
Job-SDF: A Multi-Granularity Dataset for Job Skill Demand Forecasting and Benchmarking

Xi Chen^{1*}, Chuan Qin^{2*}, Chuyu Fang³, Chao Wang¹,
Chen Zhu¹, Fuzhen Zhuang^{4,5}, Hengshu Zhu^{2†}, Hui Xiong^{6,7†}

¹University of Science and Technology of China

²Computer Network Information Center, Chinese Academy of Sciences

³Baidu Inc. ⁴Institute of Artificial Intelligence, Beihang University

⁵SKLSDE, School of Computer Science, Beihang University

⁶AI Thrust, The Hong Kong University of Science and Technology (Guangzhou)

⁷Department of Computer Science and Engineering,

The HongKong University of Science and Technology, Hong Kong SAR

chenxi0401@mail.ustc.edu.cn,

{chuanqin0426, fangchuyu2022, chadwang2012, zc3930155}@gmail.com,
zhuangfuzhen@buaa.edu.cn, zhuhengshu@gmail.com, xionghui@ust.hk

Abstract

In a rapidly evolving job market, skill demand forecasting is crucial as it enables policymakers and businesses to anticipate and adapt to changes, ensuring that workforce skills align with market needs, thereby enhancing productivity and competitiveness. Additionally, by identifying emerging skill requirements, it directs individuals towards relevant training and education opportunities, promoting continuous self-learning and development. However, the absence of comprehensive datasets presents a significant challenge, impeding research and the advancement of this field. To bridge this gap, we present *Job-SDF*, a dataset designed to train and benchmark job-skill demand forecasting models. Based on millions of public job advertisements collected from online recruitment platforms, this dataset encompasses monthly recruitment demand. Our dataset uniquely enables evaluating skill demand forecasting models at various granularities, including occupation, company, and regional levels. We benchmark a range of models on this dataset, evaluating their performance in standard scenarios, in predictions focused on lower value ranges, and in the presence of structural breaks, providing new insights for further research. Our code and dataset are publicly accessible via the <https://github.com/Job-SDF/benchmark>.

1 Introduction

Job skills encompass a range of abilities and competencies essential for performing tasks effectively in the workplace. These skills are broadly categorized into hard skills, such as technical and analytical abilities, and soft skills, including communication, teamwork, and adaptability [1]. Accurate forecasting of skill demand helps businesses and policymakers anticipate and address skill shortages and mismatches, and promotes skill development in high-demand areas, thereby supporting economic growth and stability [2, 3]. By identifying emerging skill requirements, individuals are directed towards relevant training and education opportunities, fostering continuous self-learning and development to stay competitive in the labor market [4–10]. By tracking skill demand trends,

*Equal contributions

†Corresponding Authors

employers gain deeper insight into recruits’ priorities, enhancing person-job fit. [11–21]. Moreover, forecasting informs educational and training programs, ensuring that curricula align with the labor market’s evolving needs [22–24].

Traditionally, skill demand analysis has relied on labor-intensive, survey-based methods limited to specific companies or occupations [25–27]. However, over the past decade, the rapid evolution of the internet has spurred the emergence of online recruitment platforms. These platforms have become the primary channels for job advertisements for numerous enterprises and organizations, accumulating vast amounts of job advertisement data. By leveraging this data, researchers have formulated skill demand forecasting as a time series task, utilizing various machine learning models such as autoregressive integrated moving average (ARIMA) [28], recurrent neural networks (RNNs) [29], and dynamic graph autoencoders (DyGAEs) [30], to predict future skill needs.

A major challenge impeding progress in this field is the lack of comprehensive and publicly accessible datasets. Existing studies do not provide open-source datasets, making it difficult for researchers to replicate experimental results and identify bottlenecks in current research. Furthermore, these datasets primarily focus on predicting skill demand variations across different occupations, with a notable lack of modeling and prediction at other granularities, such as companies or regions. This limitation hinders comprehensive comparisons between different models and impedes the exploration of potential downstream applications, such as human capital strategy development and regional policy formulation. Additionally, the significant variations in skill demand present further challenges. Existing studies, which rely on metrics such as Mean Squared Error (MSE), struggle to evaluate the performance of skill demand forecasting models for low-frequency skill terms. For instance, some emerging skills, such as large language models (LLMs), may initially show low demand but are crucial for the job market due to their potential to reshape existing occupations.

To this end, in this paper, we propose *Job-SDF*, a multi-granularity dataset designed for job skill demand forecasting research. Specifically, we collected millions of public job advertisements from online recruitment platforms. By extracting skill terms from job advertisement texts, we quantified the monthly skill demand at various granularities, including occupations, companies, and regions, to construct our dataset. This dataset encompasses 2,324 types of skills, 52 occupations, 521 companies, and 7 regions. We then use the Job-SDF dataset to benchmark a wide range of models for job skill demand forecasting tasks at various granularities. These models include statistical time series models (e.g., ARIMA [31]), deep learning-based methods such as RNN-based models [32, 30], Transformer-based models [33–37], MLP-based models [38, 39], as well as several state-of-the-art time-series forecasters [40, 41]. Performance is evaluated using regression metrics such as Mean Absolute Error (MAE) and Root Mean Squared Error (RMSE). Additionally, we use the Symmetric Mean Absolute Percentage Error (SMAPE) [42] and Relative Root Mean Squared Error (RRMSE) [43] metrics to account for the significantly varying nature of skill demand values, which is particularly useful for evaluating model performance in predicting lower value ranges. Moreover, we further investigate the impact of structural breaks in job skill demand time series data on model performance. The Job-SDF dataset, along with data loaders, example codes for different models, and evaluation setup, are publicly available in our GitHub repository: <https://github.com/Job-SDF/benchmark>.

2 Related Work

Skill demand forecasting can analyze how skills evolve over time, aiding experts in evaluating technological advancements [44–46], assessing wage inequality [47–49], and generating employment opportunities [50]. Furthermore, the skills required in the 21st-century workplace will differ significantly from those in previous eras [51]. Predicting skill demands benefits personal career transitions and corporate management strategies.

Recently, with the rapid accumulation of data and continuous advancements in technology, skill demand forecasting has demonstrated significant vitality. *Das et al.* proposed a method for dynamic task allocation to investigate the evolution of job task requirements over a decade of AI innovation across different salary levels [28]. Given the effectiveness of RNN in multi-step prediction, some researchers have integrated skill demand forecasting with RNN algorithms, achieving promising results [29, 32]. In addition, considering the supply-demand dynamics of the labor market concurrently, CHGH designed a joint prediction model based on the encoder-decoder architecture to achieve trend prediction for both skill supply and demand sides [30]. Moreover, to capture the dynamic information

of occupations, a pretraining-enhanced dynamic graph autoencoder has been developed to efficiently forecast skill demand at the occupational granularity [52].

However, the predominance of closed-source datasets has significantly elevated the barrier of researchers and constrained the pace of methodological advancements. While open-source skill-related datasets such as O*NET [53] and ESCO [54] provide skill taxonomies, they do not quantify skill demand. Furthermore, the current research data focuses either on macro-market skill demand predictions or analyses at a specific granularity, neglecting multi-level labor market analysis. This limitation generally hampers the transferability of the modeling approaches.

3 Job-SDF Dataset

The Job-SDF dataset is built from job advertisements collected on online recruitment platforms, encompassing dynamic job skill demand time series data at various granularities, recorded monthly. The dataset is CC BY-NC-SA 4.0 licensed, accessible via the URL <https://github.com/Job-SDF/benchmark>. We summarize the dataset construction process, task description, and dataset analysis below, with supplementary details provided in the Appendix F.

3.1 Data Collection and Processing

Job Advertisement Collection. We collected public job advertisements for 52 occupations from 521 companies on online recruitment platforms. We obtained unique records after removing identical job advertisements posted simultaneously by different companies on various platforms. Each record contains five types of information: (1) *Job Requirement*, which is a text segment that outlines the specific skills required of candidates applying for the job; (2) *Company*, which identifies the company that posted the job advertisement; (3) *Occupation*, which specifies the job advertisement’s category. Our dataset encompasses 52 detailed occupations (L2-level), such as front-end development engineer and financial investment analyst. Additionally, these 52 occupations are grouped into 14 broader categories (L1-level); (4) *Region*, which indicates the primary geographic divisions in China where the job postings are located. These regions are classified based on their geographical orientation; (5) *Posting Time*, which records the date when the job was posted, including the year, month, and day.

Job Skill Extraction. After acquiring the job advertisement data, we utilized a Named Entity Recognition (NER) model, as referenced in [55–58], to explicitly extract skill requirements from the *Job Requirement* of each advertisement. Specifically, we first annotated a dataset for training the NER model by identifying skill terms within the job requirement texts. To achieve this, we devised a set of regular expressions tailored to the characteristics of skill descriptions and used these to match skill words in job advertisements. Subsequently, we merged all matched skill words to formulate a raw skill dictionary, including their corresponding frequencies across job advertisements. We then filtered out low-frequency words and manually annotated the raw skill dictionary to create a refined skill dictionary. Along this line, we excluded unreasonable skill words matched by the regular expressions that did not appear in the refined skill dictionary, establishing an initial correspondence between the *Job Requirement* and the skill requirements.

Based on this annotated data, we trained an NER model to extract required skills from the *Job Requirement* section for all job advertisements. Experts then aggregated the skills extracted by the NER model based on their meaning and content, grouping together those with similar meanings or repeated expressions. This process resulted in a skill dictionary \mathcal{S} of 2,324 standardized skill words, mapping original skill word descriptions to standardized skill words. The skill dictionary was then used to filter and map the skill words extracted by the NER model, ultimately obtaining standardized skill requirements for each job requirement. These standardized requirements were added to the job advertisement data as a new field, *Skill Requirements*.

Job Skill Demand Estimation. Generally, the demand for different skills in the job market can be estimated by the volume of job advertisements listing these specific skills as requirements within a given time period [30]. Formally, given job advertisement data $\mathcal{P} = \{\mathcal{P}_1, \dots, \mathcal{P}_t, \dots, \mathcal{P}_T\}$, where each \mathcal{P}_t represents the job advertisements posted at timestamp t , we use $D_{s,t} = \sum_{p \in \mathcal{P}_t} \mathbf{1}(s \in p)$ to estimate the demand for skill $s \in \mathcal{S}$ at time t . $s \in p$ indicates that job advertisement p requires skill s .

Along this line, we can calculate skill demand at various granularities, such as occupation and company levels. We define the sets of L1-level occupations, L2-level occupations, companies, and




Figure 1: Data analysis on *Job-SDF*. (a) illustrates the long-tail phenomenon of skill demands under the product manager and doctor occupations. (b) illustrates the results under the Chow test for the absence (left) and presence (right) of structural breaks.

regions as \mathcal{A}^{o_1} , \mathcal{A}^{o_2} , \mathcal{A}^c , and \mathcal{A}^r , respectively. The demands for skill s at time t under granularity $i \in \{o_1, o_2, c, r\}$ is then defined as follows:

$$D_{s,t}^i = [D_{s,t,a^i}^i]_{a^i \in \mathcal{A}^i}, \quad D_{s,t,a^i}^i = \sum_{p \in \mathcal{P}_t} \mathbf{1}(s \in p) \cdot \mathbf{1}(a^i \in p), \quad (1)$$

where $a^i \in p$ represents a job advertisement p containing the attribute a^i under granularity i . Similarly, we can further define skill demands $D_{s,t}^{i,j,\dots,k}$ across multiple granularities $\{i, j, \dots, k\}$ by calculating:

$$D_{s,t,\bar{a}}^{i,j,\dots,k} = \sum_{p \in \mathcal{P}_t} \mathbf{1}(s \in p) \cdot \mathbf{1}(a^i \in p \wedge a^j \in p \wedge \dots \wedge a^k \in p), \quad (2)$$

where $\bar{a} = \{a^i, a^j, \dots, a^k\}$, $a^i \in \mathcal{A}^i$, $a^j \in \mathcal{A}^j$, \dots , $a^k \in \mathcal{A}^k$, and $D_{s,t}^{i,j,\dots,k} \in \mathbb{R}^{|\mathcal{A}^i| \times |\mathcal{A}^j| \times \dots \times |\mathcal{A}^k|}$.

3.2 Job Skill Demand Forecasting Tasks

We study model performance through job skill demand forecasting tasks at different granularities, including single and multiple levels. The primary goal of these tasks is to predict future job skill demands based on historical time series data of various skills. Formally, we have:

Definition 1 (Job Skill Demand Forecasting) *Given a granularity or a set of granularities g and the observed job skill demand series from the previous K timestamps, i.e., $\{D_{:,t-K+1}^g, \dots, D_{:,t}^g\}$, the goal of job skill demand forecasting is to learn a forecasting model \mathcal{M} to predict the demand values for the next H timestamps, denoted by $\{\hat{D}_{:,t+1}^g, \dots, \hat{D}_{:,t+H}^g\}$.*

Our dataset includes skill demand time series data for L1-level occupations, L2-level occupations, companies, regions, and their combinations. We follow a standard protocol [59] that categorizes all time-series data into training, validation, and test sets in chronological order with a ratio of 9:1:2. In the main text, we demonstrate results with K set to 6 months and consider H as 3 months to evaluate the performance of different forecasting models. More settings and results can be found in the Appendix D and our project repository. Based on the *Job-SDF* dataset, other researchers can easily adjust the parameters to suit their research objectives.

3.3 Dataset Analysis

Varying Nature of Skill Demand. The values of skill demand exhibit significant differences and generally follow a long-tail distribution. This indicates that, at a specific granularity, only a few skills have high demand, while a wide range of skills are required by a limited number of jobs. For instance,

Figure 1a presents the skill demands under the product manager and doctor occupations in December 2022. The results clearly demonstrate the varying nature of skill demand values. This suggests that relying solely on metrics like RMSE to evaluate forecasting models’ performance may overlook the prediction accuracy for low-frequency skills.

Structural Break Phenomenon. As the labor market evolves, job skills that are not widely required today may become crucial in the future, while those currently in high demand may be supplanted by others. This dynamic can induce significant changes in the statistical properties of skill demand time series at various points in time. These changes may be reflected in the mean, variance, trend, or autocorrelation structure of the series. This phenomenon is known as structural breaks. A common method for detecting structural breaks is the Chow test, which evaluates whether there are significant differences in the regression coefficients across different periods [60]. Figure 1b illustrates the application of the Chow test in detecting structural breaks in various skill demand time series. The presence of structural breaks can impact the predictive accuracy of forecasting models. Further discussion will be provided in the experimental section.

Inter-Series Correlation. Intuitively, the proposed job skill demand forecasting tasks can be categorized as multivariate time-series forecasting tasks [61]. Figure 2 shows the absolute values of the Pearson correlation coefficients of different skill demand series for the backend development engineer, salesperson, and product manager occupations. We found that the time series data for some skills exhibit significant correlation within the same occupation (i.e., product design and market analysis in product manager), as well as for the same skills across different occupations (i.e., product design in backend development engineer and product manager). This demonstrates the necessity of considering all variables as inputs for job skill demand forecasting models, as it captures the interrelationships among variables, preventing the loss of critical information when variables are considered in isolation.

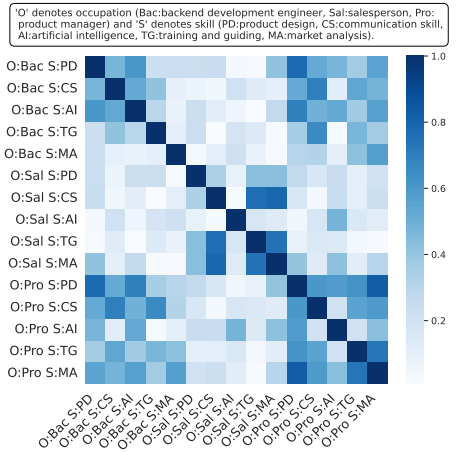


Figure 2: Pearson Correlation Coefficients.

Dataset Limitation. Recent studies [62] suggest that incorporating relationships between different variants can enhance the performance of multivariate time-series forecasting. However, due to the lack of prior knowledge, our Job-SDF dataset does not yet include a graph of relationships between skills, such as predecessor-successor relationships. Instead, we constructed a skill graph based on the co-occurrence of skills in job advertisements from the training data. This graph is included in our dataset. For detailed construction methods, please refer to the Appendix F.

4 Benchmark

4.1 Benchmark Models

We evaluated several SOTA time-series learning models using our proposed Job-SDF dataset. These models are categorized into six groups based on their underlying architectures: statistical time series models, RNN-based models, Transformer-based models, MLP-based models, Graph-based models, and Fourier-based models. The implementation details for each model are provided in the Appendix D, and the open-source model implementations are available on our GitHub repository.

Statistical Time Series Model. We first consider two statistical methodologies, namely *ARIMA* [63] and *Prophet* [64], both of which have been widely used in various contexts. The *ARIMA* model, which integrates differencing and moving averages within autoregression, has proven effective in forecasting occupational task demands [28]. *Prophet* decomposes time series data into trend, seasonality, and holiday components, allowing it to handle both linear and nonlinear trends with changepoints. However, these models often struggle to capture complex nonlinear relationships and exhibit suboptimal performance in large-scale data scenarios.

RNN-based Model. RNN-based methods are effective in capturing temporal state transitions through their recurrent structures, making them widely used in various time series forecasting tasks [65–

69]. Notably, *LSTM* have demonstrated their effectiveness in predicting changes in skill shares over time [29]. However, conventional RNNs often encounter performance degradation when handling excessively long look-back windows and forecast horizons. To address this challenge, *SegRNN* [32] introduces segment-wise iterations, which reduce the recurrence count within RNNs, thereby significantly enhancing performance in time series forecasting tasks.

