

Mind the Task Gap: Unsupervised Skill–Task Link Prediction for Workforce Upskilling^{*}

Yee Sen TAN^{1,2}, Daryl LOW^{1,*}, Eugene Chua¹, Alejandro SEIF², Leo LI^{1,2} and Lois JI^{1,2}

¹Skills Development Group, SkillsFuture Singapore, S408533, Singapore

²Artificial Intelligence Practice, Government Technology Agency, S117438, Singapore

Abstract

Rapid advances in emerging technology continue to reshape how tasks appear, disappear and mutate inside modern jobs, while formal job–skill taxonomies refresh only intermittently. To capture this dynamism, we introduce a semantically enriched knowledge graph for Singapore’s labor ecosystem that unifies 1869 job roles, 27159 tasks, 3100 skills, 28313 accredited courses and their sector context. The principal data gap is the lack of explicit Skill–Task links, so we cast their recovery as an unsupervised graph-completion problem. Using 50,000 GPT-4o pseudo-labelled pairs, a Variational Graph Autoencoder with GCN encoders attains a macro-F1 of 0.6039, substantially ahead of the semantic-similarity baseline method, demonstrating effective cold-edge discovery. To validate the pseudo labels, we randomly sampled 500 pairs and found over 90% were consistent with our internal taxonomy, although they are not a substitute for full human-verified ground truth. The enriched knowledge graph offers a more complete view of the labor market, supporting workforce planning and policy. It underpins tools like a Task–Skill radar for policymakers and a bridge-skill career-path recommender for workers, providing stakeholders with clear, data-driven insights for decision-making.

Keywords

Knowledge Graph, Labor Market, Link Prediction, Graph Neural Networks

1. Introduction

The job market is rapidly evolving, with shifting roles and growing demand for diverse skills. Technological advances accelerated by COVID-19—have intensified this trend [1, 2], giving rise to new skills such as prompt engineering, now essential in many careers with the advent of LLMs like ChatGPT and Gemini [3]. With the rapid pace at which job requirements are changing, there is a need for a unified representation of the current job-skills market.

At the same time, SkillsFuture Singapore (SSG)’s Singapore Skills Framework[4] lists more than twenty-seven thousand recognised tasks but contains no explicit mapping from those tasks to the skills that actually enable them, limiting its ability to model how skills are applied in real job functions. This gap directly prevents AI-readiness studies and conventional manpower planning alike from translating job-level insight into concrete training action.

As the job-skills market changes and with individuals seeking to upskill or reskill, the huge amount of information from various resources can still leave them uncertain about which skills are truly essential for specific occupations. In response, researchers have suggested building a knowledge graph (KG) that connects information such as skills, job titles, and online courses [1, 5]. This would create a structured representation of the relationships between competencies, roles, and learning resources, offering a comprehensive view of the skill ecosystem in the labor market.

While most existing works focus on linking skills to jobs, they often overlook the crucial role of tasks, which define what a job actually entails. There is a lack of structured connections between skills, jobs, and tasks, which limits

our understanding of how skills are applied in real job functions. As such, this research addresses the lack of existing links between skills, jobs, and tasks connections that are largely missing. To bridge this gap, we employ graph-based methodologies to better represent and infer these hidden relationships to enable a more complete and data-driven understanding of the labor market.

Our key contributions to this research are as follows:

- We aggregate multiple labor market data sources into a large-scale, semantically enriched knowledge graph, integrating job roles, tasks, skills, courses, and sectors to create a unified representation of Singapore’s labor market.
- We address the problem of inferring missing connections between tasks and skills by formulating it as a graph completion task, leveraging unsupervised link prediction techniques to enhance the knowledge graph and demonstrating superior performance over baseline approaches.
- We introduce the first labor market knowledge graph that explicitly incorporates task information interconnected with both job roles and skills, providing a more comprehensive representation of workforce dynamics, allowing stronger analysis of the labor market
- Establishing skill-to-task connections in our knowledge graph enables policy teams to use graph-based tools to uncover hidden skill gaps and target funding more effectively through data-driven audits of emerging competency needs across industries.

