Artificial Intelligence and Computing Innovations

Intelligence Science Publishing

https://www.ispublishing.com

ISSN: 3070-4944

Vo.5, No.8, 2025

Multi-Resource Workload Forecasting in Cloud Data Centers via Spatio-Temporal Graph Transformers

Feng Wang[1], Deepika Verma[2], Yuting Chen[2*]

[1]North Carolina State University, Raleigh, USA

[2]New York University, New York, USA

**Corresponding author: wangfeng.jan@gmail.com*

Abstract: Addressing the challenges of strong spatiotemporal coupling, significant cross-node propagation, and dynamic topology changes with scheduling and scaling in multi-resource load prediction for cloud data centers, this paper proposes a unified modeling framework based on the Spatio Temporal Graph Transformer. This framework organizes multi-resource observations in the form of a temporal graph and explicitly injects relational structure information. The method abstracts the data center into a dynamic graph composed of nodes and edges at each time slice. Node features aggregate multi-dimensional resource indicators such as CPU, memory, disk, and network, while relational features characterize structural cues such as link types and interaction strengths, thus unifying multi-resource and structural constraints at the input level. Spatial encoding employs a graph attention mechanism to adaptively weight and aggregate neighborhood information, strengthening the expression of critical dependency paths and suppressing noise propagation. Temporal encoding utilizes self-attention to establish global associations across time steps within a historical window, improving the ability to characterize long-range dependencies and sudden disturbances. Subsequently, a gated fusion residual mechanism adjusts the information injection intensity, balancing representation updates and robustness. The output layer adopts a structure combining a shared backbone with resource-specific prediction heads, preserving common cross-resource patterns while characterizing resource dimensional differences, achieving multi-step, multi-resource joint prediction. To verify the effectiveness of the method, a unified evaluation index system was set up, and a horizontal comparison with representative baseline methods was conducted. The results show that the proposed framework can provide more consistent multi-resource prediction signals and provide more reliable forward-looking support for capacity planning and flexible scheduling.

Keywords: Temporal graph modeling; multi-resource joint prediction; graph attention aggregation; gated residual fusion.

1 Introduction

Cloud data centers host large-scale online services and batch processing tasks, exhibiting strong non-stationarity and strong coupling between resource supply and business demands in both time and space[1,2]. With the widespread adoption of virtualization and containerization, computing instances frequently migrate and elastically scale within clusters, causing node-level load to be driven not only by local workloads but also by cross-node dependencies, network topology, and scheduling strategies, resulting in spatial propagation effects. Significant complementary and competitive relationships exist among multiple resource dimensions, such as CPU, memory, disk, and network, exhibiting asynchronous changes and sudden peaks at different service stages. Addressing this complexity, multi-resource load forecasting can provide a forward-looking basis for resource orchestration, capacity planning, and operational risk management, thereby improving resource utilization and reducing the probability of service degradation[3].

Traditional load forecasting often focuses on single-node time series, making it difficult to characterize the spatial correlations caused by cross-node interactions and to handle the dynamic dependency structures between multiple resources within a unified framework. In real-world data centers, factors such as service call chains, shared storage, and network congestion can cause load diffusion and lag across the topology, leading to incomplete local observations of the global state. Meanwhile, resource indicators suffer from scale differences, noise interference, and missing sampling, making predictions relying solely on local statistical models prone to bias and instability[4]. To clearly present the key challenges and engineering value of load forecasting, Table 1 summarizes the core difficulties and corresponding significance of multi-resource spatiotemporal forecasting, supporting the necessity and research motivation for subsequent method design.

