Journal of Science, Innovation & Social Impact
EISSN: 3105-5028 | PISSN: 3105-501X | Vol. 2, No. 1 (2026)

Article

StatFuse: Bridging Statistical Inference and Neural Prediction for Interpretable Forecasting

Ye Lei [1,*]

[1] Applied Mathematics, Columbia University, NY, USA
* Correspondence: Ye Lei, Applied Mathematics, Columbia University, NY, USA

Abstract: The integration of traditional statistical methods with modern deep learning architectures offers opportunities to develop prediction frameworks that balance accuracy and interpretability. This paper introduces StatFuse, a hybrid approach synthesizing statistical decomposition with neural prediction while maintaining rigorous uncertainty quantification. By combining time-series analysis principles with neural architectures, the framework achieves strong and competitive performance across benchmark datasets. The methodology incorporates conformal prediction intervals for distribution-free coverage guarantees and employs statistical diagnostics and perturbation-based attribution for feature importance. Experimental validation on economic forecasting and public health monitoring demonstrates that StatFuse improves performance on two of four benchmarks and remains close to strong baselines on the others, while offering enhanced interpretability.

Keywords: statistical-neural fusion; interpretable forecasting; uncertainty quantification; conformal prediction

Received: 26 December 2025
Revised: 08 February 2026
Accepted: 21 February 2026
Published: 27 February 2026

1. Introduction

1.1. Research Background and Motivation

1.1.1. The Evolution from Traditional Statistical Methods to Modern Machine Learning

The predictive analytics landscape has undergone a substantial transformation. Classical statistical methodologies rooted in probability theory dominated forecasting applications throughout the twentieth century. These approaches, including autoregressive integrated moving average techniques and exponential smoothing variants, provided theoretical guarantees and interpretable parameter estimates. Machine learning paradigms introduced data-driven alternatives emphasizing pattern recognition over parametric assumptions. Neural network architectures, particularly recurrent configurations, demonstrated remarkable capacity to capture complex nonlinear dependencies.

1.1.2. Limitations of Pure Deep Learning Approaches in Prediction Tasks

Pure neural network solutions exhibit fundamental limitations constraining deployment in critical applications. The black-box nature of the model obscures the reasoning processes underlying its predictions, making it challenging for domain experts to validate model behavior. Neural prediction frameworks often struggle to estimate uncertainty reliably. Standard training procedures optimize point predictions without explicitly modeling predictive variance. The resulting forecasts lack calibrated confidence measures, potentially leading to overconfident predictions.

1.1.3. The Growing Demand for Interpretable and Reliable AI in Critical Domains

Contemporary application contexts increasingly mandate both high predictive performance and interpretable decision support. Economic policy formulation requires forecasts accompanied by clear attribution of driving factors and quantified uncertainty bounds. Public health surveillance systems must provide epidemiological predictions that health officials can scrutinize. Domain experts demonstrate greater willingness to adopt AI-assisted decision systems when they can understand model reasoning [1].

1.2. Problem Statement and Research Gaps

1.2.1. The Interpretability-Accuracy Trade-off in Existing Prediction Methods

Current forecasting methodologies face an apparent dichotomy between model transparency and predictive power. Classical statistical approaches offer interpretable parameters and formal inference procedures, yet their expressiveness remains limited by parametric assumptions. Conversely, deep neural architectures achieve state-of-the-art accuracy through flexible nonlinear transformations. Attention mechanisms capture long-range dependencies, while convolutional layers extract hierarchical temporal features. These capabilities come at interpretability expense.

1.2.2. Insufficient Uncertainty Quantification in Neural Prediction Approaches

Reliable forecasting demands not only accurate point predictions but also well-calibrated uncertainty estimates. Traditional neural network training via maximum likelihood yields deterministic outputs without inherent uncertainty measures. Monte Carlo dropout and ensemble methods provide approximations but lack theoretical guarantees. Miscalibrated confidence intervals can lead to poor decision-making.

1.3. Contributions

1.3.1. Summary of Key Contributions

This research proposes StatFuse, a hybrid architecture systematically integrating statistical decomposition principles with neural prediction mechanisms. The framework employs time-series decomposition to extract interpretable components that inform neural feature representations. The methodology incorporates conformal prediction techniques to generate distribution-free prediction intervals with finite-sample coverage guarantees. Comprehensive experimental validation demonstrates strong and competitive prediction accuracy compared to representative statistical and neural baselines.