Transformer-based Model. Recently, Transformer-based models [70] have gained widespread recognition in long-term time series forecasting due to their global modeling capabilities. Leveraging the attention mechanism, *Reformer* [37] introduces locally sensitive hashing to approximate attention by grouping similar queries. *Informer* [33] incorporates low-rank matrices in self-attention mechanisms to accelerate computation. *Autoformer*[34] employs block decomposition and autocorrelation mechanisms to more effectively capture the intrinsic features of time series data. *FedFormer* [36] utilizes DFT-based frequency-enhanced attention, obtaining attentive weights through the spectrums of queries and keys and calculating the weighted sum in the frequency domain. To address the challenges of non-stationary time series, the *Non-stationary Transformer (NSTRansformer)* [35] introduces a sequence stabilization module and proposes a de-stationary attention mechanism. Additionally, *PatchTST* [71] is a channel-independent patch time series transformer model that features patching and channel-independence as its key design elements.

MLP-based Model. Multiple Layer Projection (MLP) has been introduced in time series forecasting, demonstrating superior performance compared to transformer-based models in both accuracy and efficiency [38]. Specifically, *DLinear* [38] uses series decomposition as a pre-processing step before linear regression. *FreTS* [72] explores a novel approach by applying MLPs in the frequency domain for time series forecasting. *TSMixer* [39] employs MLP Mixer blocks, segments input time series into fixed windows, and applies gated MLP transformations and permutations to enhance accuracy.

Graph-based Models. Graph Neural Networks (GNNs) can learn non-Euclidean relationships, making them effective for identifying associations in structured data and generating joint representations from different perspectives [73–76]. CHGH [30] uses an adaptive graph enhanced by skill co-occurrence relationships to link skill supply and demand sequences. This fusion of representations across views improves the performance of joint skill supply and demand prediction tasks. Pre-DyGAE [52] targets skill demand prediction from an occupational perspective. It builds an occupation-skill bipartite graph based on the skill demands of occupations and captures the dynamic changes in these relationships. This method allows for predicting both potential occupational skills and skill demands, leveraging a dynamic graph perspective.

Fourier-based Models. By utilizing Fourier projection, *FiLM* [40] not only captures long-term time dependencies but also effectively reduces noise in forecasting. To address the challenge of non-stationary time-series forecasting, *Koopa* [41] disentangles time-variant and time-invariant components from complex non-stationary series using a Fourier Filter and designs the Koopman Predictor to forecast dynamics.

4.2 Evaluation Metrics

To evaluate the performance of various benchmark models in job skill demand forecasting tasks, we selected two commonly used regression metrics: MAE and RMSE. MAE is calculated over H observations using the formula: $\frac{1}{H} \sum_{i=1}^H |y_i - \hat{y}_i|$, where y_i represents the ground truth value and \hat{y}_i is the predicted value. RMSE is calculated as: $\sqrt{\frac{1}{H} \sum_{i=1}^H (y_i - \hat{y}_i)^2}$. Both MAE and RMSE are scale-dependent metrics, which makes them unsuitable for comparison across different granularities. Additionally, these metrics are less sensitive to prediction errors at lower skill demand values. Therefore, we additionally applied SMAPE [42] and RRMSE [77] to assess the performance of various forecasting models. SMAPE considers both the magnitude and direction of errors, making it suitable for comparing forecasts across different scales. RRMSE measures the square root of the average of the squared percentage errors.

$$SMAPE = \frac{2}{H} * \sum_{i=1}^H \frac{|y_i - \hat{y}_i|}{|y_i| + |\hat{y}_i|}, \quad RRMSE = \sqrt{\frac{\frac{1}{H} \sum_{i=1}^H (y_i - \hat{y}_i)^2}{\sum_{i=1}^H (\hat{y}_i)^2}}. \quad (3)$$

Table 1: Performance comparison on MAE and RMSE.

Model	L1-Occupation		L2-Occupation		Region&L1-O		Region&L2-O		Company	
	MAE	RMSE	MAE	RMSE	MAE	RMSE	MAE	RMSE	MAE	RMSE
ARIMA	20.27	256.89	6.46	115.79	3.98	58.65	1.31	27.42	1.31	38.88
Prophet	29.15	356.67	8.95	161.01	5.08	72.21	1.62	33.02	1.55	41.19
LSTM	19.05	194.67	7.09	116.36	3.92	51.59	1.29	23.31	1.35	26.47
SegRNN	12.28	108.28	5.01	68.83	3.14	34.26	1.05	15.96	1.01	16.03
CHGH	22.09	261.49	7.09	116.58	3.91	51.46	1.28	23.24	1.34	26.52
Pre-DyGAE	22.98	187.90	7.04	82.97	4.24	38.62	1.37	17.39	1.24	18.24
Transformer	22.06	215.09	7.58	118.21	4.01	52.04	1.35	23.44	1.26	24.99
Autoformer	23.06	186.76	8.22	100.02	6.45	57.77	2.41	24.10	3.31	38.55
Informer	22.21	205.24	7.43	117.38	3.88	50.13	1.30	23.07	1.26	24.92
Reformer	22.11	204.35	7.46	116.60	3.91	50.95	1.25	22.81	1.54	27.37
FEDformer	22.87	181.93	7.46	88.97	4.63	43.21	1.98	21.73	2.43	26.92
NStransformer	17.36	149.46	5.75	86.24	3.45	37.09	1.15	17.45	2.13	34.83
PatchTST	14.91	141.06	5.15	78.86	3.10	35.38	1.04	16.57	1.01	19.09
DLinear	16.61	154.88	5.44	81.61	3.24	36.67	1.07	16.79	1.05	18.85
TSMixer	21.34	192.85	8.14	106.65	5.81	62.14	5.95	68.26	13.96	144.96
FreTS	16.47	167.61	6.52	106.39	3.65	47.81	1.22	21.92	1.26	25.39
FiLM	12.95	117.17	5.08	65.65	3.24	29.90	1.14	14.01	1.17	15.87
Koopa	19.91	179.30	6.05	91.87	3.53	40.73	1.15	18.71	1.08	20.18

4.3 Benchmark Results

Overall Performance. In Table 1, we present the performance of various models evaluated using two metrics: MAE and RMSE. The following conclusions can be drawn: (1) The traditional statistical method, Prophet, demonstrates relatively poor predictive performance. This may be due to seasonal and holiday factors not being the primary influencers in skill demand prediction. (2) Most Transformer-based models, including Transformer, Autoformer, Informer, and Reformer, exhibit subpar overall predictive performance. This is likely because these models are designed to address long-range temporal dependencies, which are not well-suited for the current shorter time series context. (3) In contrast, PatchTST, unlike these Transformer-based models that perform point-wise modeling of time series, segments the time series into patches and inputs them into the Transformer. This allows the model to focus on more local information. A similar idea is also employed in the SegRNN. This strategy significantly enhances the performance of these models in predicting job skill demand. (4) The performance of different linear models on our dataset varies significantly. For instance, DLinear outperforms most Transformer-based models, while TSMixer performs poorly. This discrepancy may be due to the tendency of more complex MLP-based models to overfit our dataset. (5) CHGH and Pre-DyGAE exhibit poor performance in the separate skill demand forecasting scenario, likely due to a mismatch between their model design and the context of our dataset. Specifically, CHGH relies on sequential data from the supply side of skills, which is lacking in our dataset. Conversely, Pre-DyGAE focuses more on predicting whether a skill will be required by an occupation in the future. (6) Finally, FiLM achieved the best performance in most cases, demonstrating the robustness of the denoising-based model.

Low-Demand Skill Prediction Performance. Considering the varying nature of skill demand values, we further employed SMAPE and RRMSE metrics to focus on the predictive performance of different models for low-demand skills. As shown in Table 2, the experimental results indicate the following: (1) PatchTST achieved the best SMAPE performance in most cases, validating its ability to more accurately predict the trends of low-demand skills. (2) Based on scale-independent metrics, we can compare the performance of models at different granularities. It can be observed that RRMSE exhibits a significant trend of variation across different granularities; specifically, as the granularity becomes finer, the RRMSE performance deteriorates. This indicates that predicting skill demand at finer granularities is more challenging. Additionally, FiLM shows the least variation across multiple granularities, further validating its ability to provide stable and reliable predictions under varying granularities and demand value ranges. (3) Although Koopa performs averagely on MAE and RMSE metrics, it excels in predicting low-demand skills, particularly in terms of SMAPE. Similarly, NStransformer also performs well in scenarios focusing on low-demand skill predictions. This success can be attributed to both methods being designed to handle non-stationary time series. They effectively filter noise from historical sequences and restore intrinsic non-stationary information

Table 2: Performance comparison on SMAPE and RRMSE.

Model	L1-Occupation (%)		L2-Occupation (%)		Region&L1-O (%)		Region&L2-O (%)		Company (%)	
	SMAPE	RRMSE	SMAPE	RRMSE	SMAPE	RRMSE	SMAPE	RRMSE	SMAPE	RRMSE
ARIMA	35.72	47.89	25.00	58.87	23.86	58.07	13.58	73.57	20.17	147.94
Prophet	41.22	67.78	28.35	88.47	26.75	71.60	15.07	93.04	22.31	167.77
LSTM	41.38	57.90	32.85	83.70	31.58	68.40	22.93	87.36	30.26	174.40
SegRNN	39.81	37.58	33.35	50.53	35.30	48.53	23.84	61.90	33.07	86.27
CHGH	40.27	66.05	29.60	84.10	28.11	68.42	17.42	87.45	26.72	176.70
PreDyGAE	49.87	83.67	60.54	83.60	59.32	66.56	72.67	98.09	26.21	145.73
Transformer	55.59	64.25	44.23	84.27	31.15	76.16	33.04	86.87	27.61	164.36
Autoformer	70.28	53.75	74.37	63.40	90.14	65.57	91.51	74.46	107.05	99.60
Informer	56.85	58.18	44.04	88.72	34.75	69.59	29.29	90.15	32.41	164.37
Reformer	56.58	61.35	40.58	83.70	32.21	72.87	20.86	90.85	45.25	169.87
FEDformer	69.30	54.03	69.29	60.00	73.17	52.69	81.73	70.06	94.19	97.97
NStransformer	38.11	47.19	26.30	60.73	24.98	48.89	14.55	63.29	24.20	100.78
PatchTST	34.70	51.17	24.52	58.80	25.15	44.96	13.50	67.48	19.89	115.34
DLinear	41.84	52.89	34.35	60.22	33.47	51.05	25.77	64.65	30.71	108.66
TSMixer	56.59	61.17	72.29	99.35	82.48	87.29	120.85	96.49	155.20	102.14
FreTS	39.76	54.42	30.18	80.44	28.58	66.11	17.62	85.04	27.24	174.56
FiLM	39.51	37.55	29.65	43.86	28.79	37.66	17.24	47.75	25.72	76.92
Koopa	37.84	58.30	25.72	65.34	24.41	57.81	13.98	74.00	20.43	123.96

Table 3: Performance comparison on data with structural breaks on MAE and RMSE.

Model	L1-Occupation		L2-Occupation		Region&L1-O		Region&L2-O		Company	
	MAE	RMSE	MAE	RMSE	MAE	RMSE	MAE	RMSE	MAE	RMSE
LSTM	87.30	554.46	57.95	400.22	18.99	149.53	7.91	52.38	24.40	159.02
SegRNN	61.92	390.54	43.97	276.57	15.85	114.04	6.56	37.84	17.98	112.13
CHGH	94.30	629.32	58.06	401.45	19.00	149.75	7.90	52.50	24.37	159.44
PreGyGAE	78.35	493.83	48.69	336.15	17.49	136.66	7.31	38.88	19.76	164.43
Transformer	98.66	580.58	61.73	404.17	19.37	151.12	8.45	55.46	22.41	152.27
Autoformer	107.22	533.06	67.66	350.97	26.84	156.50	12.19	63.04	44.10	208.96
Informer	98.89	570.35	59.95	402.75	19.03	146.91	7.72	49.15	22.37	151.87
Reformer	98.14	569.83	60.71	401.21	19.25	149.91	7.52	49.10	25.65	160.69
FEDformer	105.43	532.24	62.10	325.10	20.49	128.45	10.37	55.47	34.09	155.28
NStransformer	82.43	462.24	49.30	318.44	16.59	119.91	6.85	37.56	40.05	196.03
PatchTST	77.44	474.86	45.02	303.76	14.88	111.01	6.56	38.60	18.03	127.72
DLinear	81.17	485.25	46.67	307.34	15.94	118.94	6.50	37.72	18.18	124.32
TSMixer	107.47	614.93	83.60	479.39	29.99	187.08	25.83	190.29	155.10	766.58
FreTS	82.45	537.12	56.54	393.38	18.55	148.33	7.88	52.87	24.21	160.01
FiLM	62.86	404.82	42.63	260.99	14.31	101.23	6.37	32.28	18.78	110.65
Koopa	91.26	516.75	50.44	324.15	17.43	128.39	7.07	41.29	19.04	133.26

into time-dependent relationships, making them more adept at handling the fluctuating nature of low-demand skill time series data.

Performance on Skill Demand Series with Structural Breaks. As described in Section 3.3, in the dynamically changing job market, skill demand time series data exhibit structural breaks. To assess the impact of this phenomenon on different models in the skill demand forecasting task, we used the Chow test to detect structural breaks in the skill demand time series. The corresponding predictive performance of different models is presented in Tables 3 and 4. We observe the following phenomena: (1) Compared to the predictive performance on the full dataset, the performance on time series data with structural breaks is significantly worse. This finding underscores the complexity and unpredictability of skill trends that experience structural breaks. (2) FiLM has achieved results close to the overall skill demand prediction in terms of SMAPE and RRMSE metrics. This validates that FiLM can effectively mitigate the disruptive impact of structural breaks on skill demand forecasting. (3) Furthermore, while the overall predictive performance of skill demand forecasting at both the Region&L2-O and Company granularity levels is similar, significant differences emerge when forecasting skills experiencing structural breaks. This suggests that skills undergoing structural breaks display more predictable patterns at the Region&L2-O granularity level compared to the Company level, making them relatively easier to forecast.

Table 4: Performance comparison on data with structural breaks on RRMSE and SMAPE.

Model	L1-Occupation (%)		L2-Occupation (%)		Region&L1-O (%)		Region&L2-O (%)		Company (%)	
	SMAPE	RRMSE	SMAPE	RRMSE	SMAPE	RRMSE	SMAPE	RRMSE	SMAPE	RRMSE
LSTM	43.78	58.05	48.93	84.46	46.64	78.31	42.03	58.48	68.38	187.30
SegRNN	39.22	37.80	43.09	51.14	45.17	54.31	39.41	41.40	57.45	89.65
CHGH	44.91	66.31	48.90	84.87	45.43	78.32	39.79	58.89	68.36	189.91
PreDyGAE	52.35	47.15	56.56	59.31	52.06	61.22	44.13	42.31	70.26	106.88
Transformer	50.01	64.47	53.10	84.95	46.50	86.56	47.67	61.23	64.92	177.43
Autoformer	63.46	54.08	68.62	64.14	87.93	68.97	88.95	63.85	115.00	100.60
Informer	51.11	58.40	51.89	89.70	47.81	80.86	44.90	57.55	65.11	177.16
Reformer	50.79	61.59	51.51	84.53	46.86	84.15	40.81	58.59	72.36	181.36
FEDformer	62.83	54.37	64.37	60.84	72.24	58.55	80.03	54.29	103.27	100.65
NSTransformer	45.36	47.46	47.63	61.85	43.04	57.60	36.72	39.72	170.57	113.87
PatchTST	40.89	51.48	43.26	59.69	41.51	51.85	34.74	43.12	55.26	122.56
DLinear	43.14	53.20	45.25	61.13	45.26	58.80	41.15	41.71	57.65	115.24
TSMixer	54.31	61.31	76.08	99.84	85.12	95.81	117.39	93.66	160.55	102.23
FreTS	42.44	54.59	48.24	81.17	45.39	75.43	39.85	57.83	68.39	187.94
FILM	38.96	37.82	44.23	44.52	44.95	43.06	40.05	30.80	56.37	80.77
Koopa	46.45	58.59	47.13	66.28	42.60	66.20	36.24	47.48	58.98	131.77

5 Conclusion

In this work, we introduced Job-SDF, a dataset designed for training and benchmarking job-skill demand forecasting models. Compiled from millions of public job advertisements collected from online recruitment platforms, this dataset includes monthly recruitment demand for 2,324 types of skills across 52 occupations, 521 companies, and 7 regions. Using this dataset, we validated a wide range of time-series forecasting approaches, including statistical models, RNN-based models, Transformer-based models, MLP-based models, Graph-based models, and Fourier-based models. Furthermore, we conducted extensive experiments to compare the performance of various methods in predicting skill demand at different granularities. We hope that Job-SDF will facilitate further research in this field.

Acknowledgements

This work was supported in part by the National Key R&D Program of China (Grant No.2023YFF0725001), in part by the National Natural Science Foundation of China (Grant No.92370204), in part by the guangdong Basic and Applied Basic Research Foundation (Grant No.2023B1515120057), in part by Guangzhou-HKUST (GZ) Joint Funding Program (Grant No.2023A03J0008), Education Bureau of Guangzhou Municipality, in part by Nansha Postdoctoral Research Project, and in part by the National Natural Science Foundation of China (Grant No.62176014), the Fundamental Research Funds for the Central Universities.

