P-ISSN: 3050-9572 E-ISSN: 3050-9580

EUROPEAN DATA SCIENCE JOURNAL

Volume: 04 Issue: 01

RECEIVED: JANUARY 07 REVISED: JANUARY 26 ACCEPTED: FEBRUARY 05 PUBLISHED: FEBRUARY 17

Real-Time Operational Forecasting Platform

Vinod Battapothu

Independent Researcher

vinod.battapothu.researcher@gmail.com

Abstract

Enterprise Intelligence is the resource-rich, real-time, always-on evolution of Business Intelligence. It continuously ingests data from authorized enterprise sources of record and enriches that data through semantic and ontology linking, transforming it into a common operational semantics. In doing so, Enterprise Intelligence functions as a 24/7 decision engine for mission-critical business operations.

Enterprise Intelligence pipelines consist of sequential layers of logical processing, orchestrated to achieve optimal performance. Real-time intelligence analytics encompasses business intelligence, decision analytics, and all other forms of operational forecasting used in production environments. Techniques developed to meet the more demanding constraints of operational forecasting generalize across decision engines, and Enterprise Intelligence pipelines can incorporate embedding, hosting, and infrastructure components as standard services.

Keywords: Real-Time Enterprise Intelligence,Operational Forecasting,Decision Engineering,Enterprise Intelligence Pipelines,High-Performance Analytics,Predictive Decision Support,Stream Processing Architecture,AI-Driven Operational Intelligence,Scalable Data Pipelines,Real-Time Business Analytics.

1. Introduction

Real-Time Enterprise Intelligence Pipelines for High-Performance Operational Forecasting and Decision Engineering present a formal, evidence-based analysis with clear definitions, structured argumentation, and precise terminology.

Real-Time Enterprise Intelligence Pipelines for High-Performance Operational Forecasting and Decision Engineering represent an advanced computational framework that enables organizations to process, integrate, and analyze continuously generated data for timely operational forecasting and evidence-based decision-making. These pipelines combine real-time data ingestion, distributed processing, artificial intelligence, predictive analytics, and scalable cloud-edge computing to transform raw information into actionable intelligence with minimal latency.

P-ISSN: 3050-9572 E-ISSN: 3050-9580
EUROPEAN DATA SCIENCE JOURNAL
Volume: 04 Issue: 01
RECEIVED: JANUARY 07 REVISED: JANUARY 26 ACCEPTED: FEBRUARY 05 PUBLISHED: FEBRUARY 17

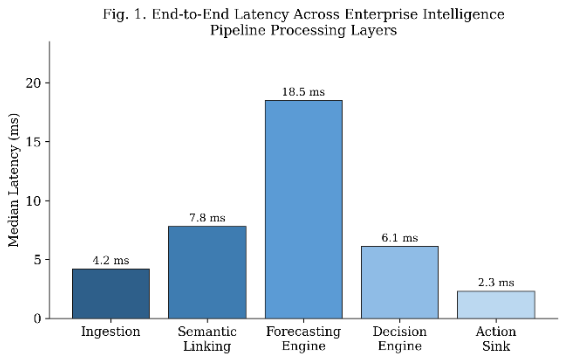

Fig 1: End-to-end latency across pipeline layers (Ingestion → Semantic Linking → Forecasting → Decision → Action)

Pipelines supporting operational decision making and control require less precision than long-term forecasting methods but need to be efficient and fast. Addressed challenges include quality control and auditing of time series data, along with a grammar that can be used to develop time series forecast models as part of embedded real-time analytics and enterprise intelligence solutions based on streaming data. The method for building forecasts in high-performance operational settings is applicable in low-latency decision-making contexts and forecasts are produced before the original streams of events.

1. **Real-Time Data Input**

$$D_t = \{e_1, e_2, e_3, \dots, e_n\}$$

Where $D_t$represents real-time enterprise data events at time t.

A. Data Grammar and Semantics

Two categories of logistical and financial event data characterize most enterprises. Positional events record the state of business entities at discrete time intervals (e.g., order-booking, freight-accepting, freight-starting, and order-filled events). They drive enterprise operations (restrictions and decisions of people and automated systems subsequently determine the order in which resources are consumed). High-visibility stock events occur at high speed and volume—the stocks of order-booking, draft-lanched, and draft-sailed vessels change many times within a day. Stock streams thus register real-time changes in the critical-value status of entities in operations such as filling, balancing, and exhaustion.