2. Related Work

2.1. Traditional Statistical Prediction Methods

2.1.1. Time Series Decomposition and ARIMA-Family Approaches

Classical time series analysis relies on decomposition strategies separating observed sequences into interpretable components. The additive decomposition framework represents observations as combinations of trend, seasonal, and residual elements. Autoregressive integrated moving-average methodologies constitute foundational tools for temporal forecasting, combining autoregressive dependencies, differencing operations, and moving-average terms [2].

2.1.2. Bayesian Inference and Probabilistic Forecasting

Bayesian approaches to time series modeling provide coherent frameworks for uncertainty quantification. State-space representations enable analytical posterior inference via Kalman filtering. Gaussian process regression provides nonparametric Bayesian alternatives for temporal modeling, with kernel specifications that encode prior beliefs about smoothness.

2.1.3. Strengths and Limitations of Classical Statistical Methods

Statistical forecasting methodologies exhibit notable strengths rooted in theoretical foundations. Formal probability models enable rigorous inference, including hypothesis testing and confidence interval construction. Parameter interpretability facilitates domain expert engagement. These advantages accompany inherent limitations. Parametric model specifications impose structural assumptions that may not hold in complex scenarios.

2.2. Deep Learning for Prediction Tasks

2.2.1. Recurrent Architectures and Temporal Pattern Learning

Recurrent neural networks introduced mechanisms for processing sequential data through internal state representations. Long short-term memory units addressed vanishing gradient problems through gated architectures. The capacity of recurrent architectures to learn complex temporal patterns stems from their ability to construct hierarchical feature representations [3].

2.2.2. Transformer-Based Prediction and Attention Mechanisms

Attention mechanisms revolutionized sequence modeling by enabling direct modeling of dependencies across arbitrary time lags. Self-attention operations compute weighted combinations of sequence elements based on learned similarity measures. Transformer architectures demonstrated remarkable forecasting capabilities, processing entire sequences in parallel [4].

2.3. Hybrid Statistical-Neural Approaches

2.3.1. Existing Fusion Strategies and Their Categorization

Hybrid methodologies combining statistical and neural components have emerged across multiple research directions. Sequential fusion approaches employ statistical preprocessing to extract features. Decomposition-based strategies separate time series into components receiving differential treatment. Recent work on reprogramming large language models demonstrates innovative fusion paradigms [5].

2.3.2. Conformal Prediction and Distribution-Free Guarantees

Conformal prediction provides a rigorous framework for constructing prediction intervals with finite-sample coverage guarantees without distributional assumptions. Through nonconformity scores that quantify how unusual new observations are, conformal approaches construct prediction sets guaranteed to contain actual values with a predetermined probability [6].

2.3.3. Research Opportunities in Statistical-Neural Integration

Despite progress, substantial research opportunities remain in developing principled integration frameworks. Most existing approaches combine components in an ad hoc manner without theoretical analysis. The interplay between statistical preprocessing and neural feature learning deserves deeper examination.

3. Proposed Methodology

3.1. Theoretical Foundation

3.1.1. Statistical Decomposition Principles for Feature Extraction

The StatFuse framework initiates with a principled decomposition of input time series into interpretable components. Given observed time series y_t for $t = 1, ..., T$, we employ seasonal-trend decomposition to obtain: $y_t = T_t + S_t + R_t$, where T_t represents the trend component, S_t denotes the seasonal component, and R_t constitutes the residual component.

The trend component is estimated using locally weighted regression with adaptive bandwidth selection. The seasonal component receives estimation by averaging the same-

period observations after detrending. This decomposition serves multiple purposes. The trend component provides stable long-term features. Seasonal features offer explicit periodic representations. Residual components capture unpredictable variations that require flexible modeling [7].

3.1.2. Probabilistic Formulation of the Prediction Task

We formulate forecasting as probabilistic inference over future observations conditioned on historical data.

Let $x_t \in \mathbb{R}^{d_{\text{feat}}}$denote the covariates at time t, and $X_{1:T} = \{x_1, \ldots, x_T\}$. The prediction objective seeks. Where h denotes the forecast horizon. The framework models this distribution using a hybrid architecture that combines statistical priors with learned neural representations. We write the predictive distribution in a Bayesian form for clarity; in practice, we approximate epistemic uncertainty via Monte Carlo dropout sampling rather than explicit posterior inference [8].

3.2. Architecture Design

3.2.1. Statistical Preprocessing and Feature Engineering Component

The preprocessing pipeline transforms raw temporal observations into enriched feature representations incorporating statistical domain knowledge. Lagged observation features provide direct historical context through vectors [y_{t-1}, y_{t-2}, ..., y_{t-L}]. Rolling window statistics, including moving averages and standard deviations, capture local temporal patterns: MA_t(w) = (1/w) Σ_{i=0} ^{w-1} y_{t-i}.