References

- [1] David Autor et al. The polarization of job opportunities in the us labor market: Implications for employment and earnings. *Center for American Progress and The Hamilton Project*, 6:11–19, 2010.
- [2] James J Heckman, Jora Stixrud, and Sergio Urzua. The effects of cognitive and noncognitive abilities on labor market outcomes and social behavior. *Journal of Labor economics*, 24(3):411–482, 2006.
- [3] Chuan Qin, Hengshu Zhu, Dazhong Shen, Ying Sun, Kaichun Yao, Peng Wang, and Hui Xiong. Automatic skill-oriented question generation and recommendation for intelligent job interviews. *ACM Transactions on Information Systems*, 42(1):1–32, 2023.
- [4] Marios Kokkodis and Panagiotis G Ipeirotis. Demand-aware career path recommendations: A reinforcement learning approach. *Management science*, 67(7):4362–4383, 2021.
- [5] Rui Zha, Chuan Qin, Le Zhang, Dazhong Shen, Tong Xu, Hengshu Zhu, and Enhong Chen. Career mobility analysis with uncertainty-aware graph autoencoders: A job title transition perspective. *IEEE Transactions on Computational Social Systems*, 11(1):1205–1215, 2023.

- [6] Rui Zha, Ying Sun, Chuan Qin, Le Zhang, Tong Xu, Hengshu Zhu, and Enhong Chen. Towards unified representation learning for career mobility analysis with trajectory hypergraph. *ACM Transactions on Information Systems*, 42(4):1–28, 2024.
- [7] Xiaoshan Yu, Chuan Qin, Qi Zhang, Chen Zhu, Haiping Ma, Xingyi Zhang, and Hengshu Zhu. Disco: A hierarchical disentangled cognitive diagnosis framework for interpretable job recommendation. *arXiv preprint arXiv:2410.07671*, 2024.
- [8] Ying Sun, Hengshu Zhu, Lu Wang, Le Zhang, and Hui Xiong. Large-scale online job search behaviors reveal labor market shifts amid covid-19. *Nature Cities*, 1(2):150–163, 2024.
- [9] Le Dai, Yu Yin, Chuan Qin, Tong Xu, Xiangnan He, Enhong Chen, and Hui Xiong. Enterprise cooperation and competition analysis with a sign-oriented preference network. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 774–782, 2020.
- [10] Hengshu Zhu, Hui Xiong, Fangshuang Tang, Qi Liu, Yong Ge, Enhong Chen, and Yanjie Fu. Days on market: Measuring liquidity in real estate markets. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 393–402, 2016.
- [11] Chuan Qin, Hengshu Zhu, Tong Xu, Chen Zhu, Liang Jiang, Enhong Chen, and Hui Xiong. Enhancing person-job fit for talent recruitment: An ability-aware neural network approach. In *The 41st international ACM SIGIR conference on research & development in information retrieval*, pages 25–34, 2018.
- [12] Chuan Qin, Hengshu Zhu, Chen Zhu, Tong Xu, Fuzhen Zhuang, Chao Ma, Jingshuai Zhang, and Hui Xiong. Duerquiz: A personalized question recommender system for intelligent job interview. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 2165–2173, 2019.
- [13] Chuan Qin, Hengshu Zhu, Tong Xu, Chen Zhu, Chao Ma, Enhong Chen, and Hui Xiong. An enhanced neural network approach to person-job fit in talent recruitment. *ACM Transactions on Information Systems (TOIS)*, 38(2):1–33, 2020.
- [14] Kaichun Yao, Jingshuai Zhang, Chuan Qin, Peng Wang, Hengshu Zhu, and Hui Xiong. Knowledge enhanced person-job fit for talent recruitment. In *2022 IEEE 38th International Conference on Data Engineering (ICDE)*, pages 3467–3480. IEEE, 2022.
- [15] Likang Wu, Zhaopeng Qiu, Zhi Zheng, Hengshu Zhu, and Enhong Chen. Exploring large language model for graph data understanding in online job recommendations. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 9178–9186, 2024.
- [16] Shuqing Bian, Xu Chen, Wayne Xin Zhao, Kun Zhou, Yupeng Hou, Yang Song, Tao Zhang, and Ji-Rong Wen. Learning to match jobs with resumes from sparse interaction data using multi-view co-teaching network. In *Proceedings of the 29th ACM International Conference on Information & Knowledge Management*, pages 65–74, 2020.
- [17] Yong Luo, Huaizheng Zhang, Yonggang Wen, and Xinwen Zhang. Resumegan: an optimized deep representation learning framework for talent-job fit via adversarial learning. In *Proceedings of the 28th ACM international conference on information and knowledge management*, pages 1101–1110, 2019.
- [18] Yoosof Mashayekhi, Nan Li, Bo Kang, Jeffrey Lijffijt, and Tijn De Bie. A challenge-based survey of e-recruitment recommendation systems. *ACM Computing Surveys*, 56(10):1–33, 2024.
- [19] Yong Luo, Huaizheng Zhang, Yongjie Wang, Yonggang Wen, and Xinwen Zhang. Resumenet: A learning-based framework for automatic resume quality assessment. In *2018 IEEE International Conference on Data Mining (ICDM)*, pages 307–316. IEEE, 2018.
- [20] Dazhong Shen, Chuan Qin, Hengshu Zhu, Tong Xu, Enhong Chen, and Hui Xiong. Joint representation learning with relation-enhanced topic models for intelligent job interview assessment. *ACM Transactions on Information Systems (TOIS)*, 40(1):1–36, 2021.

- [21] Feihu Jiang, Chuan Qin, Kaichun Yao, Chuyu Fang, Fuzhen Zhuang, Hengshu Zhu, and Hui Xiong. Enhancing question answering for enterprise knowledge bases using large language models. In *International Conference on Database Systems for Advanced Applications*, pages 273–290. Springer, 2024.
- [22] David J Deming. The growing importance of social skills in the labor market. *The Quarterly Journal of Economics*, 132(4):1593–1640, 2017.
- [23] Liyi Chen, Chuan Qin, Ying Sun, Xin Song, Tong Xu, Hengshu Zhu, and Hui Xiong. Collaboration-aware hybrid learning for knowledge development prediction. In *Proceedings of the ACM on Web Conference 2024*, pages 3976–3985, 2024.
- [24] Yunfei Zhang, Chuan Qin, Dazhong Shen, Haiping Ma, Le Zhang, Xingyi Zhang, and Hengshu Zhu. Relicid: a reliable cognitive diagnosis framework with confidence awareness. In *2023 IEEE International Conference on Data Mining (ICDM)*, pages 858–867. IEEE, 2023.
- [25] William Rasdorf, Joseph E Hummer, and Stephanie C Vereen. Data collection opportunities and challenges for skilled construction labor demand forecast modeling. *Public Works Management & Policy*, 21(1):28–52, 2016.
- [26] Joshua Healy, Kostas Mavromaras, and Peter J Sloane. Adjusting to skill shortages in australian smes. *Applied Economics*, 47(24):2470–2487, 2015.
- [27] Lutz Bellmann and Olaf Hübler. The skill shortage in german establishments before, during and after the great recession. *Jahrbücher für Nationalökonomie und Statistik*, 234(6):800–828, 2014.
- [28] Subhro Das, Sebastian Steffen, Wyatt Clarke, Prabhat Reddy, Erik Brynjolfsson, and Martin Fleming. Learning occupational task-shares dynamics for the future of work. In *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*, pages 36–42, 2020.
- [29] Maysa Malfiza Garcia de Macedo, Wyatt Clarke, Eli Lucherini, Tyler Baldwin, Dilermando Queiroz Neto, Rogerio Abreu de Paula, and Subhro Das. Practical skills demand forecasting via representation learning of temporal dynamics. In *Proceedings of the 2022 AAAI/ACM Conference on AI, Ethics, and Society*, pages 285–294, 2022.
- [30] Wenshuo Chao, Zhaopeng Qiu, Likang Wu, Zhuoning Guo, Zhi Zheng, Hengshu Zhu, and Hao Liu. A cross-view hierarchical graph learning hypernetwork for skill demand-supply joint prediction. In *AAAI*, 2024.
- [31] Adebisi A Ariyo, Adewumi O Adewumi, and Charles K Ayo. Stock price prediction using the arima model. In *2014 UKSim-AMSS 16th international conference on computer modelling and simulation*, pages 106–112. IEEE, 2014.
- [32] Shengsheng Lin, Weiwei Lin, Wentai Wu, Feiyu Zhao, Ruichao Mo, and Haotong Zhang. Segrnn: Segment recurrent neural network for long-term time series forecasting. *arXiv preprint arXiv:2308.11200*, 2023.
- [33] Haoyi Zhou, Shanghang Zhang, Jieqi Peng, Shuai Zhang, Jianxin Li, Hui Xiong, and Wancai Zhang. Informer: Beyond efficient transformer for long sequence time-series forecasting. In *Proceedings of the AAAI conference on artificial intelligence*, pages 11106–11115, 2021.
- [34] Haixu Wu, Jiehui Xu, Jianmin Wang, and Mingsheng Long. Autoformer: Decomposition transformers with auto-correlation for long-term series forecasting. *Advances in neural information processing systems*, 34:22419–22430, 2021.
- [35] Yong Liu, Haixu Wu, Jianmin Wang, and Mingsheng Long. Non-stationary transformers: Exploring the stationarity in time series forecasting. *Advances in Neural Information Processing Systems*, 35:9881–9893, 2022.
- [36] Tian Zhou, Ziqing Ma, Qingsong Wen, Xue Wang, Liang Sun, and Rong Jin. Fedformer: Frequency enhanced decomposed transformer for long-term series forecasting. In *International conference on machine learning*, pages 27268–27286. PMLR, 2022.

- [37] Nikita Kitaev, Łukasz Kaiser, and Anselm Levskaya. Reformer: The efficient transformer. *arXiv preprint arXiv:2001.04451*, 2020.
- [38] Ailing Zeng, Muxi Chen, Lei Zhang, and Qiang Xu. Are transformers effective for time series forecasting? In *Proceedings of the AAAI conference on artificial intelligence*, pages 11121–11128, 2023.
- [39] Shiyu Wang, Haixu Wu, Xiaoming Shi, Tengge Hu, Huakun Luo, Lintao Ma, James Y Zhang, and Jun Zhou. Timemixer: Decomposable multiscale mixing for time series forecasting. *arXiv preprint arXiv:2405.14616*, 2024.
- [40] Tian Zhou, Ziqing Ma, Qingsong Wen, Liang Sun, Tao Yao, Wotao Yin, Rong Jin, et al. Film: Frequency improved legendre memory model for long-term time series forecasting. *Advances in Neural Information Processing Systems*, 35:12677–12690, 2022.
- [41] Yong Liu, Chenyu Li, Jianmin Wang, and Mingsheng Long. Koopa: Learning non-stationary time series dynamics with koopman predictors. *Advances in Neural Information Processing Systems*, 36, 2024.
- [42] Saba Sareminia. A support vector based hybrid forecasting model for chaotic time series: Spare part consumption prediction. *Neural Processing Letters*, 55(3):2825–2841, 2023.
- [43] Chao Chen, Jamie Twycross, and Jonathan M Garibaldi. A new accuracy measure based on bounded relative error for time series forecasting. *PLoS one*, 12(3):e0174202, 2017.
- [44] Alan Manning. We can work it out: the impact of technological change on the demand for low-skill workers. *Scottish Journal of Political Economy*, 51(5):581–608, 2004.
- [45] John M Abowd, John C Haltiwanger, Julia Lane, Kevin L McKinney, and Kristin Sandusky. Technology and the demand for skill: an analysis of within and between firm differences, 2007.
- [46] Chuan Qin, Le Zhang, Yihang Cheng, Rui Zha, Dazhong Shen, Qi Zhang, Xi Chen, Ying Sun, Chen Zhu, Hengshu Zhu, et al. A comprehensive survey of artificial intelligence techniques for talent analytics. *arXiv preprint arXiv:2307.03195*, 2023.
- [47] Chinhui Juhn. Wage inequality and demand for skill: evidence from five decades. *ILR Review*, 52(3):424–443, 1999.
- [48] Jong-Wha Lee and Dainn Wie. Technological change, skill demand, and wage inequality: Evidence from indonesia. *World Development*, 67:238–250, 2015.
- [49] Ying Sun, Fuzhen Zhuang, Hengshu Zhu, Qi Zhang, Qing He, and Hui Xiong. Market-oriented job skill valuation with cooperative composition neural network. *Nature communications*, 12(1):1992, 2021.
- [50] George S Benson and Edward E Lawler III. Raising skill demand: Generating good jobs. *Transforming the US Workforce Development System. Urbana: Labor and Employment Relations Association*, 2011.
- [51] Margaret Hilton. *Research on future skill demands: A workshop summary*. National Academies Press, 2008.
- [52] Xi Chen, Chuan Qin, Zhigaoyuan Wang, Yihang Cheng, Chao Wang, Hengshu Zhu, and Hui Xiong. Pre-dygae: Pre-training enhanced dynamic graph autoencoder for occupational skill demand forecasting. In *Proceedings of the 33th International Joint Conference on Artificial Intelligence*, 2024.
- [53] Manuel Cifuentes, Jon Boyer, David A Lombardi, and Laura Punnett. Use of o* net as a job exposure matrix: a literature review. *American journal of industrial medicine*, 53(9):898–914, 2010.
- [54] Filippo Chiarello, Gualtiero Fantoni, Terence Hogarth, Vito Giordano, Liga Baltina, and Irene Spada. Towards esco 4.0—is the european classification of skills in line with industry 4.0? a text mining approach. *Technological Forecasting and Social Change*, 173:121177, 2021.

- [55] Chuan Qin, Kaichun Yao, Hengshu Zhu, Tong Xu, Dazhong Shen, Enhong Chen, and Hui Xiong. Towards automatic job description generation with capability-aware neural networks. *IEEE Transactions on Knowledge and Data Engineering*, 35(5):5341–5355, 2022.
- [56] Chuyu Fang, Chuan Qin, Qi Zhang, Kaichun Yao, Jingshuai Zhang, Hengshu Zhu, Fuzhen Zhuang, and Hui Xiong. Recruitpro: A pretrained language model with skill-aware prompt learning for intelligent recruitment. In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pages 3991–4002, 2023.
- [57] Kaichun Yao, Chuan Qin, Hengshu Zhu, Chao Ma, Jingshuai Zhang, Yi Du, and Hui Xiong. An interactive neural network approach to keyphrase extraction in talent recruitment. In *Proceedings of the 30th ACM International Conference on Information & Knowledge Management*, pages 2383–2393, 2021.
- [58] Feihu Jiang, Chuan Qin, Jingshuai Zhang, Kaichun Yao, Xi Chen, Dazhong Shen, Chen Zhu, Hengshu Zhu, and Hui Xiong. Towards efficient resume understanding: A multi-granularity multi-modal pre-training approach. *arXiv preprint arXiv:2404.13067*, 2024.
- [59] Zhiyuan Wang, Xovee Xu, Weifeng Zhang, Goce Trajcevski, Ting Zhong, and Fan Zhou. Learning latent seasonal-trend representations for time series forecasting. *Advances in Neural Information Processing Systems*, 35:38775–38787, 2022.
- [60] Gregory C Chow. Tests of equality between sets of coefficients in two linear regressions. *Econometrica: Journal of the Econometric Society*, pages 591–605, 1960.
- [61] Defu Cao, Yujing Wang, Juanyong Duan, Ce Zhang, Xia Zhu, Congrui Huang, Yunhai Tong, Bixiong Xu, Jing Bai, Jie Tong, et al. Spectral temporal graph neural network for multivariate time-series forecasting. *Advances in neural information processing systems*, 33:17766–17778, 2020.
- [62] Ting Guo, Feng Hou, Yan Pang, Xiaoyun Jia, Zhongwei Wang, and Ruili Wang. Learning and integration of adaptive hybrid graph structures for multivariate time series forecasting. *Information Sciences*, 648:119560, 2023.
- [63] George EP Box, Gwilym M Jenkins, Gregory C Reinsel, and Greta M Ljung. *Time series analysis: forecasting and control*. John Wiley & Sons, 2015.
- [64] Sean J Taylor and Benjamin Letham. Forecasting at scale. *The American Statistician*, 72(1):37–45, 2018.
- [65] Chen Zhu, Hengshu Zhu, Hui Xiong, Pengliang Ding, and Fang Xie. Recruitment market trend analysis with sequential latent variable models. In *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining*, pages 383–392, 2016.
- [66] Qi Zhang, Tong Xu, Hengshu Zhu, Lifu Zhang, Hui Xiong, Enhong Chen, and Qi Liu. After-shock detection with multi-scale description based neural network. In *2019 IEEE International Conference on Data Mining (ICDM)*, pages 886–895. IEEE, 2019.
- [67] Qi Zhang, Hengshu Zhu, Ying Sun, Hao Liu, Fuzhen Zhuang, and Hui Xiong. Talent demand forecasting with attentive neural sequential model. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*, pages 3906–3916, 2021.
- [68] Qi Zhang, Hengshu Zhu, Qi Liu, Enhong Chen, and Hui Xiong. Exploiting real-time search engine queries for earthquake detection: A summary of results. *ACM Transactions on Information Systems (TOIS)*, 39(3):1–32, 2021.
- [69] Dazhong Shen, Qi Zhang, Tong Xu, Hengshu Zhu, Wenjia Zhao, Zikai Yin, Peilun Zhou, Lihua Fang, Enhong Chen, and Hui Xiong. A machine learning-enhanced robust p-phase picker for real-time seismic monitoring. *arXiv preprint arXiv:1911.09275*, 2019.
- [70] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.

