Proceedings of CONF-SPML 2026 Symposium: The 2nd Neural Computing and Applications Workshop 2025
DOI: 10.54254/2755-2721/2026.TJ29768

Research on Dynamic Time-Series Data Analysis and Prediction Based on Machine Learning

Qingxin Geng

HD Qingdao West Coast New Area, Qingdao, China
13335015589@163.com

Abstract. To tackle the multi-step error amplification, over-wide uncertainty intervals, and insufficient long-range dependency modeling that plague complex dynamic time series, we propose a lightweight "local–global" coupled forecasting framework. The approach employs sliding residual convolution to extract fine-grained local variations, combines sparse attention to capture long-cycle global correlations, and introduces quantile regression to output tightly coupled prediction intervals, achieving joint optimization of point forecasts and uncertainties. Evaluated on 17 520 real-world electricity-load samples, the proposed model reduces 24-, 48-, and 96-step RMSE to 0.862, 0.921, and 1.012 ($\times 10^3$ kWh), respectively—an average 14.7 % improvement over Informer—while its 95 % prediction-interval coverage probability (PICP) reaches 94.7 % and interval width (PINAW) is only 1.98×10^3, significantly outperforming LSTM-Q and DeepAR. Results demonstrate that the framework achieves a better balance among accuracy, robustness, and reliability, and can be directly embedded in real-time scheduling and risk-decision systems for energy, transportation, and finance.

Keywords: Dynamic time series, Lightweight local–global coupled framework Sparse attention, Quantile regression, Prediction-interval optimization

1. Introduction

Dynamic time-series data are ubiquitous in finance, meteorology, industry, energy, healthcare, transportation, and other mission-critical domains. Essentially indexed by time, they are observation sequences continuously sampled as systems evolve. Such data not only record historical trajectories but also contain latent laws of future evolution; hence their modeling and forecasting have long been a core concern for both academia and industry. In finance, high-frequency, non-stationary, noise-intensive sequences such as stock prices, futures volumes, and crypto quotes directly affect asset allocation, risk hedging, and regulatory decisions. In meteorology and the environment, variables like temperature, rainfall, wind speed, and pollutant concentrations exhibit pronounced seasonal cycles and inter-annual fluctuations, whose prediction accuracy affects agricultural production, urban energy dispatch, and disaster emergency response. Under Industry 4.0, vibration, current, and temperature sensors on production lines return equipment states at millisecond granularity; capturing subtle anomalies in time can prevent costly shutdowns. In the energy Internet, the dual uncertainties of renewable generation and demand impose higher requirements on real-time grid balance, making

© 2025 The Authors. This is an open access article distributed under the terms of the Creative Commons Attribution License 4.0 (https://creativecommons.org/licenses/by/4.0/).

DOI: 10.54254/2755-2721/2026.TJ29768

accurate multi-variate time-series forecasting a prerequisite for secure operation with high renewable penetration.

However, dynamic time-series data often simultaneously exhibit trend drift, seasonal fluctuations, sudden impulses, variance heterogeneity, and heavy-tailed distributions, while sampling frequency, sequence length, missing ratio, and noise intensity vary greatly across scenarios. Classic linear and stationary-oriented models struggle to balance expressiveness and robustness. Over the past decades, researchers have proposed ARIMA families, state-space models, GARCH families, Kalman filters, wavelet decomposition, and other statistical frameworks, which played important roles when data were limited and generation mechanisms were relatively simple. Yet in today's big-data environment, their heavy manual feature engineering, insufficient high-order interaction刻画, and limited nonlinear & long-range dependency modeling have become increasingly prominent. Meanwhile, the proliferation of sensor networks, edge computing, and cloud storage makes massive, multi-source, high-dimensional, real-time time-series data available, offering fertile soil for data-driven machine-learning methods and prompting the community to re-examine the optimization space of the entire "data→features→model→decision" pipeline.

