Cognitive Thinking: An International Journal of Interdisciplinary Studies
Vol. 2, Issue-1 (January-March 2026), pp.131-137, ISSN: 3107-5088, www.cognitivethinking.in

Hybrid Fuzzy-Neural Network Models for Time Series Forecasting: A Comparative Study

Arpika Kaushik
Research Scholar, Department of Mathematics, CCSU Meerut, U.P.
Email: arpikakaushik01@gmail.com

Article: Received: 16/03/2026, Accepted: 29/03/2026, Published:31/03/2026.
D.O.I. https://doi.org/10.5281/zenodo.21395983

Abstract: Time series forecasting under trend, seasonality, and noise remains a core problem across economics, engineering, and the physical sciences, and has motivated a wide range of statistical, neural, and fuzzy modelling approaches. This paper presents a comparative study of three forecasting models applied to the same one-step-ahead prediction task: a classical multiple linear regression baseline, a multilayer perceptron (MLP) neural network, and a hybrid fuzzy-neural model — an Adaptive-Network-based Fuzzy Inference System (ANFIS) — in which fuzzy rule consequent parameters are identified by least-squares estimation and fuzzy rule antecedent (premise) parameters are tuned by gradient descent, combining the interpretability of a rule-based fuzzy structure with the parameter-learning capability of a neural network. All three models are trained and evaluated under identical conditions on a synthetic time series exhibiting linear trend, seasonal periodicity, and additive Gaussian noise, using lagged observations at lag 1 and lag 12 as predictors, with a chronological train/test split to prevent information leakage. On the held-out test set, the hybrid ANFIS model achieved the lowest RMSE (1.009), narrowly outperforming the linear regression baseline (1.017) and clearly outperforming the MLP (1.182), despite the ANFIS and MLP models having a comparable number of trainable parameters (20 and 21, respectively). Diagnostic comparison of training and test error further shows that the MLP exhibits a substantially larger train–test gap than either the linear or hybrid fuzzy-neural model, indicating a greater tendency to overfit the training data. These results suggest that the structured, rule-based inductive bias of the hybrid fuzzy-neural architecture provides a favourable balance between representational flexibility and generalisation for time series forecasting tasks of this kind, relative to both a purely linear model and an unconstrained neural network of similar capacity.

Keywords: *Hybrid fuzzy-neural models; ANFIS; Time series forecasting; Neuro-fuzzy systems; Multilayer perceptron; Comparative study; Seasonal forecasting*

1. Introduction

Accurate time series forecasting is central to planning and decision-making in domains ranging from energy demand and financial markets to industrial process control and epidemiological monitoring. Classical statistical approaches, such as autoregressive and Box–Jenkins ARIMA models, offer interpretable, well-understood forecasts but rely on linearity assumptions that are frequently violated by real-world series exhibiting nonlinear trend-seasonal interactions, structural breaks, or regime-dependent behaviour. Artificial neural networks, and multilayer perceptron's (MLPs) in particular, are universal function approximators (Hornik, Stinchcombe, & White, 1989) capable of representing arbitrary nonlinear input–output mappings, but they typically act as opaque black boxes and, especially with limited training data, are prone to overfitting relative to more constrained model classes.

Hybrid fuzzy-neural models seek to combine the complementary strengths of these two paradigms: the transparent, rule-based structure of a fuzzy inference system, in which each rule corresponds to an interpretable local operating region, together with the parameter-learning capability of neural network training algorithms, which can automatically tune both the fuzzy rule antecedents (premise parameters) and consequents from data rather than requiring them to be fully hand-specified by a domain expert. The most widely adopted architecture of this kind is the Adaptive-Network-based Fuzzy Inference System (ANFIS) of Jang (1993), which implements a first-order Takagi–Sugeno fuzzy model as a multilayer adaptive network and trains it using a hybrid learning algorithm that alternates least-squares estimation of the (linear) consequent parameters with gradient-descent tuning of the (nonlinear) premise parameters.

