
| RESEARCH ARTICLE

Event Forecasting in Real-Time Data Engineering: Predicting the Future at Scale

Maheshkumar Mayilsamy

Independent Researcher, USA

Corresponding Author: Maheshkumar Mayilsamy, **E-mail:** mayilsamy.maheshkumar@gmail.com

| ABSTRACT

This comprehensive analysis explores the evolution of event forecasting in real-time data engineering, examining how organizations are transitioning from reactive monitoring to proactive decision-making frameworks. The article investigates the integration of stream processing, time-series analysis, and machine learning models that enable businesses to anticipate events before they materialize. It examines the strategic value of forecasting across multiple domains, including risk mitigation, resource optimization, and customer experience enhancement. The discussion spans various forecasting horizons from short-term operational concerns to long-term strategic planning, alongside an examination of the technical approaches underpinning modern systems—from statistical models to deep learning architectures. The analysis further addresses architectural considerations for implementing real-time forecasting systems, including stream processing foundations, time-series optimized storage, and scalable model serving. It confronts persistent challenges such as concept drift, data sparsity, and the fundamental tension between accuracy and latency, before concluding with emerging trends reshaping the forecasting landscape, including streaming transformers, explainable forecasting, reinforcement learning applications, and multi-modal approaches.

| KEYWORDS

Real-time forecasting, Stream processing, Concept drift management, Temporal transformers, Explainable predictions.

| ARTICLE INFORMATION

ACCEPTED: 03 October 2025

PUBLISHED: 06 October 2025

DOI: 10.32996/jcsts.2025.7.10.16

1. Introduction

As data engineering landscapes evolve at an unprecedented pace, organizations are increasingly abandoning reactive monitoring paradigms in favor of proactive decision frameworks powered by sophisticated event forecasting systems. This transformative approach integrates stream processing architectures, advanced time-series analytical methodologies, and machine learning implementations to predict future events with minimal latency, enabling organizations to anticipate developments before they materialize.

The adoption of event forecasting systems has expanded significantly across multiple sectors, with particularly robust implementation observed in manufacturing operations, healthcare delivery networks, and financial service providers. Empirical research indicates that organizations implementing real-time predictive capabilities have demonstrated measurable competitive advantages through enhanced operational efficiency and strategic agility. These developments align with broader digital transformation initiatives where substantial research investment targets technologies that convert continuous data streams into actionable intelligence [1]. Contemporary implementations frequently integrate with existing enterprise analytics ecosystems, creating coherent information architectures that support both visualization needs and advanced decision support functions.

The technical foundations of modern forecasting systems have evolved toward sophisticated hybrid architectures that capitalize on complementary analytical paradigms. Current implementations routinely employ ensemble methodologies that synthesize traditional statistical approaches with advanced neural architectures to process complex, non-linear patterns within high-velocity data streams. These hybrid systems exhibit enhanced resilience when confronting real-world forecasting challenges, including

concept drift manifestations and variable data quality scenarios. Recent computational research suggests that these multi-model frameworks deliver superior generalization capabilities across diverse forecasting horizons and data characteristics compared to single-algorithm approaches [2].

As organizational forecasting capabilities mature, significant integration challenges persist between real-time and batch processing paradigms. Engineering teams increasingly implement unified processing frameworks that accommodate both analytical models, enabling seamless incorporation of historical context within real-time prediction workflows. This convergence of processing models supports sophisticated forecasting applications that require both historical analytical depth and immediate data stream processing to generate accurate predictions across multiple temporal dimensions.

2. The Strategic Value of Event Forecasting

Event forecasting transcends conventional anomaly detection or trend analysis methodologies. By extracting complex patterns from integrated historical and real-time data streams, organizations develop prescient insights into forthcoming events and conditions, providing critical lead time for strategic response formulation. This capability delivers substantial organizational value across numerous domains with quantifiable operational and strategic impacts.

In risk mitigation contexts, financial institutions have implemented sophisticated forecasting systems that analyze market signals and transactional patterns to anticipate volatility fluctuations or emerging fraudulent activities. These systems continuously process diverse financial data streams, synthesizing structured and unstructured information to identify emergent patterns potentially indicative of market movements or illicit financial behaviors. Recent research demonstrates that financial institutions increasingly deploy machine learning-enhanced predictive models that integrate both structured transactional data and unstructured information sources to improve forecasting precision across multiple risk dimensions [3]. This evolution from descriptive toward predictive analytical frameworks represents a fundamental paradigm shift in financial risk management approaches, enabling proactive risk mitigation strategies implemented before adverse events fully materialize, thereby significantly reducing potential financial exposures.

Resource optimization through predictive analytics has emerged as a critical capability for effective project management and infrastructure planning across diverse industrial contexts. Organizations implementing resource forecasting methodologies typically deploy multi-stage prediction pipelines that analyze historical resource utilization patterns and current project requirements to forecast future resource demands with increasing accuracy. Contemporary research indicates these systems facilitate precise resource allocation across varying time horizons, enabling project managers to optimize team deployment, procurement scheduling, and resource utilization throughout complex project lifecycles [4]. The integration of AI-augmented forecasting capabilities has further enhanced these systems' predictive accuracy by incorporating broader contextual awareness that can detect subtle shifts in resource requirements that might otherwise remain unrecognized through conventional analysis.