2. Related works

2.1. Knowledge Graphs

A knowledge graph (KG) is a structured representation of real-world entities and their relationships, offering a valuable external data source that enhances the model learning process [6]. With the vast amounts of data available, KGs are gaining popularity due to their ability to provide semantic and conceptual representations, which makes them

RecSys in HR’25: The 5th Workshop on Recommender Systems for Human Resources, in conjunction with the 19th ACM Conference on Recommender Systems, September 22–26, 2025, Prague, Czech Republic.

*Corresponding author

✉ ystan98@hotmail.com (Y. S. TAN); daryl_low@ssg.gov.sg (D. LOW); eugene_chua@ssg.gov.sg (E. Chua); alejandro_seif@tech.gov.sg (A. SEIF); Leo_LI@tech.gov.sg (L. LI); Lois_JI@tech.gov.sg (L. JI)
ID 0009-0007-4373-9824 (Y. S. TAN); 0009-0004-9690-6567 (D. LOW); 0009-0008-3201-2790 (E. Chua); 0000-0003-1275-5438 (A. SEIF); 0009-0008-1641-137X (L. LI); 0009-0005-4741-6112 (L. JI)



© 2025 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

fundamental for structuring knowledge [7], being applied across various fields, from finance to sentiment analysis and to labor market studies [5, 8].

Recent research has specifically highlighted the use of KGs to map connections between skills, job roles, and learning resources, providing a comprehensive view of the job-skills ecosystem in the labor market [1, 5]. One of the key advantages of graph based approaches is their ability to connect and represent data from diverse domains cohesively. Studies have also shown that incorporating additional information, such as multimodal data or data from various sources, can enhance contextual understanding and lead to more robust data representations [9, 10, 11, 12]. In the case of knowledge graphs, integrating various modalities or data from different sources improves the quality and informativeness of the knowledge representation. For job-skills knowledge graphs, this can include incorporating sector-specific information, key tasks from job postings, and other relevant data.

2.2. Knowledge Graphs for Labor Market Representation

Researchers have proposed Job-Skills Knowledge Graph (JSKG) for organizations to answer labor market related questions and analysis [5]. In their work, they used job data on a yearly basis, aggregating job posting data annually and calculating the median values of the salary. Apart from jobs data, including industry or sector data in analysis is crucial, as it not only highlights how sectoral changes impact the labor market and broader economy, but also ensures vital information is not lost by focusing solely on jobs data [13]. Without sector data, it will also be impossible to analyze the evolution of job across the industries overtime, impacting the ability to predict skill demand and workforce trends effectively.

One notable gap in the existing literature is that most research has predominantly focused on the relationship between jobs and skills, typically do not include a critical intermediate layer: tasks. Tasks serve as the functional link between job roles and the specific skills required to perform them. By omitting tasks, analyses risk oversimplifying the structure of work, as they fail to account for how skills are operationalized within specific job functions.

Furthermore, incorporating tasks into labor market models provides a more nuanced framework for analyzing changes in employment and earnings, as these are shaped by the interaction of worker skills, job tasks, and evolving technologies, thereby enabling better alignment between workforce capabilities and industry demands [14].

2.3. Link Prediction in Knowledge Graphs for Job-Skills Representation

Most labor-market graphs link jobs directly to skills, leaving tasks as unstructured notes or buried within job descriptions [15, 16]. By modelling tasks as first-class nodes and inferring Skill-Task edges, our study addresses this critical blind spot—particularly urgent in an era where AI is rapidly shifting the skill demands of many occupations.

Despite the growing use of skills taxonomies in labor market analysis, existing frameworks rarely establish explicit, structured connections between tasks and skills, which are typically left as unstructured text or embedded within

job descriptions. To address this structural gap, we frame Skill-Task linkage as a graph-based link prediction problem, which introduces unique challenges for evaluation due to the lack of labeled ground truth.