P-ISSN: 3050-9572 E-ISSN: 3050-9580
EUROPEAN DATA SCIENCE JOURNAL
Volume: 04 Issue: 01
RECEIVED: JANUARY 07 REVISED: JANUARY 26 ACCEPTED: FEBRUARY 05 PUBLISHED: FEBRUARY 17

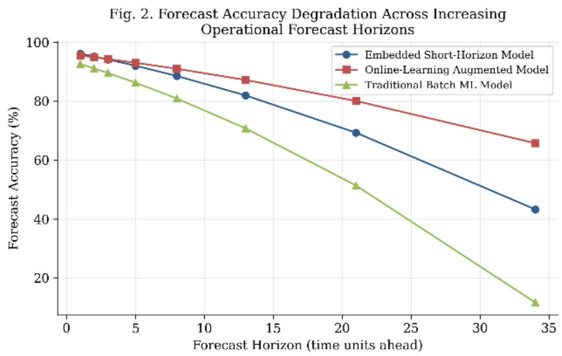

Fig 2: Forecast accuracy degradation across increasing horizons (embedded short-horizon, online-learning, batch ML)

Data models, semantics, ontologies, and coding schemas must be defined to guarantee the business/enterprise-domain specific data quality, lineage accuracy, governance integrity, and industry-interoperability requirements of all data streams and forecasted data that feed into or out of enterprise-intelligence pipelines. Three aspects deserve explicit treatment here: the canonical data sources used in enterprise-intelligence pipelines; the varieties of event-data stream or-schema-mapping definitions; and the reasons for requiring semantic consistency across enterprise-intelligence pipelines.

2. **Pipeline Processing Function**

$$P_t = f(D_t)$$

Where P_t is the processed pipeline output.

2. Conceptual Foundations of Real-Time Enterprise Intelligence

Real-Time Enterprise Intelligence Pipelines for High-Performance Operational Forecasting and Decision Engineering

Definitions and Scope. Current information technology (IT) paradigms provide exponentially increasing computational and storage capacity. Simultaneously, market demands evolve (higher uncertainty) and management practices change (increased globalization and outsourcing). Correspondingly, the open data movement enables the use of publicly available external data collections. Although external data can enhance decisions beyond the typical operational focus of the information systems of an enterprise, integrating data from different sources remains usually challenging. Moreover, internal data tend to be partially explored, stored in silos, and modelled for different purposes. An active approach to operations forecasting represents a valuable alternative or complement. A more direct impact on firm performance can be obtained by relating forecasted data with appropriate decision models to enhance the quality and consistency of key operational decisions.

P-ISSN: 3050-9572 E-ISSN: 3050-9580
EUROPEAN DATA SCIENCE JOURNAL
Volume: 04 Issue: 01
RECEIVED: JANUARY 07 REVISED: JANUARY 26 ACCEPTED: FEBRUARY 05 PUBLISHED: FEBRUARY 17

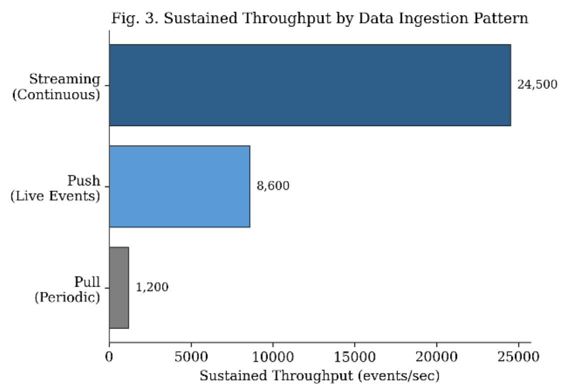

Fig 3: Sustained throughput by ingestion pattern (pull, push, streaming)

Forecasting and decision models bound AI in an enterprise environment. Forecasts embedded in a real-time workflow may directly impact operations. Such methods do not require a batch cycle nor use all available data resources. Time series models based on a single historical sequence are most likely employed, as they present low-projection computational cost and reduced time series length. Learning over time during normal operations is mandatory. High-performance models—either in terms of accuracy, computational efficiency (when embedded), or resource demand (when loaded)—should be the focus of engineering resources. Although the forecasts can be directly delivered to the corresponding operational area, additional value can be retrieved by embedding them into a decision model used to partially or fully automate the decision-making process. The integration of a suitable forecasting and decision model combination forms the core of operations decision engineering.