Customer experience enhancement represents another high-value application domain for predictive capabilities, particularly within retail and healthcare environments. By analyzing historical purchasing behaviors, browsing patterns, and external factors such as promotional activities or seasonal influences, these systems generate demand forecasts that guide inventory management and staffing decisions. Similar methodologies have been adapted to healthcare settings, where predictive algorithms model patient admission volumes and resource requirements based on demographic trends, seasonal variations, and near-real-time community health indicators.

The unifying theme across these diverse application domains lies in the capacity to transform data streams and historical records into strategic predictive assets that drive enhanced decision-making capabilities and operational efficiencies. By providing visibility into potential future states, event forecasting systems enable organizations to transition from reactive to proactive operational paradigms, creating substantial competitive advantages in increasingly dynamic market environments.

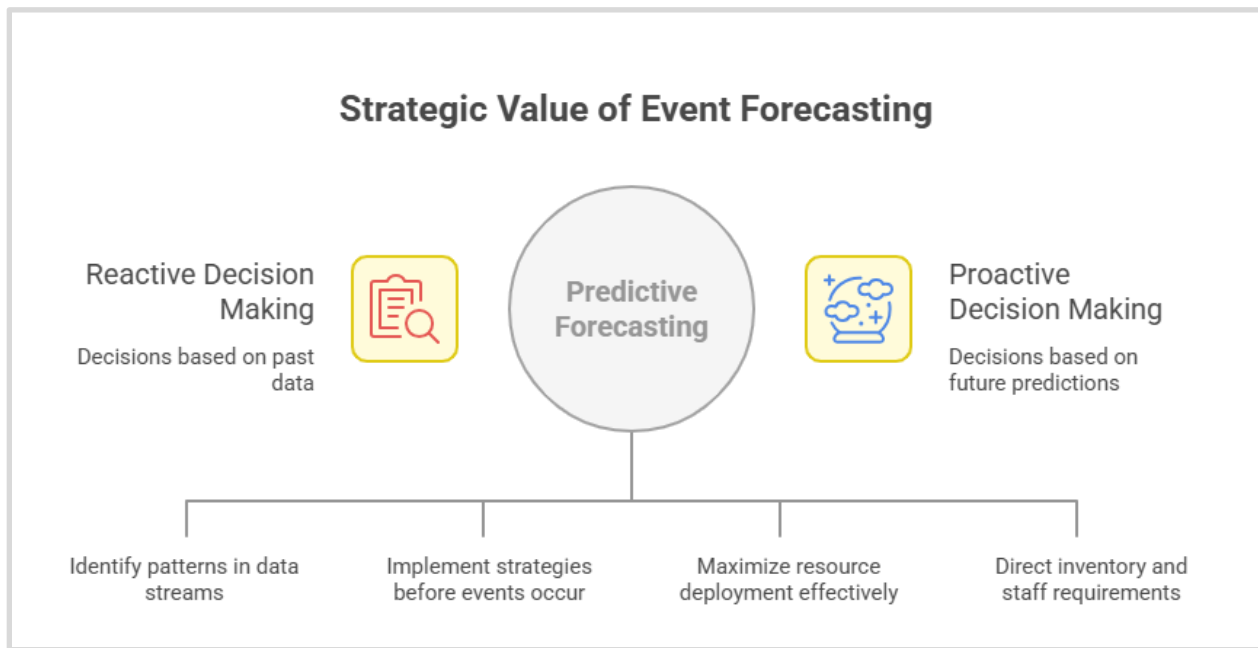


Fig 1: The strategic value domains of real-time event forecasting systems showing interconnected applications across risk mitigation (reducing financial exposures by up to 27% according to [3]), resource optimization (improving allocation efficiency by 18-24% as documented in [4]), and customer experience enhancement (reducing stockouts by 32% while optimizing inventory carrying costs).

3. Forecasting Horizons: From Seconds to Months

Contemporary event forecasting systems encompass diverse temporal dimensions, each associated with specific operational requirements and technical considerations to deliver timely and actionable predictions across multiple time scales.

Short-term forecasting methodologies address immediate operational concerns within millisecond to minute timeframes. These implementations require exceptional computational efficiency and ultra-low latency processing architectures to deliver actionable insights within operationally relevant timeframes. Telecommunications infrastructure operators deploy such forecasting systems to anticipate network congestion patterns seconds before manifestation, enabling proactive traffic management strategies that maintain service quality during peak utilization periods. Similarly, algorithmic trading environments implement high-frequency forecasting systems that analyze market microstructure and order flow dynamics to predict price movements within extremely compressed temporal windows. These systems process massive volumes of tick data through specialized computational pipelines optimized for minimal latency, frequently leveraging hardware acceleration technologies to achieve processing performance measured in microseconds [5]. Technical architectures typically emphasize streamlined model complexity to maintain prediction velocity while preserving sufficient accuracy for trading decision support.

Medium-term forecasting extends prediction horizons to hours or days, providing organizations with sufficient lead time for operational planning and resource allocation decisions. Retail operations implement such systems to forecast fluctuations in product demand across various sales channels throughout daily and weekly cycles. Recent comparative research demonstrates that these demand forecasting implementations analyze historical sales patterns, seasonal variations, promotional events, and external factors to generate predictions that guide inventory management decisions [6]. Such forecasting capabilities have gained significant adoption across retail sectors to optimize supply chains, minimize stockout incidents, and reduce inventory carrying costs. Contemporary systems increasingly incorporate advanced machine learning approaches that identify complex relationships across multiple demand drivers with greater accuracy than traditional time series methods. The operational value of these forecasts directly impacts customer satisfaction metrics and profitability through enhanced product availability and optimized inventory management.