Evaluating link prediction in knowledge graphs typically relies on well-established supervised metrics such as Precision@k, Recall@k, Mean Reciprocal Rank (MRR), and Mean Average Precision (MAP), particularly when a labeled ground truth is available [17, 18]. However, when annotated links are sparse or nonexistent, evaluation becomes more challenging. In such cases, researchers often turn to proxy methods, including semantic similarity via cosine distance in embedding space [19], or graph embedding distance metrics, where node proximity is derived from structural encodings such as TransE, Node2Vec, or GraphSAGE [20, 21]. These approaches allow for inference of link plausibility based on learned representations, even in the absence of labeled links. Earlier works also proposed graph-based autoencoder architectures such as Graph Autoencoders (GAE) and Variational Graph Autoencoders (VGAE) for unsupervised link prediction tasks [22] and training with reconstruction loss. Additionally, human-in-the-loop evaluations, such as expert assessments or structured prompts using large language models (LLMs), have gained traction as reliable complements, particularly in applied settings like labor market modeling [23].

In our context, we propose an evaluation strategy using LLM-generated pseudo labels to assess Skill-Task pairings, with cosine similarity between pairs serving as a baseline to benchmark graph-based link predictions.

3. Data Sources

In this paper, we utilize data from various sources, including SkillsFuture Singapore (SSG) and job aggregator platforms. Additionally, we have developed preprocessing tools to extract and connect crucial information across different labor market datasets.

3.1. SkillsFuture’s Singapore Skills Framework (SFw)

The SkillsFuture’s Singapore Skills Framework [4] serves as a comprehensive knowledge base that provides essential labor market information, including job roles in different sectors, lists of skills (along with relevant applications and tools), and mappings that link job roles to critical work functions and tasks. In our research, the majority of our data is sourced from this framework, specifically from the Job Role, Skills, and Task datasets. Altogether, we collected 1,869 unique job roles, 27,159 tasks, and 3,100 skills. Each job role, task, and skill is accompanied by a textual description, which forms the primary basis for our data features.

3.2. Course Data

For Course data, the data source is readily available on SkillsFuture Singapore web portal, where each course comes with a course description, with other essential information such as the course fees and duration of the course. The courses were already pre-tagged internally during creation of the courses by the course providers and verified by the in-house experts themselves. During the time of this research, the data set comprises a total of 28,313 courses.

3.3. Job Postings Data

Job postings data will be the main source of data to represent the labor market in Singapore. Job postings, often found in Job Portals like MyCareersFuture, JobStreet, LinkedIn, provides views on the labor market, though, it is only a proxy as it does not truly represent the entire labor market economy. Leveraging data acquired from a job aggregation service, we obtained over 20M job postings from the years 2019 to 2025. Each of these job posting obtained is also accompanied with the company name, job title, textual description, and salary information. Using the respective company names, we will map them to their corresponding industries based on an company-to-sector mapping, done by our in-house manual labelers well versed with the domain.

3.4. Additional Tools to link Job Market Data to Singapore Skills Framework (SFw)

Although we obtained data from various source such as job roles data, skills data, and job postings data, deriving accurate and meaningful insights from these sources individually posed several challenges. In particular, job postings often consist of a job title accompanied by an unstructured text description, with no explicit mention of the required skills or the SFw job role.

For example, a posting titled "Data Analyst" may describe responsibilities such as "Perform analysis using Tableau and provide recommendations based on data insights", yet it may omit explicit mention of the skill "Data Visualization", which is defined in the SFw.

In another instance, a job posting with the same title might mention tasks such as "fine-tuning large language models, applying deep learning, and developing ensemble methods", which actually align better with the SFw defined role of "Data Scientist".

To preprocess the data and extract insights aligned with our research objectives, we employed the skills extraction methodology from [24]. This algorithm processes textual job descriptions and maps the identified skills to those recognized in the SkillsFuture taxonomy, producing a structured list of skills standardized to the SF framework.

In addition, we used an internal tool, a job role classification tool with, which classifies job postings to their corresponding SFw recognized job roles instead of skills. These tools enabled robust alignment between job postings and data sources from the SFw, providing a coherent and enriched data set for our knowledge graph creation and subsequent analysis.

