

An Exploratory Approach to Discovering Salary-Related Wording in Job Postings in Korea

Taehyun Ha 

Future Technology Analysis Center, Korea Institute of Science and Technology Information (KISTI), Seoul, Korea
E-mail: taehyunha@kisti.re.kr

Byoung-Youl Coh 

Future Technology Analysis Center, Korea Institute of Science and Technology Information (KISTI), Seoul, Korea
E-mail: cohby@kisti.re.kr

Mingook Lee 

Future Technology Analysis Center, Korea Institute of Science and Technology Information (KISTI), Seoul, Korea
E-mail: mingook@kisti.re.kr

Bitnari Yun 

Center for Research and Development Investment and Strategy Research, Korea Institute of Science and Technology Information (KISTI), Seoul, Korea
E-mail: nryun@kisti.re.kr

Hong-Woo Chun* 

Future Technology Analysis Center, Korea Institute of Science and Technology Information (KISTI), Seoul, Korea
E-mail: hw.chun@kisti.re.kr


ABSTRACT

Online recruitment websites discuss job demands in various fields, and job postings contain detailed job specifications. Analyzing this text can elucidate the features that determine job salaries. Text embedding models can learn the contextual information in a text, and explainable artificial intelligence frameworks can be used to examine in detail how text features contribute to the models' outputs. We collected 733,625 job postings using the WORKNET API and classified them into low, mid, and high-range salary groups. A text embedding model that predicts job salaries based on the text in job postings was trained with the collected data. Then, we applied the SHapley Additive exPlanations (SHAP) framework to the trained model and discovered the significant words that determine each salary class. Several limitations and remaining words are also discussed.

Keywords: job posting, text embedding, explainable artificial intelligence, SHapley Additive exPlanations

Received: April 22, 2022
Accepted: May 17, 2022

Revised: May 17, 2022
Published: June 20, 2022

*Corresponding Author: Hong-Woo Chun
 <https://orcid.org/0000-0002-9584-7065>
E-mail: hw.chun@kisti.re.kr



All JISTaP content is Open Access, meaning it is accessible online to everyone, without fee and authors' permission. All JISTaP content is published and distributed under the terms of the Creative Commons Attribution License (<https://creativecommons.org/licenses/by/4.0/>). Under this license, authors reserve the copyright for their content; however, they permit anyone to unrestrictedly use, distribute, and reproduce the content in any medium as far as the original authors and source are cited. For any reuse, redistribution, or reproduction of a work, users must clarify the license terms under which the work was produced.

1. INTRODUCTION

Technology constantly interacts with the world. The advances in technology create new occupations, and changes in the job market yield additional advances. Although the relationship between advancements in technology and the job market is nonlinear, researchers could yield reliable predictions regarding their future based on an appropriate approximation of that relation. In other words, if researchers can identify what features of jobs are highly valued in the job market, they can capture technologies that the industry mainly focuses on and forecast changes in technology in an inductive way. This forecast would be beneficial to people who are interested in predicting future technology and responding to changes in the technology.

Researchers can utilize various features to capture changes in the job market. For example, several economic indicators such as the number of employees or full-time employees, average wages and salary, and demographics are typical features used to explain the job market. These indicators provide macro-level diagnoses for the status of the job market. However, these indicators do not provide detailed features about jobs. In other words, the indicators are limited in terms of explaining what features of jobs are highly valued and how they relate to changes in technology. Job postings are valuable sources for capturing such micro-level features of the job market. Many companies recruit employees through online websites, and postings uploaded on the websites contain detailed job descriptions such as the required and preferred job qualifications, duties and responsibilities, and salaries. Statistics suggest that approximately 86 percent of job seekers have used social media to search for jobs, and 84 percent of companies used social media for job recruiting purposes (Ku, 2021). These results suggest that analyzing the job postings uploaded on online websites can reveal highly valued features in the job market and provide a reliable forecast for changes in jobs and technologies.

Existing studies have shown that online job postings could be used to examine the characteristics of the job market (e.g., Antenucci et al., 2014; Borisyuk et al., 2017). These studies primarily focused on numerical values rather than the text found in job postings. It is difficult to extract features from text without using sophisticated embedding methods. However, advances in text embedding techniques and deep learning models have made it possible to extract features from job postings and make predictions about promising jobs. For example, Ha et al.

(2022) showed that a word embedding model could be trained by descriptions of job and patent classification codes and utilized to match two different types of information. They forecasted promising jobs based on changes in the number of patents. An advanced method utilizing text embedding techniques and deep learning models may be highly effective in terms of predicting promising jobs. Nonetheless, researchers had trouble understanding what features in the texts contributed to the prediction because the models used for prediction were too deep and difficult to interpret.

Regarding this problem, a framework of explainable artificial intelligence, the SHapley Additive exPlanations (SHAP), provides functions that interpret trained deep learning models (Lundberg & Lee, 2017). SHAP calculates Shapley values for the input variables and estimates their contributions to the trained model's output. Researchers have employed SHAP in various fields to examine significant features in prediction problems (e.g., J. Chen et al., 2021; T. Chen et al., 2019; Park et al., 2022). For our problem, we can train a prediction model using the text and salary information described in job postings and interpret the model using SHAP. The interpretation results provide the textual features that contribute to a job's salary as well as meaningful insights that elucidate the characteristics of the job market.

