about 5% of jobs have automated risks. However, the data utilized in Arntz et al. are job survey data conducted on different workers. Arntz et al. considers this data as job level data compared to occupation level data without considering the duplicates of same tasks and skills performed by different workers' job. Therefore, the results may be underestimated. That is to say, job-level studies consider each task to be independent and unique which masked the similarity between tasks. Furthermore, current studies failed to discover the semantic information hidden in text data and ignored the interactions between occupations.

2.2 Graph Neural Networks.

Recent years have witnessed a growing application in using graph neural network-based algorithms to solve graph structure problems [23,24,25]. These methods include both supervised [7] and unsupervised methods [15] [20]. Especially, Graph Convolutional Networks (GCN) [7] have achieved significant performance compared to previous methods, which is an efficient variant of Convolutional Neural Networks (CNN) [22] operating directly on graphs. In GCN, both the graph structure information and node features are also aggregated from neighbours during convolution. The graph convolution operation is defined as feature aggregations of neighbours, and through iterative convolutions, the whole graph's information can be propagated and aggregated to each part of the graph. After GCN, GraphSAGE [21] is proposed, which is a simpler but efficient inductive learning model that breaks the limitation of applying GCN in transductive learning. Furthermore, Graph Attention Networks (GAT) [26] introduces an attention mechanism into GCN and allows nodes to focus on the most relevant neighbours during training.

In this paper, a GCN-based method is first applied to automated occupation classification task through building an occupation-skill graph. Our proposed method can encode not only the features of occupation but also the interaction between occupations.

3 Methodology

3.1 Preliminaries

Generally, GCN-based models [7][21][26] consist of multiple propagation layers and all the nodes are updated simultaneously in each propagation layer. A propagation layer can be divided into two sub-layers: *aggregation* and *combination*. Let $G = (V, E)$ be a graph with node $v \in V$, edge $(v, v') \in E$, and node feature $x_v = h_v^0 \in \mathbf{R}^{d_0}$ for $v \in V$ where d_0 denotes the feature dimension of the node and $h_v^l \in \mathbf{R}^{d_l}$ denotes the hidden state of node v learned by the l -th layer of the model. For a GCN with L layers, aggregation and combination sub-layers at l -th layer ($l = 1, 2, \dots, L$) can be written as :

$$h_{N(v)}^l = \sigma(W^l \cdot \text{AGG}(\{h_{v'}^{l-1}, \forall v' \in N(v)\})) \quad (1)$$

$$h_v^l = \text{COMBINE}(h_v^{l-1}, h_{N(v)}^l) \quad (2)$$

where $N(v)$ is a set of nodes adjacent to v , AGG is a function used to aggregate embeddings from neighbor nodes of v . W^l is a trainable matrix shared among all nodes at layer l . σ is a non-linear activation function e.g., RELU . $h_{N(v)}^l$ denotes the aggregated feature of all neighbors of node v at l -th layer. COMBINE function is used to combine the embedding of node itself and the aggregated embeddings of neighbors.

For heterogeneous graph with different types of nodes and edges, it always comes with different sizes of embeddings. For each edge connecting two different nodes, we generally concatenate the edge embedding from the last propagation layer together with embeddings of the two nodes this edge links to. Similarly, for each node, besides the information from neighbor nodes, the features of edges connected to them are also collected. Therefore the aggregation sub-layer for an edge e can be defined as:

$$h_e^l = \sigma(W_E^l \cdot \text{AGG}_E^l(h_e^{l-1}, h_{U(e)}^{l-1}, h_{I(e)}^{l-1})) \quad (3)$$

where $h_{U(e)}^{l-1}$ and $h_{I(e)}^{l-1}$ denote the hidden states of two types of nodes U and I from the previous propagation layer and

$$\text{AGG}_E^l(h_e^{l-1}, h_{U(e)}^{l-1}, h_{I(e)}^{l-1}) = \text{concat}(h_e^{l-1}, h_{U(e)}^{l-1}, h_{I(e)}^{l-1}) \quad (4)$$

For the two nodes $u \in U$ and $i \in I$, the aggregated neighbor embedding $h_{N(u)}^l$ and $h_{N(i)}^l$ can be calculated as:

$$\begin{aligned} h_{N(u)}^l &= \sigma(W_U^l \cdot \text{AGG}_U^l(\text{concat}(h_i^{l-1}, h_e^{l-1}), \forall e = (u, i) \in E(u))) \\ h_{N(i)}^l &= \sigma(W_I^l \cdot \text{AGG}_I^l(\text{concat}(h_u^{l-1}, h_e^{l-1}), \forall e = (u, i) \in E(i))) \end{aligned} \quad (5)$$

Here the two types of nodes maintain different parameters (W_U^l, W_I^l) and different aggregation functions ($\text{AGG}_U^l, \text{AGG}_I^l$).