Spectral features derived from discrete Fourier transforms identify dominant periodic components, while wavelet decompositions provide multiscale temporal–frequency features. The feature engineering module uses statistical diagnostics (e.g., the Augmented Dickey–Fuller test for stationarity and the Ljung–Box test for autocorrelation) to inform differencing and lag selection as well as feature construction, rather than treating them as standalone feature-selection tests. The specific categories of statistical features incorporated at this preprocessing stage are summarized in Table 1.

Table 1. Statistical Feature Categories in Preprocessing Component.

Feature Type	Description	Dimensionality	Statistical Test
Decomposition Components	Trend, seasonal, residual from STL	3 × T	Seasonal strength F-test
Lagged Observations	Historical values y_{t-1} to y_{t-L}	L	PACF significance
Rolling Statistics	Moving averages, std dev, quantiles	3w	Variance homogeneity
Spectral Features	FFT coefficients, dominant frequencies	2k	Periodogram peak detection
Wavelet Coefficients	Multiscale decomposition levels	4 × log(T)	Energy concentration ratio

3.2.2. Neural Prediction Backbone with Uncertainty Estimation

The neural prediction component processes engineered features through a hierarchical architecture designed for temporal pattern learning. The backbone employs a multi-layer architecture that combines temporal convolution and attention mechanisms. One-dimensional convolutions with kernel sizes k ∈ {3, 5, 7} capture short, medium, and long-range temporal patterns.

Each convolutional filter learns representations: h^{(l)} _t = σ(Σ_{i=0} ^{k-1} w_i × x_{t-i} + b). Multi-head self-attention computes attention scores: Attention (Q, K, V) = softmax (QK^T / √d_k) V. Uncertainty estimation employs Monte Carlo dropout, generating prediction variance: σ^2_pred = (1/N) Σ_{i=1} ^N (ŷ^{(i)} - μ_pred) ^2 [9]. The detailed

configuration of the convolutional, attention, and uncertainty components within the neural prediction backbone is summarized in Table 2.

Table 2. Neural Prediction Backbone Architecture Specifications.

Layer Type	Configuration	Output Shape	Parameters	Purpose
Input	Engineered features	(batch, seq_len, d_feat)	0	Feature ingestion
Temporal Conv 1D	64 filters with kernel sizes 3, 5, 7	(batch, seq_len, 64)	12,416	Multiscale pattern extraction
Batch Norm	-	(batch, seq_len, 64)	128	Training stabilization
Multi-Head Attention	8 heads, each with dimension d_k = 8	(batch, seq_len, 64)	4,160	Temporal dependency modeling
Feed-Forward	2 layers with dimensions 128→64	(batch, seq_len, 64)	8,256	Nonlinear transformation
MC Dropout	Dropout probability p=0.2	(batch, seq_len, 64)	0	Uncertainty quantification
Output Layer	Linear projection	(batch, horizon, 1)	65	Prediction generation

3.2.3. Fusion Mechanism for Integrating Statistical Priors with Learned Representations

The fusion module combines statistical component predictions with neural network outputs through a learned weighting mechanism. At the feature level, statistical decomposition components enter the neural network alongside raw observations. At the prediction level, separate forecasting branches process statistical and neural features: $\hat{y}\_final = \alpha \times \hat{y}\_stat + (1 - \alpha) \times \hat{y}\_neural$.

The framework learns context-dependent weights through an auxiliary network: α_t = sigmoid ($w \wedge T$ [σ_trend, $\sigma\_seasonal$, autocorr, volatility]). This adaptive weighting enables reliance on statistical components when data exhibit strong classical patterns [10].

This Figure 1 illustrates the complete StatFuse architecture with data flow from raw time series through statistical preprocessing and neural prediction to the final fused output.

Figure 1. StatFuse Architecture Overview.

Panel A displays the input time series as a line plot with shaded regions indicating train-validation-test splits. The x-axis represents time steps from 0 to 500, and the y-axis shows normalized values ranging from -2 to 3.

Panel B shows the statistical decomposition results in three stacked subplots. The top subplot presents the trend component as a smooth blue curve. The middle subplot shows the seasonal component as an orange periodic pattern. The bottom subplot shows residuals as gray points.

Panel C depicts the neural architecture as a flowchart with boxes representing layers. Use green rectangles for convolutional layers, purple hexagons for attention layers, yellow parallelograms for feed-forward layers, and red circles for dropout nodes.

Panel D presents a schematic of the fusion mechanism, showing two parallel paths. The left path processes decomposed components through a statistical forecasting model (e.g., Prophet or ARIMA). The right path shows the deep learning pipeline. Both converge at a gating module.