Table 1. Key challenges and research significance of multi-resource load forecasting in cloud data centers

Key Challenge	Typical Manifestation	Research Significance
Spatio-temporal coupling and propagation effects	Loads diffuse along networks and dependency relations, with latency and cascading effects.	Enables modeling of cross-node influence and improves global forecasting consistency
Multi-resource collaboration and contention	CPU and memory correlations vary across task phases; network and disk are affected by shared bottlenecks	Promotes joint multi-resource forecasting and avoids resource mismatch caused by single-metric optimization
Dynamic topology and dependency changes	Instance migration and elastic scaling lead to frequent changes in the correlation structure	Improves adaptability to structural drift and enhances forecasting robustness
Non-stationary and bursty workloads	Periodicity coexists with bursts, and peak duration is uncertain	Strengthens short-term risk early warning and reduces overload and jitter
Noisy, missing, and heterogeneous sampling	Metric sampling frequencies are inconsistent, and missing values and outliers are prevalent.	Improves robustness to incomplete observations and ensures forecasting usability

Multi-resource load prediction based on the Spatio Temporal Graph Transformer emphasizes simultaneously modeling temporal dependencies and spatial relationships within a unified representation space. It adaptively aggregates cross-node information through an attention mechanism to characterize

propagation patterns[5]. Abstracting the data center into a temporal graph structure allows nodes to be viewed as computational entities or physical hosts, and edges as network connectivity, service dependencies, or statistical correlations, thus making spatial relationships learnable and dynamically updated over time. Furthermore, multi-resource metrics, as multi-channel feature inputs, can explicitly express resource collaboration and competition relationships, enabling the prediction process to consider both local fluctuations and global constraints, reducing mutual interference between resource dimensions. This type of method provides a structured inductive bias for load prediction with complex dependencies, helping to improve generalization capabilities for non-stationary scenarios and structural drift[6,7].

For engineering applications, accurate multi-resource load prediction can provide more reliable prior information for scheduling systems, enabling resource allocation to shift from passive response to proactive planning, and completing elastic scaling and load migration before peak arrival. Characterizing cross-node propagation and multi-resource coupling can also support bottleneck localization and capacity assessment, reduce energy waste caused by over-reservation, and avoid service degradation and cascading failures caused by insufficient resources. As data centers expand in scale and business models diversify, load forecasting research oriented towards spatiotemporal coupling not only has direct economic and operational value but also provides key support for building cloud infrastructure with greater adaptability and security boundaries.

2 Methodology

The core objective of multi-resource workload forecasting for cloud data centers is to jointly capture temporal evolution patterns and cross-node propagation relationships, thereby providing more reliable foresight for capacity planning and elastic scheduling. This paper also presents the overall model architecture, as shown in Figure 1.

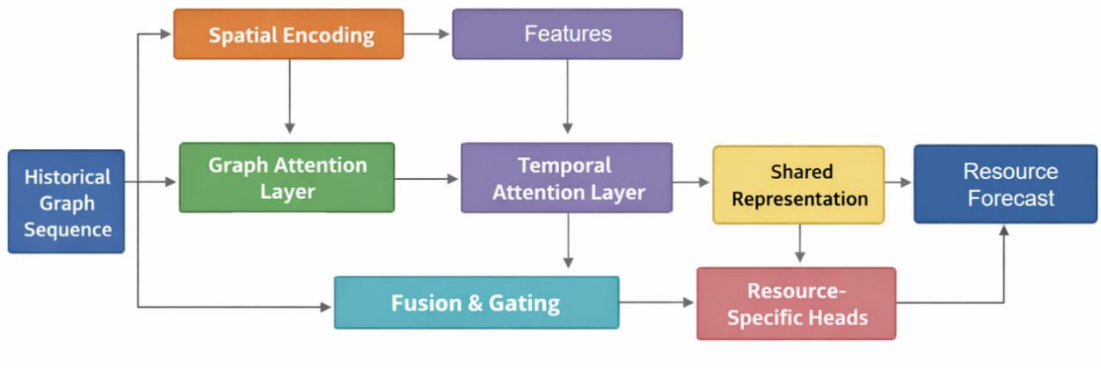


Figure 1. Overall model architecture

To this end, the data center is abstracted as a time-varying graph structure, enabling the states of service instances or physical nodes to interact under topological constraints, and further unifying multi-resource cooperation and competition within a single modeling framework. Formally, the temporal graph is defined as:

