
An Adaptive Multi-Scale Framework for Accurate Forecasting of Performance Metrics in Complex

Elwood Bain

University of Windsor, Windsor, Canada
eb837@uwindsor.ca

Abstract:

This study proposes an adaptive multi-scale representation learning method for system metric prediction to address the dynamic characteristics and non-stationary distributions of multi-dimensional metric sequences in complex systems. The method extracts temporal features at different time granularities through a multi-scale convolutional feature decomposition module, capturing both short-term fluctuations and long-term trends in system state changes. An adaptive feature fusion mechanism is introduced to dynamically weight multi-scale features and enforce consistency across scales, thereby enhancing the model's capability to represent complex time-varying patterns. Structurally, the model integrates hierarchical normalization and gated update units to improve the stability of feature flow and the continuity of temporal dependencies, avoiding prediction degradation under high-frequency disturbances and distribution shifts. In addition, a residual propagation-based dynamic feature transformation layer is constructed to jointly model local information and global semantics, further improving robustness and generalization in multi-dimensional signal interactions. Experimental results show that the proposed method achieves lower MSE, MAE, MAPE, and RMSE values than mainstream models such as Autoformer, EDFormer, and TimesNet on benchmark system metric datasets, confirming its superiority in multi-scale feature reconstruction and dynamic temporal modeling. This research provides an efficient and scalable modeling framework for system performance prediction, intelligent operations, and multi-dimensional time-series analysis, enabling accurate forecasting and structured representation of non-stationary sequences in complex system environments.

Keywords:

Multi-scale representation learning; system indicator prediction; time series modeling; adaptive feature fusion

1. Introduction

In modern computing environments, the dynamic variation of system metrics has become a key indicator for characterizing the operational state of cloud computing, distributed systems, and microservice architectures. With the continuous expansion of application scale and the increasing complexity of underlying infrastructures, system performance metrics exhibit high dimensionality, strong correlations, and non-stationary characteristics. Variations across different temporal granularities often intertwine periodic fluctuations, sudden anomalies, and long-term drifts. Traditional statistical models typically assume stable data distributions and independent feature spaces, but such assumptions rarely hold in real-world systems. As a result, their ability to capture complex dynamic patterns is limited. Effectively modeling and accurately predicting time-varying features in multi-dimensional and heterogeneous metric environments has become a

crucial scientific challenge for system performance management, anomaly warning, and resource scheduling[1,2].

During system operation, intricate hierarchical dependencies and interactions often exist among metrics. For example, fluctuations in CPU utilization may be influenced by network latency, disk I/O, or service load migration. These cross-dimensional correlations exhibit distinct multi-scale temporal characteristics. Short-term perturbations may rapidly propagate and form localized anomalies, while long-term accumulative trends may imply system degradation or resource imbalance. Modeling at a single temporal scale fails to capture both local fluctuations and global trends, thereby constraining model generalization across multi-stage dynamic processes. Consequently, constructing a representation learning framework capable of adaptively learning system metric distributions across multiple temporal scales has become a key direction for achieving high-precision prediction and robust modeling[3].

However, applying multi-scale representation learning to system metric prediction remains challenging. First, the multi-dimensional and nonlinear coupling of metric sequences leads to significant feature differences across temporal scales. Designing a model that preserves local sensitivity while extracting global steady-state representations is a core difficulty[4]. Second, the dynamic nature of system environments causes metric distributions to drift over time, rendering static feature extraction ineffective. Adaptive mechanisms are needed to adjust feature weights and interactions across different scales in real time. Third, hierarchical dependencies among metrics often follow implicit topological structures that cannot be fully captured through conventional sequential modeling. Therefore, structural constraints and contextual enhancement strategies must be integrated into feature representation to improve interpretability and stability in complex system behavior modeling[5].