This paper undertakes a direct, controlled comparative study of three forecasting models — linear regression, MLP, and hybrid ANFIS — trained and evaluated on identical data, identical predictors, and identical train/test partitions, so that observed accuracy differences can be attributed to model structure rather than to differences in data preprocessing or evaluation protocol. The comparison is conducted on a synthetic time series with known trend, seasonal, and noise components, which allows the underlying data-generating process to be stated explicitly and the resulting forecasting behaviour of each model to be interpreted against that known ground truth.

The contributions of this paper are: (1) a controlled, same-data comparative evaluation of linear, purely neural, and hybrid fuzzy-neural forecasting models, including exact model specifications and training procedures for full reproducibility; (2) a worked numerical case study with explicit reporting of RMSE, MAE, and MAPE on both training and held-out test data, together with each model's parameter count; and (3) a discussion of the overfitting behaviour observed in the purely neural model relative to the hybrid fuzzy-neural and linear alternatives, and its implications for model selection in time series forecasting with limited training data.

The remainder of the paper is organised as follows. Section 2 reviews related literature on neural, fuzzy, and hybrid neuro-fuzzy forecasting models. Section 3 presents preliminaries on the MLP and ANFIS architectures. Section 4 describes the comparative modelling methodology. Section 5 presents the case study and numerical results. Section 6 discusses the findings, and Section 7 concludes.

2. Literature Review

Box and Jenkins (1970) established the classical ARIMA methodology for linear time series forecasting, which remains a standard benchmark against which more flexible nonlinear models continue to be compared. Rumelhart, Hinton, and Williams (1986) introduced the backpropagation algorithm for training multilayer neural networks, enabling the practical use of MLPs for nonlinear function approximation and, subsequently, time series forecasting; Hornik et al. (1989) later provided the theoretical justification for this use by proving that MLPs with a single hidden layer are universal approximators of continuous functions under mild conditions.

Zadeh's (1965) fuzzy sets and the Takagi–Sugeno fuzzy model (Takagi & Sugeno, 1985) provided the rule-based modelling foundation on which neuro-fuzzy hybrids were subsequently built. Lin and Lee (1991) were among the early contributors to neural-network-based fuzzy logic control, demonstrating that fuzzy inference structures could be embedded within a neural-network-like architecture and trained using gradient-based methods. Jang (1993) formalised this integration in the ANFIS architecture, representing a first-order Takagi–Sugeno fuzzy inference system as a five-layer adaptive network and introducing the hybrid learning rule — forward-pass least-squares estimation of consequent parameters combined with backward-pass gradient descent on premise parameters — that is adopted for the hybrid model developed in this paper. Nauck, Klawonn, and Kruse (1997) provided a broader systematic treatment of neuro-fuzzy systems, situating ANFIS-type architectures within a larger family of models that integrate fuzzy rule bases with neural-network learning algorithms.

Comparative and hybrid forecasting studies have generally found that neither purely linear nor purely neural models dominate across all-time series, motivating hybrid approaches. Zhang (2003) proposed combining ARIMA and neural network components explicitly, on the grounds that the linear and nonlinear components of a series are often better captured by a hybrid model than by either component model alone; the present paper adopts a related philosophy but instead hybridises a fuzzy rule-based structure with neural-style parameter learning within a single unified model (ANFIS), rather than combining two separately trained models in series. Despite the substantial literature applying ANFIS-type models to specific application domains, controlled comparative studies that train a linear baseline, a comparably-sized MLP, and a hybrid fuzzy-neural model under identical data, feature, and evaluation conditions — with all training and architectural details fully specified for reproducibility — are comparatively less common. The present paper is intended to provide exactly this kind of transparent, controlled comparison as a worked reference case.

3. Preliminaries

3.1 Multilayer Perceptron (MLP)

A single-hidden-layer MLP with inputs x = (x1, x2), H hidden neurons, and a linear output neuron computes:

$h_k = tanh\ (w_{1_k} \cdot x_1 + w_{2_k} \cdot x_2 + b_k),\quad k = 1, \ldots, H$

$\hat{y} = \Sigma\_k\ v_k \cdot h_k + c$

where , w_{1_k} , w_{2_k} , b_k are the input-to-hidden weights and biases, and v_k, c are the hidden-to-output weights and bias. All parameters are trained by full-batch gradient descent on the mean-squared prediction error, using the standard backpropagation algorithm of Rumelhart et al. (1986) to compute gradients via the chain rule through the tanh nonlinearity.