4. Our Methodology

4.1. Knowledge Graph Construction

Using the data and information available, we construct a structured knowledge graph to represent relationships between key labor market entities, including **Skills**, **Tasks**, **Courses**, **Job Roles**, and **Sectors**. Each entity is modeled as a node with various relevant attributes (e.g., name, description, or salary information, while edges encode typed relationships with weights to represent their strength or relevance. Note that for each of these nodes, the textual description are represented as embeddings, computing using the all-base-mpnet-v2 model [25]. Note that in our




Figure 1: Overview of the knowledge graph nodes and edges model. Note that the linkage between Skill and Task will be predicted, whereas the other edges are ground truth.

research and experiments for link prediction, populate our graph with data from a single year, using snapshots taken over an arbitrary 12-month period (2023 September to 2024 September).

The graph includes several core relationships as seen in Figure 1. Such relation includes **REQUIRES**, linking job roles to skills, **TEACHES** which connects courses to the skills they impart, **REQUIRED IN** to understand the different JobRoles required in different sectors, **DELIVERS** to represent the required tasks for each job role. **SIMILAR TO** is a symmetric relation modeling semantic or functional similarity between nodes of the same type (e.g. task-to-task or skill-to-skill), derived based on its semantic similarity. Finally, the edges denoted in red dotted line, **PERFORMS**, would be our key aim for our experiments, to infer the potential edges between skill and tasks, telling us what task can be performed upon knowing each skill. Each edge is also associated with a `float` weight representing either empirical relevance derived by the SFw and in-house subject matter experts or derived by semantic similarity. Overall, our graph schema is setup to enable downstream tasks such as link prediction, graph analytics, and recommendation for job changes or courses for upskilling.

4.1.1. Deriving relationship between Skills and Job Roles

In deriving the relationship between skills and job roles, we follow the methodology proposed in [5], which combines expert-defined knowledge (done by SkillsFuture Singapore's labor economist) from the Singapore Skills Framework (SFw) with job posting data to estimate the importance of each skill for a given job. To maintain relevance and reduce noise, we retain only the top 10 highest-ranked skills per job role, as unfiltered job postings may list over 100 skills. Note that these 10 skills could be further tuned and optimized for. While the original approach aggregates and computes edge scores at the Singapore Standard Occupational Classification (SSOC) level [5], we adapt it by aggregating at the more granular level of *Job Role* instead, keeping all other aspects unchanged. The formulas used are shown below where s indicates *skill*, JR indicates *Job Role*, y indicates *year* and Δy indicates the size in years of the window to analyze:

$$R_{LM}(s, JR, y) = R \left(\sum_{y=y_{\max}-\Delta y}^{y_{\max}} \frac{(1 + y_{\max} - y)^{-\frac{1}{2}}}{R(s, JR, y)} \right) \quad (1)$$

$$W(s, JR) = \frac{\epsilon_{SFw}}{R_{SFw}(s, JR)} + \frac{1}{R_{LM}(s, JR, y)} \quad (2)$$

Once $W(s, JR)$ is calculated for each skill–JobRole pair, we rank the results based on their corresponding W values, keeping only the top 10 scores. This informs the most important required skills for each Job Role.

4.1.2. Deriving relationship between Course-Skill, JobRole-Task, edges towards Sector

The derivation of course-skill and jobrole-task relationships are straightforward, as the Singapore Skills Framework already encapsulates this information. This framework is collaboratively developed by a diverse group of stakeholders, including employers, industry associations, educational institutions, unions, and government agencies, specifically for the Singapore workforce.

For the edges pointing to Sector, the information is also self-contained, as the job dataset we receive already includes sector annotations for each job posting, provided by human annotators. Therefore, to derive the relationship between each JobRole and Sector, we compute both the count and proportion of each job role within each sector. These metrics capture the relative prevalence of a role in a sector, enabling the edge weights between JobRole and Sector to reflect both absolute frequency and sector-specific importance.

4.1.3. Deriving SIMILAR TO relationships

The derivation of all the SIMILAR_TO relationships is also straightforward. Since these edges represent similarity between nodes of the same type and each node includes a textual description, we use the all-mpnet-base-v2 model [25] to compute embeddings. Cosine similarity between these embeddings is then used to quantify how similar the skills are. We opt for cosine similarity over Jaccard because the latter is limited to exact token overlaps. For example, while "Deep Learning" and "Machine Learning" are conceptually related, Jaccard similarity may fail to capture this connection.