After aggregating the neighbors’ information, we follow the concatenation in [6] for the two types of nodes as

$$\begin{aligned} h_u^l &= \text{concat}(V_U^l \cdot h_u^{l-1}, h_{N(u)}^l) \\ h_i^l &= \text{concat}(V_I^l \cdot h_i^{l-1}, h_{N(i)}^l) \end{aligned} \quad (6)$$

Where V_U^l and V_I^l denote trainable weight matrix for node U and node I , and the h_u^l and h_i^l are the hidden states of l -th layer for the two types of nodes.

3.2 Problem Setup

Our purpose is to identify the automated risk for occupations listed in Standard Classification Code (SOC). The attributes used in this problem are task statements and skills extracted from O*NET database. According to the properties of data, each task statement is unique for each occupation while each skill data are a list of skills which are overlapped and shared among other occupations. So we can build a bipartite graph $G(O, S, E)$ where O is the set of occupation nodes, S is the set of skill nodes, and E is the set of edges between occupations and skills. An edge $e \in E$ from an occupation $o \in O$ to a skill $s \in S$ exists if o contains s . It is worth noting that the task statements data will be used as occupation nodes’ feature representations which will be introduced in section 3.4. Therefore, the task can be formulated to a node classification problem on an undirected bipartite graph with two types of attributed nodes.

3.3 Graph Convolutional Networks on Occupation-skill Graph

The occupation-skill graph is a bipartite graph with two types of nodes and one type of edge. Moreover, there is no attributes on edges in our problem, so applying a heterogeneous GCN to solve this problem can be excessive. In fact, we take a shortcut way by applying GCN to treat this bipartite graph as a homogeneous graph by setting identical dimension sizes for both nodes features. Therefore, the propagation rule at layer l is this problem can be defined as:

$$h_v^l = \text{COMBINE}(h_v^{l-1}, \sigma(W^l \cdot \text{AGG}(\{h_{v'}^{l-1}, \forall v' \in N(v)\}))) \quad (7)$$

where $v \in O \cup S$, $N(v)$ is a set of nodes adjacent to v , h_v^l denote the hidden states at l -th layer for each nodes.

3.4 Nodes Feature Representation

Text data of task statements and skills should be converted into an embedding before being encoded as nodes’ features. Here we apply two word embeddings technologies – Word2vec [27] and Doc2vec [28] to generate feature representations for occupation nodes and skill nodes. Firstly, we use Glove pretrained word embeddings to initialize the Word2vec model and retrain the model on task statements data to aggregate the context and semantic similarity information in

this domain. The output from the Word2vec model is a word embeddings table containing vectors for each word. Since the skill data are single word or grouped words, we represent each skill’s feature by simply applying the word vector or making additions of multiple word vectors from word embeddings. Then, we initialize our Doc2vec model with word embeddings from last step and train on task statements data to generate document-level embeddings as occupation nodes feature representations. In detail,

$$h_s^0 = \text{fusion}(\mathbf{w}_1, \mathbf{w}_2, \dots, \mathbf{w}_n) \quad (8)$$

$$h_o^0 = \text{doc2vec}(t_0) \quad (9)$$

where h_s^0, h_o^0 denote the initial embedding of a skill node and an occupation node. \mathbf{w}_i represents the i -th word vector in a skill generated from Word2vec. t_0 is the task statement document of an occupation. Therefore, the parameters of Word2vec and Doc2vec are trained together with others in the model described in Section 3.3.

3.5 Semi-supervised Automated Occupation Classification Model

The final embeddings of AOC-GCN are the concatenation of nodes embeddings learned from Occupation-skill Graph, so we can get

$$y = \text{softmax}(f(\text{concat}(z_o, z_s))) \quad (10)$$

where z_o and z_s denote the embeddings of O and S learned by the proposed GCN method. $f(\cdot)$ is the mapping function to map the embeddings to lower dimension space before putting into *softmax* which is the classifier used to classify automated and non-automated occupations. This method is semi-supervised because the whole graph is updated during training process which is on a partially labeled dataset. The whole pipeline of the study is shown as in Fig 2.