In this study, we gathered 733,625 job postings in Korea from WORKNET and classified them into three categories based on the terciles of the salary. We used the Efficiently Learning an Encoder that Classifies Token Replacements Accurately (ELECTRA) model to train the textual features in job postings. Furthermore, we employed an ELECTRA model that was pre-trained with Korean texts, named KoELECTRA, and fine-tuned it using text from job postings and the salary terciles as the input and output variables. SHAP interpreted the trained model and estimated how the textual features of job postings contributed to the salary classification. We identified significant words that are highly valued in terms of salary and discuss several characteristics in the current job market and suggest several implications.

2. BACKGROUND

2.1. Job Trend Analysis

Researchers are interested in using online job postings to discover trends or significant features of the job market. They have collected job postings from online recruitment websites and identified characteristics using coding tech-

niques. Daneva et al. (2017) gathered 235 job postings related to engineering job requirements from online recruitment websites in the Netherlands and manually analyzed their properties using a coding technique. The results revealed the statistical features of the job postings and suggested several implications for the engineering job market in the Netherlands. Similarly, Chang et al. (2019) collected 390 job postings from LinkedIn related to data analytics and knowledge management, and categorized them according to their entry, mid-senior, and top-level experience requirements. They coded the text in job postings and plotted word clouds to identify the characteristics and essential words used in the job postings. Hirudayaraj and Baker (2018) investigated the expectations of employers in human resource development based on 500 job postings collected from the Association for Talent Development (ATD) portal. They coded the data in terms of educational qualifications, work experience, knowledge, skills, technical skills, industry, and additional information as well as responsibilities, and they reported the jobs' features based on the coding results.

Other studies applied mathematical approaches to analyze job postings. Zhu et al. (2016) developed a sequential latent variable model named the Market Trend Latent Variable Model (MTLVM) and showed how the developed model discovers trends in the job market. They collected 257,166 job postings from a Chinese online recruitment website and trained the model using that data. The model discovered chronological changes in job-related words for several companies and suggested that job postings can be used to identify significant changes in the job market. Karakatsanis et al. (2017) considered job descriptions provided by the Occupational Information Network (O*NET) and matched them with 12,849 job postings gathered from the United States, United Kingdom, and Gulf cooperation council countries. They vectorized the job descriptions and postings using the Term Frequency - Inverse Document Frequency (TF-IDF) matrix and the Singular Value Decomposition (SVD) method and matched two different types of information. Based on the matched results, they estimated the demands of the jobs listed on O*NET.

While these studies examined limited data, other researchers have considered massive amounts of data. Scrivner et al. (2020) considered 30 million job postings related to healthcare and reported their changes between 2010 and 2018. Zhao et al. (2021) gathered 60 million job postings from CareerBuilder and vectorized them utilizing Word2Vec. They trained a convolutional neural network

with the extracted vectors and utilized it to develop a job recommendation system. In similar contexts, several studies employed text embedding models to vectorize texts in job postings and developed a job recommendation system (Kaya & Bogers, 2021; Lacic et al., 2019). Sun et al. (2021) constructed a deep neural network based on job postings in China gathered from Lagou. The model identified major skill-related words and predicted job salaries. Existing studies show that text embedding models can vectorize texts in job postings and predict trends in the job market or find similar jobs. However, the models and vectorized texts themselves do not provide information regarding how the text in job postings contributes to job salaries. In other words, researchers cannot identify the important keywords that determine salaries and have only a limited grasp of the trends in the job market. To solve this problem, the trained models must be interpreted using explainable Artificial Intelligence (XAI).

2.2. Explainable Artificial Intelligence

Recent advances in deep neural networks have contributed to solving many complex problems through many neuron layers in deep neural networks, which are nonlinearly connected and adjust their connection weights based on prediction errors. However, it is difficult to ascertain how the input values contribute to the output values because of the deep network structure. Gunning (2017) focused on this problem and developed the concept of XAI. Several techniques, such as layer-wise relevance propagation (LRP, Bach et al., 2015) and local interpretable model-agnostic explanations (LIME, Ribeiro et al., 2016) have been suggested for implementing XAI. However, some studies report instability issues in the explanation results from LRP and LIME (Alvarez-Melis & Jaakkola, 2018; Bach et al., 2015; Ward et al., 2021). Focusing on these limitations, Lundberg and Lee (2017) developed an XAI framework named SHAP. SHAP calculates Shapley values based on the input and output values from a trained model and estimates the contributions of features based on those values. SHAP applies to any trained model and is not limited to any model structures (i.e., it is model-agnostic). Some studies showed that SHAP could perform better than LIME under specific conditions (Lundberg & Lee, 2017; Ward et al., 2021).