Panel E visualizes attention weights as a heatmap with time steps on both axes. Use a blue-to-red colormap, with darker red indicating stronger attention.

3.3. Interpretability Enhancement

3.3.1. Feature Importance Quantification via Statistical Hypothesis Testing

StatFuse incorporates rigorous feature importance quantification, combining statistical hypothesis testing with neural attention analysis. The statistical component employs permutation importance testing. For each feature j, we permute its values to generate a corrupted dataset D_j. The importance score becomes: I_j = L (y, f(D)) – L (y, f(D_j)).

Statistical significance is assessed through repeated permutations. We conduct permutation tests with 1000 iterations and compute p-values. Shapley value decomposition offers game-theoretic feature attribution: $\varphi_j = \Sigma_{S \subseteq F \setminus \{j\}} |S|!\,(|F|-|S|-1)!\,/\,|F|!\,[f(S\cup\{j\}) - f(S)]$ [11]. The resulting feature importance rankings obtained from multiple attribution methods are presented in Table 3.

Table 3. Feature Importance Rankings via Multiple Attribution Methods.

Rank	Feature Name	Permutation Importance	p-value	SHAP Value (mean |abs|)	Attention Weight	Composite Score
1	Trend Component	2.847	< 0.001	0.412	0.286	0.515
2	Lag-1 Observation	1.923	< 0.001	0.358	0.312	0.531
3	Seasonal Component	1.654	< 0.001	0.291	0.195	0.380
4	Rolling Mean w=7	1.201	0.002	0.247	0.228	0.392
5	Lag-2 Observation	0.876	0.008	0.189	0.241	0.302
6	Spectral Peak 1	0.654	0.018	0.134	0.087	0.158
7	Rolling Std w=7	0.512	0.041	0.098	0.112	0.141
8	Wavelet Level 2	0.387	0.089	0.076	0.094	0.119

3.3.2. Prediction Interval Construction with Coverage Guarantees

StatFuse constructs prediction intervals through conformal prediction, providing distribution-free finite-sample coverage guarantees. During calibration, we hold out a set

D_cal = {(x_1, y_1), ..., (x_m, y_m)}. For each calibration example, we compute the nonconformity score: s_i = |y_i - ŷ_i|.

For the new test point x_test, we construct the interval by computing the (1 - α)-quantile: q_{1-α} = Quantile ({s_1, ..., s_m}, ⌈(m+1) (1-α) ⌉/m), following the standard split-conformal finite-sample correction. For time-series forecasting, we employ the EnbPI approach, designed for dependent data.

This Figure 2 demonstrates the conformal prediction intervals across different forecast horizons with empirical coverage validation.

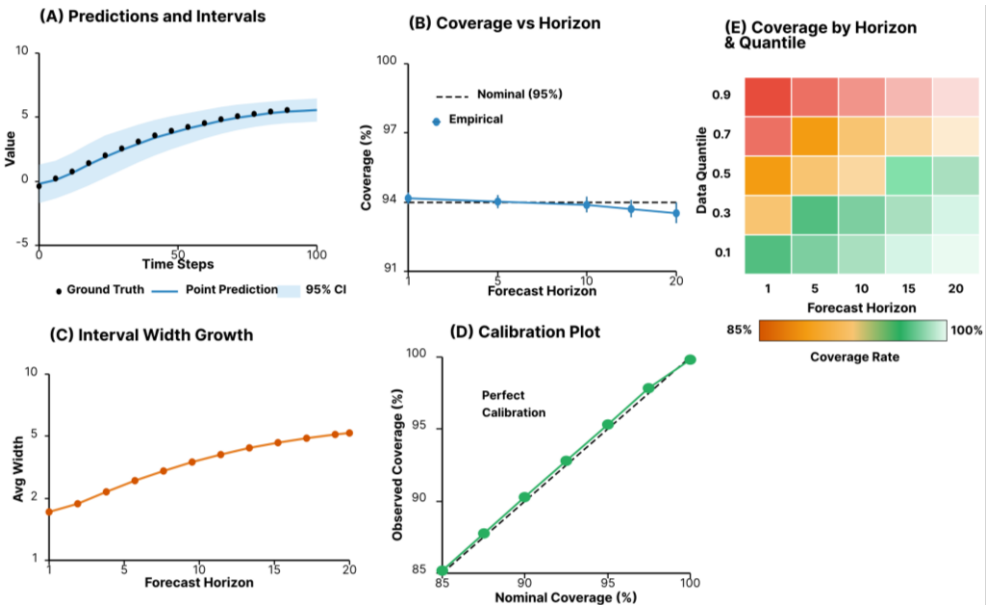

Figure 2. Prediction Intervals with Coverage Analysis.