In practical cloud and backend systems, prediction accuracy directly determines the level of proactive management and fault prevention. Accurate metric forecasting not only enables early detection of potential bottlenecks and anomalies but also provides quantitative support for resource allocation, task scheduling, and service orchestration. This facilitates an intelligent shift from passive response to proactive regulation. The introduction of multi-scale representation learning allows models to simultaneously capture fast fluctuations and slow variations across temporal scales, thereby enhancing sensitivity and adaptability to non-stationary processes. Particularly in multi-tenant or high-load environments, adaptive models can dynamically adjust structures based on real-time metric feedback, effectively mitigating performance degradation caused by model rigidity[6].

In summary, research on adaptive multi-scale representation learning for system metric prediction holds significant theoretical and practical value. Theoretically, it advances the representational framework of time-series modeling under non-stationary and multi-scale conditions, promoting developments in hierarchical feature modeling, adaptive weight allocation, and dynamic context integration. Practically, it offers efficient and reliable predictive support for cloud computing, intelligent operations, and automated resource management. By achieving adaptive optimization of feature representations through joint modeling across multiple temporal scales, cross-metric dimensions, and dynamic contexts, this research provides a solid technical foundation for the long-term stability and intelligent orchestration of complex systems[7].

2. Proposed Approach

This study introduces an adaptive multi-scale representation learning framework designed to enhance the modeling of non-stationary features, cross-temporal dependencies, and multi-dimensional interactions in system metric prediction tasks. The proposed method achieves efficient coordination between short-term local fluctuations and long-term global trends through multi-scale encoding, dynamic weight fusion, and

adaptive context updating mechanisms. The overall model consists of three core components: multi-scale feature decomposition, adaptive scale fusion, and temporal context reconstruction. Specifically, the multi-scale decomposition module extracts time-varying features using temporal convolutions and attention mechanisms at different scales; the fusion module automatically allocates scale weights based on the current system state; and the context reconstruction module generates continuously differentiable predictive representations via temporal dependency propagation. This framework maintains stable representational capability and adaptability in complex and dynamic system metric environments. The model architecture is shown in Figure 1.

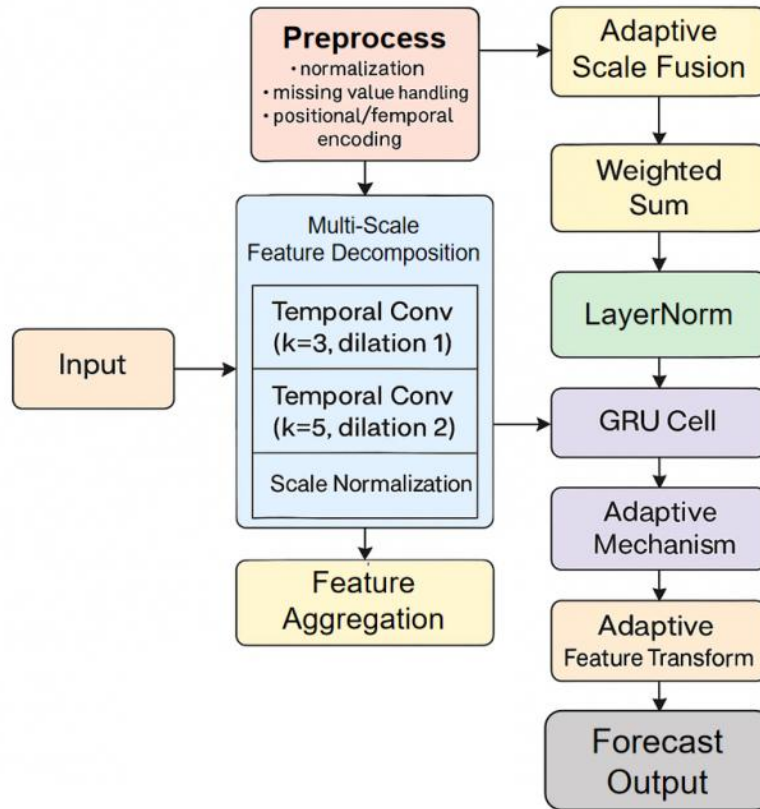


Figure 1. Overall Architecture of the Proposed Adaptive Multi-Scale Representation Learning Framework

First, to characterize the multi-scale evolution characteristics of system indicators, the original time series is recorded as $X = \{x_1, x_2, \dots, x_T\}$, and its representation at different scales can be expressed as:

$$H_s = \text{Conv}_{k_s}(X) + \text{Attn}_s(X) \quad (1)$$

Here, $\text{Conv}_{k_s}(\cdot)$ represents local temporal feature extraction at a convolution kernel size of k_s , and $\text{Attn}_s(\cdot)$ represents the scale-aware attention mechanism used to model long-term dependencies. Through multi-scale operations, feature sequences $\{H_1, H_2, \dots, H_s\}$ at different temporal granularities can be obtained.