SHAP explains a trained model based on Shapley values, which can be calculated by Equation 1. As shown in the equation, a Shapley value estimates the contribution of a specific feature based on the calculated marginal contributions of all subsets of features. Calculating Shapley

values requires a large number of computations. However, they satisfy efficiency, symmetry, dummy, and additivity conditions, and this satisfaction implies fairness when attributing the features' contributions to the model. SHAP provides instance and model-level explanations based on Shapley values for features and the means of the values. SHAP consists of interpretation functions concerning different model types such as Tree SHAP, Deep SHAP, and Kernel SHAP. These functions enable the drawing of a global surrogate for the trained model based on local explanations for the input features.

$$\phi_i(f) = \sum_{S \subseteq F \setminus i} \frac{|S|!(|F|-|S|-1)!}{|F|!} (f(S) - f(S \setminus i)) \quad (1)$$

where

Φ_i : Shapley value for feature i

f : Blackbox model

F : Set of features

S : Subset of features excluding i

$f(S)$: Contribution of S

Researchers have employed XAI frameworks to address problems in various fields. Chen et al. (2019) employed LightGBM to predict extubation failure in patients

by training the model with patient data and interpreting it to identify the significant patient features that determine this failure. Park et al. (2021) focused on features affecting the diagnosis of nuclear power plants and trained LightGBM and Gated Recurrent Unit-based AutoEncoder (GRU-AU) models with nuclear power plant data. They interpreted the trained model using the SHAP framework and showed that SHAP could be used to provide evidence for diagnoses. Chen et al. (2021) trained a multilayer neural network with Walmart retail data and interpreted it using the SHAP framework to investigate the major features that determine sales for the following month. In the context of job analysis, Choi et al. (2020) gathered 52,190 IT-related job postings from CareerBuilder and trained a Long Short-Term Memory (LSTM) model for classifying job postings. They interpreted the trained model using the LIME framework and suggested the major words that characterize IT-related jobs. Although they employed LIME rather than SHAP, the study showed that XAI frameworks can be utilized to investigate the features of job postings.

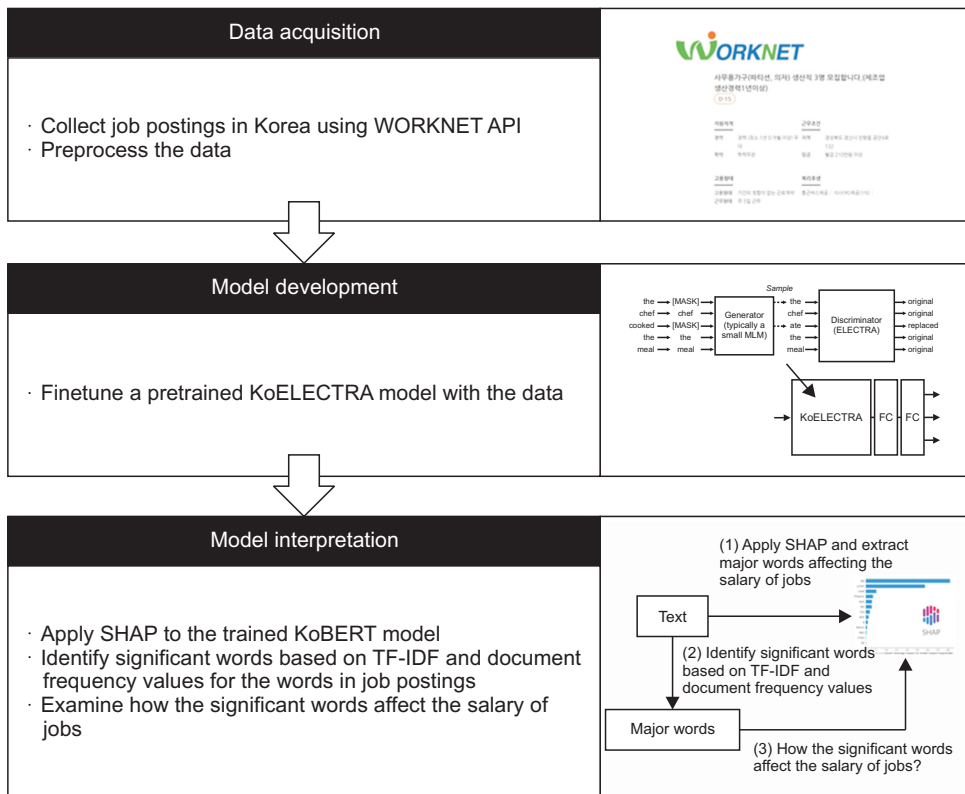


Fig. 1. Analysis flowchart. API, Application Programming Interface; SHAP, SHapley Additive exPlanations; TF-ID, Term Frequency - Inverse Document Frequency.