3. **Semantic Linking**

$$S_t = D_t + O_t$$

Where $O_t$represents ontology-based semantic mapping.

A. Definitions and Scope

Real-time enterprise intelligence is defined as the analysis and repeatable orchestration of two or more composite data pipelines in real-time, leading to business-relevant decisions made within seconds or minutes and executed with near-zero delay. The term pipeline has a wide usage in data engineering; the relevant meaning in this context is a composition of parallel data processing paths, in batch or streaming mode, whose results feed into another pipeline. Pipelines are thus distinct from typical data flows, where the output feeds data receivers, rather than subsequent processing steps. Orchestration refers to the automated and repeatable sequencing of the execution of multiple pipelines on a common data model—the latter an essential condition for rapid execution of the various steps of next-best-action and what-if scenarios in decision engineering.

P-ISSN: 3050-9572 E-ISSN: 3050-9580
EUROPEAN DATA SCIENCE JOURNAL
Volume: 04 Issue: 01
RECEIVED: JANUARY 07 REVISED: JANUARY 26 ACCEPTED: FEBRUARY 05 PUBLISHED: FEBRUARY 17

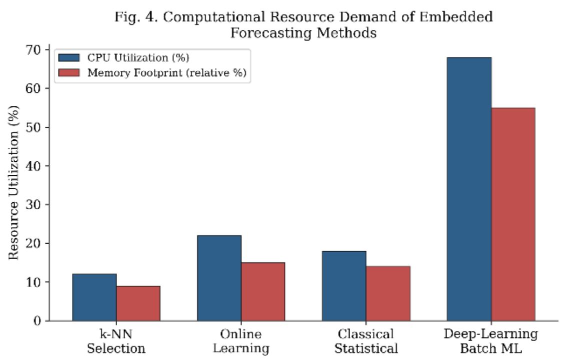

Fig 4: Computational resource demand across forecasting methods (k-NN, online learning, statistical, deep-learning batch)

The definition excludes batch-oriented data design and processing, even when implemented in a near-real-time mode with data warehouse technology. The distinction becomes clearer when considering decision engineering, defined as the determination of next-best-action and what-if scenarios in business, executed with near-zero delay and incorporated into standard operating procedures. The contrast between decision-engineering and batch methods resides both in the timeframes of the decision process and its execution, the critical factors being the number and importance of each decision in the execution of business processes. Batch data processes support visualization, predictive and prescriptive analytics, and are therefore necessary for decision support.

4. **Forecasting Output**

$$\hat{Y}_{t+1} = F(X_t)$$

Where $\hat{Y}_{t+1}$ is the predicted operational value.

3. System Architecture and Pipeline Orchestration

The conceptual foundations and high-performance forecasting methods described above are therefore combined into a high-level system architecture that provides a framework for real-time enterprise intelligence (REI). A definition of the REI pipeline and its orchestration are presented next. Pipeline orchestration implements a logical expressivity that constitutes a grammar for real-time enterprise intelligence. Four types of elements are identified: a data model; the ingestion of events and data streams; the real-time operations that process and analyse the data; and sinks for displaying and using the results. Each type is instantiated in the pipeline. The proposed method is designed primarily for operational forecasting applications, but is also applicable to operational analytics tasks such as those in decision engineering.