$$\mathcal{G}_t = (\mathcal{V}_t, \mathcal{E}_t, \boldsymbol{X}_t, \boldsymbol{R}_t) \tag{1}$$

where the node set $\mathcal{V}_t$ and edge set $\mathcal{E}_t$ jointly describe the dependency relationships at time t, the feature matrix $\boldsymbol{X}_t \in \mathbb{R}^{|\mathcal{V}_t| \times F}$ aggregates multi-resource observations such as CPU, memory, disk, and network, and the relation features $\boldsymbol{R}_t$ are used to carry structural cues, including link type, interaction intensity, or statistical correlation, to enhance the flexibility of propagation modeling. The forecasting task organizes the sequence in a sliding-window manner, learning the future H-step multi-resource workload from a historical graph sequence of length L, and thus, the input–output mapping is defined as:

$$\widehat{\boldsymbol{Y}}_{t+1:t+H} = f_\theta\left(\mathcal{G}_{t-L+1:t}\right) \tag{2}$$

The significance of this mapping lies in that the unified model f_θ controlled by parameters θ simultaneously absorbs temporal dependencies and structural dependencies, thereby reducing the sensitivity of single-node modeling to incomplete global states and enhancing cross-resource consistency.

High-order dependency graph learning for container anomaly detection demonstrates the value of preserving multi-hop operational dependencies when constructing node relations; accordingly, the proposed graph treats direct topology, service interaction, and correlation-derived links as complementary relational evidence rather than a fixed adjacency [8]. Graph-based knowledge representation and semantic alignment further motivate the use of relation embeddings to maintain a coherent feature space across heterogeneous node and edge types, allowing structural cues to remain comparable as the data-center graph evolves [9].

To mitigate topological noise and explicitly represent cross-node propagation, spatial encoding adopts a graph-attention aggregation mechanism, enabling node representations to adaptively select more interpretable dependency paths within neighborhoods. Specifically, a weighted aggregation is performed over the neighborhood $\mathcal{N}(i)$ of node i, where attention weights are jointly determined by node representations and relation features, thereby avoiding the rigidity induced by fixed adjacency. The updated form of the spatial layer is:

$$\boldsymbol{h}_{i,t}^{(l+1)} = \sum_{j \in \mathcal{N}(i)} \alpha_{ij,t}^{(l)} \boldsymbol{W}^{(l)} \boldsymbol{h}_{j,t}^{(l)} \tag{3}$$

where $\boldsymbol{h}_{i,t}^{(l)}$ denotes the node representation at time t in layer l, the linear mapping matrix $\boldsymbol{W}^{(l)}$ aligns different resource scales in the feature space, and the weight $\alpha_{ij,t}^{(l)}$ emphasizes critical dependency edges to improve the robustness of propagation modeling. To make the weights comparable within the neighborhood, a normalized scoring function is used:

$$\alpha_{ij,t}^{(l)} = softmax_j\left(\frac{\left(\boldsymbol{Q}\boldsymbol{h}_{i,t}^{(l)}\right)^\top \left(\boldsymbol{K}\left[\boldsymbol{h}_{j,t}^{(l)} \| \boldsymbol{r}_{ij,t}\right]\right)}{\sqrt{d}}\right) \tag{4}$$

In this expression, the concatenated vector $\left[\boldsymbol{h}_{j,t}^{(l)} \| \boldsymbol{r}_{ij,t}\right]$ fuses the neighbor state with the edge relation embedding $\boldsymbol{r}_{ij,t}$, and scaling by dimension d stabilizes the numerical range, so that selective aggregation remains learnable when structural drift and heterogeneous links coexist.