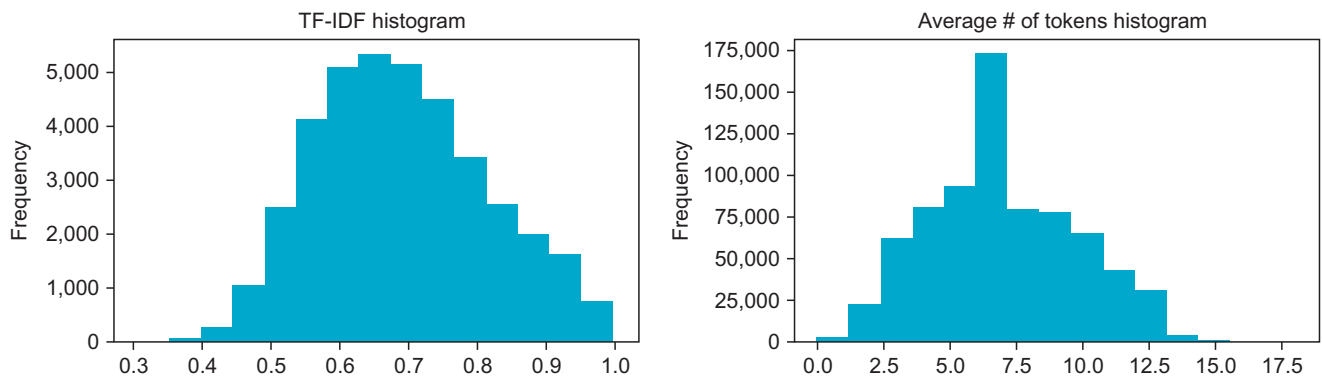


Fig. 2. Term Frequency - Inverse Document Frequency (TF-IDF) and average numbers of tokens in job postings.

3. METHOD

To examine what textual features of job postings contributed to job salaries, we collected job postings in Korea using the WORKNET Application Programming Interface (API) and trained a text embedding model that uses text from job postings and annual salaries as the input and output variables, respectively. SHAP interprets the trained model and suggests which keywords in the job postings affect the jobs' salaries. Fig. 1 shows the analysis flowchart.

3.1. Data Acquisition

We gathered 733,625 job postings from WORKNET. The job posting registration dates ranged from July 22, 2021 to March 16, 2022 and described their hourly, weekly, monthly, and annual salaries. We transformed the different types of wages into annual salaries. Then, the annual salaries are classified into three groups (high, mid, and low) based on their terciles (23,520,000 won for 2/3 and 28,800,000 won for 1/3). The mean annual salaries for the three groups were presented as 20,488,440; 25,839,360; and 35,894,140 won, and the standard deviations were 3,155,976; 1,617,983; and 131,177,700. The text in the job postings was then used to train a text embedding model. Fig. 2 shows the histograms of TF-IDF and average numbers of tokens in job postings.

3.2. Model Development

We employed a pretrained KoELECTRA model developed by Park (2020). The ELECTRA model was developed by Clark et al. (2020), who were focusing on the limitations of Bidirectional Encoder Representations from Transformers (BERT)-based models. The model can be pretrained using general texts and fine-tuned using domain-specific texts. The original study provided mod-

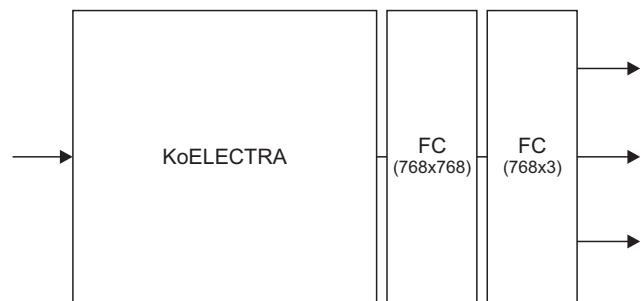


Fig. 3. KoELECTRA model structure used for fine-tuning.

els that were pretrained with multilingual texts, but their performance was not as good as the models pretrained with English texts. Thus, several researchers suggested pretrained models oriented to the Korean language, and KoELECTRA demonstrated high performances in several downstream tasks. We employed the basic version of the pretrained KoELECTRA models and fine-tuned the model for our text classification task. The data were divided into training and validation data at a ratio of 0.2. The AdamW optimizer was used for training, and the learning rate and batch size were set to 5e-5 and 32, respectively. Fig. 3 shows the model structure used for fine-tuning. We stopped the training process early based on the F1 scores for the training and validation data.

3.3. Model Interpretation

We applied an explainer for the SHAP framework to the fine-tuned model and explored how the model classifies job posting text into three salary categories. To calculate the contributions of each word, we employed a tokenizer from the Konlpy package named Mecab and used it as a masker for the input words. As discussed in the background section, calculating Shapley values requires

a massive amount of time. Thus, we sampled ten percent of the data for each category and extracted the major words that determine each salary category. In addition, the number of extracted major words was sometimes too large to review as the number of job postings considered in the analysis increased. Thus, we calculated the TF-IDF and document frequency values for the words in the job postings and identified the significant words, which we then made our primary focus. Consequently, we observed how the identified significant words in the job postings affected the jobs' salaries.