Table 1: Core Components of Enterprise Intelligence Pipelines

P-ISSN: 3050-9572
E-ISSN: 3050-9580
EUROPEAN DATA SCIENCE JOURNAL
Volume: 04 Issue: 01
RECEIVED: JANUARY 07 REVISED: JANUARY 26 ACCEPTED: FEBRUARY 05 PUBLISHED: FEBRUARY 17

Component	Purpose
Data Sources	Collect enterprise records and operational events
Semantic Linking	Converts data into common operational meaning
Pipeline Orchestration	Coordinates multiple processing layers
Forecasting Engine	Predicts short-term operational outcomes
Decision Engine	Supports next-best-action decisions

Ingestion patterns and the associated technologies are described next. The patterns are classified as pull for periodic corporate data acquisition, push for live events generated by enterprise processes, and streaming for flows from embedded monitoring devices or specialist applications such as stock exchanges. Event sources, adapters, and streaming technologies implement the patterns. A definition of data provenance specifies the requirements for data quality, latency, and fault tolerance for the various ingestion patterns. Although fault-tolerant mechanisms are needed only in the streaming case, the definition provides a common vocabulary for specifying latency targets for all ingestion patterns and for audit purposes. The definition of provenance also implies requirements for enterprise-free, enterprise-complete, and enterprise-consistent data lineage.

5. **Decision Rule**

$$A_t = g(\hat{Y}_{t+1}, C_t)$$

Where $A_t$is the recommended action and $C_t$is the current operational condition.

A. Ingestion and Streaming Data Sources

Operational forecasting and decision engineering require the integration of high-velocity time-series, event-based, and transaction-based data streams from operational environments with relatively low-velocity and medium- to low-volume data from various information systems. The ingestion of external data sources supporting enterprise-intelligence pipelines can be achieved through data streaming, appending, and bulk-copy patterns, either directly or through a combination of these techniques, depending on latency requirements, resource availability, and possibly fault-tolerance concerns.

Table 2: Types of Enterprise Event Data

Event Type	Description
Positional Events	Capture state changes of business entities
Stock Events	Track fast-changing operational volumes

P-ISSN: 3050-9572 E-ISSN: 3050-9580
EUROPEAN DATA SCIENCE JOURNAL
Volume: 04 Issue: 01
RECEIVED: JANUARY 07 REVISED: JANUARY 26 ACCEPTED: FEBRUARY 05 PUBLISHED: FEBRUARY 17

Event Type	Description
Transaction Events	Record business process activities
Time-Series Events	Support forecasting over time
Contextual Events	Add location, time, or business context

The sources of streaming data can be sorted into a few categories: time-series data from known-structure event streams; transaction flows that can be modelled as Events by the data grammar, possibly integrating a limited set of Geo-, Time-, and Contextual- data coding; and transaction data that can also be associated with a decision-intent language grammar. The ingestion of batch or periodic sources such as those from Data Grids and Data Lakes may be converted into a streaming format by incrementally updating or mirroring the Source data every time a change occurs – such as all the Data Transfer Points of an ETL application or a low-velocity Set of Data – or simply by periodically copying new records into a sink-integrating system.

6. **Latency Constraint**

$$L_t \leq L_{max}$$

Where $L_t$is processing latency and $L_{max}$is the maximum allowed delay.

4. High-Performance Forecasting Methods for Operations

Suitable methods of operational forecasting that fulfil embedded embedded context requirements for real-time, low-latency reaction. Time series forecasting – the core need of high-volume decision support – dominates operational forecasting. Nevertheless, many time series forecasting methods, especially ML-based models, have relatively high requirements on data set size and dimensions, accuracy and computational speed. As such, occasional reconciliation is necessary with higher-latency decisions or those areas with lower data volume. Feature engineering from related data sources is sometimes important.

Table 3: Ingestion Patterns in Real-Time Pipelines

Pattern	Usage
Pull	Periodic data collection from corporate systems
Push	Live events generated by business processes
Streaming	Continuous data from devices or applications
Bulk Copy	Large-volume periodic data movement

P-ISSN: 3050-9572 E-ISSN: 3050-9580
EUROPEAN DATA SCIENCE JOURNAL
Volume: 04 Issue: 01
RECEIVED: JANUARY 07 REVISED: JANUARY 26 ACCEPTED: FEBRUARY 05 PUBLISHED: FEBRUARY 17

Pattern	Usage
Appending	Incremental data updates

Time series forecasting frequently provides the core requirement of operational forecasting. Digitalisation makes time series forecasting applicable in embedded pipelines requiring ultra-low latency from event-triggered operation. With capacity planning and near-term troubleshooting decision engineering as key applications, responsive time series forecasting methods are mandated for incorporation into the corresponding forecasting pipelines.