To address the non-stationarity and burstiness of workload sequences, temporal encoding introduces position-related learnable representations to strengthen the ability to capture long-range dependencies and establishes global associations across time steps within the window via self-attention. First, the outputs of the spatial layer are stacked into a temporal sequence representation $\boldsymbol{S}_t = [\boldsymbol{H}_{t-L+1}, \ldots, \boldsymbol{H}_t]$, and time embeddings are added to distinguish the semantic roles of different time steps, thereby reducing representational confusion caused by periodic perturbations. The attention update of the temporal layer can be written as:

$$\boldsymbol{Z} = softmax\left(\frac{\boldsymbol{Q}_T \boldsymbol{K}_T^{\top}}{\sqrt{d}}\right)\boldsymbol{V}_T \tag{5}$$

where $\boldsymbol{Q}_T, \boldsymbol{K}_T, \boldsymbol{V}_T$ are obtained via linear projections of the sequence representation, enabling different time steps to form connections through content relevance to alleviate the gradient attenuation of traditional recurrent structures on long sequences. To further enhance responsiveness to extreme fluctuations, a gating mechanism is introduced in the residual path to adaptively regulate the update strength and avoid over-correction caused by noise and missing values, in the form:

This temporal formulation also draws on adaptive representation learning and temporal state modeling, which emphasizes separating persistent operational states from transient fluctuations when sequences exhibit structural drift [10]. Generalizable latent representation learning for multi-source cloud time series further supports using a shared temporal space before resource-specific decoding, so that common demand dynamics can be retained despite heterogeneous measurement scales and sources [11].

$$\boldsymbol{H}' = \boldsymbol{H} + \sigma(\boldsymbol{U}\boldsymbol{H}) \odot \boldsymbol{Z} \tag{6}$$

where the gating function $\sigma(\cdot)$ and the element-wise product $\odot$ jointly implement bit-wise filtering of the temporal update $\boldsymbol{Z}$, thereby suppressing the disruption of outliers to overall trend modeling while keeping information flow smooth.

Residual-based anomaly characterization and Transformer sequence modeling provide complementary robustness signals for this gating path: deviations between reconstructed and observed states can expose unstable operating patterns [12], while sequence-aware anomaly detection in Kubernetes clusters shows the importance of retaining contextual temporal evidence around bursts and regime changes [13].

For multi-resource joint forecasting, the output layer adopts a shared backbone with resource-specific heads, so as to balance cross-resource common patterns and resource-differentiated dynamics, thereby alleviating the bias of a single output head toward indicators with different units and scales. Specifically, the final representation is first mapped into the multi-resource prediction space, and then the future H-step sequence is output along the resource dimension, formulated as:

$$\widehat{\boldsymbol{Y}} = Head(\boldsymbol{H}') \tag{7}$$

The significance of this design is that the shared representation captures the common structures of cross-resource coupling and cross-node propagation, while the parameterization of Head enhances the

characterization of differentiated behaviors such as CPU spikes, memory lag, and network congestion, thereby providing consistent and alignable multi-dimensional predictive signals for risk warning and elastic control decisions. The training objective adopts a weighted regression loss over multiple steps and multiple resources to balance scale differences across resources, where the weight vector can be set according to historical variance or business priority and embedded into the loss to strengthen forecasting reliability for key resources, thereby achieving a unified trade-off between usability and robustness under engineering constraints.

To make the multi-resource outputs comparable for downstream decisions, the training process additionally benefits from similarity-aware calibration: logistic similarity learning provides a useful principle for aligning related representations while preserving discriminability among resource-specific behaviors [14].

3 Dataset

This paper uses Alibaba's open-source cluster trace data, cluster trace v2018, as the research dataset. This dataset originates from the operational logs of a real production data center cluster and is provided as a public project for download, along with field descriptions. Covering multiple days of continuous operation, this dataset includes server and container scheduling event information and resource usage observation records, providing a high-density, long-span, and representative time-series observation foundation for multi-resource load prediction.

At the resource dimension, this dataset provides various utilization or usage information related to computing and storage, and uses multiple tables to depict the temporal state evolution of entities such as machines, tasks, and instances, thus supporting a modeling organization method that uses nodes as machines or instances and time windows as sequence segments. Simultaneously, the dataset contains structural information of batch processing workloads and task association records, which can be used to construct graph structure constraints consistent with dependency propagation. This enables spatiotemporal graph prediction tasks to form reproducible input-output definitions and data processing flows under open-source data conditions.