- [71] Yuqi Nie, Nam H Nguyen, Phanwadee Sinthong, and Jayant Kalagnanam. A time series is worth 64 words: Long-term forecasting with transformers. *arXiv preprint arXiv:2211.14730*, 2022.
- [72] Kun Yi, Qi Zhang, Wei Fan, Shoujin Wang, Pengyang Wang, Hui He, Ning An, Defu Lian, Longbing Cao, and Zhendong Niu. Frequency-domain mlps are more effective learners in time series forecasting. *Advances in Neural Information Processing Systems*, 36, 2024.
- [73] Liyi Chen, Zhi Li, Tong Xu, Han Wu, Zhefeng Wang, Nicholas Jing Yuan, and Enhong Chen. Multi-modal siamese network for entity alignment. In *Proceedings of the 28th ACM SIGKDD conference on knowledge discovery and data mining*, pages 118–126, 2022.
- [74] Liyi Chen, Zhi Li, Weidong He, Gong Cheng, Tong Xu, Nicholas Jing Yuan, and Enhong Chen. Entity summarization via exploiting description complementarity and salience. *IEEE Transactions on Neural Networks and Learning Systems*, 34(11):8297–8309, 2023.
- [75] Dazhong Shen, Chuan Qin, Chao Wang, Hengshu Zhu, Enhong Chen, and Hui Xiong. Regularizing variational autoencoder with diversity and uncertainty awareness. *arXiv preprint arXiv:2110.12381*, 2021.
- [76] Miao Chen, Chao Wang, Chuan Qin, Tong Xu, Jianhui Ma, Enhong Chen, and Hui Xiong. A trend-aware investment target recommendation system with heterogeneous graph. In *2021 International Joint Conference on Neural Networks (IJCNN)*, pages 1–8. IEEE, 2021.
- [77] Keerti Rawal and Aijaz Ahmad. Mining latent patterns with multi-scale decomposition for electricity demand and price forecasting using modified deep graph convolutional neural networks. *Sustainable Energy, Grids and Networks*, page 101436, 2024.
- [78] Benedek Rozemberczki, Paul Scherer, Yixuan He, George Panagopoulos, Alexander Riedel, Maria Astefanoaei, Oliver Kiss, Ferenc Beres, Guzman Lopez, Nicolas Collignon, and Rik Sarkar. PyTorch Geometric Temporal: Spatiotemporal Signal Processing with Neural Machine Learning Models. In *Proceedings of the 30th ACM International Conference on Information and Knowledge Management*, page 4564–4573, 2021.
- [79] Haixu Wu, Tengge Hu, Yong Liu, Hang Zhou, Jianmin Wang, and Mingsheng Long. Timesnet: Temporal 2d-variation modeling for general time series analysis. In *International Conference on Learning Representations*, 2023.
- [80] Aldo Pareja, Giacomo Domeniconi, Jie Chen, Tengfei Ma, Toyotaro Suzumura, Hiroki Kanezashi, Tim Kaler, Tao Schardl, and Charles Leiserson. Evolvegc: Evolving graph convolutional networks for dynamic graphs. In *Proceedings of the AAAI conference on artificial intelligence*, pages 5363–5370, 2020.
- [81] Youngjoo Seo, Michaël Defferrard, Pierre Vandergheynst, and Xavier Bresson. Structured sequence modeling with graph convolutional recurrent networks. In *Neural Information Processing: 25th International Conference, ICONIP 2018, Siem Reap, Cambodia, December 13-16, 2018, Proceedings, Part I 25*, pages 362–373. Springer, 2018.
- [82] Ling Zhao, Yujiao Song, Chao Zhang, Yu Liu, Pu Wang, Tao Lin, Min Deng, and Haifeng Li. T-GCN: A Temporal Graph Convolutional Network for Traffic Prediction. *IEEE Transactions on Intelligent Transportation Systems*, pages 3848–3858, 2020.
- [83] Jinyin Chen, Xueke Wang, and Xuanheng Xu. Gc- lstm: Graph convolution embedded lstm for dynamic network link prediction. *Applied Intelligence*, pages 1–16, 2022.
- [84] Aynaz Taheri, Kevin Gimpel, and Tanya Berger-Wolf. Learning to represent the evolution of dynamic graphs with recurrent models. In *Companion Proceedings of The 2019 World Wide Web Conference, WWW '19*, New York, NY, USA, 2019. Association for Computing Machinery.

A Ethics Statement

In accordance with the Code of Ethics as outlined by NeurIPS ¹, we present the following discussion on the ethical considerations and potential societal impact of our research.

Privacy. Our data is collected from online platforms, encompassing job advertisements, which is public by various companies. This collection method ensures that all information utilized is already available to the public, having been disclosed by companies for the purpose of recruitment. Consequently, our analysis inherently respects privacy norms, as it involves no proprietary or confidential data, relying solely on publicly shared statistics and general information.

Consent. We have extensively collected job advertisement data from online platforms where the information is openly and transparently available. Our analysis is based on the statistical data derived from these job advertisements, rather than the original textual information, significantly minimizing the risk of privacy breaches. Therefore, the use of this data is permissible.

Deprecated Datasets. Not applicable.

Copyright and Fair Use. The library used in our research is publicly available and distributed under the MIT License. We list the used assets along with their licenses as follows:

- PyTorch Geometric Temporal [78] consists of various dynamic and temporal geometric deep learning, embedding, and spatio-temporal regression methods from a variety of published research papers. This library is open-source and available at https://github.com/benedekroemberczki/pytorch_geometric_temporal, which is publicly available and distributed under the MIT License.
- Time Series Library [79] is an open-source library for deep learning researchers, especially for deep time series analysis. This library is available at <https://github.com/thuml/Time-Series-Library>, which is publicly available and distributed under the MIT License.

Representative Evaluation Practice. Job-SDF has extensively covered skill demand data at multiple granular levels, including regions, occupations, and companies. The detailed nature of the data proves invaluable in advancing tasks related to skill demand prediction. It significantly assists job seekers in identifying employment opportunities, employers in recruiting suitable talent, and educational institutions in tailoring their curricula to meet market demands. Moreover, this comprehensive dataset not only facilitates the analysis of skill trends but also serves as a platform for verifying time-series algorithms within a specialized application domain.

Safety. Our research does not employ technologies that directly inflict harm on individuals.

Security. Since the data sources are all publicly available, there is no risk of security incidents.

Discrimination. Our data is strictly limited to market analysis and does not judge the superiority or inferiority of skills themselves, thus it does not involve any issues of discrimination.

Deception & Harassment. We firmly believe that our data cannot be used for deceptive practices or harassment activities.

Human Rights. Not applicable.

Bias and Fairness. Since our data is collected from the public job postings of a subset of companies, there may be statistical biases in skill demands across different industries or regions. While it may not represent the entirety, it effectively captures the main trends in skill demands.

¹<https://neurips.cc/public/EthicsGuidelines>

B Code and Data Availability

The dataset is CC BY-NC-SA 4.0 licensed, accessible via the URL <https://github.com/Job-SDF/benchmark>. We also provide a website for the project: <https://job-sdf.github.io/> with information on how to use the data and code.

C Computational Resource

Due to inherent design and size constraints of the models combined with varying data sizes at different granularities, the deployment environments for each model are distinct. The CHGH model, which requires over 80GB of memory, is exclusively deployed on CPU platforms to accommodate its substantial resource demands. In contrast, the PreDyGAE model operates solely on GPU infrastructure, leveraging the computational efficiencies of the NVIDIA A800 GPUs. For other models, deployment strategies are tailored according to the granularity of the data. Experiments at the labor market, regions, L1 occupations, L2 occupations, and Region & L1 occupations granularities are conducted on GPUs, capitalizing on the enhanced processing capabilities of these units for handling moderate data volumes. However, at the granularities of Region & L2 and company, where data volumes are significantly larger, deployment shifts to CPUs. Overall, the training time of different models are shown in Table 5.

Table 5: Training time (minute) of different models for job skill demand forecasting.

Model	Market	Region	L1-O	L2-O	R&L1-O	R&L2-O	Company
LSTM	0-0.5	0-0.5	0-0.5	0-0.5	0-0.5	37.7	39.0
SegRNN	0-0.5	0-0.5	0-0.5	0-0.5	0-0.5	342.8	458.2
CHGH	17.7	132.8	170.2	258.3	1300.3	490.6	6604.2
PreDyGAE	1-10	16.5	30.0	48.1	48.1	88.2	126.2
Transformer	0-0.5	0-0.5	0-0.5	0.5-1	0.5-1	128.2	166.5
Autoformer	0-0.5	0-0.5	0-0.5	0-0.5	0.5-1	304.3	325.0
Informer	0-0.5	0-0.5	0-0.5	0-0.5	0-0.5	133.8	171.7
Reformer	0-0.5	0-0.5	0-0.5	0-0.5	0-0.5	36.0	52.5
FEDformer	0-0.5	0-0.5	0-0.5	0-0.5	0-0.5	193.8	198.5
NStransformer	0-0.5	0.5-1	0.5-1	0-0.5	0.5-1	128.5	195.7
PatchTST	0-0.5	0-0.5	0-0.5	0-0.5	0.5-1	1202.8	2558.0
DLinear	0-0.5	0-0.5	0.5-1	0-0.5	0-0.5	20.0	39.1
TSMixer	0-0.5	0-0.5	0-0.5	0-0.5	0-0.5	24.0	97.0
FreTS	0-0.5	0-0.5	0-0.5	0-0.5	0-0.5	85.0	200.0
FiLM	0-0.5	0.5-1	0.5-1	1-10	1-10	598.0	1464.7
Koopa	0-0.5	0-0.5	0-0.5	0-0.5	0-0.5	38.3	68.5

D Additional Experimental Results

Due to page limitations in the main text, we have included additional experimental content in the appendix. First, we present the results of repeated trials of the benchmark models discussed in the main text in the first subsection. Subsequently, we focus on the performance of existing benchmark models in predicting demand for low-frequency skills. Further, we have constructed a co-occurrence relationship between skills as prior knowledge based on the training set and employed various Graph Neural Network (GNN)-based multivariate time series forecasting methods in the task of job skill demand forecasting, demonstrating promising results. Finally, considering that the skill demand proportion may be more meaningful than the skill demand volume in certain contexts, we have constructed a dataset for skill demand proportion and showcased the performance of benchmark models on this task.

D.1 Repeated Experiments on Job Skill Demand Forecasting

To demonstrate the robustness and reliability of our experimental results, we first repeated the experiments multiple times as described in the main text. Additionally, we extended our analysis to include experiments across the entire labor market and at various regional granularities.

Implement Details. We utilized the Time-series-Library² to implement some of the models. The hyperparameters were uniformly set as follows: a learning rate of 0.0001, 20 epochs of training, a hidden layer dimension of 2048, the GELU activation function, and MSE loss as the loss function. Early stopping was employed to prevent overfitting by terminating training early when necessary. Data from the first 24 months were used for pre-training, and the model was fine-tuned using the next 6 months to capture trend changes. Finally, the model was used to infer skill demands for the last 6 months. All other hyperparameters were kept consistent with those in the original paper. To ensure the reliability of our findings, we repeated these experiments four times, using random seeds set to 0, 1, 2, and 3, respectively.

Overall Performance Table 6 displays the mean and standard deviation results of repeated experiments on the benchmark models for the job skill demand forecasting task as presented in the main text. Initially, we supplement the experimental results at the overall labor market level and regional granularity, where the RMSE averages over 1000. In cases of coarser granularity, due to the larger base of demand values, the prediction deviations are significant.

Performance on Skill Demand Series with Structural Breaks Table 7 presents the results of repeated experiments on forecasting skill demand sequences that have undergone structural breaks. Initially, the overall errors are quite pronounced, underscoring the challenge of accurately predicting these skills. Moreover, FiLM performs well on most metrics, which further verifies its robustness.

D.2 Job Skill Demand Forecasting for Low-Frequency Skills

In multigranular skill demand sequences, a significant number of skills remain inactive or in low frequency over extended periods. These skills might continue to have low demand in the future (indicating low importance), or they might suddenly gain interest from certain professions or companies, leading to rapid growth. In this study, we define low-frequency skills as those that appear fewer than twice in the time slices of the training set. Predicting the demand for these skills is challenging because their data points are predominantly zero during training, resulting in a lack of effective observational data. Therefore, we specifically present the demand prediction results of the existing benchmark models for these low-frequency skills.

Results. We continued to test the demand prediction effect on low-frequency skills using the benchmark models described in the main text, and the results are shown in Table 8. From this, we can draw the following conclusions: Firstly, there is a significant increase in the error on the RRMSE metric, indicating that low-demand skills are difficult to predict accurately. Secondly, Koopa has the best predictive performance in this scenario. We also found that the performance of SegRNN significantly decreases, suggesting that SegRNN’s segment learning approach is not suitable for predicting low-frequency skill demands due to a lack of effective observational data, rendering the learning segments meaningless.

D.3 Skill Co-occurrence Graph Enhanced Job Skill Demand Forecasting

In the task of job skill demand forecasting, fully leveraging the inter-relationships among different skills is beneficial for downstream tasks. Therefore, we construct a prior graph with co-occurrence frequency from the training data to include as a dataset component. Given a set of granularities i, j, \dots, k , we constructed the skill co-occurrence graph as $\mathcal{G}^{i,j,\dots,k} = (\mathcal{V}^{i,j,\dots,k}, \mathcal{E}^{i,j,\dots,k})$, where $\mathcal{V}^{i,j,\dots,k}$ is the extended skill set under the multiple granularities. The edge weight $e_{v,v'} \in \mathcal{E}^{i,j,\dots,k}$ between nodes v and v' is determined by the co-occurrence frequency of the node pair v, v' in the job advertisement data for training \mathcal{P}^{train} . Specifically, given $v = (a^i, a^j, \dots, a^k, s)$ and $v' = (a^{i'}, a^{j'}, \dots, a^{k'}, s')$, $e_{v,v'}$ is calculated as:

$$e_{v,v'} = \sum_{p \in \mathcal{P}^{train}} \prod_{x \in \{a^i, a^j, \dots, a^k, a^{i'}, a^{j'}, \dots, a^{k'}, s, s'\}} \mathbf{1}_p(x \in p). \quad (4)$$

This information will serve as prior knowledge, reflecting global inter-skill dependency patterns.

²<https://github.com/thuml/Time-Series-Library>

Table 6: Overall performance comparisons on repeated experiments.