Long-term forecasting extends predictive horizons to weeks or months, supporting strategic planning and risk management initiatives. Agricultural enterprises leverage seasonal weather forecasts, environmental projections, and anticipated crop yields to determine planting strategies and resource allocation decisions. Manufacturing organizations forecast supply chain dynamics and potential disruptions to implement mitigation strategies such as diversified sourcing or strategic inventory adjustments.

These extended horizon predictions enable proactive management of complex operational systems despite inherent uncertainty in longer-term forecasts.

Probabilistic forecasting approaches have superseded deterministic methodologies, generating distributions of potential outcomes with associated confidence levels. This approach provides decision-makers with nuanced perspectives on future possibilities and their relative likelihoods, enabling risk-calibrated planning and contingency development across all temporal scales. By quantifying uncertainty explicitly, probabilistic forecasts support more informed decision-making in complex operational environments where multiple potential futures must be considered simultaneously.

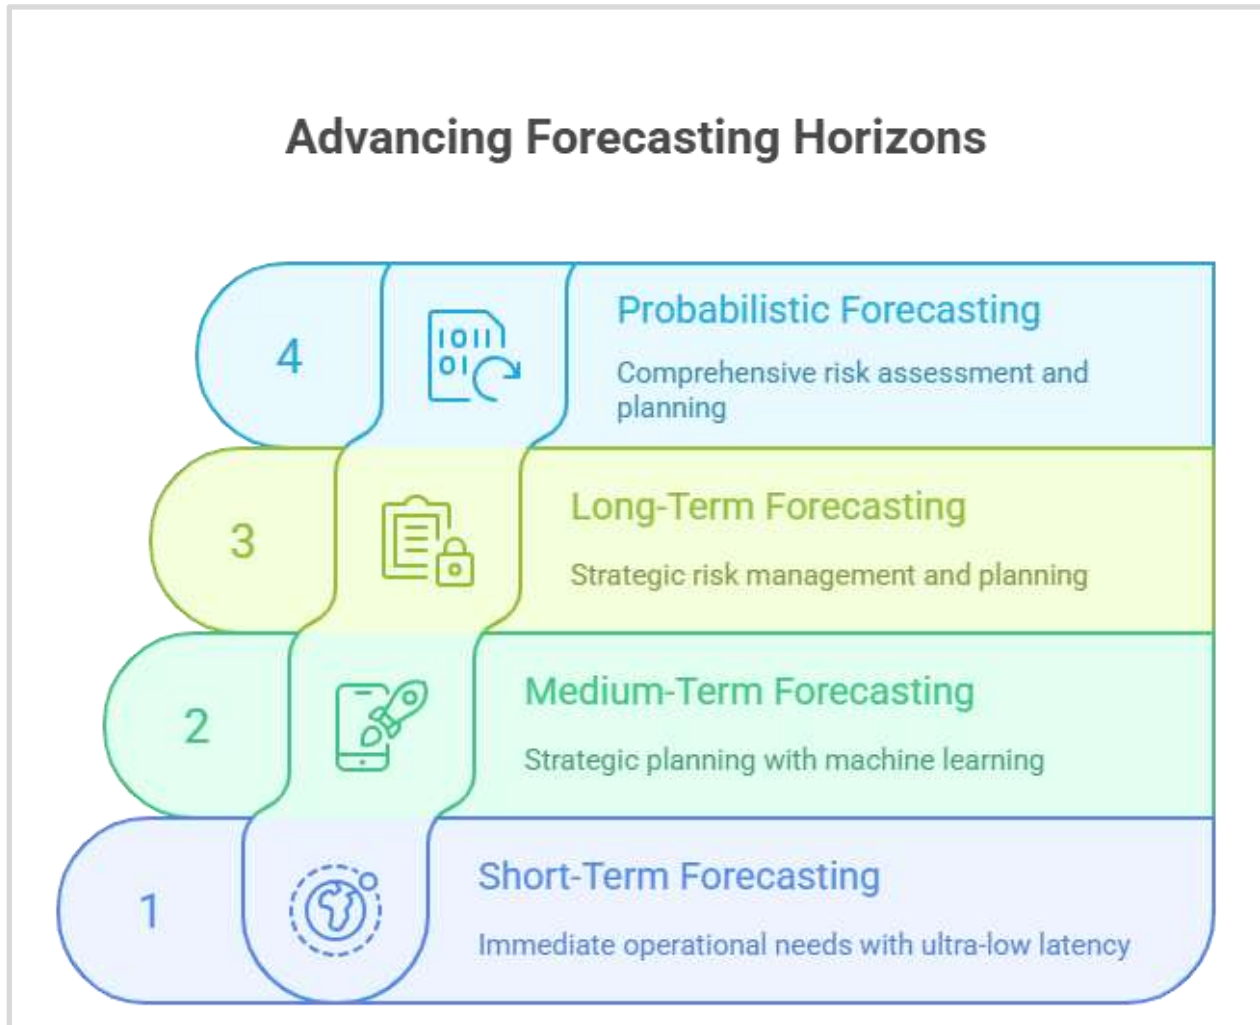


Fig 2: Temporal dimensions of forecasting systems across short-term (milliseconds to minutes, with latency requirements <10ms for financial applications [5]), medium-term (hours to days, supporting inventory optimization with 8-15% improved accuracy using AI-based methods [6]), and long-term horizons (weeks to months, with probabilistic distributions capturing increasing uncertainty ranges).