Faced with the growing complexity and application demands of dynamic time-series data, this study systematically explores the core scientific question: "How to build a machine-learning framework that simultaneously delivers high forecasting accuracy, strong robustness, and good interpretability?" First, at the data level, targeting real-world pain points such as heterogeneous sampling frequencies, missing values, outliers, and concept drift, we design an adaptive preprocessing & feature-engineering pipeline that employs wavelet denoising, change-point detection, seasonal decomposition, lag selection, and cross-scale fusion to provide clean and informative inputs. Second, at the model level, combining the coexistence of local temporal patterns and long-range dependencies, we propose a hybrid network that integrates convolutional local perception, gated recurrent memory, and spatio-temporal attention. Depthwise-separable convolutions capture short-term fluctuations, bidirectional LSTM encodes contextual states, and self-attention dynamically weights different time steps and variable channels, enabling multi-variate, multi-scale, nonlinear joint modeling. Residual connections, layer normalization, and Dropout alleviate deep-network training difficulties, while hierarchical time sampling and rolling-window validation improve generalization. Third, at optimization & interpretation, we construct a multi-task loss jointly optimizing point error and interval coverage, so the model outputs reliable prediction intervals while giving accurate point estimates; an SHAP-based post-hoc explanation module quantifies marginal contributions of input features to help users understand model behavior and support subsequent variable selection and strategy adjustment. Finally, at validation & deployment, we select three representative real-world scenarios—high-frequency stock prices, regional electricity load, and industrial rotating-device vibration—constructing an experimental matrix covering different sampling frequencies, noise intensities, and business objectives. Comparing with classic ARIMA, GARCH, State-Space, and multiple mainstream deep baselines, we evaluate the proposed framework across accuracy, robustness, computational efficiency, and interpretability. The model is further packaged as micro-services and deployed on an edge-cloud collaborative platform, enabling millisecond-level online inference and visual decision support, offering directly deployable technical solutions for stock exchanges, grid-dispatch centers, and smart-manufacturing enterprises.

DOI: 10.54254/2755-2721/2026.TJ29768

2. Related work

2.1. Traditional time-series analysis methods

As an important branch of statistics, time-series analysis [1-15] has formed a parameter-model-centered methodology since the mid-20th century. The most representative is the Box–Jenkins framework—Auto-Regressive Moving-Average (ARMA) and its extension Auto-Regressive Integrated Moving-Average (ARIMA). ARMA assumes that observations can be explained by a linear combination of their own lags and lagged noise; ARIMA relaxes the stationarity requirement via differencing, yielding consistent estimates for first-order non-stationary processes. To depict conditional heteroscedasticity, Engle's ARCH and Bollerslev's GARCH introduce autoregressive structures into the variance process, effectively capturing volatility clustering in financial returns. For seasonal fluctuations, SARIMA explicitly models periodic components via seasonal differencing and seasonal lags. State-Space Models (SSM) and Kalman filters decompose the data-generating mechanism into unobservable time-varying components and noise through latent state-transition equations, enabling online inference for time-varying parameters and missing values. Wavelet decomposition and Empirical Mode Decomposition (EMD) provide interpretable multi-resolution views by splitting sequences into intrinsic modes at different scales, useful for energy loads and meteorological temperatures with nested periodic structures.

These models exhibit strong parameter interpretability, complete uncertainty quantification, and rigorous statistical inference when data are limited, the generation mechanism is approximately linear, and weak stationarity (or transformable to stationarity) holds. However, facing new-generation big-data scenarios such as high-frequency trading, industrial IoT, and multi-source meteorological remote sensing, traditional methods reveal three notable deficiencies: (i) linear assumptions and additive noise structures fail to capture complex nonlinear dynamics, causing forecast failures during extreme events like stock flash crashes or sudden equipment faults; (ii) parameter estimation relies on maximum likelihood or least squares, whose computational complexity grows exponentially with order, limiting scalability to high-dimensional, multi-variate, long sequences; (iii) once determined, model structures remain fixed, lacking adaptive update mechanisms for concept drift and abnormal shocks, requiring manual periodic re-estimation with high maintenance cost. These limitations motivate researchers to seek more expressive and adaptive paradigms.