	Model	Market	Region	L1-O	L2-O	R&L1-O	R&L2-O	Company
MAE	LSTM	314.54±0.57	49.92±0.0	24.43±4.91	8.04±0.87	4.44±0.47	1.45±0.15	1.49±0.13
	SegRNN	190.05±0.37	35.92±0.47	16.37±3.73	5.81±0.73	3.68±0.5	1.32±0.25	1.23±0.2
	CHGH	315.47±0.04	50.03±0.02	25.62±3.23	8.04±0.87	4.43±0.48	1.51±0.7	1.54±0.73
	PreDyGAE	189.95±0.01	35.84±0.09	21.72±1.15	6.81±0.21	4.08±0.15	1.85±0.43	1.85±0.56
	Transformer	340.72±0.79	54.98±0.13	27.43±4.9	8.93±1.24	4.92±0.83	1.75±0.36	1.69±0.39
	Autoformer	465.86±2.78	60.9±0.64	31.97±8.13	9.97±1.6	6.68±0.21	2.6±0.17	2.85±0.42
	Informr	340.76±1.6	55.08±0.35	27.54±4.87	8.81±1.26	4.87±0.91	1.73±0.39	1.69±0.39
	Reformer	344.9±2.62	54.83±0.2	27.35±4.79	8.88±1.3	4.9±0.9	1.71±0.42	1.8±0.24
	FEDformer	469.46±4.52	60.37±0.19	31.87±8.22	9.54±1.9	5.83±1.1	2.41±0.39	2.52±0.08
	NStransformer	208.23±3.43	37.39±0.46	19.13±1.61	6.29±0.5	3.79±0.31	1.26±0.1	1.64±0.44
	PatchTST	204.94±12.41	36.33±2.23	18.69±3.45	6.12±0.89	3.69±0.54	1.23±0.18	1.21±0.18
	DLinear	201.05±1.95	35.45±0.56	18.23±1.48	5.89±0.41	3.55±0.28	1.21±0.12	1.18±0.11
	TSMixer	517.95±1.49	67.56±11.06	30.83±8.66	10.54±2.19	5.95±0.12	3.58±2.16	6.75±6.58
	FreTS	310.62±4.57	49.52±0.73	23.1±6.05	7.74±1.11	4.3±0.59	1.41±0.17	1.45±0.17
	FiLM	201.7±24.86	36.18±2.29	16.79±3.51	5.9±0.75	3.7±0.42	1.28±0.13	1.31±0.12
	Koopa	205.12±4.95	35.95±0.5	19.9±0.01	6.19±0.13	3.68±0.13	1.2±0.05	1.15±0.06
RMSE	LSTM	1799.7±0.49	341.97±0.07	308.58±103.99	148.32±29.18	66.87±13.95	30.01±6.11	35.35±8.1
	SegRNN	941.59±1.08	194.73±1.59	154.02±41.75	78.84±9.14	38.6±3.97	18.21±2.05	21.05±4.58
	CHGH	1803.56±0.73	344.15±0.11	335.97±67.99	148.94±29.54	67.12±14.29	19.3±12.73	20.61±14.53
	PreDyGAE	925.27±0.01	193.94±0.04	186.36±1.4	84.55±1.44	40.37±1.6	18.79±1.27	22.05±3.48
	Transformer	1829.87±2.96	360.87±6.13	318.17±94.1	151.16±30.08	67.66±14.26	30.42±6.37	35.12±9.25
	Autoformer	2177.05±87.58	313.84±8.55	270.26±76.23	122.29±20.33	62.56±4.38	26.57±2.26	35.47±2.81
	Informr	1854.81±3.0	359.97±7.72	314.97±100.17	148.73±28.62	67.01±15.41	30.14±6.45	35.1±3.9
	Reformer	1890.32±13.72	357.61±1.23	314.03±100.12	152.43±32.71	67.58±15.18	30.21±6.76	36.14±8.01
	FEDformer	2137.44±58.69	321.53±2.51	268.73±79.24	113.47±22.37	56.36±12.0	25.64±3.57	30.9±3.64
	NStransformer	1121.07±25.29	236.56±12.38	198.03±44.33	103.2±15.48	45.99±8.13	21.42±3.62	32.8±1.86
	PatchTST	1098.82±43.54	220.21±7.5	204.56±57.97	97.9±17.38	45.32±9.07	20.88±3.94	25.48±5.83
	DLinear	1107.0±28.41	222.14±4.31	211.02±51.25	99.81±16.61	46.04±8.55	21.12±3.95	25.7±6.25
	TSMixer	2578.82±27.27	438.06±67.68	350.22±143.66	172.63±60.23	78.88±15.28	48.76±17.8	82.91±56.64
	FreTS	1763.69±25.03	336.65±4.04	289.29±111.07	140.54±31.17	63.72±14.52	28.68±6.17	34.0±7.86
	FiLM	1071.14±223.85	216.77±38.28	173.2±51.15	83.31±16.12	39.26±8.54	18.23±3.85	22.53±6.08
	Koopa	1147.88±31.89	228.36±7.6	227.46±43.96	106.03±12.93	49.37±7.89	22.53±3.49	27.06±6.28
SMAPE(%)	LSTM	42.33±0.04	49.78±0.03	42.31±0.85	33.14±0.27	31.79±0.19	22.67±0.24	30.42±0.14
	SegRNN	39.67±0.26	51.14±0.53	42.16±2.14	34.62±1.16	36.02±0.65	30.62±6.19	32.77±0.27
	CHGH	42.21±0.01	49.03±0.03	40.83±0.51	29.98±0.35	28.42±0.28	46.97±9.54	30.69±14.64
	PreDyGAE	42.1±0.14	64.37±1.24	75.3±23.21	59.69±0.77	46.51±11.69	76.6±3.58	60.36±31.18
	Transformer	49.86±0.46	60.59±0.02	58.06±3.26	52.7±7.73	47.37±14.81	47.03±12.77	45.69±16.5
	Autoformer	73.47±0.36	79.08±0.19	79.44±8.36	79.16±4.38	86.16±3.64	86.32±4.73	91.12±14.54
	Informr	49.2±0.2	60.45±0.03	58.59±1.59	52.61±7.82	48.84±12.86	45.52±14.82	47.61±13.88
	Reformer	48.69±0.17	60.48±0.17	58.49±1.74	51.26±9.75	47.8±14.24	42.15±19.43	52.76±6.85
	FEDformer	73.57±0.53	78.11±0.55	79.26±9.09	76.98±7.02	78.54±4.9	82.05±0.3	86.29±7.21
	NStransformer	38.88±0.09	47.74±0.14	38.94±0.7	27.49±1.08	26.43±1.33	15.47±0.84	23.75±0.41
	PatchTST	37.56±2.15	46.52±2.49	37.82±2.85	26.98±2.25	26.7±1.42	15.2±1.55	22.29±2.19
	DLinear	37.7±1.82	49.27±3.85	41.79±0.05	34.65±0.28	34.07±0.54	26.6±0.75	31.42±0.65
	TSMixer	73.63±1.83	72.27±9.41	61.01±4.04	68.37±3.58	65.19±15.78	79.28±37.95	95.62±54.39
	FreTS	42.81±0.56	50.46±0.62	41.53±1.62	31.26±0.98	29.76±1.08	18.89±1.16	28.37±1.03
	FiLM	37.79±1.75	46.85±4.01	41.03±1.38	30.26±0.56	29.49±0.64	17.65±0.37	26.19±0.42
	Koopa	36.26±0.39	44.7±0.95	36.69±1.05	25.59±0.12	24.51±0.09	14.13±0.14	20.62±0.18
RRMSE(%)	LSTM	52.7±0.09	52.87±0.03	73.45±14.2	97.42±12.52	80.24±10.81	102.19±13.54	210.9±33.32
	SegRNN	20.55±0.05	22.49±0.39	32.2±4.91	41.44±8.3	39.72±8.04	49.27±11.53	71.42±13.56
	CHGH	52.92±0.03	53.46±0.02	77.26±10.23	98.13±12.8	80.86±11.35	54.98±47.9	100.68±96.78
	PreDyGAE	20.07±0.0	22.65±0.01	51.16±29.68	54.87±26.22	47.52±17.38	64.98±30.23	96.4±45.03
	Transformer	52.49±0.18	55.57±1.14	75.98±10.71	98.09±12.61	82.63±5.9	101.34±13.21	200.07±32.6
	Autoformer	49.81±1.73	38.47±1.72	54.53±0.71	64.59±1.09	61.08±4.1	71.01±3.15	103.89±3.92
	Informr	54.11±0.2	55.59±2.17	71.42±12.08	97.76±8.25	81.04±10.45	102.15±10.95	198.51±31.17
	Reformer	55.75±0.54	54.83±0.62	76.12±13.48	98.47±13.48	81.65±8.01	103.29±11.35	206.04±33.02
	FEDformer	48.25±2.89	39.48±0.85	56.52±2.27	60.81±0.74	55.93±2.95	69.07±0.9	102.11±3.78
	NStransformer	26.75±0.76	30.42±1.94	43.64±3.24	59.15±1.44	49.08±0.17	62.42±0.79	107.45±6.09
	PatchTST	26.38±1.12	28.31±1.01	47.51±3.34	56.73±1.89	47.96±2.74	64.14±3.05	108.92±5.86
	DLinear	26.6±1.35	28.58±1.32	49.04±3.52	58.69±1.4	51.47±0.38	64.72±0.06	110.8±1.95
	TSMixer	74.69±3.7	69.2±15.69	86.56±23.18	118.56±17.54	99.17±10.85	100.88±4.01	136.0±30.91
	FreTS	51.01±0.94	51.53±0.82	69.56±13.82	92.39±10.91	76.55±9.53	97.37±11.25	200.81±23.97
	FiLM	25.72±7.38	27.83±7.19	36.13±1.29	43.32±0.5	38.93±1.16	48.51±0.7	79.66±2.5
	Koopa	27.81±1.3	29.66±1.55	52.74±5.07	62.29±2.78	56.41±1.28	70.91±2.82	122.3±1.52

Benchmark Models To fully utilize the prior information from the co-occurrence graph, we introduce several GNN-based methods for multivariate time series prediction. These methods leverage GNNs to extract the influences between different variables, effectively capturing the relationships among various time series. The specific models are as follows:

- **EvolveGCN** [80]: EvolveGCN introduces a recurrent mechanism to update the network parameters, as GCN parameters, for capturing the dynamism of the graphs. Two methods are introduced: EvolveGCNH, which learns the weight matrix of the graph at each time

Table 7: Performance comparisons on skill demand series with structural breaks.

	Model	Market	Region	L1-O	L2-O	R&L1-O	R&L2-O	Company
MAE	LSTM	423.99±0.56	109.63±0.04	101.68±13.13	63.39±4.97	21.02±1.85	8.55±0.59	26.23±1.67
	SegRNN	256.67±0.54	76.09±1.01	68.08±5.62	44.93±0.88	16.12±0.24	7.07±0.47	18.91±0.85
	CHGH	425.17±0.04	109.82±0.06	104.41±9.23	63.54±5.0	18.11±1.93	15.16±4.33	19.75±13.35
	PreDyGAE	296.27±0.01	85.05±5.06	84.11±8.87	56.28±7.2	16.57±0.84	9.51±0.19	22.9±0.13
	Transformer	460.45±2.31	120.62±0.4	113.11±13.19	69.45±7.05	22.4±2.76	9.42±0.88	26.93±4.12
	Autoformer	632.29±2.76	131.92±1.52	131.01±21.72	75.4±7.07	25.82±0.93	11.76±0.39	36.65±6.8
	Informer	462.88±1.99	120.89±0.68	113.36±13.21	68.13±7.46	22.22±2.91	9.1±1.26	26.82±4.06
	Reformer	465.72±3.72	120.58±0.34	112.51±13.12	69.02±7.59	22.4±2.88	9.09±1.44	28.16±2.29
	FEDformer	645.02±6.45	131.05±0.27	130.18±22.59	72.12±9.14	23.06±2.34	11.01±0.59	32.89±1.1
	NStransformer	285.18±6.1	81.54±1.2	80.2±2.03	49.66±0.33	16.88±0.26	7.07±0.2	28.34±10.69
	PatchTST	282.34±17.19	79.1±4.75	80.46±2.76	48.45±3.13	16.29±1.29	7.01±0.41	19.65±1.48
	DLinear	277.33±2.32	77.47±0.99	77.64±3.23	46.5±0.16	16.03±0.08	6.64±0.12	18.73±0.51
	TSMixer	698.27±6.79	146.22±23.62	130.91±21.4	88.39±4.38	28.49±1.37	16.46±8.55	79.02±69.45
	FreTS	419.07±6.85	108.63±1.75	98.46±14.62	62.17±5.14	20.5±1.82	8.47±0.54	25.89±1.53
	FiLM	277.66±34.84	78.58±6.24	69.56±6.12	45.37±2.5	15.26±0.87	6.79±0.39	19.86±0.98
	Koopa	282.85±6.89	78.75±1.42	83.25±7.31	48.75±1.55	16.92±0.47	6.93±0.13	19.16±0.11
RMSE	LSTM	2148.33±0.79	534.61±0.67	704.41±136.89	466.21±60.24	178.29±26.25	58.54±5.62	192.72±30.76
	SegRNN	1130.19±0.81	304.2±2.76	387.1±3.14	264.53±10.99	110.27±3.44	36.38±1.34	120.84±7.95
	CHGH	2153.77±0.87	537.97±0.18	735.74±97.15	468.27±61.0	179.17±26.85	41.0±28.76	203.78±87.33
	PreDyGAE	1510.41±0.01	403.02±0.06	529.39±58.83	388.54±43.46	159.45±15.71	36.48±2.19	162.0±20.48
	Transformer	2154.53±30.61	562.07±12.46	717.62±125.1	474.06±63.8	179.2±25.64	61.55±5.56	191.91±36.18
	Autoformer	2566.07±102.19	489.69±10.56	621.06±80.33	387.57±33.41	155.94±0.51	58.1±4.51	187.41±19.68
	Informer	2212.83±2.36	559.86±11.51	715.12±132.16	467.05±58.69	178.41±28.75	57.53±7.65	191.85±36.5
	Reformer	2248.94±15.68	558.75±0.8	713.91±131.53	478.6±70.65	179.85±27.33	59.48±9.48	195.6±31.87
	FEDformer	2551.16±77.11	501.84±4.46	621.78±81.74	363.84±35.37	149.67±19.37	57.89±2.21	166.63±10.36
	NStransformer	1337.16±47.86	373.7±20.75	473.97±10.71	334.17±14.36	128.79±8.11	39.55±1.82	176.86±17.5
	PatchTST	1328.54±49.25	347.07±11.24	499.82±22.78	321.16±15.89	125.31±13.05	39.74±1.04	144.13±14.98
	DLinear	1341.55±32.44	350.8±6.91	506.06±18.99	325.15±16.25	130.09±10.18	38.99±1.16	144.53±18.45
	TSMixer	3040.93±86.91	678.15±105.71	816.57±184.07	582.2±93.85	211.98±22.73	115.88±67.93	436.64±301.19
	FreTS	2105.11±31.62	526.0±6.74	679.7±130.16	451.97±53.49	173.34±22.83	57.76±4.47	188.29±25.82
	FiLM	1298.06±265.59	341.89±60.42	426.14±19.46	276.0±13.7	111.46±9.34	34.32±1.86	128.54±16.33
	Koopa	1386.85±35.21	360.64±11.92	532.64±14.5	338.26±12.88	138.63±9.35	41.39±0.09	152.48±17.55
SMAPE(%)	LSTM	37.23±0.05	45.67±0.02	45.95±1.99	50.34±1.28	47.36±0.66	42.47±0.4	70.15±1.62
	SegRNN	32.88±0.2	43.35±0.39	41.54±2.12	46.46±3.08	47.11±1.77	44.29±4.45	59.48±1.85
	CHGH	37.21±0.01	45.55±0.03	46.29±1.26	50.3±1.28	46.22±0.72	45.92±21.79	77.34±37.44
	PreDyGAE	32.0±0.02	43.25±0.18	49.33±2.76	51.02±5.06	50.59±1.35	63.2±17.41	70.72±0.42
	Transformer	44.18±0.71	53.31±0.19	53.13±2.85	56.88±3.45	55.61±3.32	55.47±7.12	73.21±7.57
	Autoformer	67.46±0.27	68.4±0.03	70.67±6.58	71.92±3.02	83.21±4.31	83.71±4.79	99.62±14.04
	Informer	43.4±0.31	53.28±0.16	53.54±2.22	56.42±4.14	56.23±7.69	54.33±8.61	73.34±7.52
	Reformer	43.1±0.01	53.09±0.1	53.42±2.4	56.23±4.31	55.84±8.2	52.75±10.9	76.27±3.57
	FEDformer	68.1±0.29	67.4±0.48	70.57±7.06	70.35±5.46	76.14±3.56	79.91±0.11	95.19±7.37
	NStransformer	32.85±0.14	42.35±0.05	44.17±1.08	47.32±0.29	44.23±1.08	38.15±1.31	103.89±60.87
	PatchTST	31.96±2.0	41.14±2.29	42.63±1.59	45.79±2.31	43.83±2.11	37.55±2.57	58.04±2.54
	DLinear	31.38±1.08	40.56±1.37	41.8±1.23	44.83±0.38	45.42±0.14	41.61±0.42	58.22±0.52
	TSMixer	68.55±1.22	67.06±11.47	61.4±6.47	75.25±0.75	72.63±11.4	83.29±31.13	109.31±46.77
	FreTS	37.47±0.7	46.12±0.91	45.36±2.66	50.03±1.63	46.82±1.3	41.25±1.28	70.0±1.47
	FiLM	32.01±0.16	41.26±0.46	41.12±1.97	45.56±1.21	45.99±0.95	40.91±0.79	58.02±1.51
	Koopa	31.01±0.02	39.92±0.13	42.95±3.2	45.33±1.65	41.88±0.65	35.91±0.3	57.19±1.63
RRMSE(%)	LSTM	53.94±0.16	53.9±0.12	73.77±14.35	98.43±12.75	93.46±13.83	65.39±6.3	227.07±36.31
	SegRNN	21.23±0.03	22.98±0.43	32.33±4.99	41.84±8.49	44.18±9.25	34.49±6.31	73.51±14.73
	CHGH	54.21±0.03	54.5±0.02	77.63±10.33	99.15±13.03	93.96±14.27	53.56±32.26	105.96±104.02
	PreDyGAE	30.72±0.01	33.15±0.01	46.59±9.64	55.32±12.77	57.56±12.47	45.53±6.19	91.47±23.2
	Transformer	52.69±0.67	56.3±1.62	76.3±10.8	99.08±12.9	94.57±7.31	67.56±5.78	215.96±35.18
	Autoformer	51.07±0.19	39.44±1.34	54.75±0.61	65.4±1.15	64.75±3.85	54.46±8.57	106.51±5.39
	Informer	55.42±0.04	56.22±2.17	71.7±12.14	98.81±8.31	93.29±11.35	64.2±6.07	213.07±32.78
	Reformer	56.91±0.64	55.93±0.41	76.53±13.64	99.47±13.64	95.12±10.02	67.05±7.72	221.45±36.6
	FEDformer	49.35±2.5	40.99±1.28	56.8±2.22	61.57±0.66	63.27±4.31	54.5±0.19	105.93±4.82
	NStransformer	27.47±1.23	31.55±2.22	43.9±3.25	60.16±1.54	56.5±1.0	39.75±0.02	116.26±2.18
	PatchTST	27.53±1.1	29.31±1.0	47.82±3.34	57.57±1.93	55.03±2.9	40.96±1.97	114.76±7.12
	DLinear	27.83±1.38	29.66±1.38	49.37±3.5	59.59±1.4	59.48±0.62	40.74±0.89	116.83±1.45
	TSMixer	76.26±6.24	69.83±15.63	86.8±23.27	119.17±17.64	113.43±16.08	77.61±14.65	141.79±36.12
	FreTS	52.23±1.04	52.51±0.92	69.87±13.95	93.32±11.09	88.52±11.95	63.64±5.3	215.64±25.29
	FiLM	26.91±7.65	28.84±7.48	36.38±1.32	43.93±0.54	44.44±1.26	31.45±0.6	83.15±2.17
	Koopa	28.98±1.27	30.77±1.6	53.08±5.03	63.24±2.77	65.24±0.87	44.14±3.04	129.36±2.2

step as a hidden state, and EvolveGCNO, which directly employs the weight evolution as a hidden state output, decoupled from node embedding.