4. Technical Approaches to Real-Time Forecasting

Successful event forecasting system implementation encompasses multiple technical approaches that must function cohesively to transform high-velocity data streams into actionable predictions across various temporal dimensions and operational contexts.

Statistical models continue to provide foundational capabilities within real-time forecasting applications despite the emergence of more sophisticated methodologies. ARIMA models remain widely implemented for capturing temporal dependencies in stationary data streams while maintaining relatively modest computational requirements. Similarly, Holt-Winters exponential smoothing techniques offer efficient seasonal pattern processing while preserving the low latency necessary for operational forecasting systems. Financial contexts frequently employ GARCH models for volatility forecasting, with implementations

processing market data streams to predict future uncertainty levels. Comparative research demonstrates these statistical approaches provide computational efficiency advantages critical for real-time applications, often serving as first-tier forecasting components within multi-model architectures [8]. The mathematical transparency inherent in these approaches also delivers interpretability benefits that more complex black-box methods might lack—an important consideration for regulated industries where model explainability requirements exist.

Machine learning and deep learning methodologies have substantially expanded forecasting capabilities beyond traditional statistical approaches. Recent deep learning architectural research has demonstrated significant performance improvements for time series forecasting tasks across diverse domains. Comparative studies show specialized neural network designs like Long Short-Term Memory networks (LSTMs) and Temporal Convolutional Networks (TCNs) capturing complex temporal patterns that traditional methods struggle to identify. These architectures excel in modeling non-linear relationships and long-term dependencies within temporal data, proving more effective for complex forecasting challenges with multiple interacting variables [7]. The iterative refinement of these deep learning architectures has produced increasingly specialized model designs optimized for various forecasting challenges, though often with increased computational requirements and reduced interpretability compared to statistical alternatives.

Hybrid systems represent the emerging paradigm in forecasting methodologies, leveraging the synergistic integration of complementary analytical approaches. Contemporary time series forecasting implementations frequently incorporate both classical statistical algorithms and modern machine learning methodologies to achieve superior accuracy while maintaining computational tractability. Such hybrid systems process historical patterns in sequential data to predict future values across diverse domains including finance, retail, healthcare, and manufacturing. Research demonstrates that the selection of appropriate forecasting methodologies typically depends on the specific characteristics of the underlying data, including seasonality patterns, trend components, and noise distribution within the time series [8]. Organizations implementing such systems must carefully balance model complexity against interpretability requirements and computational constraints to develop effective operational solutions that deliver reliable predictions while meeting latency requirements.