4 Experimental Results and Analysis

To systematically characterize the method positioning and comparison boundary of the proposed model in the scenario of multi-resource load prediction in cloud data centers, representative spatiotemporal graph modeling methods for cloud workload and resource usage prediction in recent years are selected as references, and a horizontal comparison table is given under a unified evaluation index system, to make a comprehensive comparison from the dimensions of error magnitude, relative deviation and fitting consistency. The experimental results are shown in Table 2.

Table 2. Experimental results compared with other models

Method	MAE	RMSE	MAPE	RMAE
Li et al.[15]	0.52	0.78	6.45	0.60
Ding et al.[16]	0.49	0.74	6.10	0.57
Karthikeyan et al.[17]	0.47	0.71	5.96	0.55
Li et al.[18]	0.45	0.69	5.72	0.53

Shen et al.[19]	0.43	0.66	5.48	0.51
Gao et al.[20]	0.41	0.63	5.22	0.49
Xue et al.[21]	0.40	0.62	5.14	0.48
Ours	0.36	0.57	4.78	0.44

The results shown in Table 2 demonstrate that the proposed method exhibits superior overall performance in error control and stability for multi-resource load prediction. In the joint modeling of spatial and temporal dependencies, this method more fully absorbs cross-node propagation information while maintaining robust characterization of local non-stationary fluctuations, thus achieving more consistent prediction quality across different resource dimensions. The synergistic mechanism of shared representations and resource-specific outputs further reduces interference from dimensional differences and coupling noise, making the overall estimate closer to the actual load pattern and reducing extreme biases.

From an indicator perspective, both error correlation and relative deviation metrics show superior performance, indicating that the model is not only more accurate in an average sense but also maintains more reliable fitting consistency under scale normalization evaluation. The gated fusion structure provides adaptive adjustment in information injection intensity, preserving effective patterns while suppressing anomalous disturbances, thereby improving the usability of multi-step prediction. In summary, this method is more valuable for applications in resource orchestration and capacity planning, providing more stable forward signals for subsequent decision-making modules.

The length of the time window determines the range of historical context that the model can utilize and directly affects its ability to characterize periodic patterns and sudden disturbances. A window that is too short is prone to losing lag information propagating across nodes, while a window that is too long may introduce outdated states and structural drift noise. Based on this, different historical window lengths were set to examine the sensitivity of prediction errors to the range of time context, thus providing a basis for window selection in practical deployments. The experimental results are shown in Figure 2.

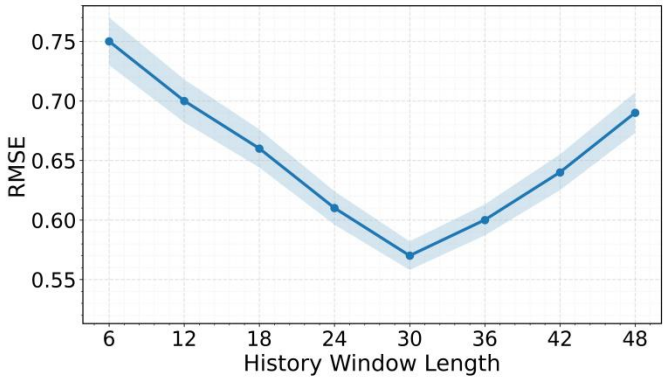


Figure 2. Experimental results on the sensitivity of different time window lengths to RMSE

This figure illustrates the clear controllability of the proposed method's choice of temporal context range. A suitable historical window can strike a balance between preserving key dependency information and suppressing redundant noise. With a shorter window, the lagged effects of cross-node propagation and the phased characteristics of multi-resource coupling are difficult to capture fully,

resulting in limited error control space. As the window gradually expands, spatiotemporal dependencies are more completely incorporated into the representation learning process, and the model shows greater stability in jointly characterizing periodic patterns and sudden disturbances, demonstrating the effective modeling capability of the proposed spatiotemporal graph Transformer for long-range associations.