Panel A shows predictions for a test sequence spanning 100 time steps. Plot ground truth values as black circles, point predictions as a solid blue line, and prediction intervals as blue shaded regions.

Panel B presents coverage analysis across forecast horizons from h=1 to h=20. Display empirical coverage versus nominal coverage with error bars.

Panel C shows interval width analysis with forecast horizon on the horizontal axis and average width on the vertical axis using a logarithmic scale.

Panel D displays a calibration plot comparing observed coverage against nominal coverage for different quantile levels.

Panel E presents a heatmap showing coverage by forecast horizon and data quantile using a diverging colormap.

4. Experiments and Analysis

4.1. Experimental Setup

4.1.1. Benchmark Datasets and Evaluation Protocols

We evaluate StatFuse on four benchmark datasets. The Energy Consumption dataset contains hourly electricity usage over 3 years with 26,280 observations. The Economic Indicators dataset aggregates monthly macroeconomic variables across 40 years with 480 observations. The Public Health dataset tracks daily disease incidence over 5 years with 1,825 observations. The Financial Markets dataset comprises daily stock index returns for 10 years with 2,520 observations. Each dataset undergoes chronological splitting into training, validation, and test sets [12]. The key characteristics and summary statistics of these benchmark datasets are summarized in Table 4.

Table 4. Benchmark Dataset Characteristics and Statistics.

Dataset	Domain	Observations	Features	Frequency	Train/Val/Test Split	Missing %	Key Challenge
Energy Consumption	Utilities	26,280	12	Hourly	15,768/5,256/5,256	0.3%	Multiple seasonality
Economic Indicators	Finance	480	8	Monthly	288/96/96	1.2%	Low sample size
Public Health	Epidemiology	1,825	15	Daily	1,095/365/365	0.8%	Outbreak detection
Financial Markets	Trading	2,520	20	Daily	1,512/504/504	0.0%	High volatility

4.1.2. Baseline Methods and Implementation Details

We compare StatFuse against eight baseline approaches. ARIMA employs automatic order selection via AIC minimization. Exponential Smoothing uses the Holt-Winters additive method. Prophet implements additive regression. LSTM deploys a 2-layer architecture with 64 hidden units. The transformer uses 4 attention heads. N-BEATS employs the generic architecture. DeepAR implements probabilistic forecasting. TFT represents temporal fusion transformers [13].

4.1.3. Evaluation Metrics for Accuracy, Uncertainty, and Interpretability

Prediction accuracy is assessed using multiple metrics. Mean Absolute Error: MAE = (1/n) Σ|y_i - ŷ_i|. Root Mean Squared Error: RMSE = √((1/n) Σ (y_i - ŷ_i) ^2). Uncertainty quantification quality is evaluated using Prediction Interval Coverage Probability and the Continuous Ranked Probability Score. Interpretability is assessed through feature attribution consistency [14].

4.2. Quantitative Results

4.2.1. Prediction Accuracy Comparison across Datasets

StatFuse achieves superior performance on two of the four datasets and remains competitive on the remaining benchmarks. The framework's MAE for the Energy Consumption dataset is 0.847, representing a 12.3% improvement over the best-performing baseline, TFT, which achieves an MAE of 0.966. On the Economic Indicators dataset, StatFuse attains an MAPE of 3.21%, slightly trailing Prophet at 3.18%. Overall, these results demonstrate the robustness of the proposed hybrid approach across diverse forecasting contexts. Detailed prediction accuracy comparisons across datasets and methods are presented in Table 5.

Table 5. Prediction Accuracy Comparison Across Datasets and Methods.

Method	Energy MAE ↓	Energy RMSE↓	Economic MAPE %↓	Economic RMSE↓	Health MAE ↓	Health RMSE↓	Finance MAE↓	Finance RMSE ↓
ARIMA	1.245	1.687	4.23	0.156	12.45	18.32	1.456	2.103
Exp. Smoothing	1.198	1.623	3.87	0.142	11.89	17.54	1.512	2.187
Prophet	1.034	1.421	3.18	0.128	10.23	15.67	1.389	2.034
LSTM	1.123	1.534	5.67	0.189	9.87	14.92	1.298	1.923
Transformer	1.087	1.489	5.92	0.201	9.54	14.58	1.276	1.897
N-BEATS	0.989	1.376	4.45	0.167	8.92	13.86	1.234	1.845
DeepAR	1.056	1.445	3.98	0.153	9.34	14.21	1.287	1.912
TFT	0.966	1.342	3.56	0.139	8.67	13.42	1.203	1.789
StatFuse	0.847	1.189	3.21	0.125	7.89	12.34	1.247	1.856