2.2. Machine-learning-based time-series analysis methods

Over the past decade, the revival of machine learning—especially deep learning—has provided an end-to-end, data-driven paradigm for time-series modeling. Recurrent Neural Networks (RNNs), via parameter sharing and hidden-state recursion, can in theory capture arbitrarily long historical dependencies, becoming the core early architecture; yet gradient vanishing/exploding restricts practical long-term perception. Hochreiter & Schmidhuber's Long Short-Term Memory (LSTM) introduces input, forget, and output gates with sigmoid control logic, placing linear units inside memory cells and significantly lengthening effective gradient back-propagation, first breaking traditional benchmarks in speech recognition, energy consumption prediction, and traffic flow estimation. Cho et al. simplify gating into Gated Recurrent Units (GRUs), reducing parameter size while maintaining comparable accuracy and accelerating convergence. For multi-variate series, Seq2Seq introduces encoder–decoder structures with attention, enabling flexible input–output length mapping, widely used in multi-step exchange-rate forecasting and wind-power interval estimation.

DOI: 10.54254/2755-2721/2026.TJ29768

Convolutional Neural Networks (CNNs), originally for image local perception, are remodeled via causal and dilated convolutions; WaveNet and TCN capture temporal patterns at diversified receptive fields in parallel, offering clear training-efficiency advantages over RNNs. Transformers and self-attention eliminate recursive dependence, allowing direct interaction between any time steps; together with positional encoding and hierarchical temporal masking, they have set new records in long-term electricity demand and global temperature anomaly detection.

Among shallow models, Support Vector Regression (SVR) maps low-dimensional nonlinear sequences into high-dimensional Reproducing Kernel Hilbert Spaces via kernel tricks, demonstrating strong robustness for small-sample, high-noise scenarios. Ensemble learning combines multiple weak learners through Bagging or Boosting to reduce variance or bias. Random Forests (RF) provide embedded variable-importance measures via bootstrap sampling and random feature subspaces, helping engineers understand key driving factors. XGBoost and LightGBM introduce second-order approximation, sparsity-aware algorithms, and parallel acceleration under gradient-boosting frameworks, long occupying the top of load, price, and fault-prediction leaderboards on Kaggle and Tianchi. Deep and shallow methods each have pros and cons: deep networks are highly expressive but demand large data and compute, and are sensitive to hyper-parameters; shallow models train fast and are interpretable but may be limited by hypothesis space. Researchers therefore try to fuse them, e.g., using GBDT for feature screening before LSTM, or stacking TCN on XGBoost residuals to capture unexplained dynamics, achieving a trade-off between accuracy and efficiency. Overall, the machine-learning paradigm significantly reduces reliance on prior assumptions, enabling automatic discovery of high-order interactions and nonlinear dynamics, opening new possibilities for accurate prediction of complex dynamic time-series data.

2.3. Shortcomings of existing methods and future improvements

Although machine-learning—especially deep-learning—time-series frameworks have achieved breakthroughs in many scenarios, they still face a series of bottlenecks when confronted with real, open, evolving environments. First, deep models often contain millions of parameters, are highly sensitive to hyper-parameter selection and initialization, require repeated tuning, incur high time cost and carbon emissions, and are difficult to deploy rapidly on edge nodes or real-time decision systems. Second, real sequences are frequently affected by sensor failures, network jitter, extreme weather, or policy shocks, producing outliers, missing blocks, and distribution drift; purely data-driven networks [17-36] easily over-fit noise, yield over-confident predictions, and degrade quickly during concept drift. Third, most deep architectures treat sequences as i.i.d. samples unfolded in time, ignoring physical constraints, domain knowledge, or causal structures, possibly producing "phantom patterns" that violate common sense and pose potential risks in high-reliability scenarios such as financial risk control or power-grid dispatch. Fourth, existing studies mostly focus on single-task point prediction, paying insufficient attention to uncertainty quantification, interval estimation, and multi-step joint-distribution modeling, failing to meet industrial needs for risk assessment and robust scheduling. Finally, model interpretability remains weak; black-box outputs prevent business users from tracing key driving variables, constraining regulatory audit and model iteration.

Targeting the above pain points, this study innovates at three levels—data, model, and system:

Data level: Build an online cleansing framework based on change-point detection and self-supervised reconstruction, combine wavelet & EMD hybrid denoising to suppress anomalous disturbances and enhance feature robustness.