7. **Pipeline Throughput**

$$T = \frac{N}{\Delta t}$$

Where $N$is the number of processed events in time interval Δt.

A. Time Series Forecasting in Embedded Pipelines

The analysis defines a class of forecasting problems with special constraints and performance requirements: predicting attributes of operational processes and records a few time units ahead. Embedding times-series forecasting in real-time analytic or decision pipelines removes the long and unpredictable latencies of traditional forecasting frameworks. Designs are provided for low-latency, embedded time-series forecasting for both short- and medium-term horizons. Short-horizon forecasts (a few time units ahead) are supplied by methods with the lowest latency and computational load, ideally on the reacting forces that target near-time processes. Medium-horizon forecasts (a small number of time units ahead) are constructed from larger samples and can handle more complex and longer-lead processes. Feature-engineering pipelines extend the applicability of predictive models for different time-horizon requirements. Learning from real-time data simplifies and improves medium-horizon forecasting.

Table 4: Forecasting Requirements

Requirement	Importance
Low Latency	Enables fast operational response
Computational Efficiency	Reduces processing load
Online Learning	Improves models during live operations
Short-Horizon Prediction	Supports immediate decisions
Feature Engineering	Enhances forecast accuracy

P-ISSN: 3050-9572 E-ISSN: 3050-9580
EUROPEAN DATA SCIENCE JOURNAL
Volume: 04 Issue: 01
RECEIVED: JANUARY 07 REVISED: JANUARY 26 ACCEPTED: FEBRUARY 05 PUBLISHED: FEBRUARY 17

Time-series forecasting methods are generally heavy processes. They require long training periods, large historical data samples for physical or statistical modeling, or extensive feature engineering for machine-learning approaches. Because forecast timelines are remote with respect to the current data sample, the methods call for explicit management of lag times and parallel operations. For most forecast methods, the optimizer reacts by committing resources to the decision path that is “last out the door.”

8. **Forecast Error**

$$E_t = Y_t - \hat{Y}_t$$

Where $Y_t$is the actual value and $\hat{Y}_t$is the forecasted value.

5. Conclusion

Real-time enterprise intelligence pipelines offer a practical operational forecasting and decision-engineering framework. Foundations, a decision-support approach, high-performance-forecasting methods, and adaptive online decision-engineering paradigms form a comprehensive enterprise-intelligence concept. Embedded operational analytics and decision-support forecasting rely on diverse analytics engines that operate within low-latency and resource-constrained environments.

Enterprise-intelligence pipelines exemplify a data-analytic process designed to provide timely, complete, accurate, and faithful answers to operation-decision questions. Real-time data sources supply massive amounts of business events that are processed, enriched, temporally joined, and combined with external data. Condition-action decision-engineering rules, enablement-status-computation rules, and event-test rules address key questions. Among those questions are whether any problems exist with the key performance indicators that matter to the business and whether action should be taken to profitably mitigate an imminent or in-process problem.

References

1. Chen, Y., Kang, Y., Chen, Y., & Wang, Z. (2020). Probabilistic forecasting with temporal convolutional neural network. Neurocomputing, 399, 491–501.
2. Salinas, D., Flunkert, V., Gasthaus, J., & Januschowski, T. (2020). DeepAR: Probabilistic forecasting with autoregressive recurrent networks. International Journal of Forecasting, 36(3), 1181–1191.
3. Lim, B., & Zohren, S. (2021). Time-series forecasting with deep learning: A survey. Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences, 379(2194), 20200209.
4. Mashetty, S., Malempati, M., Paleti, S., Adusupalli, B., & Singireddy, J. (2025). A Multidisciplinary Framework for AI and Data-Driven Transformation in Taxation, Insurance, Mortgage Financing, and Financial Advisory: Integrating Cloud Computing, Deep Learning, and Agentic AI for Community-Centric Economic Development. Insurance, Mortgage Financing, and Financial Advisory: Integrating Cloud Computing, Deep Learning, and Agentic AI for Community-Centric Economic Development.