4 Experiment

4.1 Dataset and Metrics

Experiments are conducted on occupation lists, task statements and skills in Standard Classification Code (SOC) from O*NET database. The label information utilized in the experiments is generated from expert-based assessments on O*NET occupations conducted by [2][6] with automated occupations being labeled as 1 and non-automated occupations being labeled as 0. However, the original label information has defects that the number of two classes are unbalanced, and some occupations are too close e.g., web developer and software programmer. We preprocess the data and finally select 112 labeled occupations with 56 automated occupations and 56 non-automated occupations which also have distance between each other. Then, all the data will form an undirected graph structure with 910 occupation nodes, 135 skill nodes and 13222 edges. Training, validation and test set are randomly split with a ratio of 8:1:1. Given the prediction and ground truth, we evaluate the model’s performance using accuracy, precision, recall and F1 score.

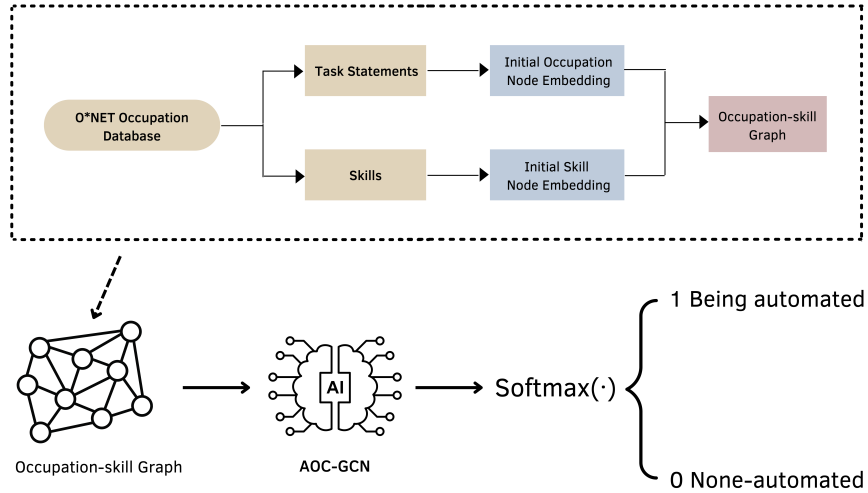


Fig. 2: The whole System Pipeline

4.2 Baselines

We compare our model with traditional classification models. Specifically, we design features of each occupation by concatenate the embeddings generated from Section 3.4 for both task statements and skills. Then we use zero padding to make sure all the occupations have identical feature dimension sizes before putting into baseline models. The chosen baseline models are as following:

- **Decision Tree.** DT splits the dataset as a tree based on a set of rules and conditions which is a supervised learning algorithm that can be used for both classification and regression [32].
- **Random Forest.** Random Forest is an ensemble method which combines the output of multiple decision trees to reach a single result [30].
- **Adaptive Boost.** AdaBoost (Adaptive Boosting) is a boosting technique that aims at combining multiple weak classifiers to build one strong classifier and is widely used in both classification and regression problem [29].
- **Light Gradient Boosting Machine.** LightGBM is a gradient boosting framework based on decision trees to increase the efficiency of the model and reduce memory usage [31].

4.3 Results and Analysis

Evaluation. We evaluate the results of AOC-GCN model on O*NET dataset with above metrics and compare them with baseline models. The results of the comparison are shown in Table 1 below.

Table 1: Preprocessed results of O*NET dataset

Method	Accuracy	Precision	Recall	F1
Decision Tree	0.7143	0.6667	0.8571	0.7500
Random Forest	0.7500	0.7692	0.7143	0.7407
LightGBM	0.7500	0.6842	0.9286	0.7879
AdaBoost	0.8571	0.8125	0.9286	0.8667
AOC-GCN	0.9091	1.0000	0.8750	0.9333

We find that our AOC-GCN method performs significantly better than other baseline models, since our model not only uses external node features, but also capture the graph structural information which is the interactions between nodes. Besides, baseline models using raw features only already achieve good results, the GCN’s combination of both semantic and structural information of relations successfully attains the best performance. Furthermore, our proposed model ensures high model precision to avoid disturbing non-automated occupations.