As the window continues to increase, the marginal benefit of historical information begins to weaken. Some expired states and structural drift noise are introduced into the attention aggregation path, making the prediction process more sensitive to non-stationary changes. The gating fusion mechanism alleviates the injection of irrelevant information to some extent, but still faces the risks of information dilution and attention dispersion in extremely long contexts. Overall, the proposed method performs better within a reasonable window range, maintaining a stable trade-off between engineering availability and robustness, and providing more reliable forward signals for subsequent capacity planning and flexible scheduling.

5 Conclusion

This paper addresses the critical issue of multi-resource load prediction in cloud data centers, constructing a unified modeling framework oriented towards spatiotemporal coupling and structural propagation. By abstracting the data center as a graph structure that evolves over time and jointly characterizing cross-node dependencies and time-series patterns within a unified representation space, the model can simultaneously absorb topological information and multi-resource collaborative relationships, thus providing a more consistent representational basis for prospective estimation of multidimensional loads. This method emphasizes maintaining a stable structural inductive bias in real-world environments where dynamic dependencies and non-stationary fluctuations coexist, ensuring that the prediction process focuses not only on local fitting of a single indicator but also on the overall modeling goal of global consistency and multi-resource collaboration.

At the methodological level, spatial encoding uses an attention mechanism to adaptively select key dependency paths, thereby mitigating the rigidity of fixed adjacencies and improving the ability to characterize topological changes and propagation lags. Temporal encoding further establishes global correlations across time steps within a window, enabling the model to form a compatible temporal understanding between long-range dependencies and short-term disturbances. The design of gating fusion and shared representation combined with resource-specific headers enables multi-resource prediction to utilize common structural information while preserving the dynamic differences in resource dimensions, thereby reducing the interference of dimensional differences and coupling noise on the learning process. The overall framework unifies structural dependencies, temporal dependencies, and multi-resource coupling into a single optimization objective in an end-to-end manner, providing a modeling paradigm for cloud load prediction that is closer to the actual operating mechanism.

From an application value perspective, the prediction capabilities oriented towards multi-resource and cross-node propagation provide more reliable prior information for capacity planning, elastic scaling, and scheduling orchestration, helping to shift resource management from passive response to proactive planning. More consistent multi-dimensional prediction signals can support early identification of bottleneck risks and timely correction of resource mismatches, thereby reducing the probability of

service degradation and cascading failures caused by overload, while also reducing energy consumption and cost waste caused by excessive reservation. This research provides a reproducible technical path for building adaptive, scalable, and secure cloud infrastructure, and offers transferable modeling ideas for intelligent operation and maintenance and automated reliability engineering for complex dependent systems.

Future work can be further deepened in three directions: data, structure, and decision-making closed loop. First, introducing finer-grained entity and relationship definitions into broader cloud-native scenarios will enable graph structures to simultaneously cover multiple dependencies, such as service calls, network bandwidth, and storage sharing, thereby improving the ability to express complex system behaviors. Second, addressing the issues of noise loss and sampling heterogeneity in real-world operations, more robust representation learning and uncertainty characterization mechanisms can be explored to provide more credible confidence constraints for prediction outputs. Third, a closer linkage between prediction models and resource orchestration strategies will be established to construct an end-to-end closed-loop optimization framework from prediction to scheduling, supporting more efficient elastic control, fault prevention, and globally optimal resource allocation under cost constraints. These directions will further expand the application boundaries of spatiotemporal graph prediction in cloud computing and distributed systems, and drive the evolution of intelligent cloud infrastructure towards higher reliability and stronger autonomy.

References

[1] Miao W W, Zeng Z, Zhang M, et al. Workload prediction in edge computing based on graph neural network[C]//2021 IEEE Intl Conf on Parallel & Distributed Processing with Applications, Big Data & Cloud Computing, Sustainable Computing & Communications, Social Computing & Networking (ISPA/BDCloud/SocialCom/SustainCom). IEEE, 2021: 1663-1666.