4. RESULTS

The model was trained on 45 epochs, and the F1 scores for the training and validation data were 0.913 and 0.710, respectively. We identified 586 words that presented TF-IDF values greater than 0.5 and document frequency values larger than 1,000. Then, we applied a SHAP explainer to the sample data and explored how the identified 586 words affected each salary category. Fig. 4 shows two example cases. The explainer highlights a category used by the model to classify the input text. A category is highlighted more strongly if the model classifies the input

text with higher confidence. The words in the text are also highlighted in red or blue coloring. Red highlighting indicates that a word contributed positively to the classification result, whereas blue highlighting represents a negative contribution by a word.

In example 1, “조리사 및 청소원 채용” was classified into the low-range salary category. In the prediction for the low-range salary category, “청소원” (cleaner) and “조리사” (cook) were highlighted in red. This result means that “청소원” and “조리사” positively contributed to the low-range salary prediction, and jobs related to these words are generally paid a low salary. In example 2, “정보시스템 BPR/ISP 및 빅데이터 분석 컨설턴트 모집” was classified into the mid-range salary category. In the prediction for the mid-range salary category, “빅데이터 분석” (big data analysis) and “ISP (Information Strategy Planning)” were highlighted in red, whereas “BPR (Business Process Reengineering)” and “정보시스템” (information system) were highlighted in blue. On the other hand, the prediction for the high-range salary category reverses the highlight colors. This result means that job postings containing “빅데이터 분석” and “ISP” are likely to suggest mid-range salaries, whereas “BPR” and “정보시스템” are more likely to provide high-range salaries.

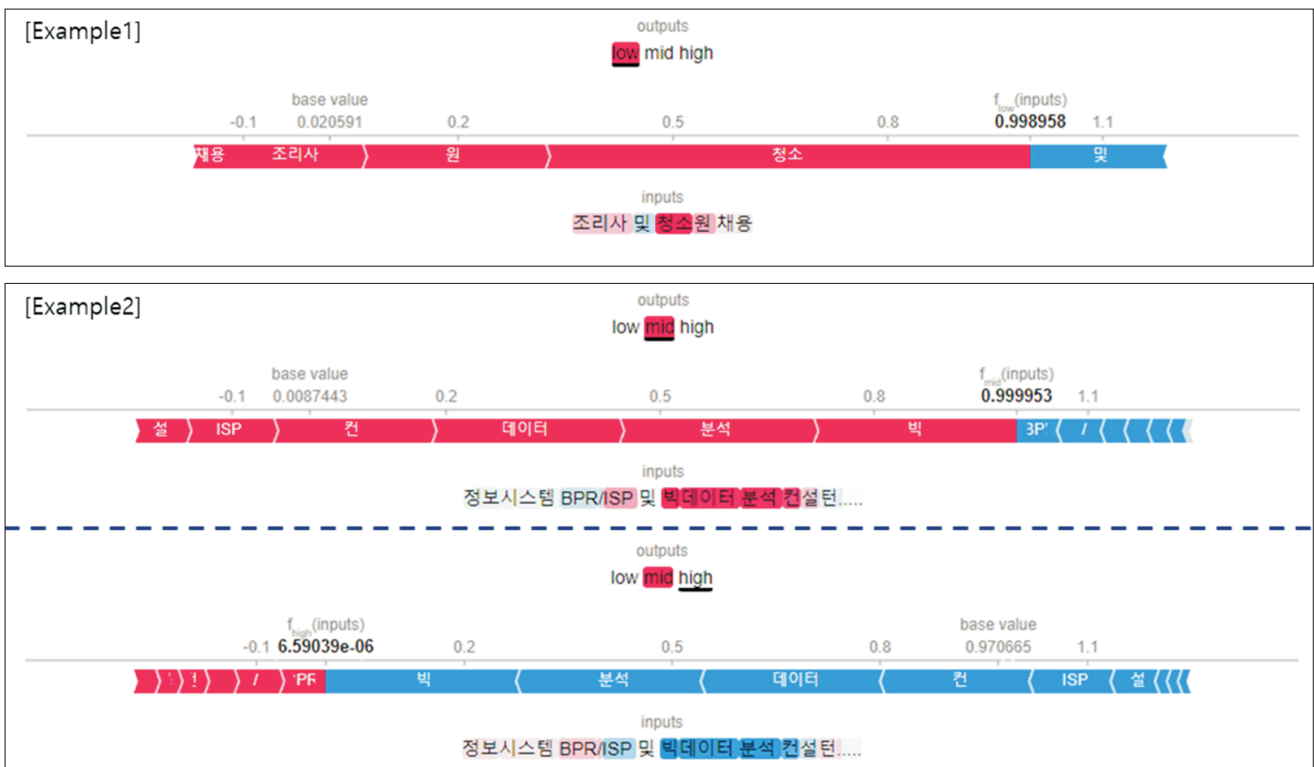


Fig. 4. Example cases of the model interpretation.