Parameter Experiments. In this section, we explore the effects on the model performance of different parameters. The key parameters used here are the dimensions of initial node embedding sizes and the dimension sizes of GCN layer. The dimension sizes of node embeddings range from 50 to 300 and the dimension sizes of GCN layer range from 16 to 512.

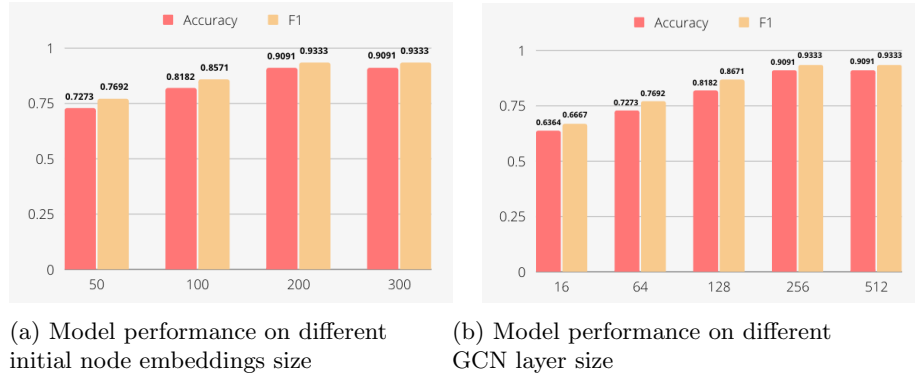


Fig. 3: Parameters Effect On Model Performance

We can see from the Fig 3 that with the increase of sizes of initial node embeddings from 50 to 200, both the accuracy and F1 scores rise and reach the maximum value. When the size is greater than 200, both of the values keep stable. The same pattern can be found in GCN layer dimension sizes, when the size increases from 16 to 256, both of the accuracy and F1 increase. When the size

is greater than 256, the two scores keep unchanged. Therefore, we can conclude that the model gets the best performance when the initial node embeddings size is greater than 200 and the GCN layer dimension size is greater than 256.

Graph Embeddings Visualization In this section, we visualize the embeddings generated from the GCN layer using t-Distributed Stochastic Neighbor Embedding (t-SNE) [33] which is a feature reduction technique mitigating the effects of the ‘‘Curse of Dimensionality’’. We also conduct an unsupervised K-means clustering on the dataset and visualize the result using a dimensionality reduction algorithm Principal component analysis (PCA) [34]. Through comparison, we can find out the difference between applying unsupervised learning and semi-supervised learning. The results are showing in Fig 4 and 5.




Fig. 4: t-SNE projection of GCN embeddings

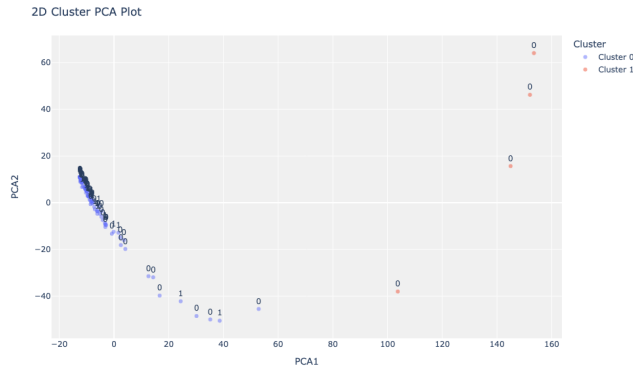


Fig. 5: PCA projection of K-means clustering

We discover that there is a distinct boundary between label 1 and label 0 which are colored in blue and red in t-SNE projection of GCN embeddings in Fig 4. However, there is no distinguishable separation of two K-means clusterings colored in red and blue from PCA projection Fig 5, in which labels are annotated in the graph as 1 and 0. It means that GCN semi-supervised method achieved much better performance compared to the unsupervised training by using only a small portion of labelled training data.

5 Results and Discussion

In this section, we apply the model with the best performance to conduct predictions on the rest of 798 unlabelled occupation nodes to get the probabilities of their automated risk. Part of the results is shown in Table 2.