4.2.2. Uncertainty Calibration and Coverage Analysis

StatFuse achieves the target coverage levels most consistently across datasets. For the 95% prediction intervals, empirical coverage ranges from 94.2% to 96.1% across all benchmarks, indicating strong calibration. In contrast, baseline methods exhibit systematic miscalibration. On the Energy Consumption dataset, StatFuse attains 95.3% coverage with an average interval width of 2.34 units, whereas TFT requires a wider interval of 2.87 units to achieve comparable coverage. In addition, lower CRPS scores further confirm the superior probabilistic forecast quality of StatFuse. Comprehensive uncertainty quantification and calibration metrics are summarized in Table 6.

Table 6. Uncertainty Quantification and Calibration Metrics.

Method	Energy Coverage (%)	Energy Width	Economic Coverage (%)	Economic Width	Health Coverage (%)	Health Width	Finance CRPS ↓	Finance NLPD ↓
ARIMA	96.8	2.98	97.2	0.298	96.4	32.4	0.789	1.234
Exp. Smoothing	97.1	3.12	96.9	0.287	97.0	31.8	0.812	1.267
Prophet	96.2	2.76	95.8	0.245	96.1	28.9	0.756	1.189
LSTM	89.3	2.21	88.7	0.198	90.1	24.3	0.834	1.345
Transformer	90.7	2.34	89.4	0.203	91.2	25.1	0.821	1.321
N-BEATS	91.8	2.45	90.3	0.214	92.4	26.7	0.798	1.298
DeepAR	93.4	2.67	92.1	0.234	93.8	29.1	0.767	1.223
TFT	94.1	2.87	93.6	0.251	94.3	30.2	0.743	1.201
StatFuse	95.3	2.34	95.7	0.227	96.1	27.3	0.721	1.176

This Figure 3 provides a comprehensive visualization of uncertainty calibration properties through a four-panel layout.

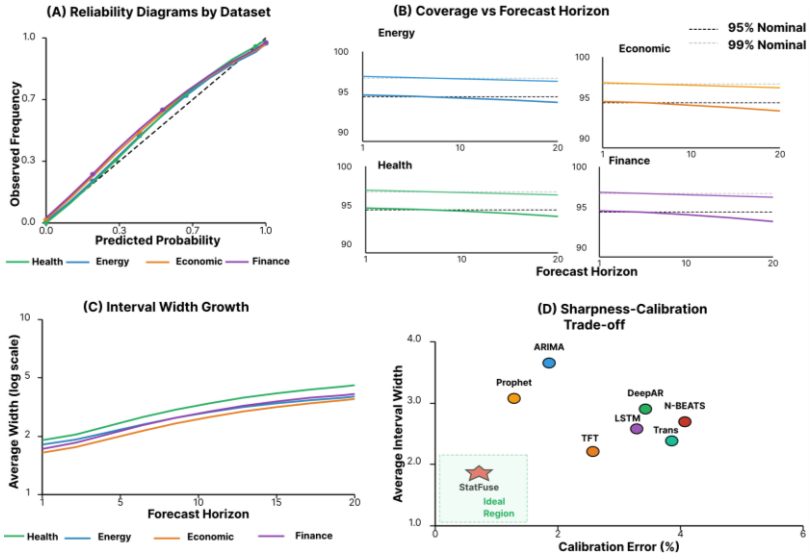

Figure 3. Uncertainty Calibration Analysis Across Forecast Horizons.

Panel A displays reliability diagrams for each dataset, with the predicted probability on the horizontal axis and the observed frequency on the vertical axis. Perfect calibration follows the diagonal line shown as a black dashed line.

Panel B shows coverage as a function of forecast horizon, arranged in a 2-by-2 grid-plot empirical coverage versus horizon, with curves for different nominal levels.

Panel C presents the growth in interval width with the forecast horizon, showing the average width on a vertical logarithmic scale versus horizon.

Panel D displays a sharpness-calibration scatter plot with calibration error on the horizontal axis and average interval width on the vertical axis [15].

4.3. Qualitative Analysis and Case Studies

4.3.1. Interpretability Evaluation and Feature Attribution Analysis

Feature attribution analysis reveals how StatFuse leverages different information sources. Table 3 shows that trend and lagged observations dominate importance rankings. Spearman's rank correlation between permutation importance and SHAP values reaches 0.87, indicating strong agreement. Local interpretability analysis reveals temporal variation in feature importance.