4.2. Knowledge Graph Completion

After constructing the knowledge graph, we infer missing Skill–Task links to enrich the graph and support workforce analytics, recommendations, and labor market predictions. We leverage unsupervised techniques, particularly Graph Autoencoders (GAE) and Variational Graph Autoencoders (VGAE), which have proven effective in learning meaningful latent representations without requiring labeled edges. These models encode nodes into a continuous latent space by capturing both node features and graph topology, and then decode these embeddings to reconstruct the original graph structure, optimizing a reconstruction loss that encourages the accurate prediction of edges.

We will integrate three different encoder architectures, to investigate their respective impacts on embedding quality and link prediction performance. Specifically, the following:

- **GraphSAGE** aggregates neighborhood information through sampling strategies, extended with attention mechanisms to dynamically weigh neighbor importance.
- **Graph Attention Networks (GAT)** employ self-attention to assign learnable weights to neighbors, enhancing the model’s ability to focus on relevant graph regions.

- **Graph Convolutional Networks (GCN)** utilize spectral convolutions to propagate and aggregate node features over graph structures.

To address the lack of labeled edges in job-skills graphs, we generate pseudo labels using GPT-4o, which classifies whether a given skill is relevant to a specific task, as illustrated below.

LLM Skill–Task Classification Prompt

Given the following **Skill** and **Task**, determine whether the skill is relevant to the task and whether it can help accomplish it.

- Respond with **Yes** or **No** for relevance.
- Provide a brief explanation (1–2 sentences) supporting your answer.

Skill: [Insert skill here]

Task: [Insert task here]

This enables a quantitative assessment of link plausibility between Skill and Task nodes using standard classification metrics (Accuracy, Precision, Recall, Macro F1 score). Given the large number of Skill→Task pairs (>80M), generating pseudo labels for the entire set would be prohibitively expensive. Therefore, we randomly sampled 50,000 pairs for pseudo-labeling and evaluation. We note that these pseudo labels are used solely for relative method comparison and do not represent absolute correctness, since a different LLM could provide entirely different result. However, in a smaller validation sample of 500 pairs, we found that over 90% of the pseudo labels aligned with our expectations after reviewing, suggesting that the labels are reasonably reliable.

5. Experiment Setup and Evaluation

Our experiment will involve comparing the performance of GAE and VGAE, as well as the different encoders (GraphSage, GraphSage with attention, Graph Attention Networks, and Graph Convolutionary Networks). Additionally, as a baseline, we incorporate a semantic similarity method based on the all-mpnet-base-v2 [25], computing cosine similarity between Skill and Task embeddings. This baseline facilitates direct comparison against graph-based approaches.

For all experiments, the dataset edges were split into training and test sets using a 20% test ratio with `train_test_split_edges`. The model was trained on a single T4 GPU with 768-dimensional embeddings using the AdamW optimizer (learning rate 0.001, weight decay 1×10^{-4}) and using reconstruction loss to measure how well the model reconstructs and predict from the learned representation. Edge dropout ($p = 0.3$) was applied for regularization, and training employed early stopping with a patience of 20 epochs and a maximum of 2000 epochs.

For evaluation, we use the same standard metrics across all experiments, namely, Accuracy, Precision, Recall, and Macro F1 score on the same fixed dataset of 50,000 samples which is pseudo-labelled by GPT-4o. To address class imbalance, we also tune each model to maximize the Macro F1 score, thereby promoting balanced performance across classes. Note that we did not consider the Area Under the Curve (AUC) scores due to its limitations in link prediction and inapplicability for imbalanced tasks [26]. Furthermore,

Table 1Performance Comparison of Graph-Based Methods. *Note:* Values are reported after optimizing for Macro F1 score.