Table 2: Top 10 Occupations in High Risk of Being Automated

Occupation	Risk
File Clerks	0.7002
Real Estate Brokers	0.6968
Word Processors and Typists	0.6943
Payroll and Timekeeping Clerks	0.6941
Data Entry Keyers	0.6940
Transportation Engineers	0.6938
Credit Analysts	0.6938
Insurance Appraisers, Auto Damage	0.6938
Insurance Underwriters	0.6938
Tax Examiners and Collectors, and Revenue Agents	0.6938

Our results show that the maximum probability of automated risk is 0.7002 which is the “File Clerks”. We follow the definition in [2] and set 69% as the cut-off point so that if the probability is greater than 69% then this occupation has a risk of being automated. Then we obtain 230 occupations and a percentage of 25.3% among all occupations are at risk of being automated. This result is lower than the experiments conducted by [2][9][10] which are about 50% occupations are at risk and is also higher than [3][6] which conducted a job-level estimation of 9% automated risk. Considering the defects in occupation-level and job-level automated risk estimations, our result is quite reasonable to reach a median value between these two estimations since it predicts the occupation-level risk using granular task and skills data.

To further validate our results, we compare them with “U.S. Bureau of Labor Statistics (BLS) Occupations with the largest job declines, 2020 and projected 2030”. It projected out 29 occupations that have declining employments in 2030. The outcome is visualized in Figure 6.

We assume that automated risk greater than 50% would cause employment declining and we can find that 76% of the BLS projected declining occupations have automated risks higher than 50% except for “customer service representatives”, “computer programmers”, “inspectors, testers, sorters, samplers and




Fig. 6: BLS Projected Declining Occupations in 2030

weighters” which have an automated risk probability of 0.1%, 8.15% and 7.5% that are obviously automated safe but are projected declining in 2030. This discrepancy can be caused by rules used in expert assessment of the ground truth data. According to [2], experts assess the risks based on three attributes of an occupation which are “perception and manipulation”, “creative intelligence tasks” and “social intelligence tasks”. i.e., “customer service representatives” have tasks including persuading and dispute resolution which requires higher social intelligence and perception, therefore, they are not easy to be automated according to the analysis in [2]. BLS projection made predictions based on employment change in time series. And employment decline may be caused by various reasons e.g., supply and demand change [6], other than automation technologies which mean the declining employment of “customer service representatives” can be caused by the demand change, i.e. some the customers’ basic demands can be satisfied by AI dialogue agents. That is to say, expert assessment based on occupations’ attributes only can be biased since some occupations may still decline even if they have tasks and skills that automation technologies cannot perform because of the supply and demand change. Overall, the performance and evaluation demonstrate the effectiveness of our AOC-GCN method in occupation’s automated risk prediction.

6 Limitations and Future Work

We acknowledge several limitations in our method and the results presented in this paper. Firstly, data limitations hindered the performance of our model. O*NET is the only globally admitted database in labour market, however, it only classified 1016 occupations, and this data size limited the performance of classification using machine-learning method. Moreover, each occupation is an aggregation of several jobs which still obfuscates the difference between jobs. Secondly, lack of ground truth data is also a barrier that inhibits predictions of automated risk. Currently, most of the ground truths are expert assessments

based on occupation’s attributes which are inevitably subjective to some extent. Expert assessments based on occupations’ attributes can be biased without considering the supply and demand change. Moreover, there is no job-level or skill-level ground truth data which impedes predictions on a more granular level. Thirdly, predicting on static information cannot capture the technological change and labor trends since the employment trends and changing demand for specific tasks and skills might change faster than static O*NET data can capture. Therefore, our future work will focus on automated risk identifications on job-level which requires more granular ground truth data such as which skills or tasks decide a specific occupation being automated [19]. Moreover, the study will be conducted on real-world datasets that capture technological change, employment trends and supply and demand change.

7 Conclusion

In this paper, we applied the AOC-GCN model to the Occupation-skill graph to identify the automated risk of a specific occupation. The initial node embeddings are generated from Word2vec and Doc2vec. GCN learns the graph structure and interactions between occupations. Extensive experiments and comparisons with real-world projections demonstrate the effectiveness of our method. Our study paves the way for understanding where to concentrate upskilling efforts in the following decades [1]. In the future, we will further extend our model to solve more granular-level tasks in labour market.