P-ISSN: 3050-9572 E-ISSN: 3050-9580
EUROPEAN DATA SCIENCE JOURNAL
Volume: 04 Issue: 01
RECEIVED: JANUARY 07 REVISED: JANUARY 26 ACCEPTED: FEBRUARY 05 PUBLISHED: FEBRUARY 17

5. Zhou, H., Zhang, S., Peng, J., Zhang, S., Li, J., Xiong, H., & Zhang, W. (2021). Informer: Beyond efficient Transformer for long sequence time-series forecasting. Proceedings of the AAAI Conference on Artificial Intelligence, 35(12), 11106–11115.
6. Wu, H., Xu, J., Wang, J., & Long, M. (2021). Autoformer: Decomposition Transformers with auto-correlation for long-term series forecasting. Advances in Neural Information Processing Systems, 34, 22419–22430.
7. Al-Gabalawy, M., Hosny, N. S., & Adly, A. R. (2021). Probabilistic forecasting for energy time series considering uncertainties based on deep learning algorithms. Electric Power Systems Research, 196, 107216.
8. Marcjasz, G., Serafin, T., & Weron, R. (2021). Forecasting electricity prices: A comparative analysis of deep learning and classical models. Energy Economics, 101, 105402.
9. Kummari, D. N., Singireddy, J., Sheelam, G. K., Nandan, B. P., Pandiri, L., Lakkarasu, P., & Dwaraka. (2025, August). Generative AI Models for Process Optimization in Semiconductor Wafer Design and Yield Prediction. In International Conference on Artificial Intelligence: Theory and Applications (pp. 220-233). Cham: Springer Nature Switzerland.
10. Wen, Q., Zhou, T., Zhang, C., Chen, W., Ma, Z., Yan, J., & Sun, L. (2022). Transformers in time series: A survey. Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, 2022, 1–8.
11. Zhou, T., Ma, Z., Wen, Q., Wang, X., Sun, L., & Jin, R. (2022). FEDformer: Frequency enhanced decomposed Transformer for long-term series forecasting. Proceedings of the 39th International Conference on Machine Learning, 162, 27268–27286.
12. Pan, B. (2022). Improving seasonal forecast using probabilistic deep learning. Journal of Advances in Modeling Earth Systems, 14(3), e2021MS002766.
13. Nie, Y., Nguyen, N. H., Sinthong, P., & Kalagnanam, J. (2023). A time series is worth 64 words: Long-term forecasting with Transformers. International Conference on Learning Representations.
14. Adusupalli, B., Malempati, M., Paleti, S., Mashetty, S., & Singireddy, J. (2025). Integrated financial ecosystems: AI-driven innovations in taxation, insurance, mortgage analytics, and community investment through cloud, big data, and advanced data engineering. Journal of Information Systems Engineering and Management, 10, 1103-1117.
15. Wu, H., Hu, T., Liu, Y., Zhou, H., Wang, J., & Long, M. (2023). TimesNet: Temporal 2D-variation modeling for general time series analysis. International Conference on Learning Representations.
16. Jin, M., Wang, S., Ma, L., Chu, Z., Zhang, J. Y., Shi, X., Chen, P.-Y., Liang, Y., Li, Y.-F., Pan, S., & Wen, Q. (2024). Time-LLM: Time series forecasting by reprogramming large language models. International Conference on Learning Representations.
17. Liu, Y., Hu, T., Zhang, H., Wu, H., Wang, S., Ma, L., & Long, M. (2024). iTransformer: Inverted Transformers are effective for time series forecasting. International Conference on Learning Representations.
18. Wang, S., Wu, H., Shi, X., Hu, T., Luo, H., Ma, L., Zhang, J. Y., & Zhou, J. (2024). TimeMixer: Decomposable multiscale mixing for time series forecasting. International Conference on Learning Representations.
19. Mashetty, S. (2025). LEVERAGING DEEP LEARNING, NEURAL NETWORKS, AND DATA ENGINEERING FOR INTELLIGENT MORTGAGE LOAN VALIDATION. INTERNATIONAL JOURNAL OF SOCIAL SCIENCE & INTERDISCIPLINARY RESEARCH ISSN: 2277-3630 Impact factor: 8.036, 14(04), 51-65.
20. Woo, G., Liu, C., Kumar, A., Xiong, C., Savarese, S., & Sahoo, D. (2024). Unified training of universal time series forecasting Transformers. Proceedings of the 41st International Conference on Machine Learning, 235, 53140–53164.
21. Liu, Y., Zhang, H., Li, C., Huang, X., Wang, J., & Long, M. (2024). Timer: Generative pre-trained Transformers are large time series models. Proceedings of the 41st International Conference on Machine Learning, 235, 32369–32399.
22. Ansari, A. F., Stella, L., Turkmen, C., Zhang, X., Mercado, P., Shen, H., Shchur, O., Rangapuram, S. S., Pineda Arango, S., Kapoor, S., Zschiegner, J., Maddix, D. C., Mahoney, M. W., Torkkola, K., Wilson, A. G., Bohlke-Schneider, M., & Wang, Y. (2024). Chronos: Learning the language of time series. Transactions on Machine Learning Research.