Table 1 shows the top and bottom ten words identified as significantly affecting the low, mid, and high-range salaries of jobs. 미화원 (Cleaner), 데이케어 (Daycare), 파트타임 (Part-time), 작업자 (Worker), 도우미 (Helper), 검사원 (Inspector), 미화 (Cleaning), 포장 (Packaging), 연장 (Extension), and 식당 (Restaurant) had the most positive effect when predicting a low-range salary, whereas 목욕 (Bath), 방문 (Visit), 기술직 (Technical work), 물리 (Physical), 간호사 (Nurse), 폐기물 (Waste), 재 (Resident), 단체

(Group), 종 (Following), and 입주 (Resident) had the most negative effect. 경리사무 (Accounting), 주주야야 (2 days/2 nights shift), 영양사 (Nutritionist), 경리 (Accounting), 보안 (Security), 경비원 (Security guard), 조 (Group/Serving), 점검 (Inspection), 창고 (Warehouse), and 편집 (Editing) were the most positively impactful when predicting a mid-range salary. At the same time, 연장, 도우미, 대원 (Crew), 식당, 오후 (Afternoon), 영아 (Infant), 미화원, 데이케어, 입주, and 목욕 were the most negatively impactful. 입

Table 1. Top and bottom ten influential words

Category	Top ten ranked words	Mean Shapley value	Bottom ten ranked words	Mean Shapley value
Low-range salary	미화원(Cleaner)	0.580	목욕(Bath)	-0.086
	데이케어(Daycare)	0.482	방문(Visit)	-0.088
	파트타임(Part-time)	0.477	기술직(Technical work)	-0.090
	작업자(Worker)	0.448	물리(Physical)	-0.091
	도우미(Helper)	0.398	간호사(Nurse)	-0.100
	검사원(Inspector)	0.394	폐기물(Waste)	-0.117
	미화(Cleaning)	0.393	재(Resident)	-0.136
	포장(Packaging)	0.361	단체(Group)	-0.150
	연장(Extension)	0.358	종(Following)	-0.161
	식당(Restaurant)	0.355	입주(Resident)	-0.238
Mid-range salary	경리사무(Accounting)	0.506	연장(Extension)	-0.142
	주주야야(2 days/2 nights shift)	0.430	도우미(Helper)	-0.146
	영양사(Nutritionist)	0.414	대원(Crew)	-0.146
	경리(Accounting)	0.380	식당(Restaurant)	-0.155
	보안(Security)	0.354	오후(Afternoon)	-0.164
	경비원(Security guard)	0.345	영아(Infant)	-0.171
	조(Group/Serving)	0.330	미화원(Cleaner)	-0.186
	점검(Inspection)	0.329	데이케어(Daycare)	-0.218
	창고(Warehouse)	0.313	입주(Resident)	-0.272
	편집(Editing)	0.305	목욕(Bath)	-0.276
High-range salary	입주(Resident)	0.510	미화(Cleaning)	-0.368
	목욕(Bath)	0.362	사무직원(Clerical worker)	-0.385
	금형(Mold)	0.192	미화원(Cleaner)	-0.395
	과장(Manager)	0.173	경리(Accounting)	-0.406
	폐기물(Waste)	0.163	파트타임(Part-time)	-0.414
	간호사(Nurse)	0.158	조(Group/Serving)	-0.420
	대원(Crew)	0.157	작업자(Worker)	-0.425
	톤(Ton)	0.154	검사원(Inspector)	-0.428
	용(For)	0.150	간호조무사(Nurse's aide)	-0.528
	안전관(Safety manager)	0.145	경리사무(Accounting)	-0.547

주, 목욕, 금형 (Mold), 과장 (Manager), 폐기물, 간호사, 대원, 톤 (Ton), 용 (For), and 안전관 (Safety manager) had the most positive effect when predicting a high-range salary. In contrast, 미화, 사무직원 (Clerical worker), 미화원, 경리, 파트타임, 조 (Group/Serving), 작업자, 검사원, 간호조무사 (Nurse's aide), and 경리사무 had the most negative effect.

5. DISCUSSION

We collected approximately 730 thousand job postings in Korea using the WORKNET API and classified them into three groups based on the terciles of the salaries described in the postings. We employed a pretrained KoELECTRA model and fine-tuned it by using the text in job postings as input values and the salary categories as output values. The job posting data was split into training and validation data at a ratio of 0.2, and the fine-tuned model showed F1 values of 0.913 and 0.710 for the training and validation data, respectively. We interpreted the trained model using SHAP and the Shapley values it provides. The Mecab tokenizer was used to mask the input words and calculate their Shapley values. For the 586 words identified based on the TF-IDF and document frequency values, SHAP provided information regarding how the words related to the three categories of job salaries.

In the low-range salary classification, simple labor-related words were the most frequently used. However, some simple labor-related words such as 목욕 (bath), 입주 (resident), and 폐기물 (waste) contributed negatively to the low-range salary classification. Instead, these words were the most contributed words to the high-range salary classification. This result indicates that employers offered high salaries for hard labor jobs even if they were simple. We were also able to find 연장 (extension) from job postings in the low-range salary category, which implies that many job postings failed to find applicants due to the low-range salary and the job postings were extended. Accounting, shift work, and security-related words were the most frequently used words in the mid-range salary classification, and hard labor, nursing, and technology-related words were the most frequently used words in the high-range salary classification. Inspection-related words (검사원, inspector, 점검, inspection, and 안전관, safety manager) were found in three categories. This result suggests that salaries for inspection-related jobs are diverse but likely to be high if the job is highly technical (i.e., safety manager).