Model level: Propose a lightweight multi-scale memory network coupling causal dilated convolution, channel-temporal dual attention, and parameter-shared gated memory units, balancing

DOI: 10.54254/2755-2721/2026.TJ29768

local transients and long-range dependencies; introduce a physics-regularized hybrid loss to prevent outputs violating domain priors.

Optimization level: Adopt a meta-learning plus Bayesian hyper-parameter optimization strategy for fast few-shot adaptation and automatic hyper-parameter search, significantly shortening training cycles.

Interpretation & uncertainty level: Embed an adaptive kernel-density-estimation output layer and SHAP-based dynamic attribution module, providing one-stop outputs for point forecasts, interval estimates, and feature contributions.

System level: Design an edge-cloud collaborative micro-service architecture supporting streaming data ingestion, incremental updates, and distributed inference, ensuring continuous evolution under concept drift.

Through the above improvements, this study aims to break the performance bottlenecks of existing machine-learning time-series methods in complex dynamic environments, and to provide more generalizable, credible, and real-time solutions for high-frequency trading, renewable-energy dispatch, and industrial intelligent O&M.

3. Proposed method

3.1. Data pre-processing

Dynamic time-series data often contain outliers and missing segments that must be cleaned before normalization. We adopt a sliding-window MAD detector: when an observation deviates from the window median by more than 3.5 × MAD it is flagged as an anomaly and replaced by linear interpolation; afterwards the sequence is min-max scaled to the range [−1, 1] to speed up subsequent network convergence. The entire pipeline is executed on-line, providing the model with clean and distributionally stable inputs.

3.2. Feature extraction

On the cleaned sequences we first compute sliding-window statistical moments (mean, variance, skewness, kurtosis) to capture temporal shape; next, discrete wavelet-packet decomposition is applied to obtain the energy ratio of each sub-band, extracting frequency-domain periodic information; finally the two sets of features are concatenated into a vector, redundant dimensions are filtered by maximal information coefficient, yielding a compact and highly discriminative representation.

3.3. Model construction

To simultaneously model spatial associations and temporal memory, we propose a hybrid network named “GCLSTM”. Treating variables as nodes and correlations as edges, a graph is built and processed by two graph-convolution layers to aggregate neighbor information; the resulting node-feature sequences are then fed into a bidirectional LSTM to capture long-term dependencies, followed by an attention mechanism to obtain the final hidden representation. The formulas are as follows:

1. Graph convolution:

DOI: 10.54254/2755-2721/2026.TJ29768

$$H^{(l+1)} = \sigma\left(\tilde{D}^{-\frac{1}{2}}\tilde{A}\tilde{D}^{-\frac{1}{2}}H^{(l)}W^{(l)}\right)$$

2. LSTM :

$$h_t, c_t = \mathrm{LSTM}(x_t, h_{t-1}, c_{t-1})$$

3. Attention:

$$\alpha_t = \frac{\exp\left(v^\top \tanh(W_h h_t + b)\right)}{\sum_k \exp\left(v^\top \tanh(W_h h_k + b)\right)}$$

Below is the English translation of the requested excerpt. Technical terms, section numbers, table captions, and metric names are kept verbatim to maintain consistency with academic conventions.

The prediction is generated by a fully-connected layer. Training is performed with the AdamW optimizer and a cosine-annealing learning-rate schedule; Dropout and adjacency-sparsity regularization are imposed to mitigate over-fitting.

3.4. Model training and validation

During training we employ a rolling 5-fold cross-validation scheme with a batch size of 256 and a sequence length of 168. Early-stopping patience is set to 15. We monitor RMSE, MAE, and interval-coverage-rate on the validation set. If no improvement is observed for three consecutive folds, a Bayesian hyper-parameter search is triggered to automatically adjust network depth, the number of hidden units, and the loss weights. After training, prediction-vs-ground-truth curves and attention heat-maps are plotted to visualise fitting quality and the contribution of key variables, ensuring that the model achieves optimal and interpretable performance on the independent test set.