P-ISSN: 3050-9572 E-ISSN: 3050-9580
EUROPEAN DATA SCIENCE JOURNAL
Volume: 04 Issue: 01
RECEIVED: JANUARY 07 REVISED: JANUARY 26 ACCEPTED: FEBRUARY 05 PUBLISHED: FEBRUARY 17

23. Sen, R., & Zhou, Y. (2024). A decoder-only foundation model for time-series forecasting. Proceedings of the 41st International Conference on Machine Learning.
24. Oakley, J., Conlan, C., Demirci, G. V., Sfyridis, A., & Ferhatosmanoglu, H. (2024). Foresight plus: Serverless spatio-temporal traffic forecasting. GeoInformatica, 28, 649–677.
25. Wang, Y., Zhang, Y., & Zhang, X. (2024). Deep learning-based multivariate time-series forecasting for large-scale operational systems. Applied Sciences, 14.
26. Liu, Y., Hu, T., Zhang, H., Wu, H., Wang, S., Ma, L., & Long, M. (2024). Time-series forecasting with inverted Transformers and multivariate representations. International Conference on Learning Representations.
27. Wang, S., Li, J., Shi, X., Ye, Z., Mo, B., Lin, W., Shengtong, J., Chu, Z., & Jin, M. (2025). TimeMixer++: A general time series pattern machine for universal predictive analysis. International Conference on Learning Representations.
28. Liu, Y., Zhang, H., Li, C., Huang, X., Wang, J., & Long, M. (2025). Timer-XL: Long-context Transformers for unified time series forecasting. International Conference on Learning Representations.
29. Recharla, M. (2024). Antioxidants, Biological Markers, Catalase, Glutathione Peroxidase, Chronic Periodontitis, Saliva, Smokeless tobacco, Smoker. Frontiers in Health Informatics, 13(8), 4999.
30. Liu, Y., Zhang, H., Li, C., Huang, X., Wang, J., & Long, M. (2025). Sundial: A family of large-scale generative time-series foundation models. Proceedings of the 42nd International Conference on Machine Learning.
31. Ansari, A. F., Shchur, O., Küken, J., Auer, A., Han, B., Mercado, P., Rangapuram, S. S., Shen, H., Stella, L., Zhang, X., Goswami, M., Kapoor, S., Maddix, D. C., Guerron, P., Hu, T., Yin, J., Desai, P. M., Wang, H., Rangwala, H., Karypis, G., & Wang, Y. (2025). Chronos-2: From univariate to universal forecasting. arXiv preprint.
32. Wen, Q., Zhou, T., Zhang, C., Chen, W., Ma, Z., Yan, J., & Sun, L. (2023). Transformers in time series: A survey. Proceedings of the 32nd International Joint Conference on Artificial Intelligence, 6775–6783.
33. Rasul, K., Ashok, A., Reddy, A. S., Bhatt, C., Lalla, R., Lensch, H., Pineda Arango, S., Kapoor, S., & Le, D. (2024). Lag-Llama: Towards foundation models for probabilistic time series forecasting. Proceedings of the 2024 International Conference on Learning Representations.
34. Al-Gabalawy, M., Hosny, N. S., & Adly, A. R. (2021). Deep learning approaches for uncertainty-aware time-series forecasting in energy systems. Electric Power Systems Research, 196, 107216.
35. Oakley, J., Conlan, C., Demirci, G. V., Sfyridis, A., & Ferhatosmanoglu, H. (2024). Real-time spatio-temporal forecasting using serverless computing architectures. GeoInformatica, 28, 649–677.