References

1. Ahadi, A., Kitto, K., Rizoiu, M.A., Musial, K. (2022) Skills Taught vs Skills Sought: Using Skills Analytics to Identify the Gaps between Curriculum and Job Markets. In: International Conference on Educational Data Mining. pp. 538–542. Durham, UK <https://doi.org/10.5281/zenodo.6853121>
2. Frey, Carl & Osborne, Michael. (2013). The Future of Employment: How Susceptible Are Jobs to Computerisation?. Oxford Martin. 114.<https://doi.org/10.1016/j.techfore.2016.08.019>.
3. Arntz, Melanie & Gregory, Terry & Zierahn, Ulrich. (2017). Revisiting the Risk of Automation. Economics Letters. 159. 157-160.<https://doi.org/10.1016/j.econlet.2017.07.001>.
4. Frank, Morgan & Autor, David & Bessen, James & Brynjolfsson, Erik & Cebrian, Manuel & Deming, David & Feldman, Maryann & Groh, Matthew & Lobo, José & Moro, Esteban & Wang, Dashun & Youn, Hyejin & Rahwan, Iyad. (2019). Toward understanding the impact of artificial intelligence on labor. Proceedings of the National Academy of Sciences of the United States of America. 116. 6531-6539.<https://doi.org/10.1073/pnas.1900949116>.
5. Cortes, Matias & Jaimovich, Nir & Siu, Henry. (2017). Disappearing Routine Jobs: Who, How, and Why?. Journal of Monetary Economics. 91.<https://doi.org/10.1016/j.jmoneco.2017.09.006>.

6. Vermeulen, Ben & Pyka, Andreas & Saviotti, Pier-Paolo & Kesselhut, Jan. (2018). The Impact of Automation on Employment: Just the Usual Structural Change?. Sustainability. 10.<https://doi.org/10.3390/su10051661>.
7. Kipf, Thomas & Welling, Max. (2016). Semi-Supervised Classification with Graph Convolutional Networks. <https://doi.org/https://doi.org/10.48550/arXiv.1609.02907>
8. Nakamura, A. O., Shaw, K. L., Freeman, R. B., Nakamura, E., & Pyman, A. (2009). Jobs online. In Studies of labor market intermediation (pp. 27-65). University of Chicago Press. <https://doi.org/10.7208/chicago/9780226032900.003.0002>
9. Brzeski C, Burk C (2015) Die Roboter kommen: Folgen der Automatisierung für den deutschen Arbeitsmarkt [The Robots are coming: Consequences of Automation for the German Labor Market]. *INGDiBa Econ Res* 30:1–7. German.
10. Pajarinen M, Rouvinen P, Ekeland A (2015) Computerization threatens one-third of Finnish and Norwegian employment. *ETLA Brief* 34. Available at <https://www.etla.fi/wp-content/uploads/ETLA-Muistio-Brief-34.pdf>.
11. Bowles J (July 24, 2014) The computerisation of European jobs. Bruegel blog. Available at <https://www.bruegel.org/2014/07/the-computerisation-of-european-jobs>.
12. Arntz, M.; Gregory, T.; Zierahn U. The Risk of Automation for Jobs in OECD Countries: A Comparative Analysis; OECD Social, Employment and Migration Working Papers; OECD Publishing: Paris, France, 2016. <https://doi.org/https://doi.org/10.1787/1815199X>
13. Acemoglu D, Restrepo P (2017) Robots and jobs: Evidence from US labor markets. NBER Working Paper 23285 (National Bureau of Economic Research, Cambridge, MA). <https://doi.org/10.3386/w23285>
14. Li, Ao & Qin, Zhou & Liu, Runshi & Yang, Yiqun & Li, Dong. (2019). Spam Review Detection with Graph Convolutional Networks. 2703-2711. <https://doi.org/10.1145/3357384.3357820>.
15. Aditya Grover and Jure Leskovec. 2016. node2vec: Scalable feature learning for networks. In Proceedings of the 22nd ACM international conference on Knowledge discovery and data mining, SIGKDD. ACM, 855–864. <https://doi.org/1607.00653>
16. Dawson, N., Molitorisz, S., Rizoiu, M.A., Fray, P.: Layoffs, inequity and COVID-19: A longitudinal study of the journalism jobs crisis in Australia from 2012 to 2020. *Journalism* p. 146488492199628 (feb 2021). <https://doi.org/10.1177/1464884921996286>
17. Dawson, N., Rizoiu, M.A., Johnston, B., Williams, M.A.: Adaptively selecting occupations to detect skill shortages from online job ads. In: 2019 IEEE International Conference on Big Data (Big Data). pp. 1637–1643. IEEE, Los Angeles, CA, USA (dec 2019). <https://doi.org/10.1109/BigData47090.2019.9005967>
18. Dawson, N., Rizoiu, M.A., Johnston, B., Williams, M.A.: Predicting Skill Shortages in Labor Markets: A Machine Learning Approach. In: Proceedings - 2020 IEEE International Conference on Big Data, Big Data 2020. pp. 3052–3061. University of Technology Sydney, IEEE (dec 2020). <https://doi.org/10.1109/BigData50022.2020.9377773>
19. Dawson, N., Williams, M.A., Rizoiu, M.A.: Skill-driven recommendations for job transition pathways. *PLOS ONE* **16**(8), e0254722 (aug 2021). <https://doi.org/10.1371/journal.pone.0254722>
20. Bryan Perozzi, Rami Al-Rfou, and Steven Skiena. 2014. Deepwalk: Online learning of social representations. In Proceedings of the 20th ACM international conference on Knowledge discovery and data mining, SIGKDD. ACM, 701–710. <https://doi.org/1403.6652>