- **GConvGRU** [81]: This model integrates convolutional neural networks (CNNs) on graphs to identify spatial structures and recurrent neural networks (RNNs) to detect dynamic patterns. Two architectures, GConvGRU and GConvLSTM, are explored for the Graph Convolutional Recurrent Network (GCRN).
- **TGCN** [82]: The temporal graph convolutional network (T-GCN) model, which is in combination with the graph convolutional network (GCN) and gated recurrent unit (GRU). Specifically, the GCN is used to learn complex topological structures to capture spatial

Table 8: Performance comparisons on skill demand series with low-frequency.

	Model	Market	Region	L1-O	L2-O	R&L1-O	R&L2-O	Company
MAE	LSTM	33.69±0.44	26.88±0.05	16.49±0.01	12.11±0.04	12.28±0.02	8.92±0.1	12.6±0.02
	SegRNN	51.5±0.57	46.12±0.28	25.65±0.39	18.03±0.03	22.08±1.92	19.1±3.22	19.37±0.88
	CHGH	32.37±0.0	24.99±0.02	14.2±0.04	9.32±0.01	9.45±0.0	67.89±10.03	77.01±5.58
	PreDyGAE	113.67±1.73	157.17±18.68	343.32±0.15	64.46±13.49	24.16±0.13	60.41±55.15	68.89±62.89
	Transformer	54.84±1.99	52.95±0.09	46.1±0.13	45.47±0.07	45.64±0.02	45.1±0.01	45.43±0.0
	Autoformer	132.98±11.42	106.13±1.65	110.06±0.19	97.66±1.23	98.28±0.46	95.26±1.1	90.23±0.54
	Informer	54.42±1.83	52.76±0.13	46.1±0.0	45.52±0.01	45.69±0.02	45.07±0.01	45.47±0.0
	Reformer	52.72±0.3	53.25±0.15	46.19±0.07	45.58±0.02	45.65±0.01	45.08±0.03	45.49±0.03
	FEDformer	133.07±6.09	101.47±1.47	109.51±1.3	96.43±1.0	95.83±1.48	94.18±0.48	90.71±0.47
	NStransformer	57.96±1.83	34.24±0.3	17.99±0.37	10.84±0.04	11.46±0.02	6.45±0.02	11.9±0.0
	PatchTST	47.25±6.07	33.21±3.3	17.31±1.85	10.39±1.03	10.91±1.1	6.1±0.67	11.4±1.23
	DLinear	51.7±7.31	42.08±8.83	29.64±9.87	25.18±10.9	25.6±10.9	22.53±11.72	25.79±10.77
	TSMixer	102.23±14.26	64.77±10.8	51.89±9.94	52.61±1.14	40.06±7.0	34.02±2.2	41.32±2.03
	FreTS	38.55±0.83	28.89±0.96	16.6±0.58	11.19±0.47	11.46±0.51	7.44±0.47	11.94±0.48
	FILM	51.42±20.47	34.97±10.99	18.23±5.69	10.9±3.11	11.47±3.38	6.43±1.96	11.95±3.58
	Koopa	43.3±4.56	29.66±2.37	15.52±1.38	9.49±0.73	9.95±0.82	5.53±0.48	10.35±0.86
RMSE	LSTM	99.83±0.77	306.02±0.33	168.19±0.01	140.57±0.0	247.24±0.02	125.6±0.01	272.77±0.04
	SegRNN	141.21±0.31	452.49±0.09	233.37±0.06	146.34±0.22	216.77±0.01	117.81±0.63	227.59±0.0
	CHGH	100.44±0.02	306.28±0.01	168.23±0.01	140.54±0.0	247.21±0.01	134.20±0.0	208.01±17.09
	PreDyGAE	169.76±1.46	542.42±24.84	771.72±50.15	195.02±16.23	218.94±0.65	119.76±109.33	188.17±171.77
	Transformer	112.3±5.16	311.19±1.43	176.74±0.52	150.63±0.53	252.9±0.1	137.39±0.13	278.45±0.13
	Autoformer	249.03±36.15	364.32±9.52	231.67±3.95	184.29±1.11	266.1±0.64	172.38±1.84	281.53±0.17
	Informer	111.59±6.84	312.15±0.35	176.18±0.19	150.46±0.16	252.84±0.01	137.07±0.12	279.56±0.36
	Reformer	107.11±3.02	314.04±1.77	176.73±0.57	150.62±0.23	252.85±0.14	137.02±0.12	279.79±0.61
	FEDformer	233.86±32.3	346.25±1.54	227.64±2.23	182.54±0.06	264.12±0.64	171.43±0.76	280.57±1.13
	NStransformer	159.4±10.66	329.86±4.31	182.83±8.95	136.01±2.69	249.63±0.44	129.74±5.67	255.48±4.76
	PatchTST	126.71±17.24	339.92±16.66	182.49±8.84	134.57±2.89	233.14±1.45	120.94±1.7	253.85±3.35
	DLinear	109.06±5.5	320.09±2.51	173.51±2.67	133.82±2.37	235.23±1.35	123.01±2.84	254.7±1.91
	TSMixer	186.01±18.76	359.57±28.14	199.07±10.44	160.52±5.4	261.65±6.67	139.7±3.15	296.15±1.83
	FreTS	104.43±1.38	310.94±2.43	170.36±1.47	139.96±0.74	239.77±4.49	122.58±2.07	266.82±3.88
	FILM	129.11±32.28	343.63±34.15	185.3±16.68	135.8±1.72	235.13±10.73	122.84±2.06	256.54±7.91
	Koopa	111.96±8.52	317.8±3.82	171.28±3.23	134.11±0.67	242.79±2.91	124.64±0.36	260.87±0.46
SMAPE(%)	LSTM	26.1±0.39	19.43±0.17	15.18±0.01	12.88±0.07	13.01±0.02	11.46±0.17	13.01±0.03
	SegRNN	33.71±0.71	29.55±0.4	22.04±0.56	18.44±0.11	24.16±2.97	24.8±4.63	19.83±1.32
	CHGH	23.61±0.0	15.99±0.04	11.08±0.09	7.8±0.03	7.86±0.01	9.61±1.08	8.07±2.03
	PreDyGAE	83.59±0.72	82.33±4.38	114.25±1.42	48.0±7.33	23.97±0.1	47.34±43.22	48.31±44.1
	Transformer	49.05±1.26	48.23±0.0	48.68±0.14	50.78±0.03	50.92±0.01	52.6±0.01	50.28±0.0
	Autoformer	80.95±1.5	76.73±0.44	81.69±0.03	80.14±0.55	80.43±0.21	81.25±0.49	76.73±0.25
	Informer	48.87±0.78	48.07±0.25	48.69±0.03	50.83±0.01	50.96±0.02	52.57±0.01	50.29±0.0
	Reformer	48.32±1.06	48.33±0.02	48.76±0.08	50.87±0.03	50.93±0.0	52.59±0.02	50.28±0.02
	FEDformer	82.56±1.23	75.32±0.52	81.88±0.57	79.7±0.39	79.55±0.62	80.83±0.21	76.96±0.21
	NStransformer	32.62±0.69	18.84±0.31	11.59±0.02	7.3±0.01	7.81±0.01	4.76±0.0	7.76±0.0
	PatchTST	29.48±1.53	18.23±1.46	11.15±1.04	7.01±0.71	7.51±0.76	4.56±0.52	7.46±0.8
	DLinear	42.19±8.16	35.1±10.8	31.64±12.55	30.55±14.04	30.93±14.03	30.35±15.14	30.52±13.82
	TSMixer	68.79±5.48	51.84±6.45	48.47±6.43	54.61±0.92	45.87±6.93	43.34±2.84	45.68±1.95
	FreTS	29.63±0.33	19.71±0.62	14.01±0.56	10.6±0.64	10.86±0.67	8.45±0.75	11.09±0.61
	FILM	30.74±10.37	18.56±6.37	11.42±4.06	7.19±2.55	7.7±2.79	4.69±1.69	7.64±2.84
	Koopa	28.27±2.34	16.67±1.38	10.15±0.9	6.35±0.57	6.79±0.63	4.09±0.39	6.73±0.63
RRMSE(%)	LSTM	432.63±21.64	1364.7±202.93	1057.49±28.76	1468.93±15.86	2581.13±33.65	1485.07±28.13	1828.92±33.42
	SegRNN	114.7±0.21	125.97±0.16	126.24±0.18	146.81±0.73	130.23±0.45	130.24±0.87	133.03±0.54
	CHGH	452.99±0.65	1662.7±2.11	1394.35±4.28	1824.0±1.1	3436.23±14.46	289.09±14.07	406.07±8.08
	PreDyGAE	104.17±0.99	216.23±2.0	201.73±0.3	220.44±4.62	229.42±0.22	165.21±59.52	272.3±66.0
	Transformer	172.07±30.78	513.13±29.31	320.11±12.98	269.97±3.45	468.03±0.41	242.99±4.1	466.5±1.48
	Autoformer	89.47±0.4	161.16±10.32	136.15±6.04	134.91±1.17	195.47±1.88	134.43±1.31	191.17±2.0
	Informer	193.88±30.22	538.02±31.16	316.88±4.27	275.64±0.03	470.15±2.27	246.67±0.09	462.58±1.78
	Reformer	200.52±32.93	536.21±13.15	324.86±11.04	276.66±0.14	470.18±0.3	247.32±0.89	460.86±4.08
	FEDformer	96.72±6.76	176.64±0.39	136.62±1.69	136.93±0.31	199.25±1.04	134.8±0.73	189.29±1.49
	NStransformer	90.14±7.16	219.18±39.19	193.48±36.91	218.73±0.68	406.29±7.04	272.29±31.63	259.15±8.34
	PatchTST	95.85±5.0	210.2±36.63	199.12±34.35	266.55±38.3	391.36±46.33	298.48±41.01	278.87±33.7
	DLinear	90.72±2.96	241.41±14.46	219.68±17.61	281.05±37.87	431.86±48.87	296.1±49.89	300.39±25.41
	TSMixer	125.99±9.23	200.39±22.34	211.03±36.38	192.85±35.49	418.03±16.7	275.29±1.9	293.38±24.47
	FreTS	222.55±30.4	479.45±20.87	462.19±33.86	720.43±51.2	962.73±47.52	758.24±41.67	809.76±76.22
	FILM	108.05±34.05	444.72±401.85	360.89±298.05	402.84±297.56	652.7±546.99	417.27±297.85	364.47±235.25
	Koopa	93.83±0.11	287.42±57.85	264.45±56.56	345.86±79.42	582.52±154.25	368.66±64.94	345.87±66.81

dependence and the gated recurrent unit is used to learn dynamic changes of traffic data to capture temporal dependence.

- **GCLSTM** [83]: GCLSTM is an end-to-end model integrating a Graph Convolution Network (GCN) embedded Long Short-Term Memory network (LSTM) for dynamic network link prediction. The GCN captures local structural properties, while the LSTM learns temporal features across snapshots of a dynamic network.
- **DyGrEncoder** [84]: This approach combines a sequence-to-sequence encoder-decoder model with gated graph neural networks (GGNNs) and long short-term memory networks

Table 9: Performance comparisons on skill demand series with GNN-based methods.

	Model	Market	Region	L1-O	L2-O	R&L1-O	R&L2-O	Company
MAE	EvolveGCNH	1053.18±804.38	30.07±22.97	16.51±12.61	5.43±4.15	2.92±2.23	1.11±0.85	0.96±0.73
	EvolveGCNO	151.13±115.43	27.01±20.63	15.4±11.76	5.04±3.85	2.95±2.25	1.21±0.92	0.89±0.68
	GConvGRU	570.31±435.58	92.96±71.0	46.43±35.46	11.45±8.75	5.47±4.18	1.4±1.07	1.04±0.79
	TGCN	741.05±565.99	96.43±73.65	55.41±42.32	13.21±10.09	6.28±4.8	1.63±1.24	1.11±0.85
	GCLSTM	729.48±557.15	57.13±43.63	25.84±19.74	11.49±8.78	5.67±4.33	1.61±1.23	1.07±0.82
	GConvLSTM	741.45±566.29	93.61±71.5	46.57±35.57	11.19±8.55	5.57±4.25	1.49±1.14	1.09±0.83
	DyGrEncoder	732.53±559.48	92.21±70.43	46.96±35.87	12.26±9.36	5.57±4.26	1.47±1.13	1.07±0.82
	RMSE	EvolveGCNH	4998.7±3817.82	166.34±127.04	178.55±136.37	76.14±58.15	35.12±26.82	16.97±12.96
EvolveGCNO		709.19±541.65	143.72±109.77	164.16±125.38	71.57±54.67	33.21±25.37	15.04±11.49	17.93±13.7
GConvGRU		2968.75±2267.42	579.66±442.72	442.53±337.99	170.09±129.91	81.95±62.59	31.75±24.25	30.65±23.41
TGCN		3163.46±2416.13	581.49±444.12	450.79±344.29	172.13±131.47	83.33±63.64	32.19±24.59	30.77±23.5
GCLSTM		3159.61±2413.19	497.97±380.33	371.13±283.45	170.23±130.01	82.66±63.13	32.25±24.63	30.81±23.53
GConvLSTM		3166.41±2418.38	581.26±443.94	442.72±338.13	169.43±129.4	82.38±62.92	32.29±24.66	30.88±23.58
DyGrEncoder		3161.61±2414.72	579.78±442.81	443.49±338.72	171.91±131.3	82.27±62.83	32.02±24.46	30.70±23.49
SMAPE(%)		EvolveGCNH	76.69±58.58	41.67±31.83	39.77±30.38	33.11±25.29	28.99±22.14	18.79±14.35
	EvolveGCNO	45.32±34.61	56.75±43.35	46.34±35.39	35.09±26.8	29.33±22.4	33.19±25.35	25.99±19.85
	GConvGRU	63.93±48.82	62.41±47.67	38.65±29.52	31.09±23.74	23.67±18.08	12.93±9.88	20.5±15.66
	TGCN	71.13±54.32	51.37±39.23	56.23±42.95	37.21±28.42	25.83±19.73	21.74±16.6	23.35±17.83
	GCLSTM	60.49±46.2	38.68±29.54	43.6±33.3	28.15±21.5	24.34±18.59	18.57±14.18	21.74±16.6
	GConvLSTM	64.79±49.49	44.77±34.2	35.9±27.42	26.21±20.02	21.67±16.55	13.95±10.66	23.05±17.61
	DyGrEncoder	61.55±47.01	44.74±34.17	38.12±29.11	26.94±20.58	26.49±20.23	18.42±14.07	23.75±18.14
	RRMSE(%)	EvolveGCNH	68.68±52.46	19.57±14.95	32.99±25.2	36.53±27.9	32.1±24.52	41.01±31.32
EvolveGCNO		16.21±12.38	17.85±13.63	28.93±22.09	34.6±26.43	30.4±23.22	32.27±24.64	58.07±44.35
GConvGRU		418.84±319.89	989.03±755.38	1111.69±849.06	701.09±535.46	414.17±316.33	307.07±234.53	385.69±294.58
TGCN		3478.62±2656.84	1037.13±792.12	112164.83±85667.3	948.04±724.08	501.07±382.7	338.75±258.73	404.5±308.94
GCLSTM		3083.62±2355.15	176.53±134.83	159.96±122.17	721.9±551.36	467.49±357.05	334.26±255.3	425.46±324.95
GConvLSTM		4479.9±3421.58	1123.32±857.95	1101.0±840.9	640.31±489.05	449.12±343.02	382.23±291.94	451.71±345.0
DyGrEncoder		3386.49±2586.47	1010.93±772.11	1247.51±952.8	1004.62±767.29	450.57±344.13	331.31±253.04	418.67±319.77

Table 10: Performance comparisons on skill demand series with structural breaks with GNN-based methods.