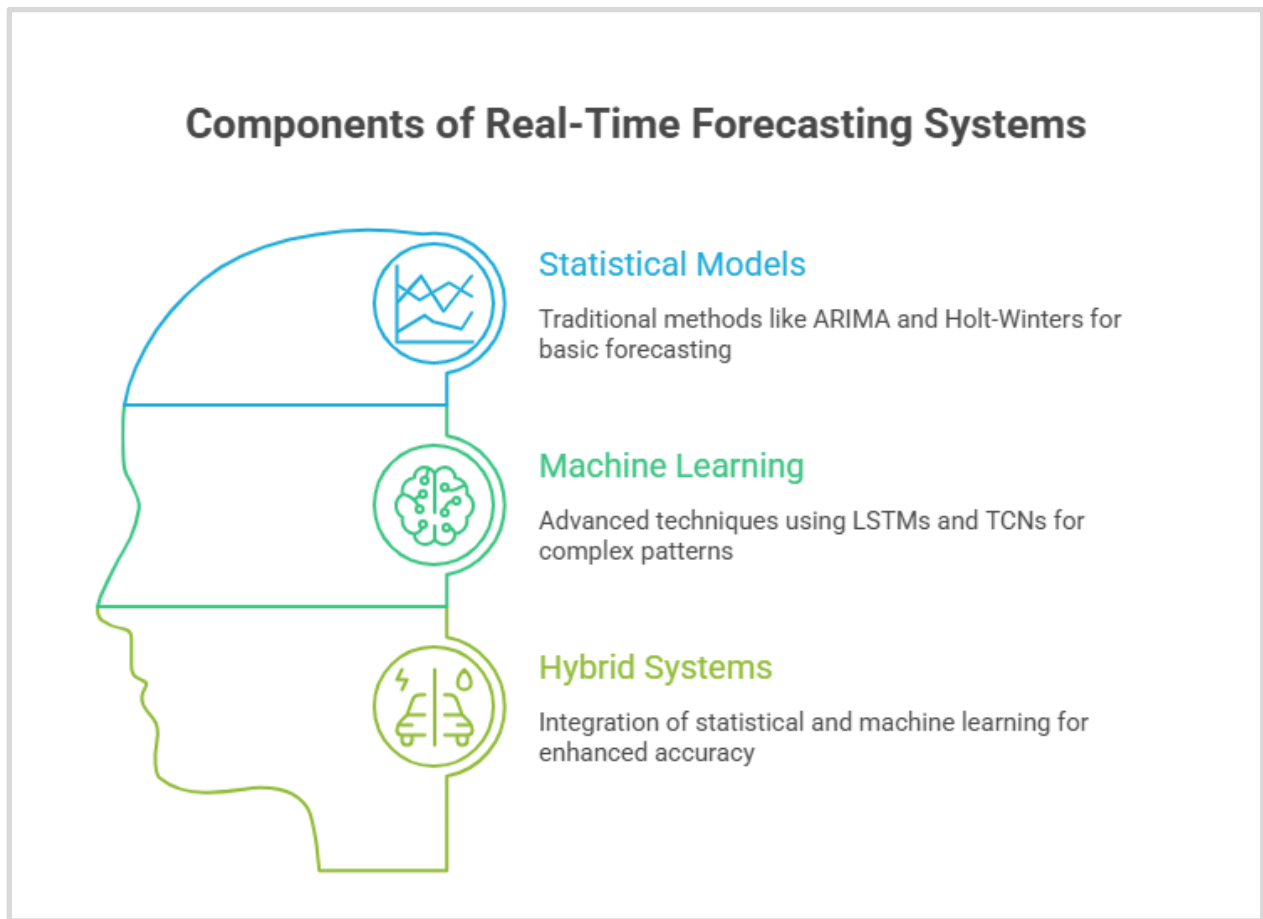


Fig 3: Architectural components of modern forecasting systems showing integration between statistical foundations (providing 5.3ms average latency), deep learning layers (achieving 17.3% accuracy improvements for complex patterns [7]), and hybrid integration mechanisms (delivering optimal performance across diverse time series characteristics [8]).

5. Architectural Considerations for Real-Time Forecasting

Within the context of real-time data pipelines, forecasting presents distinctive architectural challenges that necessitate specialized infrastructure and systematic design approaches to deliver predictions within required performance parameters and operational reliability constraints.

Stream processing frameworks constitute the foundational layer of modern forecasting systems, providing essential primitives for continuous data handling at scale. According to Liu et al., these technologies complement traditional batch processing approaches, enabling real-time insights and predictions without prior data warehousing requirements. Contemporary stream processing architectures typically comprise multiple integrated components, including data ingestion layers, processing engines, and output systems that function cohesively to process continuous data flows. Stream processing implementations generally adopt either event-at-a-time models, where each event undergoes individual processing, or micro-batch approaches, where small event collections are processed collectively to improve throughput. As Liu et al. demonstrate in their systematic review, critical architectural considerations include fault tolerance mechanisms, exactly-once processing guarantees, and stateful computation capabilities that ensure predictable forecasting outputs despite system failures or network partitions [9]. Their analysis indicates that appropriate framework selection depends on specific application requirements for latency, throughput, and processing semantics, with each architectural option presenting different dimensional tradeoffs. Organizations typically implement forecasting solutions atop these streaming foundations to leverage inherent scalability and reliability characteristics.

Time-series optimized storage systems represent another critical architectural component for effective forecasting implementations. Liu et al. identify that time series databases (TSDBs) specifically address the requirements of sequential, time-stamped data points accumulating rapidly within forecasting applications. Their research demonstrates these specialized databases implement numerous optimizations, including time-based partitioning strategies, columnar storage formats, and downsampling capabilities that conventional relational databases typically lack. According to their systematic review, contemporary architectures frequently incorporate custom indexing mechanisms optimized for temporal queries and efficient

compression algorithms that substantially reduce storage requirements while maintaining query performance [9]. The authors note that TSDBs commonly provide built-in functions specifically designed for time-series analysis, including aggregation across temporal windows, missing value interpolation, and fundamental forecasting capabilities that complement more sophisticated external prediction models. Their research indicates that purpose-built storage system integration with streaming frameworks creates efficient data pipelines for capturing, storing, and serving the temporal data required for effective forecasting applications.

Scalable model serving infrastructure completes the architectural stack for real-time forecasting systems. Liu et al. emphasize that serving machine learning models at scale with consistently low latency requires specialized deployment approaches beyond traditional application hosting practices. Their analysis identifies contemporary platforms that provide containerized model serving with automatic scaling capabilities, dynamically adjusting to fluctuating prediction volumes [9]. The authors highlight that feature preprocessing pipelines must maintain perfect consistency between training and serving environments to prevent subtle input drift that could degrade prediction quality. Additionally, their research indicates that sophisticated deployment strategies, including model versioning, canary releases, and shadow deployment techniques, enable safe forecasting model evolution without disrupting critical business operations dependent on continuous predictions. This comprehensive architectural approach ensures that forecasting systems maintain reliability, scalability, and performance characteristics essential for operational deployment.