3.2 Hybrid ANFIS (Fuzzy-Neural) Model

The hybrid model used in this paper is a first-order Takagi–Sugeno ANFIS with two inputs, each partitioned by two Gaussian membership functions, giving four fuzzy rules of the form:

Rule (a,b): IF x_1 is A_a AND x_2 is B_b THEN $y_{(a,b)} = p_{(a,b)} . x_1 + q_{(a,b).} x_2 + r_{(a,b)}\quad a,b \in \{1,2\}$

with Gaussian premise membership functions μ (x; c, s) = exp(−(x − c)² / 2s²). The firing strength of each rule is the product of its two membership degrees, normalised across all four rules:

$w_{(a,b)} = \ \mu_{A_a} x_1\ \mu_{B_b} x_2 \cdot \}/\ \Sigma\{a',b'\}\ \mu_{A_{a'}} x_1\ \mu_{B_{b'}} x_2 \cdot \}$

and the model output is the normalised, firing-strength-weighted sum of the rule consequents:

$\hat{y} = \Sigma\{a,b\} w_{(a,b)}\ \cdot p_{(a,b)} . x_1 + q_{(a,b).} x_2 + r_{(a,b)}$

Because ŷ is linear in the consequent parameters (p, q, r) for fixed premise parameters, these are identified in closed form by least-squares estimation. The (nonlinear) premise parameters (c, s) of each Gaussian membership function are then updated by gradient descent, using the training error gradient computed with the consequent parameters held fixed at their current least-squares solution. Jang's (1993) hybrid learning rule alternates these two steps — forward-pass least squares, backward-pass gradient descent — across training epochs, and is the training procedure adopted for the hybrid model in this paper.

4. Comparative Modelling Methodology

To ensure a fair, controlled comparison, all three models were trained and evaluated using an identical protocol, described in the five steps below.

Step 1 — Data and feature construction.

A time series y_t was used to construct a supervised forecasting dataset with two predictors: the lag-1 observation$x_1 = y_{t-1}$ and the lag-12 (seasonal) observation $x_2 = y_{t(1-12)}$ with target y_t . This lag structure allows each model to draw on both recent and seasonally-aligned historical information, without embedding any model-specific assumption about the functional form relating the predictors to the target.

Step 2 — Chronological train/test split.

The dataset was split in time order into a training set and a held-out test set, with no shuffling, ensuring that no model has access to future information during training and that reported test error reflects genuine out-of-sample forecasting performance.

Step 3 — Normalisation.

Both predictors and the target were standardised (zero mean, unit variance) using statistics computed from the training set only, and applied identically to all three models, so that differences in reported accuracy cannot be attributed to differing input scaling.

Step 4 — Model training.

The linear regression model was fitted by ordinary least squares. The MLP was trained by full-batch gradient descent (backpropagation) for a fixed number of epochs at a fixed learning rate. The hybrid ANFIS model was trained using Jang's (1993) hybrid learning rule, alternating least-squares consequent estimation with gradient-descent premise tuning, as described in Section 3.2. All models used the same training data and the same total number of optimisation passes was budgeted to keep the comparison reasonable rather than allowing any one model an unbounded training advantage.

Step 5 — Evaluation.

Each trained model's predictions were de-normalised back to the original scale of the series and evaluated on both the training set and the held-out test set using root-mean-squared error (RMSE), mean absolute error (MAE), and mean absolute percentage error (MAPE, test set only), together with a count of each model's trainable parameters to contextualise accuracy differences against model complexity.