	Model	Market	Region	L1-O	L2-O	R&L1-O	R&L2-O	Company
MAE	EvolveGCNH	1385.2±1057.96	61.79±47.19	58.05±44.33	36.04±27.53	10.69±8.17	5.56±4.25	13.25±10.12
	EvolveGCNO	195.87±149.6	53.86±41.14	52.46±40.07	32.76±25.02	10.85±8.28	5.21±3.98	12.22±9.33
	GConvGRU	749.83±572.69	204.83±156.44	189.48±144.72	92.29±70.49	26.94±20.58	9.23±7.05	17.26±13.18
	TGCN	990.21±756.28	210.7±160.92	222.48±169.92	100.93±77.09	29.27±22.36	9.93±7.58	17.87±13.65
	GCLSTM	977.98±746.94	128.33±98.02	102.44±78.24	92.71±70.81	27.94±21.34	10.25±7.83	17.81±13.61
	GConvLSTM	993.55±758.84	207.86±158.76	190.31±145.35	90.33±68.99	27.51±21.01	9.93±7.58	18.0±13.75
	DyGrEncoder	982.05±750.06	204.89±156.49	191.91±146.57	98.78±75.44	27.41±20.93	9.6±7.33	17.51±13.38
	RMSE	EvolveGCNH	5799.81±4429.68	261.83±199.97	371.99±284.11	227.88±174.05	87.98±67.2	29.87±22.82
EvolveGCNO		863.11±659.21	226.98±173.36	341.34±260.7	213.99±163.44	82.25±62.82	25.03±19.11	93.36±71.3
GConvGRU		3463.23±2645.08	890.59±680.2	922.56±704.62	508.61±388.46	198.01±151.24	69.96±53.43	158.49±121.05
TGCN		3688.76±2817.34	893.17±682.17	938.93±717.12	514.25±392.77	201.03±153.54	71.44±54.56	159.07±121.49
GCLSTM		3684.35±2813.97	768.16±586.69	776.42±593.0	508.99±388.74	199.63±152.47	71.71±54.77	159.26±121.64
GConvLSTM		3692.17±2819.94	892.93±681.98	922.93±704.9	506.7±387.0	198.99±151.98	71.93±54.94	159.59±121.89
DyGrEncoder		3686.65±2815.72	890.76±680.33	924.49±706.09	513.77±392.4	198.73±151.79	70.95±54.19	159.01±121.44
SMAPE(%)		EvolveGCNH	77.85±59.46	31.5±24.06	30.67±23.43	33.05±25.24	33.43±25.54	29.61±22.61
	EvolveGCNO	31.65±24.17	33.88±25.88	30.29±23.14	31.57±24.11	34.17±26.1	37.43±28.59	37.69±28.78
	GConvGRU	49.4±37.73	44.31±33.84	44.6±34.06	38.91±29.72	34.25±26.16	27.57±21.06	37.43±28.59
	TGCN	71.56±54.65	49.2±37.58	120.31±91.89	48.37±36.94	36.76±28.08	32.09±24.51	39.02±29.8
	GCLSTM	63.9±48.8	31.28±23.89	28.84±22.03	39.16±29.91	35.03±26.76	32.18±24.58	37.87±28.93
	GConvLSTM	68.73±52.49	45.39±34.66	44.43±33.94	38.47±29.38	33.75±25.77	28.62±21.86	37.96±28.99
	DyGrEncoder	64.22±49.05	43.72±33.39	44.73±34.16	43.73±33.4	35.73±27.29	30.42±23.23	37.41±28.57
	RRMSE(%)	EvolveGCNH	68.88±52.61	20.26±15.47	33.11±25.29	36.88±28.17	35.03±26.76	26.91±20.55
EvolveGCNO		17.01±12.99	18.47±14.11	28.97±22.12	34.88±26.64	32.23±24.61	20.17±15.41	59.82±45.69
GConvGRU		426.16±325.49	1072.11±818.84	1215.75±928.55	763.66±583.25	485.64±370.91	255.02±194.77	427.8±326.74
TGCN		3708.55±2832.45	1159.31±885.44	156814.93±119769.38	1119.75±855.22	621.0±474.3	290.69±222.02	451.93±345.17
GCLSTM		3210.27±2451.88	184.79±141.14	164.56±125.68	786.29±600.54	553.67±422.87	291.57±222.69	475.27±362.99
GConvLSTM		4672.71±3568.84	1217.46±929.85	1206.73±921.66	695.01±530.83	527.58±402.95	326.17±249.12	505.82±386.33
DyGrEncoder		3526.47±2693.38	1093.03±834.81	1365.47±1042.9	1110.66±848.28	525.05±401.02	279.39±213.39	465.82±355.78

(LSTMs). The encoder captures temporal dynamics in an evolving graph, and the decoder reconstructs the dynamics using the encoded representation.

We implement these benchmark models using the PyG library ³ and demonstrate the effectiveness of these GNN-based methods in skill demand forecasting.

Results We have implemented a series of graph-based multivariate time series forecasting methods based on the co-occurrence graph and verified their experimental effects under the three scenarios discussed above. Firstly, Table 9 presents the overall performance of the methods based on the co-occurrence graph for skill demand forecasting. It is observed that the prediction accuracy significantly

³https://github.com/benedekrozemberczki/pytorch_geometric_temporal

Table 11: Performance comparisons on low-frequency skill demand series with GNN-based methods.

	Model	Market	Region	L1-O	L2-O	R&L1-O	R&L2-O	Company
MAE	EvolveGCNH	35.53 \pm 27.13	2.28 \pm 1.74	1.47 \pm 1.13	0.55 \pm 0.42	0.31 \pm 0.23	0.14 \pm 0.11	0.19 \pm 0.15
	EvolveGCNO	27.87 \pm 21.29	3.35 \pm 2.56	1.69 \pm 1.29	0.57 \pm 0.43	0.31 \pm 0.24	0.27 \pm 0.21	0.17 \pm 0.13
	GConvGRU	63.05 \pm 48.16	3.06 \pm 2.34	0.25 \pm 0.19	0.18 \pm 0.14	0.11 \pm 0.09	0.05 \pm 0.04	0.11 \pm 0.09
	TGCN	7.37 \pm 5.63	2.16 \pm 1.65	0.13 \pm 0.1	0.61 \pm 0.46	0.26 \pm 0.2	0.15 \pm 0.12	0.15 \pm 0.11
	GCLSTM	0.62 \pm 0.47	0.41 \pm 0.32	1.37 \pm 1.04	0.15 \pm 0.11	0.12 \pm 0.09	0.09 \pm 0.07	0.12 \pm 0.09
	GConvLSTM	0.79 \pm 0.6	0.43 \pm 0.33	0.2 \pm 0.15	0.13 \pm 0.1	0.1 \pm 0.08	0.06 \pm 0.05	0.13 \pm 0.1
	DyGrEncoder	1.04 \pm 0.79	0.47 \pm 0.36	0.24 \pm 0.18	0.13 \pm 0.1	0.15 \pm 0.11	0.09 \pm 0.07	0.14 \pm 0.11
	RMSE	EvolveGCNH	144.53 \pm 110.39	11.08 \pm 8.46	11.65 \pm 8.9	5.52 \pm 4.22	3.48 \pm 2.66	1.7 \pm 1.3
EvolveGCNO		83.5 \pm 63.77	10.36 \pm 7.91	12.82 \pm 9.79	5.74 \pm 4.38	3.8 \pm 2.9	1.95 \pm 1.49	1.65 \pm 1.26
GConvGRU		63.06 \pm 48.16	3.65 \pm 2.79	1.19 \pm 0.91	0.87 \pm 0.66	1.59 \pm 1.22	0.81 \pm 0.62	1.72 \pm 1.31
TGCN		21.13 \pm 16.14	8.89 \pm 6.79	1.13 \pm 0.86	3.12 \pm 2.38	2.32 \pm 1.77	1.11 \pm 0.85	1.78 \pm 1.36
GCLSTM		1.63 \pm 1.24	2.48 \pm 1.89	2.8 \pm 2.14	0.86 \pm 0.66	1.61 \pm 1.23	0.83 \pm 0.63	1.73 \pm 1.32
GConvLSTM		1.23 \pm 0.94	2.1 \pm 1.6	1.14 \pm 0.87	0.85 \pm 0.65	1.6 \pm 1.22	0.81 \pm 0.62	1.74 \pm 1.33
DyGrEncoder		1.49 \pm 1.14	2.12 \pm 1.62	1.15 \pm 0.88	0.86 \pm 0.66	1.59 \pm 1.22	0.81 \pm 0.62	1.73 \pm 1.32
SMAPE(%)		EvolveGCNH	47.39 \pm 36.19	42.25 \pm 32.27	34.5 \pm 26.35	25.49 \pm 19.47	20.27 \pm 15.48	12.66 \pm 9.67
	EvolveGCNO	115.04 \pm 87.86	89.67 \pm 68.48	48.24 \pm 36.84	28.83 \pm 22.02	20.7 \pm 15.81	29.63 \pm 22.63	19.09 \pm 14.58
	GConvGRU	131.13 \pm 100.15	105.97 \pm 80.93	26.95 \pm 20.58	23.04 \pm 17.6	13.04 \pm 9.96	6.02 \pm 4.6	11.92 \pm 9.1
	TGCN	83.42 \pm 63.71	48.87 \pm 37.32	11.74 \pm 8.97	28.63 \pm 21.87	14.45 \pm 11.04	16.13 \pm 12.32	15.27 \pm 11.67
	GCLSTM	42.8 \pm 32.69	36.75 \pm 28.07	43.27 \pm 33.05	17.57 \pm 13.42	14.03 \pm 10.72	12.17 \pm 9.3	13.44 \pm 10.26
	GConvLSTM	53.85 \pm 41.13	36.93 \pm 28.2	21.67 \pm 16.55	14.71 \pm 11.24	10.07 \pm 7.69	7.15 \pm 5.46	14.91 \pm 11.39
	DyGrEncoder	62.53 \pm 47.76	39.84 \pm 30.43	26.20 \pm 20.06	15.61 \pm 11.92	17.33 \pm 13.23	12.71 \pm 9.7	16.23 \pm 12.39
	RRMSE(%)	EvolveGCNH	66.67 \pm 50.92	67.56 \pm 51.6	66.87 \pm 51.08	67.29 \pm 51.39	73.11 \pm 55.84	73.29 \pm 55.97
EvolveGCNO		66.67 \pm 50.92	67.59 \pm 51.63	66.81 \pm 51.03	67.22 \pm 51.34	69.99 \pm 53.43	71.57 \pm 54.66	119.39 \pm 91.18
GConvGRU		66.46 \pm 50.76	75.73 \pm 57.84	122.73 \pm 93.74	153.14 \pm 116.96	297.91 \pm 227.54	242.77 \pm 185.42	241.43 \pm 184.39
TGCN		66.51 \pm 50.8	68.02 \pm 51.95	676.51 \pm 516.7	68.73 \pm 52.5	89.88 \pm 68.65	92.87 \pm 70.93	190.19 \pm 145.26
GCLSTM		60.79 \pm 46.43	103.27 \pm 78.88	70.15 \pm 53.58	169.24 \pm 129.26	359.25 \pm 274.38	214.34 \pm 163.7	272.3 \pm 207.97
GConvLSTM		65.93 \pm 50.36	189.03 \pm 144.38	156.39 \pm 119.44	166.51 \pm 127.18	345.76 \pm 264.08	287.94 \pm 219.92	285.4 \pm 217.98
DyGrEncoder		60.35 \pm 46.09	171.99 \pm 131.36	167.79 \pm 128.15	171.6 \pm 131.06	325.47 \pm 248.58	279.07 \pm 213.14	269.12 \pm 205.54

declines across the overall labor market. However, as the granularity of the forecast becomes finer, the model performance improves, and at a finer granularity, the EvolveGCN method outperforms the state-of-the-art (SOTA) methods mentioned in the main text considerably. We analyze that the finer the granularity, the more accurately the co-occurrence graph reflects the associations between skills, while coarser granularity might introduce excessive noise leading to decreased model performance. The fine-grained co-occurrence graph accurately reflects the interrelationships between skills at different granularities, which aids in enhancing the model’s prediction accuracy. Secondly, we find significant differences in the RRMSE metric among these methods, with EvolveGCN showing superior performance because it can learn the evolution of GCN parameter weights over time, thus capturing the evolving dependencies among edges. Therefore, based on the provided co-occurrence graph, it can effectively learn the evolution of skill relationships, which is beneficial for dynamic prediction of skill demand.

For skill demand forecasting in scenarios involving structural breaks, as shown in Table 10, the improvements in the methods based on the co-occurrence graph are greater than those in the overall skill demand forecasting task. This suggests that skills experiencing structural breaks have strong interconnections, and the co-occurrence graph helps the model to identify the patterns of skill demand sequences that are likely to undergo structural breaks, thus further enhancing the prediction effectiveness for this category of skills.

In the task of predicting low-frequency skills, as shown in Table 11, methods like GConvLSTM significantly outperform EvolveGCN. This is due to the sparse observable data for these skills, which leads to sparse connectivity edges on the co-occurrence graph.

D.4 Job Skill Demand Proportion Forecasting

In the main text, we discussed the issue of skill demand prediction. However, consider a scenario where the number of skill postings for a particular occupation is very low, leading to a low demand for that occupational skills. Nevertheless, these skills might constitute a significant portion of the profession’s core competencies. Therefore, using skill demand alone may not adequately measure the importance of these skills within the occupation. To address this, we introduce an extended dataset that includes the skill demand proportion. We define the skill demand proportion as:

$$R_{s,t}^i = [R_{s,t,a^i}^i]_{a^i \in \mathcal{A}^i}, \quad R_{s,t,a^i}^i = \frac{\sum_{p \in \mathcal{P}_t} \mathbf{1}(s \in p) \cdot \mathbf{1}(a^i \in p)}{\sum_{p \in \mathcal{P}_t} \mathbf{1}(a^i \in p)}, \quad (5)$$

Table 12: Performance comparisons on skill demand proportion forecasting.

	Model	Market	Region	L1-O	L2-O	R&L1-O	R&L2-O	Company
MAE(%)	LSTM	0.14±0.0	0.49±0.01	1.21±0.03	2.26±0.05	2.43±0.03	3.51±0.06	2.69±0.03
	SegRNN	0.15±0.02	1.25±0.02	2.63±0.11	4.2±0.04	6.6±1.12	10.16±2.44	5.52±0.53
	CHGH	0.1±0.01	0.19±0.0	0.35±0.0	0.64±0.01	0.7±0.0	0.91±0.07	1.03±0.01
	PreDyGAE	0.08±0.01	0.09±0.0	0.12±0.02	0.23±0.02	0.18±0.0	0.27±0.15	0.29±0.16
	Transformer	0.62±0.02	3.89±0.02	10.87±0.05	20.58±0.01	22.02±0.01	31.26±0.02	23.47±0.01
	Autoformer	1.29±0.14	8.84±0.15	24.5±0.22	43.99±0.94	47.53±0.67	66.32±1.02	50.02±0.66
	Informer	0.55±0.02	3.83±0.03	10.86±0.03	20.6±0.01	22.04±0.01	31.23±0.01	23.48±0.0
	Reformer	0.56±0.01	3.85±0.01	10.9±0.03	20.65±0.02	22.01±0.01	31.25±0.01	23.48±0.0
	FEDformer	1.42±0.04	8.58±0.18	25.11±0.59	44.15±0.6	47.07±0.66	65.95±0.23	50.0±0.07
	NStransformer	0.06±0.0	0.09±0.0	0.11±0.0	0.19±0.0	0.24±0.0	0.34±0.0	0.34±0.0
	PatchTST	0.06±0.0	0.09±0.01	0.11±0.02	0.18±0.04	0.23±0.05	0.33±0.07	0.33±0.06
	DLinear	0.27±0.15	1.65±1.07	4.63±3.09	8.76±5.86	9.4±6.26	13.33±8.89	10.05±6.65
	TSMixer	0.85±0.03	3.15±0.69	7.57±1.43	18.88±0.49	15.12±3.85	18.76±2.85	17.77±1.58
	FreTS	0.12±0.0	0.31±0.04	0.69±0.11	1.26±0.21	1.37±0.22	1.94±0.32	1.54±0.24
	FiLM	0.06±0.01	0.09±0.01	0.12±0.01	0.19±0.02	0.25±0.03	0.35±0.05	0.34±0.05
Koopa	0.06±0.0	0.08±0.0	0.09±0.0	0.15±0.01	0.2±0.01	0.27±0.02	0.28±0.01	
RMSE(%)	LSTM	0.61±0.01	1.74±0.04	2.89±0.08	4.06±0.09	4.16±0.05	5.21±0.1	4.43±0.05
	SegRNN	1.32±0.17	4.98±0.11	6.49±0.22	7.7±0.09	11.23±1.78	14.58±3.23	9.18±0.83
	CHGH	0.36±0.01	0.49±0.0	0.7±0.0	1.12±0.02	1.14±0.0	2.11±0.32	2.31±0.07
	PreDyGAE	0.23±0.0	0.29±0.0	0.36±0.0	0.85±0.0	0.91±0.0	1.15±1.05	1.29±0.72
	Transformer	6.22±0.12	16.65±0.08	27.9±0.14	38.45±0.03	39.77±0.01	47.4±0.02	40.97±0.01
	Autoformer	15.11±2.51	38.94±0.61	63.72±0.49	82.56±1.83	86.21±1.22	100.99±1.58	87.83±1.13
	Informer	5.63±0.15	16.41±0.13	27.89±0.05	38.5±0.03	39.79±0.01	47.36±0.01	40.98±0.0
	Reformer	5.62±0.05	16.45±0.05	27.96±0.11	38.58±0.02	39.76±0.03	47.38±0.02	40.98±0.02
	FEDformer	17.24±0.68	37.52±0.77	64.93±1.42	82.85±1.16	85.38±1.21	100.35±0.34	87.74±0.1
	NStransformer	0.21±0.0	0.28±0.0	0.35±0.0	0.77±0.01	0.86±0.0	1.69±0.0	1.24±0.0
	PatchTST	0.21±0.01	0.28±0.02	0.35±0.03	0.75±0.04	0.84±0.06	1.65±0.1	1.22±0.07
	DLinear	2.21±1.49	6.05±4.1	10.27±6.97	14.16±9.59	14.64±9.91	17.51±11.77	15.11±10.19
	TSMixer	10.5±0.04	14.32±2.43	21.09±3.64	34.92±1.37	25.27±5.66	26.48±3.26	31.07±2.82
	FreTS	0.47±0.02	0.96±0.07	1.56±0.12	2.22±0.17	2.29±0.17	2.99±0.2	2.54±0.17
	FiLM	0.21±0.0	0.28±0.0	0.35±0.02	0.76±0.04	0.85±0.06	1.67±0.11	1.22±0.05
Koopa	0.21±0.0	0.27±0.0	0.33±0.01	0.72±0.01	0.79±0.02	1.57±0.04	1.17±0.01	
SMAPE(%)	LSTM	0.26±0.0	0.92±0.02	2.25±0.05	4.21±0.09	4.52±0.05	6.5±0.11	5.0±0.05
	SegRNN	0.27±0.03	2.09±0.03	4.55±0.17	7.38±0.06	11.1±1.63	16.79±3.54	9.58±0.83
	CHGH	0.19±0.01	0.36±0.0	0.69±0.0	1.24±0.02	1.36±0.0	1.78±0.01	1.77±0.06
	PreDyGAE	0.09±0.0	0.19±0.0	0.22±0.0	0.28±0.0	0.39±0.0	0.64±0.27	0.57±0.3
	Transformer	0.81±0.02	4.78±0.01	13.28±0.04	25.1±0.0	26.89±0.01	38.12±0.01	28.71±0.01
	Autoformer	1.13±0.02	7.42±0.05	20.74±0.1	38.35±0.38	41.28±0.29	58.08±0.43	43.84±0.29
	Informer	0.75±0.02	4.74±0.02	13.27±0.03	25.12±0.0	26.9±0.01	38.1±0.01	28.72±0.0
	Reformer	0.75±0.01	4.75±0.01	13.31±0.02	25.16±0.02	26.87±0.01	38.12±0.01	28.72±0.0
	FEDformer	1.17±0.01	7.35±0.08	21.02±0.24	38.42±0.24	41.09±0.28	57.94±0.1	43.84±0.03
	NStransformer	0.11±0.0	0.17±0.0	0.22±0.0	0.35±0.0	0.45±0.0	0.61±0.0	0.63±0.0
	PatchTST	0.11±0.01	0.17±0.02	0.21±0.05	0.34±0.09	0.44±0.1	0.59±0.14	0.61±0.11
	DLinear	0.43±0.19	2.58±1.39	7.19±4.0	13.59±7.59	14.59±8.11	20.67±11.52	15.62±8.62
	TSMixer	0.92±0.03	4.05±0.67	10.03±1.36	23.97±1.21	21.46±4.26	27.6±3.5	23.83±1.45
	FreTS	0.22±0.01	0.59±0.07	1.31±0.21	2.4±0.41	2.63±0.43	3.69±0.62	2.95±0.46
	FiLM	0.11±0.02	0.17±0.02	0.22±0.02	0.36±0.03	0.46±0.06	0.63±0.09	0.64±0.09
Koopa	0.1±0.0	0.15±0.01	0.18±0.01	0.28±0.01	0.37±0.02	0.49±0.03	0.52±0.03	
RRMSE(%)	LSTM	43.76±0.35	81.48±0.6	85.53±0.64	91.34±0.35	93.23±0.15	97.48±0.09	96.6±0.07
	SegRNN	72.41±4.5	96.92±0.17	96.4±0.21	96.99±0.11	98.79±0.46	99.24±0.25	98.15±0.28
	CHGH	27.4±0.7	37.66±0.07	37.45±0.03	53.25±0.58	58.32±0.03	80.91±5.07	70.89±2.0
	PreDyGAE	20.64±0.0	28.83±0.0	40.02±0.01	40.12±0.0	47.5±0.0	76.9±42.81	66.68±33.49
	Transformer	97.99±0.07	99.73±0.01	99.8±0.0	99.89±0.0	99.92±0.0	99.97±0.0	99.96±0.0
	Autoformer	99.63±0.12	99.95±0.0	99.96±0.0	99.97±0.0	99.98±0.0	99.99±0.0	99.98±0.0
	Informer	97.59±0.12	99.72±0.0	99.8±0.0	99.89±0.0	99.92±0.0	99.97±0.0	99.96±0.0
	Reformer	97.57±0.03	99.72±0.01	99.81±0.0	99.89±0.0	99.92±0.0	99.97±0.0	99.96±0.0
	FEDformer	99.75±0.02	99.95±0.0	99.96±0.0	99.97±0.0	99.98±0.0	99.99±0.0	99.98±0.0
	NStransformer	16.44±0.44	22.17±0.31	19.66±0.16	37.34±0.26	45.81±0.26	75.98±0.13	61.84±0.13
	PatchTST	16.32±0.88	21.85±1.31	19.5±1.44	36.66±1.59	44.97±1.95	75.59±0.49	61.55±1.34
	DLinear	78.33±15.19	95.41±3.54	96.69±2.57	97.94±1.58	98.55±1.11	99.37±0.44	98.87±0.84
	TSMixer	99.29±0.02	99.61±0.13	99.64±0.11	99.86±0.02	99.8±0.11	99.85±0.0	99.91±0.05
	FreTS	34.81±1.43	61.32±3.03	66.35±3.08	77.25±2.53	81.41±2.37	92.32±1.5	90.05±1.64
	FiLM	16.49±0.39	21.81±0.76	19.48±0.05	36.44±0.97	44.79±2.15	75.91±10.46	61.27±8.05
Koopa	16.47±0.04	21.47±0.18	18.71±0.07	35.58±0.07	43.65±0.38	75.85±2.58	61.26±1.98	