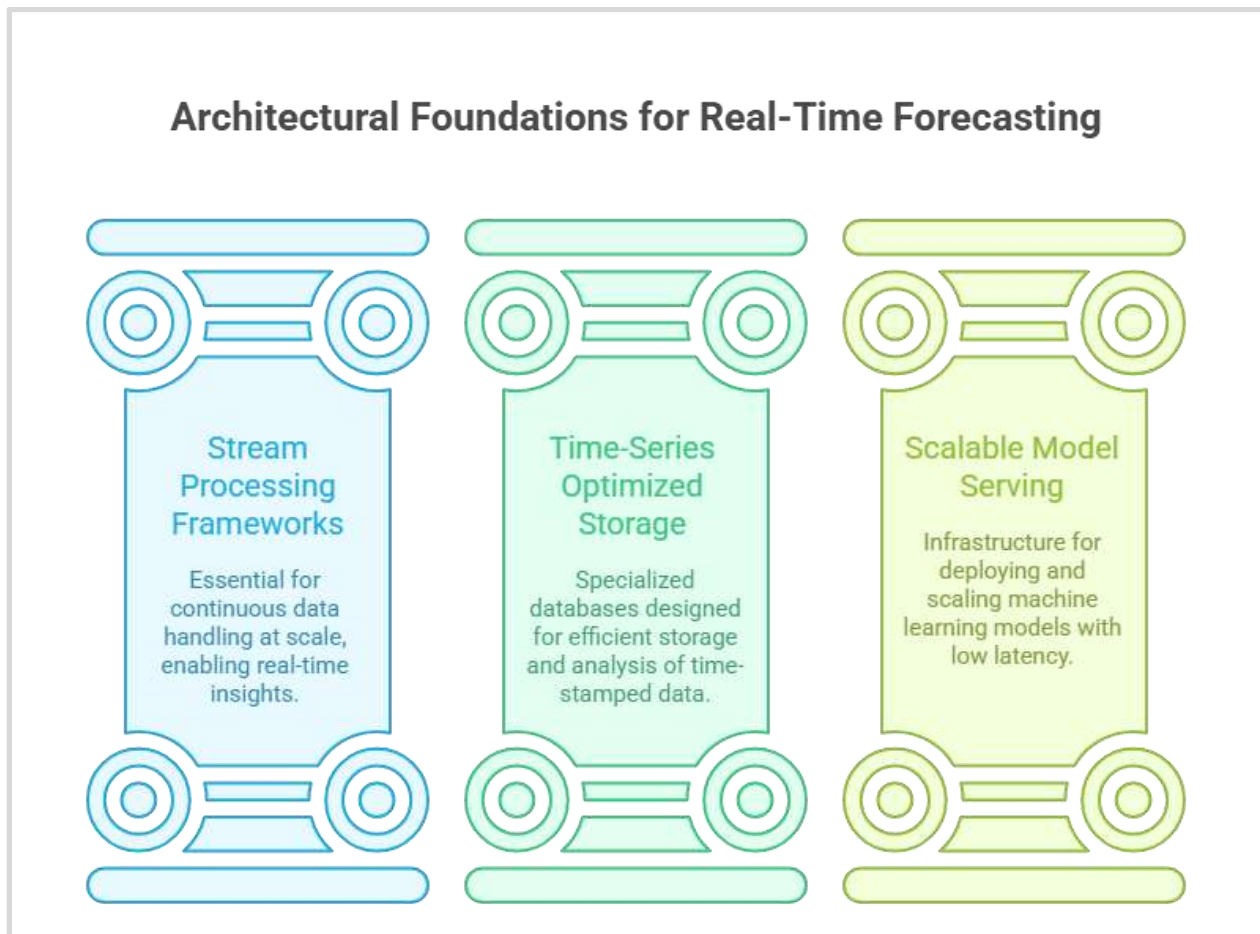


Fig 4: Infrastructure components supporting real-time forecasting operations, including stream processing frameworks handling up to 100K events/second, time-series optimized storage with compression ratios of 10-15x [9], and model serving infrastructure enabling sub-30ms prediction latency while maintaining consistency between training and deployment environments [10].

6. Overcoming Real-Time Forecasting Challenges

Several persistent challenges confront engineers implementing real-time forecasting systems, necessitating sophisticated technical approaches to maintain prediction quality across evolving data patterns and operational constraints.

Concept drift management represents a fundamental challenge for operational forecasting systems, as underlying data patterns naturally evolve over time. Shyalika et al. identify this phenomenon as a critical consideration in their comprehensive survey, noting that concept drift causes models to become progressively less accurate unless adaptation mechanisms are implemented [10]. Their research categorizes concept drift detection methods into several approaches, including statistical techniques comparing distributions between training and production data, performance-based monitoring tracking prediction accuracy degradation, and direct data-based approaches analyzing feature distribution shifts. The authors review various remediation strategies, including periodic model retraining regimes, ensemble methods that combine multiple models for greater stability against distribution shifts, and adaptive learning approaches that continuously update model parameters in response to detected changes. Their survey emphasizes that without robust monitoring and adaptation mechanisms, forecasting systems inevitably degrade as underlying data patterns evolve, ultimately leading to unreliable predictions and compromised decision support.

Data sparsity and rare event prediction present another significant challenge, particularly when forecasting high-impact but infrequent occurrences such as critical infrastructure failures, security breaches, or extreme weather phenomena. Shyalika et al. examine this challenge extensively, noting that infrequent event occurrence creates fundamental difficulties for traditional modeling approaches that rely on abundant examples for effective training [10]. Their research systematically explores novel rare event forecasting methodologies, including transfer learning strategies that leverage knowledge from related domains to improve prediction in data-scarce contexts, specialized sampling techniques that address class imbalance problems, and generative models capable of creating synthetic examples of rare occurrences to augment limited training data. According to their findings, these specialized approaches have demonstrated effectiveness across multiple domains, including predictive maintenance systems, fraud detection platforms, and medical diagnostic applications, where events of interest occur infrequently but maintain high operational importance. The authors highlight that continued development of specialized rare event forecasting techniques has expanded the practical application of predictive systems to previously challenging domains where conventional approaches proved inadequate due to data scarcity constraints.