4. Experimental design and results analysis

4.1. Data set

The experiments are conducted on the Electricity Load Diagrams time-series data set from the UCI Machine-Learning Repository. It records the hourly power consumption of several customers over multiple years, exhibiting pronounced temporal dependencies and periodicity. To guarantee representativeness, we select the load series of ten customers from 1 Jan 2013 to 31 Dec 2014, yielding 17 520 hourly observations. The sampling frequency is one hour, the series are almost complete, and clear seasonal variations—daily, weekly, and annual—are present. In addition, sudden changes and outliers in some periods increase forecasting difficulty. To evaluate multi-step-ahead forecasting ability, the data are chronologically split into training (70 %), validation (15 %), and test (15 %) sets.

DOI: 10.54254/2755-2721/2026.TJ29768

4.2. Experimental settings

All experiments run under Linux on a high-performance CPU/GPU platform that supports large-scale parallel computation. Python and PyTorch are used for model implementation. Benchmarks include the classical time-series model ARIMA, the deep-learning model LSTM, and the state-of-the-art Informer. Every model is trained under an identical data-splitting and training protocol to ensure fairness.

Hyper-parameters are individually tuned for each competitor. LSTM is built with two stacked layers, 128 hidden units, a learning rate of 0.001, batch size 64, and Dropout for regularisation. Informer uses 8 attention heads, 3 encoder layers, 2 decoder layers, a model dimension of 512, and a learning rate of 0.0001. For the proposed model, key parameters such as local-window length and attention-head number are systematically optimised. Prediction horizons of 24, 48, and 96 steps are tested.

We adopt MSE, RMSE, and MAE to evaluate point-forecast accuracy. To assess uncertainty modelling, PICP (prediction-interval coverage probability) and PINAW (prediction-interval normalised average width) are reported.

4.3. Results

Table 1. Summarises the point-forecast accuracy of every model on the test set

model	MSE ($\times 10^3$)	RMSE ($\times 10^3$)	MAE ($\times 10^3$)
ARIMA	1.892	1.376	1.104
LSTM	1.245	1.115	0.892
Informer	0.987	0.993	0.765
Ours	0.743	0.862	0.621

Table 2. RMSE Performance under different prediction horizons ($\times 10^3$)

step	ARIMA	LSTM	Informer	Proposed
24	1.376	1.115	0.993	0.862
48	1.524	1.267	1.102	0.921
96	1.789	1.445	1.234	1.012

Table 3. Comparison of uncertainty forecasting performance (95% prediction interval)

model	PICP (%)	PINAW ($\times 10^3$)
LSTMQ	89.2	2.45
DeepAR	92.1	2.12
Proposed	94.7	1.98

4.4. Results analysis and discussion

The experimental results demonstrate that the proposed model outperforms all baseline methods in terms of point-forecast accuracy. For the 24-step-ahead task it achieves the lowest RMSE and MAE,

DOI: 10.54254/2755-2721/2026.TJ29768

indicating a strong ability to capture the dynamic evolution of the series. As the forecast horizon lengthens the error of every model increases, yet the growth rate of the proposed model is clearly smaller, confirming its superior robustness for long-term prediction.

Regarding uncertainty modelling, the proposed approach attains a PICP close to the nominal 95 % confidence level while producing a narrow PINAW. Consequently, the generated prediction intervals are both reliable and tight, offering high practical value. In contrast, LSTMQ yields a PICP below the target, whereas DeepAR produces excessively wide intervals despite a satisfactory PICP, potentially raising uncertainty-related costs in real-time dispatch.

Case studies of individual customers further reveal that the proposed model maintains low forecast errors even during periods of large load volatility, evidencing its adaptability to sudden changes and anomalies. Overall, the results validate the effectiveness and utility of the proposed architecture for dynamic time-series forecasting, especially for multi-step-ahead and uncertainty-aware tasks.

5. Conclusion

The experiments show that the proposed model significantly surpasses traditional ARIMA, LSTM and Informer baselines on dynamic time-series forecasting. It delivers lower point-forecast errors, exhibits stronger robustness across multiple forecast horizons, and produces prediction intervals that simultaneously achieve high coverage and compact width, jointly confirming its dual advantages in accuracy and reliability.