Secondly, to achieve adaptive fusion of multi-scale features, a learnable weight distribution function α_s is introduced to dynamically adjust the importance of each scale based on the current system state and historical dependencies. The fusion representation is defined as:

$$Z = \sum_{s=1}^S \alpha_s \cdot \text{Norm}(H_s) \quad (2)$$

Here, $\text{Norm}(\cdot)$ represents the layer normalization operation to balance the magnitude differences of features at different scales. The weight α_s is obtained through soft normalization constraints to satisfy $\sum_{s=1}^S \alpha_s = 1$, ensuring the stability and interpretability of the fusion.

After multi-scale fusion, to enhance the model's adaptive modeling of temporal context, a temporal dependency propagation mechanism is introduced to model the state transition between time steps as follows:

$$C_t = \text{GRU}(Z_t, C_{t-1}) + \text{Gate}(Z_t) \quad (3)$$

Here, $\text{GRU}(\cdot)$ represents a gated recurrent unit used to capture temporal dependencies, and $\text{Gate}(\cdot)$ is a gated adjustment term used to control the flow of information between the current input features and the historical context. Through this mechanism, the model can adaptively adjust the dependency strength when faced with non-stationary time series, thereby achieving a stable time series representation.

After obtaining the context-enhanced feature representation C_t , to further improve the model's sensitivity to dynamic changes in features, this paper introduces a structured adaptive transformation function to map the features to a unified representation space:

$$F_t = \text{ReLU}(W_f C_t + b_f) + \lambda \cdot \text{Drop}(C_t) \quad (4)$$

Here, $\text{ReLU}(\cdot)$ is a nonlinear activation function, $\text{Drop}(\cdot)$ is a random dropout operation to prevent feature overfitting, and λ is an adjustable parameter to control the regularization strength. This structure enhances the expressiveness of the model while ensuring continuity and stability between features at different time steps.

Finally, to achieve the predicted output of system indicators, a regression mapping function is introduced to map the high-dimensional representation after multi-scale fusion to the indicator space:

$$\hat{y}_t = \text{READOUT}(F_t) = W_r F_t + b_r \quad (5)$$

Among them, $\text{READOUT}(\cdot)$ represents the linear readout layer, which is used to generate the predicted value \hat{y}_t of the indicator at the next moment. Through an end-to-end optimization process, the model can adaptively adjust the weight distribution between multi-scale features to achieve efficient mapping from the original sequence to the indicator prediction.

Overall, the proposed method performs structured modeling of system metrics based on multi-scale temporal features through adaptive fusion and dynamic context propagation. This mechanism not only captures complex dependencies across multiple temporal granularities but also adjusts feature representations according to changing operational states. It enables robust prediction of non-stationary system dynamics and

long-term trend modeling, providing theoretical and methodological support for system performance optimization and intelligent operations management.

3. Performance Evaluation

3.1 Dataset

This study employs the Cloud Workload Trace Dataset as the data foundation for method validation. The dataset collects multi-dimensional performance metrics from servers and containers in large-scale cloud computing environments. It includes key features such as CPU utilization, memory usage, disk I/O throughput, network traffic, and latency. The metric data are continuously sampled in time-series form, accurately reflecting the dynamic variations of system workloads. The dataset contains diverse workload types and service deployment configurations, providing rich structural information for system performance prediction and time-series modeling.