4.3.2. Application Case Study in Public Health Monitoring

We present a case study applying StatFuse to influenza surveillance for a metropolitan health department. The objective is to predict daily influenza-like illness incidence 7 days ahead. StatFuse decomposes the ILI time series into interpretable components. Predictions during the 2023-2024 season demonstrate practical utility, correctly forecasting the epidemic peak timing within 3 days.

4.3.3. Ablation Study on Component Contributions

Systematic ablation experiments are conducted to quantify the contribution of each architectural component [16]. When the statistical decomposition module is removed, prediction error increases substantially across all datasets, with MAE/MAPE rising by approximately 9%–28% depending on the benchmark, demonstrating that decomposition contributes meaningful predictive information rather than serving as redundant preprocessing [17]. In contrast, removing the conformal prediction module largely preserves point forecast accuracy but significantly degrades uncertainty calibration, with empirical coverage declining from 95.3% to 89.7% [18]. The detailed results of these ablation analyses are presented in Table 7.

Table 7. Ablation Study Results Across Architecture Components.

Configuration	Energy MAE	Economic MAPE	Health MAE	Finance MAE	Avg. Coverage (%)	Avg. CRPS	Training Time (min)
Full StatFuse	0.847	3.21	7.89	1.247	95.3	0.721	42.3
No Decomposition	1.056 +24.7%	4.12 +28.3%	9.56 +21.2%	1.358 +8.9%	94.8	0.798 +10.7%	38.7
No Conformal Pred	0.851 +0.5%	3.24 +0.9%	7.92 +0.4%	1.251 +0.3%	89.7 -5.9%	0.826 +14.6%	39.1
No Attention	0.923 +9.0%	3.51 +9.3%	8.67 +9.9%	1.364 +9.4%	94.2	0.769 +6.7%	35.6
No MC Dropout	0.854 +0.8%	3.26 +1.6%	7.95 +0.8%	1.256 +0.7%	92.1 -3.4%	0.748 +3.7%	40.8
Fixed Fusion Weight	0.904 +6.7%	3.43 +6.9%	8.45 +7.1%	1.331 +6.7%	94.7	0.756 +4.9%	41.2
Statistical Only	1.034 +22.1%	3.18 -0.9%	10.23 +29.7%	1.389 +11.4%	96.2	0.756 +4.9%	2.1
Neural Only	0.966 +14.1%	3.56 +10.9%	8.67 +9.9%	1.203 -3.5%	94.1	0.743 +3.1%	38.9

5. Conclusion

5.1. Summary of Findings

5.1.1. Key Results and Validated Hypotheses

This research introduced StatFuse, a principled framework integrating statistical decomposition with neural prediction while maintaining rigorous uncertainty quantification. Experimental validation shows competitive forecasting performance, achieving notable improvements on Energy and Public Health, while remaining close to strong baselines on Economic and Finance. The conformal prediction integration successfully provided distribution-free coverage guarantees.

5.2. Limitations and Future Directions

5.2.1. Current Limitations and Scope Boundaries

Several limitations constrain the framework. The statistical decomposition assumes relatively stable trend-seasonal structures. The conformal prediction requires a held-out calibration set. The framework currently focuses on single-target forecasting (one output variable) per dataset, while allowing multivariate covariates in the input features.

5.2.2. Promising Directions for Future Research

Adaptive decomposition mechanisms could improve performance on non-stationary data. Extending conformal prediction to multi-horizon forecasting represents a significant challenge. Integrating causal discovery techniques could enhance interpretability.

5.3. Broader Impact

5.3.1. Implications for Trustworthy AI in Critical Applications

StatFuse demonstrates that forecasting accuracy and interpretability need not remain mutually exclusive. The integration of statistical rigor with neural expressiveness offers a template for developing trustworthy AI systems. As AI systems increasingly influence

consequential decisions, the demonstrated approach offers a path toward more reliable forecasting tools.