21. Will Hamilton, Zhitao Ying, and Jure Leskovec. 2017. Inductive representation learning on large graphs. In *Advances in Neural Information Processing Systems*, NIPS. 1024–1034. <https://doi.org/1706.02216>
22. LeCun, Y., Haffner, P., Bottou, L., & Bengio, Y. (1999). Object recognition with gradient-based learning. In *Shape, contour and grouping in computer vision* (pp. 319–345). Springer, Berlin, Heidelberg. <http://yann.lecun.com/exdb/publis/pdf/lecun-99.pdf>
23. Mihaita, A.S., Li, H., He, Z., Rizoiu, M.A.: Motorway Traffic Flow Prediction using Advanced Deep Learning. In: 2019 IEEE Intelligent Transportation Systems Conference (ITSC). pp. 1683–1690. IEEE, Auckland, New Zealand (oct 2019). <https://doi.org/10.1109/ITSC.2019.8916852>
24. Mihaita, A.S., Liu, Z., Cai, C., Rizoiu, M.A.: Arterial incident duration prediction using a bi-level framework of extreme gradient-tree boosting. In: *Proceedings of the 26th ITS World Congress*. pp. 1–12. Singapore (may 2019)
25. Mihaita, A.S., Papachatgis, Z., Rizoiu, M.A.: Graph modelling approaches for motorway traffic flow prediction. In: 2020 IEEE 23rd International Conference on Intelligent Transportation Systems (ITSC). pp. 1–8. IEEE, Rhodes, Greece (sep 2020). <https://doi.org/10.1109/ITSC45102.2020.9294744>
26. Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua Bengio. 2018. Graph Attention Networks. *International Conference on Learning Representations, ICLR (2018)*. <https://doi.org/1710.10903>
27. Mikolov, Tomas; et al. (2013). "Efficient Estimation of Word Representations in Vector Space" <https://doi.org/1301.3781>
28. Le, Quoc & Mikolov, Tomas. (2014). Distributed Representations of Sentences and Documents. 31st International Conference on Machine Learning, ICML 2014. 4. <https://doi.org/1405.4053>
29. Freund, Y., & Schapire, R.E. (1996). Experiments with a New Boosting Algorithm. *ICML*. <https://dl.acm.org/doi/10.5555/3091696.3091715>
30. Ho, Tin Kam (1995). Random Decision Forests (PDF). *Proceedings of the 3rd International Conference on Document Analysis and Recognition*, Montreal, QC, 14–16 August 1995. pp. 278–282. Archived from the original (PDF) on 17 April 2016. Retrieved 5 June 2016. <https://doi.org/10.1109/ICDAR.1995.598994>
31. Meng, Qi. (2018). LightGBM: A Highly Efficient Gradient Boosting Decision Tree. <https://proceedings.neurips.cc/paper/2017/file/6449f44a102fde848669bdd9eb6b76fa-Paper.pdf>
32. Quinlan, J. R. (1986). "Induction of decision trees" (PDF). *Machine Learning*. 1: 81–106. <https://doi.org/10.1007/BF00116251>. S2CID 189902138.
33. Roweis, Sam; Hinton, Geoffrey (January 2002). Stochastic neighbor embedding (PDF). *Neural Information Processing Systems*. <https://www.cs.toronto.edu/~hinton/absps/sne.pdf>
34. Hotelling H (1933) Analysis of a complex of statistical variables into principal components. *J Educ Psychol* 24:417–441, 498–520 <https://doi.org/https://doi.org/10.1037/h0071325>
35. Beaudry P, Green DA, Sand BM (2016) The great reversal in the demand for skill and cognitive tasks. *J Labor Econ* 34:S199–S247.