where $a^i \in p$ represents a job advertisement p containing the attribute a^i under granularity i . Similarly, we can further define skill demand proportions $R_{s,t}^{i,j,\dots,k}$ across multiple granularities $\{i, j, \dots, k\}$ by

calculating:

$$R_{s,t,\bar{a}}^{i,j,\dots,k} = \frac{\sum_{p \in \mathcal{P}_t} \mathbf{1}(s \in p) \cdot \mathbf{1}(a^i \in p \wedge a^j \in p \wedge \dots \wedge a^k \in p)}{\sum_{p \in \mathcal{P}_t} \mathbf{1}(a^i \in p \wedge a^j \in p \wedge \dots \wedge a^k \in p)}, \quad (6)$$

where $\bar{a} = \{a^i, a^j, \dots, a^k\}$, $a^i \in \mathcal{A}^i$, $a^j \in \mathcal{A}^j$, \dots , $a^k \in \mathcal{A}^k$, and $R_{s,t}^{i,j,\dots,k} \in \mathbb{R}^{|\mathcal{A}^i| \times |\mathcal{A}^j| \times \dots \times |\mathcal{A}^k|}$.

Results We continue to utilize the benchmark models described in the main text for this task, and the results, as shown in Table 12, lead to the following conclusions: Firstly, the best-performing model on the task of forecasting the proportion of skill demand is Koopa. This model, integrating time series decomposition and Fourier transformations, effectively captures the distribution changes in demand proportions. Secondly, there is a significant variation in performance across models in this task. For example, models like DLinear perform poorly on this task, though they are reasonably effective in skill demand forecasting. We analyze that predicting percentages is distinct from forecasting skill demand, as percentage predictions are also influenced by the demand for other skills at the same granularity. Therefore, simple linear models are not advantageous for capturing the complex interrelations and influences among multiple pieces of information.

E Data Structure and Components

Our dataset comprises five components for each granularity level: job skill demand sequences, job skill demand proportion sequences, ID mapping index, the indexes of skills with structural breaks, and skill co-occurrence graph. Each component is structured as follows: (1) *Job Skill Demand Sequences*: These are presented in tabular files, where each row represents a specific skill, and each column corresponds to a different time slice (month). Each cell within the table contains a numerical value that reflects the demand for the respective skill during that month. (ii) *Job Skill Demand Proportion Sequences*: This component is also formatted in tabular files similar to the skill demand sequences. However, each cell in these tables displays a value between 0 and 1, representing the proportion of demand defined in Eq 6. This provides a normalized view of skill demand across different granularities. (iii) *ID mapping index*: In the dataset, various elements such as regions, occupations, companies, and skills are represented using unique identifiers (IDs) for the convenience of experimentation and analysis. An index table is provided that maps each ID to the actual names of regions, occupations, and skills, facilitating clear and effective data interpretation. The names of companies, however, are withheld due to potential privacy concerns. This selective anonymization ensures that while company-related data remains confidential, researchers can still access necessary details about other variables. During the review period, we open part of id-to-entity mappings to illustrate its use and structure. (iv) *Indexes of skills with structural breaks*: In the provided dataset, data concerning skills that have experienced structural breaks are organized in JSON format. Each granularity level is represented by a separate JSON file, which contains a list of indexes. These indexes correspond to the skills that have undergone structural breaks and can be directly mapped to the skill indexes in the skill demand sequences. The purpose of supplying this data is to facilitate research on the demand trends of skills that have exhibited structural breaks, enabling a detailed analysis of their demand dynamics over time. (v) *Skill Co-occurrence Graph*: This data is provided as a set of triples (skill ID_1, skill ID_2, frequency of co-occurrence), forming a collection that outlines the co-occurrence relationships between skills. Each triple indicates how frequently two skills are mentioned or required together within the job advertisements in the training data, serving as a prior knowledge graph to enhance predictive modeling by capturing relationships between skills.

F Datasheet for Datasets

F.1 Motivation

- **For what purpose was the dataset created?** The dataset in this paper was created specifically for the task of skill demand forecasting. The creation of this dataset aims to address two primary objectives: first, to fill the existing gap in publicly available datasets for skill demand forecasting; second, to offer a dataset that encompasses multiple levels of granularity. This diversity in granularity enables the possibility of predicting skill trends at various levels of detail. By providing such a comprehensive dataset, researchers can gain a more thorough understanding of and ability to predict changes in skill demand, thereby

facilitating job seekers in finding suitable positions and aids recruiters in hiring the right talent. Furthermore, this dataset can support the development of education, training, and employment policies.

- **Who created the dataset (e.g., which team, research group) and on behalf of which entity (e.g., company, institution, organization)?** Our dataset is jointly developed by a collaborative effort from the following affiliations: University of Science and Technology of China, Computer Network Information Center, Chinese Academy of Sciences, Institute of Artificial Intelligence, Beihang University, and Artificial Intelligence Thrust, The Hong Kong University of Science and Technology (Guangzhou).
- **Who funded the creation of the dataset?** This work was supported by University of Science and Technology of China, Computer Network Information Center, Chinese Academy of Sciences, and The Hong Kong University of Science and Technology (Guangzhou).

E.2 Composition

- **What do the instances that comprise the dataset represent (e.g., documents, photos, people, countries)?** See the Section E.
- **How many instances are there in total (of each type, if appropriate)?** We collected millions of public job advertisements, covering 521 companies, 52 occupations, and 7 regions. We further extracted 2335 skills and tracked the monthly demand for these skills, thus ultimately forming sequences of demand for 2335 skills over 36 months, broken down by company, occupation, and region.
- **Does the dataset contain all possible instances or is it a sample (not necessarily random) of instances from a larger set?** Yes, the dataset encompasses an exhaustive collection of job recruitment demands posted by various companies on online platforms. In the subsequent data processing stages, we employed rule-based matching and Named Entity Recognition (NER) models to extract as many relevant skill demands as possible from each job advertisement. Additionally, we filtered out certain implausible skill terms based on skill frequency analysis and manual review. Therefore, the dataset includes all possible instances except for those terms deemed unreasonable and filtered out during the review process.
- **What data does each instance consist of?** See the Subsection 3.1 and Section E.
- **Is there a label or target associated with each instance?** Yes, given that our dataset is a time series collection, it is structured for unsupervised learning. This means there are no explicit labels or target variables associated with each instance. Instead, the data is used to observe and analyze patterns over time, particularly in skill demand and relationships between skills, without predefined outcomes or classifications.
- **Is any information missing from individual instances?** N/A
- **Are relationships between individual instances made explicit (e.g., user movie ratings, social network links)?** We supplemented our analysis with a co-occurrence graph to explicitly represent the relationships between skills based on their co-occurrence in various contexts.
- **Are there recommended data splits (e.g., training, development/validation, testing)?** Yes, we recommended the data split in Subsection 3.2.
- **Are there any errors, sources of noise, or redundancies in the dataset?** Yes. See the Subsection 3.1
- **Is the dataset self-contained, or does it link to or otherwise rely on external resources (e.g., websites, tweets, other datasets)?** The dataset is self-contained for the tasks described in the paper, which is collected from online websites. a) We can not ensure the linked resources exist for a long time and consistent. b) There is no official archive related to the linked resource. c) There are no restrictions.
- **Does the dataset contain data that might be considered confidential (e.g., data that is protected by legal privilege or by doctor-patient confidentiality, data that includes the content of individuals' non-public communications)?** No.
- **Does the dataset contain data that, if viewed directly, might be offensive, insulting, threatening, or might otherwise cause anxiety?** No.

F.3 Collection Process

- **How was the data associated with each instance acquired?** See the Section 3.1.
- **What mechanisms or procedures were used to collect the data?** See Section 3.1.
- **If the dataset is a sample from a larger set, what was the sampling strategy (e.g., deterministic, probabilistic with specific sampling probabilities)?** No.
- **Who was involved in the data collection process (e.g., students, crowdworkers, contractors) and how were they compensated (e.g., how much were crowdworkers paid)?** We have collaborated with industry experts who have relevant experience to assist in data annotation for skill dictionary construction. Their participation is voluntary and unpaid.
- **Over what timeframe was the data collected?** The raw data, consisting of job advertisements, was collected over a period in the past three years.
- **Were any ethical review processes conducted (e.g., by an institutional review board)?** No.

F.4 Preprocessing/cleaning/labeling

- **Was any preprocessing/cleaning/labeling of the data done (e.g., discretization or bucketing, tokenization, part-of-speech tagging, SIFT feature extraction, removal of instances, processing of missing values)?** Yes. See the Subsection 3.1.
- **Was the “raw” data saved in addition to the preprocessed/cleaned/labeled data (e.g., to support unanticipated future uses)?** No.
- **Is the software that was used to preprocess/clean/label the data available?** No.

F.5 Uses

- **Has the dataset been used for any tasks already?** No.
- **Is there a repository that links to any or all papers or systems that use the dataset?** No.
- **What (other) tasks could the dataset be used for?** N/A
- **Is there anything about the composition of the dataset or the way it was collected and preprocessed/cleaned/labeled that might impact future uses?** N/A.
- **Are there tasks for which the dataset should not be used?** N/A.

F.6 Distribution

- **Will the dataset be distributed to third parties outside of the entity (e.g., company, institution, organization) on behalf of which the dataset was created?** No. Our dataset will be managed and maintained by our research group.
- **How will the dataset will be distributed (e.g., tarball on website, API, GitHub)?** The evaluation dataset is released to the public and hosted on GitHub.
- **When will the dataset be distributed?** It has been released now.
- **Will the dataset be distributed under a copyright or other intellectual property (IP) license, and/or under applicable terms of use (ToU)?** Our dataset will be distributed under the CC BY-SA 4.0 license.
- **Have any third parties imposed IP-based or other restrictions on the data associated with the instances?** No.
- **Do any export controls or other regulatory restrictions apply to the dataset or to individual instances?** No.

F.7 Maintenance

- **Who will be supporting/hosting/maintaining the dataset?** The authors.
- **How can the owner/curator/manager of the dataset be contacted (e.g., email address)?** By the email address.

- **Is there an erratum?** N/A.
- **Will the dataset be updated (e.g., to correct labeling errors, add new instances, delete instances)?** Yes, the dataset is expected to be updated with additional job advertisements collected over extended periods, potentially on a biannual basis. These updates will focus on increasing our data volume over time. All updates will be made available and communicated to dataset consumers via GitHub.
- **If the dataset relates to people, are there applicable limits on the retention of the data associated with the instances (e.g., were the individuals in question told that their data would be retained for a fixed period of time and then deleted)?** N/A.
- **Will older versions of the dataset continue to be supported/hosted/maintained?** We plan to maintain the newest version only.
- **If others want to extend/augment/build on/contribute to the dataset, is there a mechanism for them to do so?** Yes. Contact the authors of the paper.

Checklist

The checklist follows the references. Please read the checklist guidelines carefully for information on how to answer these questions. For each question, change the default **[TODO]** to **[Yes]**, **[No]**, or **[N/A]**. You are strongly encouraged to include a **justification to your answer**, either by referencing the appropriate section of your paper or providing a brief inline description. For example:

- Did you include the license to the code and datasets? **[Yes]** See Section X.
- Did you include the license to the code and datasets? **[N/A]**

Please do not modify the questions and only use the provided macros for your answers. Note that the Checklist section does not count towards the page limit. In your paper, please delete this instructions block and only keep the Checklist section heading above along with the questions/answers below.

1. For all authors...
 - (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? **[Yes]** Please refer to the Abstract and Introduction sections.
 - (b) Did you describe the limitations of your work? **[Yes]** Please refer to dataset limitation in Section 3.3.
 - (c) Did you discuss any potential negative societal impacts of your work? **[Yes]** Please refer to the Appendix A.
 - (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? **[Yes]** This work adheres to the NeurIPS Code of Ethics in each respect.
2. If you are including theoretical results...
 - (a) Did you state the full set of assumptions of all theoretical results? **[N/A]**
 - (b) Did you include complete proofs of all theoretical results? **[N/A]**
3. If you ran experiments (e.g. for benchmarks)...
 - (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? **[Yes]** All the code, data, and instructions needed to reproduce the main experimental results can be found at <https://github.com/Job-SDF/benchmark>.
 - (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? **[Yes]** We have provided the detailed implementation descriptions in the Appendix.
 - (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? **[Yes]** Due to the page limit, we report error bars in the Appendix.
 - (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? **[Yes]** The experimental environment is reported in the Appendix.

4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
 - (a) If your work uses existing assets, did you cite the creators? [\[Yes\]](#) We have cited the original paper that produced the code package.
 - (b) Did you mention the license of the assets? [\[Yes\]](#) We have stated the detailed version and license of each asset. See the Appendix for the details.
 - (c) Did you include any new assets either in the supplemental material or as a URL? [\[Yes\]](#) We have released a new dataset and the codes for the benchmark models.
 - (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [\[Yes\]](#) Please refer to the Appendix.
 - (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [\[Yes\]](#) The proposed dataset does not contain personally identifiable information or offensive content. See the Appendix for the details.
5. If you used crowdsourcing or conducted research with human subjects...
 - (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [\[N/A\]](#)
 - (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [\[N/A\]](#)
 - (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [\[N/A\]](#)