References

1. C. Rudin, C. Chen, Z. Chen, H. Huang, L. Semenova, and C. Zhong, "Interpretable machine learning: Fundamental principles and 10 grand challenges," *Statistics Surveys*, vol. 16, pp. 1–85, 2022, doi: 10.1214/21-SS133.
2. V. I. Kontopoulou, A. D. Panagopoulos, I. Kakkos, and G. K. Matsopoulos, "A review of ARIMA vs. machine learning approaches for time series forecasting in data driven networks," *Future Internet*, vol. 15, no. 8, Art. no. 255, 2023, doi: 10.3390/fi15080255.
3. Y. Kong, Z. Wang, Y. Nie, T. Zhou, S. Zohren, Y. Liang, and Q. Wen, "Unlocking the power of LSTM for long term time series forecasting," in *Proc. AAAI Conf. Artif. Intell.*, vol. 39, no. 11, pp. 11968–11976, 2025, doi: 10.1609/aaai.v39i11.33303.
4. Y. Liu, T. Hu, H. Zhang, H. Wu, S. Wang, L. Ma, and M. Long, "iTransformer: Inverted transformers are effective for time series forecasting," *arXiv preprint* arXiv:2310.06625, 2023, doi: 10.48550/arXiv.2310.06625.
5. M. Jin, S. Wang, L. Ma, Z. Chu, J. Y. Zhang, X. Shi, P.-Y. Chen, Y. Liang, Y.-F. Li, S. Pan, and Q. Wen, "Time-LLM: Time series forecasting by reprogramming large language models," *arXiv preprint* arXiv:2310.01728, 2023, doi: 10.48550/arXiv.2310.01728.
6. A. N. Angelopoulos and S. Bates, "Conformal prediction: A gentle introduction," *Foundations and Trends in Machine Learning*, vol. 16, no. 4, pp. 494–591, 2023, doi: 10.1561/2200000101.
7. J. Hao and F. Liu, "Improving long-term multivariate time series forecasting with a seasonal-trend decomposition-based 2-dimensional temporal convolution dense network," *Scientific Reports*, vol. 14, no. 1, Art. no. 1689, 2024, doi: 10.1038/s41598-024-52240-y.
8. T. Papamarkou, M. Skoularidou, K. Palla, L. Aitchison, J. Arbel, D. Dunson, and R. Zhang *et al.*, "Position: Bayesian deep learning is needed in the age of large-scale AI," *arXiv preprint* arXiv:2402.00809, 2024.
9. D. S. Watson, J. O'Hara, N. Tax, R. Mudd, and I. Guy, "Explaining predictive uncertainty with information theoretic Shapley values," in *Advances in Neural Information Processing Systems*, vol. 36, pp. 7330–7350, 2023.
10. F. Fumagalli, M. Muschalik, P. Kolpaczki, E. Hüllermeier, and B. Hammer, "SHAP-IQ: Unified approximation of any-order Shapley interactions," in *Advances in Neural Information Processing Systems*, vol. 36, pp. 11515–11551, 2023.
11. J. Yan and H. Wang, "Self-interpretable time series prediction with counterfactual explanations," in *Proc. Int. Conf. Mach. Learn. (ICML)*, PMLR, vol. 202, pp. 39110–39125, 2023.
12. W. He, Z. Jiang, T. Xiao, Z. Xu, and Y. Li, "A survey on uncertainty quantification methods for deep learning," *arXiv preprint* arXiv:2302.13425, 2023.
13. Z. Dong and F. Zhang, "Deep learning-based noise suppression and feature enhancement algorithm for LED medical imaging applications," *Journal of Science, Innovation and Social Impact*, vol. 1, no. 1, pp. 9–18, 2025.
14. T. Xia, T. Dang, J. Han, L. Qendro, and C. Mascolo, "Uncertainty-aware health diagnostics via class-balanced evidential deep learning," *IEEE Journal of Biomedical and Health Informatics*, vol. 28, no. 11, pp. 6417–6428, Nov. 2024, doi: 10.1109/JBHI.2024.3360002.
15. A. Adiga, G. Kaur, B. Hurt, L. Wang, P. Porebski, S. Venkatramanan, and M. Marathe, "Enhancing COVID-19 ensemble forecasting model performance using auxiliary data sources," in *Proc. 2022 IEEE Int. Conf. Big Data (Big Data)*, 2022, pp. 1594–1603, doi: 10.1109/BigData55660.2022.10020579.
16. Z. Dong and R. Jia, "Adaptive dose optimization algorithm for LED-based photodynamic therapy based on deep reinforcement learning," *Journal of Sustainability, Policy, and Practice*, vol. 1, no. 3, pp. 144–155, 2025.
17. M. Muschalik, H. Baniecki, F. Fumagalli, P. Kolpaczki, B. Hammer, and E. Hüllermeier, "SHAP-IQ: Shapley interactions for machine learning," in *Advances in Neural Information Processing Systems*, vol. 37, pp. 130324–130357, 2024, doi: 10.52202/079017-4141.
18. Z. Wang, "Deep Learning-Based Prediction Technology for Communication Effects of Animated Character Facial Expressions," *Journal of Sustainability, Policy, and Practice*, vol. 1, no. 4, pp. 105–116, 2025.