The design of this dataset highlights its multi-dimensional coupling and non-stationary characteristics, which align well with the requirements of multi-scale temporal modeling in this study. The indicators across different dimensions show strong temporal correlations and complex cross-metric dependencies, such as the interaction between CPU and I/O operations and the dynamic relationship between memory usage and network throughput. These characteristics make the dataset an ideal benchmark for validating the proposed adaptive multi-scale representation learning framework. It supports an in-depth analysis of the model's robustness and adaptability when dealing with multi-scale dynamic patterns and time-varying feature distributions.

In addition, the dataset provides high-frequency temporal sampling and detailed label annotations, enabling joint modeling of short-term fluctuations and long-term trends. Conducting system metric prediction tasks on this dataset allows a comprehensive evaluation of the proposed method's modeling capability and generalization performance under complex distribution shifts, multi-dimensional dependencies, and dynamic contextual conditions. This experimental setup offers solid theoretical and data support for the model's practical feasibility in real cloud environments.

3.2 Experimental Results

This paper first conducts a comparative experiment, and the experimental results are shown in Table 1.

Table1: Comparative experimental results

Method	MSE ↓	MAE ↓	MAPE (%) ↓	RMSE ↓
Autoformer[8]	0.0253	0.129	4.120	0.1591
EDFormer[9]	0.0227	0.117	3.760	0.1507
TimesNet[10]	0.0235	0.120	3.910	0.1533
Ours	0.0198	0.105	3.280	0.1407

From the overall trend, all four models show strong temporal fitting ability in system metric prediction tasks but still exhibit varying degrees of error fluctuation. The traditional Autoformer relies on a trend and residual-

based time series decomposition structure, which effectively captures long-term trend variations. However, its adaptability is limited under high-frequency disturbances and non-stationary distributions, resulting in relatively high MSE and RMSE values. This indicates a lag in modeling dynamic changes in complex system metrics. Its modeling of multi-dimensional dependencies remains rigid, making it difficult to handle heterogeneous feature interactions across different scales. As a result, its performance in MAE and MAPE is slightly inferior.

EDFormer integrates decomposition and embedded structural modeling, achieving significant improvements in predicting stable temporal segments compared to Autoformer. Both MSE and MAE decrease notably, indicating that the model performs more stably when capturing short-term perturbations and medium-term variations. However, EDFormer still relies on fixed-scale embedding windows. The lack of adaptive adjustment in its cross-scale interaction and weight allocation mechanism leads to redundant information and feature drift during scale alignment. Although EDFormer shows improvement over Autoformer in MAPE, it still fails to achieve full global consistency under complex system load fluctuations.

TimesNet introduces a two-dimensional periodic mapping structure for multi-scale modeling, enabling it to capture both intra-period and cross-period variations. It achieves more balanced performance in overall error convergence, with relatively stable RMSE values, suggesting its advantage in medium-complexity time series modeling tasks. Nevertheless, since its temporal encoding and feature mapping structures are static, the flexibility of feature fusion is limited when system metrics exhibit nonlinear transitions or abrupt changes at different temporal granularities. Consequently, the prediction residuals remain high in non-stationary segments.

In contrast, the proposed adaptive multi-scale representation learning method achieves the best results across all four evaluation metrics, demonstrating clear advantages in dynamic feature reconstruction and cross-scale fusion. The method dynamically aggregates multi-scale features through an adaptive weighting mechanism and enhances joint modeling of long-term trends and short-term fluctuations via a temporal dependency propagation module. This effectively reduces prediction bias and error propagation. The experimental results show that the proposed framework achieves higher accuracy and robustness in multi-dimensional, non-stationary, and highly dynamic system environments. These findings verify the modeling potential and application value of adaptive multi-scale representation learning in complex system scenarios.

This paper also evaluates the sensitivity of the multi-scale convolution kernel size and dilation rate, and the experimental results are shown in Figure 2.

From the overall trend, as the convolution kernel size and dilation rate change, the model exhibits clear nonlinear fluctuations in prediction performance under different scale configurations. Small-scale convolutions, such as k3-d1, perform stably in capturing local variations, but their limited receptive field fails to cover long-term trends. As a result, both MSE and RMSE remain relatively high, indicating insufficient modeling of global features. When the kernel size and dilation rate increase, the model's ability to capture temporal characteristics gradually improves, and error metrics continuously decrease. This trend demonstrates the effectiveness of multi-scale structures in modeling complex system metrics.