Method	Encoder	Auto Encoder	Accuracy	Precision	Recall	Macro F1
Baseline	all-base-mpnet-v2 [25]	-	0.6687	0.9913	0.0168	0.0331
Graph Based	GraphSAGE[27] + Attention	VGAE	0.7229	0.5931	0.5792	0.5861
		GAE	0.6964	0.5600	0.4830	0.5187
Graph Based	GraphSAGE[27]	VGAE	0.6911	0.5367	0.6425	0.5848
		GAE	0.6602	0.4985	0.5740	0.5336
Graph Based	GAT[28]	VGAE	0.7294	0.6067	0.5709	0.5883
		GAE	0.6834	0.5286	0.6031	0.5634
Graph Based	GCN [29]	VGAE	0.7345	0.6102	0.5977	0.6039
		GAE	0.6498	0.4806	0.4226	0.4497

AUC focuses solely on ranking and averages performance across all thresholds, ignoring probability calibration, practical decision boundaries, and unequal costs of false positives and false negatives. Given the highly imbalanced nature of link prediction and the need for accurate, actionable probability estimates, we instead rely on the proposed metrics to better capture real world performance and decision making relevance. The full results of these experiments, based on the most optimal thresholds, are summarized in Table 1.

The results indicate that the baseline method, which computes cosine similarity between each Skill and Task pair, achieves very high precision. However, this comes at the cost of extremely low recall, resulting in a low overall Macro F1 score of 0.0331. While this baseline method excels in identifying only highly confident positive matches (hence the extremely high precision), it fails to capture a large proportion of true positives, as indicated by the very low recall. This indicates that the baseline is overly conservative in predicting positive matches, resulting in a poor F1 score despite high accuracy (0.6687). The outcome illustrates the typical pitfall of high precision but negligible recall, which is particularly undesirable in imbalanced classification problems where missing positives is costly.

For graph-based approaches, all experiments showed that it substantially outperform the baseline in terms of Macro F1 score, demonstrating their effectiveness for link prediction given its ability to understand and connect the relationships between skills and tasks, probably due to the middle node, the job role node, which connects skills and tasks in the middle. Notably, architectures incorporating VGAE consistently outperform their deterministic GAE counterparts. For example, GCN-VGAE achieves an F1 score of 0.6039 compared to 0.4497 for GCN-GAE; GAT-VGAE scores 0.5883 vs. 0.5634 for GAT-GAE; and GraphSAGE + Attention-VGAE reaches 0.5861 compared to 0.5187 for the corresponding GAE model. This consistent improvement suggests that VGAE’s probabilistic framework capturing uncertainty in latent representations provides better regularization and generalization, especially in the presence of class imbalance.

Overall, GCN-VGAE delivered the strongest overall performance, with the highest accuracy (0.7345), precision (0.6102), and F1 (0.6039), while GAT-VGAE and GraphSAGE + Attention-VGAE also performed well, though gains from attention mechanisms were modest.

6. Downstream Real-World Applications of the Enriched Graph

Establishing links between tasks and skills within our knowledge graph potentially enables a range of downstream applications for labor market stakeholders and policy analysts at SSG. Such graph-enabled tools may facilitate the translation of structural insights into actionable strategies for workforce planning, career guidance, and training optimization.

The knowledge graph enables several applications: a Task–Skill Radar highlights tasks with weak skill coverage and relevant subsidized courses; a Bridge-Skill Career Paths recommender identifies high-confidence skill sequences for career transitions, including AI-specific skills; a Course-Portfolio Optimiser flags courses misaligned with current task clusters; and a Task Centrality Monitor alerts agencies when critical tasks lack sufficient skill coverage. These tools support targeted reskilling, portfolio alignment, and policy interventions, complementing national initiatives like Skills-Future subsidies and Career Conversion Programmes, while allowing granular monitoring of emerging trends such as AI-driven task shifts.

In sum, the enriched knowledge graph represents not only a technical advancement in labor market modeling but also a practical foundation for informed decision-making across the workforce ecosystem.

7. Conclusion

In this study, we proposed a graphical method to represent Singapore’s labor market, and showed how a graphical method can be used for enriching a labor market knowledge graph, specifically, by using an unsupervised link prediction method. By leveraging a well defined schema grounded in real world workforce data, we constructed a rich knowledge graph and applied both GAE and VGAE across several encoder architectures (GCN, GAT, GraphSAGE).