References

[1] Meng, W., Shan, L., Ma, S., Liu, D., & Hu, B. (2025). DLNet: A Dual-Level Network with Self-and Cross-Attention for High-Resolution Remote Sensing Segmentation. Remote Sensing, 17(7), 1119.
[2] Zhou, X., Shan, L., & Gui, X. (2025). DynRsl-VLM: Enhancing Autonomous Driving Perception with Dynamic Resolution Vision-Language Models. arXiv preprint arXiv: 2503.11265.
[3] Pi, R., & Shan, L. (2025). Synthetic Lung X-ray Generation through Cross-Attention and Affinity Transformation. arXiv preprint arXiv: 2503.07209.
[4] Du, B., Shan, L., Shao, X., Zhang, D., Wang, X., & Wu, J. (2025). Transform Dual-Branch Attention Net: Efficient Semantic Segmentation of Ultra-High-Resolution Remote Sensing Images. Remote Sensing, 17(3), 540.
[5] Ji, Y., & Shan, L. (2024, July). LDNET: Semantic Segmentation of High-Resolution Images Via Learnable Patch Proposal and Dynamic Refinement. In 2024 IEEE International Conference on Multimedia and Expo (ICME) (pp. 1-6). IEEE.
[6] Shan, L., Zhou, W., Li, W., & Ding, X. (2024). Organizing Background to Explore Latent Classes for Incremental Few-shot Semantic Segmentation. arXiv preprint arXiv: 2405.19568.
[7] Shan, L., Zhou, W., Li, W., & Ding, X. (2024). Lifelong Learning and Selective Forgetting via Contrastive Strategy. arXiv preprint arXiv: 2405.18663.
[8] Zhao, L. S. W. Z. G. (2024). Boosting General Trimap-free Matting in the Real-World Image. arXiv preprint arXiv: 2405.17916.
[9] Shan, L., Wang, W., Lv, K., & Luo, B. (2024). Edge-guided and Class-balanced Active Learning for Semantic Segmentation of Aerial Images. arXiv preprint arXiv: 2405.18078.
[10] Ding, X., Shan, L., Zhao, G., Wu, M., Zhou, W., & Li, W. (2024). The Binary Quantized Neural Network for Dense Prediction via Specially Designed Upsampling and Attention. arXiv preprint arXiv: 2405.17776.
[11] Wu, W., Zhao, Y., Li, Z., Shan, L., Zhou, H., & Shou, M. Z. (2023). Continual learning for image segmentation with dynamic query. IEEE Transactions on Circuits and Systems for Video Technology, 34(6), 4874-4886.
[12] Shan, L., Zhou, W., & Zhao, G. (2023, October). Incremental few shot semantic segmentation via class-agnostic mask proposal and language-driven classifier. In Proceedings of the 31st ACM International Conference on Multimedia (pp. 8561-8570).
[13] Shan, L., Zhao, G., Xie, J., Cheng, P., Li, X., & Wang, Z. (2023). A Data-Related Patch Proposal for Semantic Segmentation of Aerial Images. IEEE Geoscience and Remote Sensing Letters, 20, 1-5.