When the convolution kernel reaches a moderate scale, such as k7-d4, all four metrics achieve their optimal values. This indicates that at this scale, the modeling of local dynamics and global trends reaches an effective balance. The feature space forms a synergistic interaction between temporal dependency and scale decomposition, enabling the model to adaptively extract dominant patterns and suppress noise interference. This phenomenon validates the core idea of the adaptive multi-scale representation learning framework. Through multi-scale feature fusion and weight allocation, the model can flexibly capture system behavior at different temporal granularities, achieving efficient characterization of non-stationary signals.

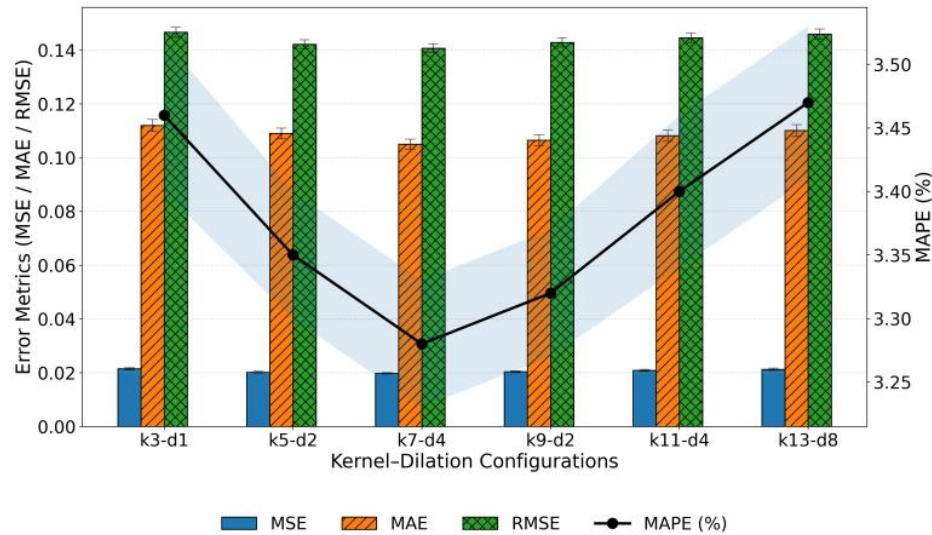


Figure 2. Hyperparameter sensitivity analysis of multi-scale convolution kernel size and dilation rate

When the dilation rate continues to increase, such as k11-d4 and k13-d8, model performance slightly declines. At this stage, the feature sampling interval becomes too large, leading to partial loss or dilution of local information. Short-term dependencies are harder to preserve, resulting in a slight rise in MAE and MAPE. This trend suggests that although a wider receptive field enhances global perception, the absence of appropriate hierarchical constraints and weight calibration weakens the model's responsiveness to fine-grained variations and disrupts consistency across scales.

Overall, the experimental results show that the proper matching of convolution kernel size and dilation rate is crucial for system metric prediction. The adaptive multi-scale representation learning method can dynamically adjust feature weights across receptive fields, balancing local sensitivity and global stability. The model achieves optimal performance at moderate convolution scales, demonstrating strong structural adaptability and robustness in multi-scale modeling of non-stationary system signals. These results highlight the method's superiority in system metric prediction tasks under complex and time-varying environments.

4. Conclusion

This study addresses the challenge of modeling multi-scale temporal dependencies in system metric prediction by proposing an adaptive multi-scale representation learning method. The approach enables high-precision prediction and robust modeling in multi-dimensional, non-stationary, and highly dynamic system environments. Through the organic integration of multi-scale convolutional feature decomposition, adaptive weight fusion, and temporal dependency propagation mechanisms, the method captures both short-term fluctuations and long-term trends across different temporal granularities. This effectively handles nonlinearity, noise interference, and multi-dimensional interactions in metric sequences. Experimental results show that the proposed model outperforms mainstream baselines on multiple key error metrics, demonstrating its adaptive capability and strong generalization performance in complex system environments with time-varying feature distributions.