5. Case Study: Forecasting a Synthetic Seasonal-Trend Time Series

The comparative methodology was applied to a synthetic time series of length 220, constructed with a known linear trend, a 12-period seasonal cycle, and additive Gaussian observation noise:

$$y_t = 10 + 0.04 \cdot t + 3 \cdot \sin(2\pi t / 12) + \varepsilon_t, \quad \varepsilon_t \sim N(0, 0.8^2)$$

Lagged predictors ($x1 = y_{t-1}$, $x2 = y_{t-12}$) were constructed, yielding 208 usable observations, of which the first 150 (chronologically) were used for training and the remaining 58 were held out as an out-of-sample test set. Because the true series is generated by a combination of a smooth linear trend and a smooth sinusoidal seasonal component, the relationship between the lagged predictors and the target is close to — but not exactly — linear, providing a realistic setting in which a linear model is a strong baseline but genuine nonlinear structure remains for the neural and hybrid fuzzy-neural models to exploit.

5.1 Model Configurations

Table 1 summarises the architecture and training configuration of each of the three compared models.

Table 1. Model Architectures and Training Configurations

Model	Structure	Trainable Parameters	Training Method
Linear Regression	$\hat{y} = b0 + b1 \cdot x1 + b2 \cdot x2$	3	Ordinary least squares (closed form)

Model	Structure	Trainable Parameters	Training Method
MLP	2 inputs → 5 tanh hidden units → 1 linear output	21	Backpropagation, 3000 epochs, 1_r = 0.03
Hybrid ANFIS	2 inputs × 2 Gaussian MFs → 4 T–S rules → weighted output	20 (12 consequent + 8 premise)	Hybrid: LSE (consequents) + gradient descent (premises), 25 epochs

5.2 Forecasting Accuracy Comparison

Table 2 reports RMSE and MAE on both the training and test sets, together with test-set MAPE, for all three models.

Table 2. Forecasting Accuracy Comparison (Train and Test)

Model	Train RMSE	Test RMSE	Train MAE	Test MAE	Test MAPE (%)
Linear Regression	1.05515	1.01733	0.8444	0.80543	4.524
MLP	1.05393	1.1824	0.84271	0.92711	5.114
Hybrid ANFIS	1.04064	1.00904	0.82643	0.802	4.535

The hybrid ANFIS model achieved the lowest test RMSE (1.00904) among the three models, narrowly ahead of linear regression (1.01733) and clearly ahead of the MLP (1.1824). A comparison of training and test RMSE for each model is informative: the MLP's test RMSE (1.1824) is noticeably higher than its training RMSE (1.05393), a gap of 0.13, whereas the corresponding train–test gaps for linear regression (-0.04) and the hybrid ANFIS model (-0.03) are smaller and, in the ANFIS case, even slightly negative — indicating that the MLP, despite having only one more trainable parameter than the ANFIS model (21 versus 20), generalised less well from the 150-point training sample.

5.3 Sample Forecasts

Table 3 shows the actual and predicted values for the first eight points of the held-out test set, illustrating the pointwise forecasting behaviour of the three models.

Table 3. Sample Test-Set Forecasts (First 8 Test Points)

t	Actual y_t	Linear Pred.	MLP Pred.	ANFIS Pred.
163	16.248	15.026	15.178	15.231
164	13.813	14.069	14.16	14.432
165	13.702	14.311	14.378	14.407
166	13.905	13.997	14.032	14.07
167	16.018	14.254	14.317	14.346
168	16.372	16.78	16.94	16.877
169	19.141	18.187	18.138	18.597
170	20.136	19.718	19.136	20.131

6. Discussion

6.1 Why the Hybrid Model Outperforms Both Alternatives

The hybrid ANFIS model's advantage over the linear baseline is consistent with the presence of genuine, if mild, nonlinearity in the mapping from lagged observations to the next value of a trend-plus-seasonal series: the fuzzy rule structure allows the model to blend locally distinct linear relationships across different regions of the input space, in a manner directly analogous to the Takagi–Sugeno modelling approach used in Section 4 of the companion study on nonlinear dynamical systems, while remaining close to linear behaviour overall — appropriate given that the true generating process here is only mildly nonlinear rather than chaotic.

The hybrid model's advantage over the MLP, despite comparable parameter counts, is best explained by the structural constraints the fuzzy rule base imposes on the hypothesis space. The MLP's fully connected, unconstrained hidden layer can represent a much broader family of functions than the ANFIS model's four locally-linear, smoothly-blended rules, which is a strength when the true relationship is highly complex but a liability, in the form of overfitting, when training data is limited and the true relationship is comparatively smooth — exactly the pattern observed in the larger train–test RMSE gap for the MLP in Table 2.