Our link prediction experiments showed that graph-based models outperformed the baseline in Macro F1 score, highlighting their strength under class imbalance. The GCN-VGAE architecture achieved the best overall results. Including the job role node as an intermediary improved relational learning, helping models better link skills to tasks. These findings demonstrate the value of graph-based approaches for labor market intelligence and set the foundation for further analysis to support data-driven policy decisions.

7.1. Limitations and Future Works

The evaluation was performed on a sample of 50,000 Skill→Task pairs, which constitutes less than 5% of the entire dataset encompassing all Skill-Task combinations. Future work could increase the number of samples and employ stratified sampling to balance representation across sectors, or analyze specific sectors individually to better reflect sector-specific patterns.

Another key limitation of our study is that evaluation relies on GPT-4o-generated pseudo labels rather than human-annotated ground truth. Consequently, our results reflect relative method performance under proxy labels and do not provide an absolute measure of correctness. Moving forward, we aim to incorporate human-annotated ground truth labels between skill and task nodes, enabling the exploration of supervised learning approaches that may yield improved performance. Additionally, increasing the sample size in future evaluations will help ensure greater robustness and better generalization of the findings.

Lastly, future work could improve the baseline's low recall by treating the number of retained skills per job as a hyperparameter to tune. Additionally, its performance may depend on the choice of LLM for generating pseudo-ground truth, suggesting that experimenting with different models or hybrid labeling approaches could capture more true positives without sacrificing precision.

8. Acknowledgments

The authors acknowledge the support of SkillsFuture Singapore and Government Technology Agency of Singapore for their support in this research.

Declaration on Generative AI

During the preparation of this work, the author(s) used GPT-4o for evaluation purposes as highlighted above. The author(s) also used ChatGPT Free Tier for grammar and spelling checks for this paper, which the author(s) further reviewed, and take full responsibility for the publication's content.