[2] Luo Y, Gao M, Yu Z, et al. Integrating system state into spatio temporal graph neural network for microservice workload prediction[C]//Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining. 2024: 5521-5531.

[3] Theodoropoulos T, Makris A, Kontopoulos I, et al. Graph neural networks for representing multivariate resource usage: A multiplayer mobile gaming case-study[J]. International Journal of Information Management Data Insights, 2023, 3(1): 100158.

[4] Yu J, Gao M, Li Y, et al. Workflow performance prediction based on graph structure aware deep attention neural network[J]. Journal of Industrial Information Integration, 2022, 27: 100337.

[5] Ba A, Harsha P, Subramanian C. Leveraging interpretability in the transformer to automate the proactive scaling of cloud resources[J]. arXiv preprint arXiv:2409.03103, 2024.

[6] Xu Y, Ge J, Tang H, et al. System states forecasting of microservices with dynamic spatio-temporal data[J]. arXiv preprint arXiv:2408.07894, 2024.

[7] Park J, Choi B, Lee C, et al. GRAF: Graph neural network-based proactive resource allocation for SLO-oriented microservices[C]//Proceedings of the 17th International Conference on Emerging Networking EXperiments and Technologies (CoNEXT). ACM, 2021.

[8] Hsieh C Y. High-Order Dependency Graph Learning for Container Anomaly Detection in Microservice Environments[J]. 2025.

[9] Zhou J. Graph-Based Knowledge Representation and Semantic Alignment for Large Language Model-Driven Entity Recognition[J]. 2025.

[10] Zhan Y. Adaptive Representation Learning and Temporal State Modeling for Robust Anomaly Detection in Distributed Systems[J]. 2024.

[11] Huang W. Generalizable Latent Representation Learning for Multi-Source Time Series Forecasting in Cloud Systems[J]. 2024.

[12] Ren L, Li S, Liu Z. Variational Autoencoder-Based Query Plan Anomaly Detection and Cost Deviation Warning for Database Systems[J]. 2024.

[13] Chang F. Intelligent Anomaly Detection in Kubernetes Clusters Through Transformer-Based Sequence Modeling[J]. 2024.

[14] Hu R. Unified Semantic Matching with Logistic Similarity Learning for Cold-Start Advertising Recommendation[J]. 2024.

[15] Li Y, Yuan H, Fu Z, et al. Elastic: edge workload forecasting based on collaborative cloud-edge deep learning[C]//Proceedings of the ACM Web Conference 2023. 2023: 3056-3066.

[16] Ding Z, Feng B, Yu W. FELT: Large-Scale Cloud Workload Prediction Through Adaptive Feature-Enhanced and Similarity-Aware Transformer[J]. Tsinghua Science and Technology, 2025.

[17] Karthikeyan R, Abdul Samad S R, Balamurugan V, et al. Workload Prediction in Cloud Data Centers Using Complex-Valued Spatio-Temporal Graph Convolutional Neural Network Optimized With Gazelle Optimization Algorithm[J]. Transactions on Emerging Telecommunications Technologies, 2025, 36(3): e70078.

[18] Li J, Yang D, Guo H, et al. Forecasting resource usage pattern changes in clouds via contrast graph-evolution learning[J]. Future Generation Computer Systems, 2024, 154: 373-383.

[19] Shen F, Wang J, Zhang Z, et al. Long-term multivariate time series forecasting in data centers based on multi-factor separation evolutionary spatial–temporal graph neural networks[J]. Knowledge-Based Systems, 2023, 280: 110997.

[20] Gao M, Xu K, Gao X, et al. Spatial-temporal heterogenous graph contrastive learning for microservice workload prediction[C]//Proceedings of the AAAI Conference on Artificial Intelligence. 2025, 39(11): 11681-11689.

[21] Xue Z, Zi Y, Qi N, et al. Multi-level service performance forecasting via spatiotemporal graph neural networks[C]//2025 5th International Conference on Intelligent Communications and Computing (ICICC). IEEE, 2025: 202-206.