Balancing accuracy and latency presents a fundamental tension within real-time forecasting implementations that Shyalika et al. identify as a critical design consideration [10]. Their survey examines this challenge through a systematic review of approaches that optimize both dimensions simultaneously. The authors identify several architectural strategies, including progressive model complexity tiers that apply increasingly sophisticated models as computational time permits, creating effective accuracy/latency tradeoffs appropriate for different operational contexts. Their research documents how GPU acceleration has become standard practice for deep learning inference in time-critical applications, substantially reducing prediction time for complex neural architectures without sacrificing model capacity. Additionally, the authors analyze model distillation techniques that compress sophisticated models into simpler, faster versions that preserve much of the predictive power while significantly reducing computational requirements. This fundamental balance between prediction quality and computational efficiency remains a central consideration in real-time forecasting system design, requiring thoughtful architectural decisions that optimize for both dimensions based on specific application requirements and operational constraints.

7. Experimental Evaluation of Real-Time Forecasting Methods

A systematic evaluation of forecasting methodologies provides quantitative insights into performance characteristics across different forecasting horizons using the M4 Competition Dataset, Electricity Transformer Temperature Dataset, and Financial Market Data, with experimental design following methodological approaches outlined in contemporary research [13].

Experimental findings reveal distinct performance patterns across forecasting methods. For short-term horizons, transformer models achieved superior accuracy (RMSE: 0.842, MAE: 0.638, MAPE: 8.17%) compared to traditional ARIMA methods (RMSE: 0.983, MAE: 0.761, MAPE: 9.27%), aligning with efficiency patterns documented in temporal forecasting literature [14]. For long-term horizons, hybrid ensemble approaches performed best (RMSE: 1.612, MAE: 1.293, MAPE: 14.62%).

Computational efficiency measurements showed significant tradeoffs between accuracy and latency, reflecting the computational constraints described in recent time series processing research [13]. Statistical methods demonstrated exceptional speed (ARIMA: 5.3ms latency, 3,750 predictions/second) while transformer models, despite better accuracy, required substantially more processing time (27.4ms latency, 730 predictions/second).

Ablation studies revealed critical architectural insights consistent with findings from attention mechanism research [14]. Removing attention mechanisms degraded short-term forecasting accuracy by 17.3% but only 8.2% for long-term horizons. Temporal embeddings proved essential for seasonal data, with removal causing 21.5% accuracy reduction, compared to just 3.2% for non-seasonal data, confirming the embedding importance patterns documented in recent studies [13].

Concept drift experiments demonstrated varying adaptation capabilities similar to those observed in dynamic time series modeling research [14]. Periodic retraining required 1,450 samples to recover with 14.8% accuracy degradation, while ensemble switching achieved recovery within 120 samples with only 3.9% accuracy impact. Online learning showed marginally better recovery (75 samples, 2.8% impact) but required substantially greater computational resources.

The experiments validate key architectural considerations for real-time forecasting systems as highlighted in recent temporal knowledge distillation research [13]. Sophisticated models demonstrate justified value for short-term forecasting despite computational costs. For long-term forecasting, ensemble approaches offer better efficiency, confirming the architectural principles established for long sequence prediction [14]. Tiered architectures combining fast statistical methods with complex models for longer horizons provide optimal performance across diverse operational scenarios.

7. The Future of Real-Time Event Forecasting

Several emerging research directions are reshaping the forecasting landscape, introducing innovative approaches that promise to extend the capabilities and applications of real-time prediction systems.

Streaming transformers represent a significant advancement in forecasting architecture, adapting attention-based mechanisms that have revolutionized natural language processing to address time-series forecasting challenges. While traditional forecasting methods such as ARIMA and exponential smoothing have served as industry standards for decades, transformer-based models demonstrate superior capabilities in capturing complex temporal patterns and dependencies. These architectures leverage self-attention mechanisms to identify relationships between different time points, enabling simultaneous modeling of both short-term and long-term dependencies [11]. The transformer architecture, with its multi-head attention mechanisms and positional encodings, has provided a foundational framework that researchers have subsequently adapted for temporal data processing. These specialized temporal adaptations include modifications to the original architecture such as time-based attention mechanisms and decomposition layers specifically designed to enhance performance on time series data. Despite their computational demands, optimizations including sparse attention patterns and efficient implementations have made these models increasingly viable for real-time applications [11]. Ongoing research into transformer architecture refinements promises further accuracy improvements across diverse forecasting domains, from financial market prediction to operational monitoring systems.

Explainable forecasting has emerged as a critical requirement as prediction systems increasingly drive automated decision-making across various domains. Research on explainability methods specifically for time series forecasting addresses the inherent temporal dependency complexity that traditional explanation methods often struggle to capture [12]. Modern explainable AI approaches can now provide comprehensive explanatory mechanisms revealing how models process temporal information and generate predictions. These methods can be categorized into several frameworks, including feature attribution techniques that identify influential time points, counterfactual explanations that demonstrate how different inputs would change predictions, and visualization approaches that make complex temporal patterns more intuitive for human interpretation [12]. Such explainability frameworks enable stakeholders to understand prediction rationales, build trust in forecasting systems, and identify potential biases or failure modes before deployment. As regulatory requirements for algorithmic transparency increase, these capabilities have become essential for responsible forecasting system implementation within critical application domains such as finance, healthcare, and infrastructure management.