The developed approach can help recruiters write more compelling postings. Job seekers are often interested in

postings that suggest high salaries and can use a filtering function to search the postings. However, constructing a description that can attract many applicants is a separate issue. Recruiters post jobs with similar descriptions on online recruitment websites, but the outcomes are different. Various factors such as the company's name value and a job website's features can affect these outcomes. However, if recruiters focus on what is within their control, they can examine what expressions improve or deteriorate a description from the salary perspective. On the other hand, researchers can develop a job recommendation system that utilizes the fine-tuned KoELECTRA model. This model reflects the contextual information in job postings and can vectorize them. A system can then match the vectorized job postings with an input text that a user has typed, which could be more effective than traditional job search systems.

We found that a text embedding-based approach can discover significant words that determine job salaries, but there are also some limitations. First, the analysis results can consist of terms that are too general and provide limited insight into technological changes in the job market. We focused on a subset of words in the job postings as indicated by TF-IDF and document frequency values because calculating Shapley values consumes considerable time, and not all of the words' Shapley values were necessary to capture significant changes in the job market. However, if one is interested in technological changes in the job market, s/he must focus only on certain words and examine how they contribute to job salaries. Second, some meanings of the words were not apparent due to the characteristics of the Korean language. For example, we observed that 조 in job postings significantly related to the mid-range salary category. However, the word is a homonym and can be interpreted as either "group" or "serving." In addition, a word can be tokenized differently depending on the tokenizers. We employed Mecab, one of the most frequently used tokenizers for Korean language processing. However, this tokenizer is dictionary-based and cannot adequately capture proper nouns. In this case, learning-based tokenizers such as Soynlp (<https://github.com/lovit/soynlp>) may be a good alternative.

Data completion and representativity issues also remain. WORKNET API provides detailed information for recruitment in Korea. However, some missing information and misdescribed values in the data may have affected the results. In particular, the average job salary presented in WORKNET (27,319,058 won) was quite lower than the average salary announced by the Korea Ministry

of Employment and Labor in 2020 (34,392,000 won). This implies that the data we used for the analysis was biased toward the low- and mid-salary jobs. The limitation of data could be addressed by considering the additional sources of data. With an extended data collection period, chronological changes in the job market can be tracked and analyzed.

6. CONCLUSION

In this study, we examined textual features determining job salaries using KoELECTRA and SHAP. The results showed that text embedding models can be used to predict job salaries and discover significant words that determine job salaries. However, we were not able to link the words with technological changes in the job market because the words were too general and provided limited insights. Future studies can consider extended data and advanced natural language processing techniques to discover meaningful words and elucidate changes in jobs and technologies effectively.

CONFLICTS OF INTEREST

No potential conflict of interest relevant to this article was reported.

REFERENCES

Alvarez-Melis, D., & Jaakkola, T. S. (2018). *On the robustness of interpretability methods*. arXiv. <https://doi.org/10.48550/arXiv.1806.08049>.

Antenucci, D., Cafarella, M., Levenstein, M., Ré, C., & Shapiro, M. D. (2014). *Using social media to measure labor market flows*. <https://www.nber.org/papers/w20010>.

Bach, S., Binder, A., Montavon, G., Klauschen, F., Müller, K. R., & Samek, W. (2015). On pixel-wise explanations for non-linear classifier decisions by layer-wise relevance propagation. *PloS One*, 10(7), e0130140. <https://doi.org/10.1371/journal.pone.0130140>.

Borisjuk, F., Zhang, L., & Kenthapadi, K. (2017, August 13-17). LiJAR: A system for job application redistribution towards efficient career marketplace. *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (pp. 1397-1406). ACM.

Chang, H. C., Wang, C. Y., & Hawamdeh, S. (2019). Emerging trends in data analytics and knowledge management job market: Extending KSA framework. *Journal of Knowledge Management*, 23(4), 664-686. [https://doi.org/10.1108/JKM-](https://doi.org/10.1108/JKM-02-2018-0088)

[02-2018-0088](https://doi.org/10.1108/JKM-02-2018-0088).

Chen, J., Koju, W., Xu, S., & Liu, Z. (2021, March 26-28). Sales forecasting using deep neural network and SHAP techniques. *Proceedings of the 2021 IEEE 2nd International Conference on Big Data, Artificial Intelligence and Internet of Things Engineering (ICBAIE)* (pp. 135-138). IEEE.

Chen, T., Xu, J., Ying, H., Chen, X., Feng, R., Fang, X., Gao, H., & Wu, J. (2019). Prediction of extubation failure for intensive care unit patients using light gradient boosting machine. *IEEE Access*, 7, 150960-150968. <https://doi.org/10.1109/ACCESS.2019.2946980>.