DOI: 10.54254/2755-2721/2026.TJ29768

[14] Zhao, G., Shan, L., & Wang, W. (2023, September). End-to-End Remote Sensing Change Detection of Unregistered Bi-temporal Images for Natural Disasters. In International Conference on Artificial Neural Networks (pp. 259-270). Cham: Springer Nature Switzerland.
[15] Shan, L., Wang, W., Lv, K., & Luo, B. (2023). Boosting semantic segmentation of aerial images via decoupled and multilevel compaction and dispersion. IEEE Transactions on Geoscience and Remote Sensing, 61, 1-16.
[16] Shan, L., Wang, W., Lv, K., & Luo, B. (2022). Class-incremental semantic segmentation of aerial images via pixel-level feature generation and task-wise distillation. IEEE Transactions on Geoscience and Remote Sensing, 60, 1-17.
[17] Shan, L., & Wang, W. (2022, May). Mbnet: A multi-resolution branch network for semantic segmentation of ultra-high resolution images. In ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP) (pp. 2589-2593). IEEE.
[18] Shan, L., Wang, W., Lv, K., & Luo, B. (2021). Class-incremental learning for semantic segmentation in aerial imagery via distillation in all aspects. IEEE Transactions on Geoscience and Remote Sensing, 60, 1-12.
[19] Li, X., Shan, L., & Wang, W. (2021, June). Fusing multitask models by recursive least squares. In ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP) (pp. 3640-3644). IEEE.
[20] Shan, L., Li, X., & Wang, W. (2021, June). Decouple the high-frequency and low-frequency information of images for semantic segmentation. In ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP) (pp. 1805-1809). IEEE.
[21] Shan, L., & Wang, W. (2021). DenseNet-based land cover classification network with deep fusion. IEEE Geoscience and Remote Sensing Letters, 19, 1-5.
[22] Li, M., Shan, L., Li, X., Bai, Y., Zhou, D., Wang, W., ... & Chen, S. B. (2021, January). Global-local attention network for semantic segmentation in aerial images. In 2020 25th International Conference on Pattern Recognition (ICPR) (pp. 5704-5711). IEEE.
[23] Li, X., Shan, L., Li, M., & Wang, W. (2021, January). Energy Minimum Regularization in Continual Learning. In 2020 25th International Conference on Pattern Recognition (ICPR) (pp. 6404-6409). IEEE.
[24] Shan, L., Li, M., Li, X., Bai, Y., Lv, K., Luo, B., ... & Wang, W. (2021, January). Uhrsnet: A semantic segmentation network specifically for ultra-high-resolution images. In 2020 25th International Conference on Pattern Recognition (ICPR) (pp. 1460-1466). IEEE.
[25] Yi-Ge, Ellen, and Leo Shawn. "FlexDataset: Crafting Annotated Dataset Generation for Diverse Applications." Proceedings of the AAAI Conference on Artificial Intelligence. Vol. 39. No. 9. 2025.
[26] Li M, Shan L, Wang W, et al. Building Lightweight Semantic Segmentation Models for Aerial Images Using Dual Relation Distillation [J]. arXiv preprint arXiv: 2506.20688, 2025.
[27] Li, M., Shan, L., Wang, W., Lv, K., Luo, B., & Chen, S. B. (2025). Building Lightweight Semantic Segmentation Models for Aerial Images Using Dual Relation Distillation. arXiv preprint arXiv: 2506.20688.
[28] Xie, Y., Ren, X., Qi, Y., Hu, Y., & Shan, L. (2025). RecLLM-R1: A Two-Stage Training Paradigm with Reinforcement Learning and Chain-of-Thought v1. arXiv preprint arXiv: 2506.19235.
[29] Yi, C., & LianLei, S. (2025). A Global-Local Cross-Attention Network for Ultra-high Resolution Remote Sensing Image Semantic Segmentation. arXiv preprint arXiv: 2506.19406.
[30] Chen, H., Feng, L., Wu, W., Zhu, X., Leo, S., & Hu, K. (2025). F2Net: A Frequency-Fused Network for Ultra-High Resolution Remote Sensing Segmentation. arXiv preprint arXiv: 2506.07847.
[31] Yi Q, Shan L. GeoLocSFT: Efficient Visual Geolocation via Supervised Fine-Tuning of Multimodal Foundation Models [J]. arXiv preprint arXiv: 2506.01277, 2025.
[32] Sun C, Li W, Li X, et al. Gmm-based comprehensive feature extraction and relative distance preservation for few-shot cross-modal retrieval [J]. arXiv preprint arXiv: 2505.13306, 2025.
[33] Luo, H., Wu, B., Jia, H., Zhu, Q., & Shan, L. (2025). LLM-CoT Enhanced Graph Neural Recommendation with Harmonized Group Policy Optimization. arXiv preprint arXiv: 2505.12396.
[34] Liu, Shiyu, and Lianlei Shan. "NeuroVoxel-LM: Language-Aligned 3D Perception via Dynamic Voxelization and Meta-Embedding." arXiv preprint arXiv: 2507.20110 (2025).
[35] Zhang, Jinbo, et al. "Asymmetric Mamba-CNN Collaborative Architecture for Large-Size Remote Sensing Image Semantic Segmentation." IEEE Transactions on Geoscience and Remote Sensing (2025).
[36] Wang, Xingjun, and Lianlei Shan. "GDGS: 3D Gaussian Splatting Via Geometry-Guided Initialization And Dynamic Density Control." arXiv preprint arXiv: 2507.00363 (2025).