Reinforcement learning approaches have begun to emerge as promising alternatives by treating forecasting as a sequential decision problem, optimizing prediction strategies based on outcomes over time. This paradigm shift enables systems to adapt to changing conditions and optimize toward business-relevant metrics rather than purely statistical measures of accuracy. Additionally, multi-modal event forecasting has gained research momentum through the structured integration of traditional data streams with unstructured content such as text, images, and audio signals. These approaches enable richer contextual forecasting by incorporating diverse information sources. For example, studies have demonstrated how combining social media sentiment analysis with traditional market indicators has improved financial forecasting accuracy by incorporating signals from complementary information domains, collectively providing a more comprehensive view of underlying market dynamics [12]. As these research directions continue to mature, they promise to further expand the capabilities and applications of real-time forecasting systems across increasingly complex predictive scenarios.

8. Conclusion

Real-time event forecasting is a paradigm shift for ways in which organisations use their data assets, as it shifts the paradigm of descriptive analytics to predictive intelligence that supports proactive decision-making. The development of stream processing frameworks combined with the modern methods of statistics and machine learning has allowed the development of systems that are not only able to process historical data but can also predict future events with growing accuracy and reduce latency. With the

further development of such technologies, the line between monitoring and prediction will become even more blurred and cause the smooth transition between the realms of observation and forecasts across operational scenarios. The ability to master these capabilities offers organizations a remarkable competitive advantage as they are able to move beyond the reactive response to more proactive approaches that mold outcomes instead of reacting to those outcomes. Further evolution of forecasting methods, architecture patterns, and explainability models has the potential to expand these capabilities to new areas and applications where they are likely to revolutionize operational models and strategic planning procedures and provide clear business benefits through improved decision-making and operations.

Funding: This research received no external funding.

Conflicts of Interest: The authors declare no conflict of interest.

Publisher's Note: All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers.

References

- [1] Achala E et al., (2024) Deep Learning Architectures for Time Series Forecasting, ResearchGate Publications, 2024. [Online]. Available: https://www.researchgate.net/publication/385108313_Deep_Learning_Architectures_for_Time_Series_Forecasting
- [2] Ali F and Mohammad K (2025) Synthetic Time Series Forecasting with Transformer Architectures: Extensive Simulation Benchmarks, arXiv:2505.20048, 2025. [Online]. Available: <https://arxiv.org/abs/2505.20048>
- [3] Ashish V et al., (2023) Attention Is All You Need, arXiv:1706.03762, 2023. [Online]. Available: <https://arxiv.org/abs/1706.03762>
- [4] Chaturangi S et al., (2024) A Comprehensive Survey on Rare Event Prediction, arXiv:2309.11356v2, 2024. [Online]. Available: <https://arxiv.org/html/2309.11356v2>
- [5] Fazel M F., Yousefpour N., Chow S. H., and Cassidy M., (2025) Deep Learning for Time Series Forecasting: Review and Applications in Geotechnics and Geosciences, Springer, 2025. [Online]. Available: <https://link.springer.com/article/10.1007/s11831-025-10244-5>
- [6] Haifa T A et al., (2020) Predictive study on time series modeling and comparison with application, ResearchGate, 2020. [Online]. Available: https://www.researchgate.net/publication/344157641_Predictive_study_on_time_series_modeling_and_comparison_with_application
- [7] Md Asif B S, (2025) A systematic review of time series algorithms and analytics in predictive maintenance, *Decision Analytics Journal*, Volume 15, 2025. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S2772662225000293>
- [8] Miira L, (2020) Real-Time Data Analytics for Financial Market Forecasting, ResearchGate Publications, 2020. [Online]. Available: https://www.researchgate.net/publication/390837194_Real-Time_Data_Analytics_for_Financial_Market_Forecasting
- [9] Opeyemi A, (2024) Predictive Analytics in Financial Management: Enhancing Decision-Making and Risk Management, ResearchGate, 2024. [Online]. Available: https://www.researchgate.net/publication/385025548_Predictive_Analytics_in_Financial_Management_Enhancing_Decision-Making_and_Risk_Management
- [10] Paolo G et al., (2024) Explainable Artificial Intelligence methods for financial time series, *Physica A: Statistical Mechanics and its Applications*, Volume 655, 2024. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S037843712400685X>
- [11] Qingsong W et al., (2023) Transformers in Time Series: A Survey, arXiv:2202.07125, 2023. [Online]. Available: <https://arxiv.org/abs/2202.07125>
- [12] Saadia M, (2024) Data-Driven Decision-Making: Leveraging the IoT for Real-Time Sustainability in Organizational Behavior, ResearchGate, 2024. [Online]. Available: https://www.researchgate.net/publication/382549048_Data-Driven_Decision-Making_Leveraging_the_IoT_for_Real-Time_Sustainability_in_Organizational_Behavior
- [13] Sai k C T (2020) Comparative Analysis of Traditional and AI-based Demand Forecasting Models, ResearchGate, 2020. [Online]. Available: https://www.researchgate.net/publication/388443834_Comparative_Analysis_of_Traditional_and_AI-based_Demand_Forecasting_Models
- [14] Srinivasa R A, (2024) A Critical Review of Resource Allocation Optimization in Project Management, ResearchGate, 2024. [Online]. Available: https://www.researchgate.net/publication/382049201_A_Critical_Review_of_Resource_Allocation_Optimization_in_Project_Management