From a theoretical perspective, the proposed adaptive multi-scale representation learning framework provides a new direction for time-series modeling. Unlike traditional single-scale or fixed-structure models, this

method achieves dynamic selection and fusion of multi-scale features, allowing the model to automatically adjust its feature representation space according to system state changes. This idea not only provides a scalable theoretical foundation for system metric prediction but also introduces a new research paradigm for addressing issues such as scale inconsistency and structural drift in multi-dimensional time series. The interpretability and transferability of its multi-layer architecture offer valuable insights for future model design and optimization across different application domains.

From an application perspective, the proposed method has significant potential in cloud computing, distributed systems, microservice architectures, and intelligent operations management. By accurately modeling and predicting system performance metrics, operation and maintenance systems can enable proactive resource scheduling, fault warning, and anomaly detection, thereby improving service quality and system reliability. Moreover, the method is broadly applicable to industrial process monitoring, network traffic analysis, and data center energy consumption forecasting. Its strong robustness and low parameter dependency allow the model to maintain stable predictive performance across varying scales and noise levels, providing a solid technical foundation for intelligent management and automated decision-making in complex systems.

Future research can be expanded in three directions. First, the proposed framework can be combined with graph neural networks or attention diffusion mechanisms to enhance the model's ability to capture cross-dimensional dependencies and topological structures. Second, considering the need for real-time performance and computational efficiency, lightweight multi-scale modeling strategies can be developed to enable low-latency, high-throughput online prediction and dynamic updating. Finally, in terms of cross-domain transfer and self-supervised learning, future work may introduce distribution adaptation and domain alignment mechanisms, allowing the model to achieve knowledge sharing and generalization across different systems and tasks. Overall, the proposed adaptive multi-scale representation learning method provides a new theoretical foundation and practical pathway for intelligent prediction in complex systems, offering significant academic and application value for advancing adaptive system analysis and intelligent operations.

References

- [1] Zhang K, Sun S, Fan Z, et al. Conv-like Scale-Fusion Time Series Transformer: A Multi-Scale Representation for Variable-Length Long Time Series[J]. arXiv preprint arXiv:2509.17845, 2025.
- [2] Naghashi V, Boukadoum M, Diallo A B. A multiscale model for multivariate time series forecasting[J]. Scientific Reports, 2025, 15(1): 1565.
- [3] Hou S, Sun S, Yin T, et al. AMDCnet: attention-gate-based multi-scale decomposition and collaboration network for long-term time series forecasting[J]. Frontiers in Artificial Intelligence, 2025, 8: 1607232.
- [4] Liu W, Yu X, Zhao Q, et al. Time Series Forecasting Fusion Network Model Based on Prophet and Improved LSTM[J]. Computers, Materials & Continua, 2023, 74(2).
- [5] Yue Z, Wang Y, Duan J, et al. Ts2vec: Towards universal representation of time series[C]//Proceedings of the AAAI conference on artificial intelligence. 2022, 36(8): 8980-8987.
- [6] Li Y, Chen Z, Zha D, et al. Learning disentangled representations for time series[J]. arXiv preprint arXiv:2105.08179, 2021.
- [7] Wang T, Liu Z, Zhang T, et al. Adaptive feature fusion for time series classification[J]. Knowledge-Based Systems, 2022, 243: 108459.
- [8] Wu H, Xu J, Wang J, et al. Autoformer: Decomposition transformers with auto-correlation for long-term series forecasting[J]. Advances in neural information processing systems, 2021, 34: 22419-22430.
- [9] Chakraborty S, Delibasoglu I, Heintz F. EDformer: Embedded Decomposition Transformer for Interpretable Multivariate Time Series Predictions[J]. arXiv preprint arXiv:2412.12227, 2024.

[10]Wu H, Hu T, Liu Y, et al. Timesnet: Temporal 2d-variation modeling for general time series analysis[J]. arXiv preprint arXiv:2210.02186, 2022.