References

- [1] Y. Fettach, A. Bahaj, M. Ghogho, Jobedkg: An uncertain knowledge graph-based approach for recommending online courses and predicting in-demand skills based on career choices, *Engineering Applications of Artificial Intelligence* 131 (2024) 107779.
- [2] L. Borghans, B. Ter Weel, B. A. Weinberg, People skills and the labor-market outcomes of underrepresented groups, *Ilr Review* 67 (2014) 287–334.
- [3] P. Bansal, Prompt engineering importance and applicability with generative ai, *Journal of Computer and Communications* 12 (2024) 14–23.
- [4] SkillsFuture Singapore, Skills framework, n.d. URL: <https://www.skillsfuture.gov.sg/skills-framework>, accessed: 22-Jul-2025.
- [5] A. Seif, S. Toh, H. K. Lee, A dynamic jobs-skills knowledge graph (2024).
- [6] S. Ji, S. Pan, E. Cambria, P. Marttinen, S. Y. Philip, A survey on knowledge graphs: Representation, acquisition, and applications, *IEEE transactions on neural networks and learning systems* 33 (2021) 494–514.
- [7] X. Zou, A survey on application of knowledge graph, in: *Journal of Physics: Conference Series*, volume 1487, IOP Publishing, 2020, p. 012016.
- [8] Z. Hu, G. Yan, Y. S. Tan, H. Pen, Stock price prediction using graph attention networks with sentiment analysis, in: *Pacific-Asia Conference on Knowledge Discovery and Data Mining*, Springer, 2025, pp. 399–408.
- [9] J. Lee, Y. Wang, J. Li, M. Zhang, Multimodal reasoning with multimodal knowledge graph, *arXiv preprint arXiv:2406.02030* (2024).
- [10] Y. S. Tan, N. Teo, E. Ghe, J. Fong, Z. Wang, Video sentiment analysis for child safety, in: *2023 International Conference on Data Mining, ICDM: Shanghai, December, IEEE, 2023*, pp. 783–790.
- [11] Y. S. Tan, Z. Wang, E. Cambria, Explainable multimodal sentiment analysis of social media visual content for child safety, *IEEE Intelligent Systems* 40 (2025) 64–72.
- [12] Y. Lu, W. Zhao, N. Sun, J. Wang, Enhancing multimodal knowledge graph representation learning through triple contrastive learning, in: *Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence, 2024*, pp. 5963–5971.
- [13] A. Pacana, D. Siwiec, An analysis of the changes of employment in industry, *Humanities and Social Sciences* 26 (2019) 121–130.
- [14] D. Acemoglu, D. Autor, Chapter 12-skills, tasks and technologies: Implications for employment and earnings (d. card & o. ashenfelter, eds.), Elsevier. [https://doi.org/10.1016/S0169-7218\(11\)\(2011\)02410-5](https://doi.org/10.1016/S0169-7218(11)(2011)02410-5).
- [15] E. Atalay, P. Phongthiengtham, S. Sotelo, D. Tannenbaum, The evolution of work in the united states, *American Economic Journal: Applied Economics* 12 (2020) 1–34.
- [16] M. Sigelman, B. Taska, L. O’Kane, J. Nitschke, R. Strack, J. Baier, F. Breitling, Á. Kotsis, Shifting skills, moving targets, and remaking the workforce, *Burning Glass Institute*. [https://static1.squarespace.com/static/61\(2022\)7](https://static1.squarespace.com/static/61(2022)7).
- [17] A. Bordes, N. Usunier, A. Garcia-Duran, J. Weston, O. Yakhnenko, Translating embeddings for modeling multi-relational data, in: *Advances in neural information processing systems*, volume 26, 2013.
- [18] M. Nickel, K. Murphy, V. Tresp, E. Gabrilovich, A review of relational machine learning for knowledge graphs, *Proceedings of the IEEE* 104 (2016) 11–33.
- [19] N. Reimers, I. Gurevych, Sentence-bert: Sentence embeddings using siamese bert-networks, in: *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2019, pp. 3982–3992.
- [20] A. Grover, J. Leskovec, node2vec: Scalable feature learning for networks, in: *Proceedings of the 22nd ACM SIGKDD international conference on Knowledge discovery and data mining*, 2016, pp. 855–864.
- [21] Q. Wang, Z. Mao, B. Wang, L. Guo, Knowledge graph embedding: A survey of approaches and applications, *IEEE Transactions on Knowledge and Data Engineering* 29 (2017) 2724–2743.
- [22] T. N. Kipf, M. Welling, Variational graph auto-encoders, *arXiv preprint arXiv:1611.07308* (2016).
- [23] R. Bommasani, D. A. Hudson, E. Adeli, R. Altman,

- S. Arora, M. von Arx, et al., On the opportunities and risks of foundation models, arXiv preprint arXiv:2108.07258 (2021).
- [24] Y. S. Tan, D. Low, S. Toh, Z. Wang, Enhancing skills extraction for domain-specific taxonomy mapping, in: Pacific-Asia Conference on Knowledge Discovery and Data Mining, Springer, 2025, pp. 324–335.
 - [25] N. Reimers, I. Gurevych, Sentence-bert: Sentence embeddings using siamese bert-networks, in: Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing, Association for Computational Linguistics, 2019. URL: <https://arxiv.org/abs/1908.10084>.
 - [26] J. Li, H. Shomer, H. Mao, S. Zeng, Y. Ma, N. Shah, J. Tang, D. Yin, Evaluating graph neural networks for link prediction: Current pitfalls and new benchmarking, Advances in Neural Information Processing Systems 36 (2023) 3853–3866.
 - [27] W. Hamilton, Z. Ying, J. Leskovec, Inductive representation learning on large graphs, Advances in neural information processing systems 30 (2017).
 - [28] P. Veličković, G. Cucurull, A. Casanova, A. Romero, P. Lio, Y. Bengio, Graph attention networks, arXiv preprint arXiv:1710.10903 (2017).
 - [29] T. N. Kipf, M. Welling, Semi-supervised classification with graph convolutional networks, arXiv preprint arXiv:1609.02907 (2016).