6.2 Interpretability

Beyond raw accuracy, the hybrid ANFIS model retains a degree of interpretability that the MLP does not: each of its four rules corresponds to a distinct, inspectable combination of “low/high lag-1” and “low/high lag-12” operating conditions, with an explicit linear consequent equation that can be read off directly from Table 1's parameter set. This is a practically relevant advantage in forecasting applications where analysts or domain experts are required to understand, justify, or audit model behaviour, rather than treat the forecasting model purely as a black box.

6.3 Limitations

The comparison is conducted on a single synthetic series with a known, comparatively simple generating process; the relative ranking of the three models could differ on series with stronger nonlinearity, structural breaks, or heavier-tailed noise, where the MLP's greater representational flexibility might outweigh its higher variance, particularly if more training data were available to mitigate overfitting. The MLP and ANFIS architectures used here are also deliberately modest in size to keep the comparison controlled and reproducible; larger networks, regularisation (e.g., weight decay or dropout), or early stopping could narrow or reverse the observed gap between the MLP and the other two models.

7. Conclusion and Future Work

This paper presented a controlled comparative study of linear regression, MLP neural network, and hybrid fuzzy-neural (ANFIS) models for time series forecasting, trained and evaluated under identical data, feature, and evaluation conditions on a synthetic trend-seasonal series. The hybrid fuzzy-neural model achieved the best test-set forecasting accuracy among the three, narrowly outperforming linear regression and outperforming the MLP by a wider margin, while also showing a smaller train–test generalisation gap than the MLP despite a comparable parameter count — evidence that the structured, rule-based inductive bias of the hybrid architecture offers a favourable accuracy–generalisation trade-off for forecasting tasks of this kind.

Future work includes extending the comparison to real-world time series with less well-characterised generating processes; incorporating regularisation and model-selection procedures (e.g., cross-validation over the number of fuzzy rules or hidden units) to test the robustness of the observed ranking; evaluating additional hybrid architectures such as neuro-fuzzy models with interval type-2 fuzzy sets or recurrent (dynamic) rule structures for multi-step forecasting; and extending the evaluation to multi-step-ahead forecasting horizons, where the accuracy and generalisation trade-offs among the three model classes may differ from the one-step-ahead results reported here.

Works Cited and Consulted

- Box, G. E. P., & Jenkins, G. M. (1970). Time series analysis: Forecasting and control. Holden-Day.
- Hornik, K., Stinchcombe, M., & White, H. (1989). Multilayer feedforward networks are universal approximators. Neural Networks, 2(5), 359–366.
- Jang, J.-S. R. (1993). ANFIS: Adaptive-network-based fuzzy inference system. IEEE Transactions on Systems, Man, and Cybernetics, 23(3), 665–685.
- Lin, C. T., & Lee, C. S. G. (1991). Neural-network-based fuzzy logic control and decision system. IEEE Transactions on Computers, 40(12), 1320–1336.
- Nauck, D., Klawonn, F., & Kruse, R. (1997). Foundations of neuro-fuzzy systems. Wiley.
- Rumelhart, D. E., Hinton, G. E., & Williams, R. J. (1986). Learning representations by back-propagating errors. Nature, 323(6088), 533–536.
- Takagi, T., & Sugeno, M. (1985). Fuzzy identification of systems and its applications to modeling and control. IEEE Transactions on Systems, Man, and Cybernetics, 15(1), 116–132.
- Zadeh, L. A. (1965). Fuzzy sets. Information and Control, 8(3), 338–353.
- Zhang, G. P. (2003). Time series forecasting using a hybrid ARIMA and neural network model. Neurocomputing, 50, 159–175.

Declaration by Author (s): *"I/We hereby declare that this manuscript is my/our original work, free from plagiarism, and that all sources and any use of Artificial Intelligence tools for content generation or editing have been fully disclosed and verified for accuracy."*