Choi, I. H., Kim, Y. S., & Lee, C. K. (2020, September 17-19). A study of the classification of IT jobs using LSTM and LIME. *Proceedings of the 9th International Conference on Smart Media and Applications* (pp. 248-252). ACM.

Clark, K., Luong, M. T., Le, Q. V., & Manning, C. D. (2020). *ELECTRA: Pre-training text encoders as discriminators rather than generators*. arXiv. <https://doi.org/10.48550/arXiv.2003.10555>.

Daneva, M., Wang, C., & Hoener, P. (2017, November 9-10). What the job market wants from requirements engineers? An empirical analysis of online job ads from the Netherlands. *Proceedings of the 11th ACM/IEEE International Symposium on Empirical Software Engineering and Measurement* (pp. 448-453). IEEE.

Gunning, D. (2017). *Explainable artificial intelligence (XAI)*. Defense Advanced Research Projects Agency (DARPA).

Ha, T., Lee, M., Yun, B., & Coh, B. Y. (2022). Job forecasting based on the patent information: A word embedding-based approach. *IEEE Access*, 10, 7223-7233. <https://doi.org/10.1109/ACCESS.2022.3141910>.

Hirudayaraj, M., & Baker, R. (2018). HRD competencies: Analysis of employer expectations from online job postings. *European Journal of Training and Development*, 42(9), 577-596. <https://eric.ed.gov/?id=EJ1198597>.

Karakatsanis, I., AlKhader, W., MacCrorry, F., Alibasic, A., Omar, M. A., Aung, Z., & Woon, W. L. (2017). Data mining approach to monitoring the requirements of the job market: A case study. *Information Systems*, 65, 1-6. <https://doi.org/10.1016/j.is.2016.10.009>.

Kaya, M., & Bogers, T. (2021, September 27-October 1). Effectiveness of job title based embeddings on résumé to job ad recommendation. *Proceedings of the 2021 Workshop on Recommender Systems for Human Resources, RECSYS IN HR 2021* (pp. 1-7). CEUR Workshop Proceedings.

Ku, D. (2021). *Social recruiting: Everything you need to know for 2022*. <https://www.postbeyond.com/blog/social-recruiting/>.

Lacic, E., Reiter-Haas, M., Duricic, T., Slawicek, V., & Lex, E.

- (2019, September 16-20). Should we embed? A study on the online performance of utilizing embeddings for real-time job recommendations. *Proceedings of the 13th ACM Conference on Recommender Systems* (pp. 496-500). ACM.
- Lundberg, S. M., & Lee, S. I. (2017, December 4-9). A unified approach to interpreting model predictions. In I. Guyon, U. von Luxburg, S. Bengio, H. M. Wallach, R. Fergus, S. V. N. Vishwanathan, & R. Garnett (Eds.), *Proceedings of the Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems* (pp. 4768-4777). NIPS.
- Park, J. (2020). *KoELECTRA: Pretrained ELECTRA model for Korean*. https://github.com/monologg/KoELECTRA/blob/master/README_EN.md.
- Park, J. H., Jo, H. S., Lee, S. H., Oh, S. W., & Na, M. G. (2022). A reliable intelligent diagnostic assistant for nuclear power plants using explainable artificial intelligence of GRU-AE, LightGBM and SHAP. *Nuclear Engineering and Technology*, 54, 1271-1287. <https://doi.org/10.1016/j.net.2021.10.024>.
- Ribeiro, M. T., Singh, S., & Guestrin, C. (2016, August 13-17). "Why should I trust you?": Explaining the predictions of any classifier. *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (pp. 1135-1144). ACM.
- Scrivner, O., Nguyen, T., Simon, K., Middaugh, E., Taska, B., & Börner, K. (2020). Job postings in the substance use disorder treatment related sector during the first five years of Medicaid expansion. *PLoS One*, 15(1), e0228394. <https://doi.org/10.1371/journal.pone.0228394>.
- Sun, Y., Zhuang, F., Zhu, H., Zhang, Q., He, Q., & Xiong, H. (2021). Market-oriented job skill valuation with cooperative composition neural network. *Nature Communications*, 12(1), 1992. <https://doi.org/10.1038/s41467-021-22215-y>.
- Ward, I. R., Wang, L., Lu, J., Bennamoun, M., Dwivedi, G., & Sanfilippo, F. M. (2021). Explainable artificial intelligence for pharmacovigilance: What features are important when predicting adverse outcomes? *Computer Methods and Programs in Biomedicine*, 212, 106415. <https://doi.org/10.1016/j.cmpb.2021.106415>.
- Zhao, J., Wang, J., Sigdel, M., Zhang, B., Hoang, P., Liu, M., & Korayem, M. (2021). *Embedding-based recommender system for job to candidate matching on scale*. arXiv. <https://doi.org/10.48550/arXiv.2107.00221>.
- Zhu, C., Zhu, H., Xiong, H., Ding, P., & Xie, F. (2016, August 13-17). Recruitment market trend analysis with sequential latent variable models. *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (pp. 383-392). ACM.