

MHBS 2026 ICSRI Fellows: Research Proposal

1 Non-Stationary Spectral Decomposition Network (NS-SDN) For Econometric Forecasting

Nikhil Sunder

Student | ICSRI Research Fellow | Open-Source Developer

B.S.B.A. Quantitative Economics, Finance & Minor in Math

University of Miami Herbert Business School

email: nss106@miami.edu

phone: (305) 409-3108

links: [Github](#) | [PyPI](#) | [Anaconda](#) | [LinkedIn](#) | [ORCID](#) | [Zenodo](#) | [Handshake](#)

2 Introduction

Recent work in neural implicit representations has shown that sinusoidal nonlinearities can represent highly oscillatory structure with strong gradient flow properties (Sitzmann et al., 2020). In particular, Sinusoidal Representation Networks (SIREN) replace conventional activations (e.g., sigmoid or tanh) with $\sin(\cdot)$, enabling compact representation of periodic structure. This is well-aligned with domains such as signal processing and physics-informed modeling, where harmonic structure is intrinsic.

This proposal asks whether sinusoidal function representations can be adapted to econometric time series forecasting, where observables often exhibit both secular trend and cyclical variation, and where cycles may be nonstationary (time-varying amplitude, time-varying frequency, phase shifts after shocks) (Huang et al., 1998; Lubik & Matthes, 2015). The goal is a forecasting architecture that explicitly decomposes a series into interpretable components and adapts these components over time via a latent state.

3 Design Concept

3.1 Econometric decomposition target

Assume an observed scalar time series $y(t)$ admits an additive decomposition into trend, cyclical structure, and idiosyncratic error (Harvey, 1989; Beveridge & Nelson, 1981):

$$y(t) = T(t) + C(t) + \epsilon(t)$$

Rather than imposing fixed-frequency stochastic cycles (Harvey, 1989), NS-SDN models allow time-varying frequency and amplitude consistent with the Hilbert–Huang framework (Huang et al., 1998).

4 Network Overview

4.1 From stationary sinusoid to nonstationary spectral sum

A single stationary sinusoid is inadequate for macro/financial data because amplitude/frequency are rarely constant. NS-SDN therefore models the signal as:

$$f(t) = B(t) + \sum_{k=1}^K A_k(t) \sin(\theta_k(t) + \varphi_k(t)),$$

Where: - $B(t)$ is the trend/low-frequency component, - $A_k(t)$ is the (positive) time-varying amplitude (“envelope”) of component k , - $\theta_k(t)$ is the cumulative phase of component k , - $\varphi_k(t)$ is a time-varying phase offset (“horizontal shift”) that allows alignment changes without forcing distortion into amplitude/frequency.

The defining feature is that $B(t)$, $A_k(t)$, $\theta_k(t)$, $\varphi_k(t)$ are functions of discrete time through a learned latent state.

5 Discrete-Time State-Space NS-SDN

Let the time grid be $t_0 < t_1 < \dots < t_N$, with $\Delta t_n := t_n - t_{n-1}$ for $n \geq 1$. Let the observed data be $y_n := y(t_n)$ and the model output be $\hat{y}_n := f(t_n)$.

5.1 Latent state with transition equation

NS-SDN is formulated as a nonlinear state-space model (Durbin & Koopman, 2012), with time-varying parameters analogous to TVP-VAR specifications (Lubik & Matthes, 2015):

State transition

$$h_n = F_\psi(h_{n-1}, x_n, \Delta t_n) \quad h_n \in \mathbb{R}^d,$$

where x_n is the information available at time t_n (e.g., $x_n = y_n$ for univariate autoregressive forecasting, or a vector of predictors in the multivariate case). This transition can be instantiated as a GRU/LSTM cell, or as a time-aware update of the form $h_n = h_{n-1} + \Delta t_n \cdot g_\psi(h_{n-1}, x_n)$, which is a discrete-time analog of continuous-time dynamics.

Emission (spectral synthesis)

Given h_n , the model emits $\hat{y}_n$ via nonstationary spectral synthesis.

5.2 Emission components

Trend component

$$B_n := B(t_n) = w_B^\top h_n + b_B.$$

Amplitude (positive envelope)

$$A_n = \text{softplus}(W_A h_n + b_A) \in \mathbb{R}_{>0}^K, \quad A_{n,k} = \text{softplus}(w_{A,k}^\top h_n + b_{A,k}).$$

Instantaneous angular frequency

$$\omega_n = \omega_{\text{base}} + W_\omega h_n + b_\omega \in \mathbb{R}^K, \quad \omega_{n,k} = \omega_{\text{base},k} + w_{\omega,k}^\top h_n + b_{\omega,k}.$$

(Design note: ω_n may be left unconstrained in $\mathbb{R}^K$, or constrained positive via softplus if interpretability/identifiability is desired. This update mirrors the definition of instantaneous frequency in the Hilbert–Huang framework (Huang et al., 1998).)

Cumulative phase (discrete-time integration)

$$\theta_{0,k} \text{ learned, } \quad \theta_{n,k} = \theta_{n-1,k} + \omega_{n,k} \Delta t_n \quad (n \geq 1).$$

Phase offset (horizontal shift from state)

$$\varphi_n = W_\varphi h_n + b_\varphi \in \mathbb{R}^K, \quad \varphi_{n,k} = w_{\varphi,k}^\top h_n + b_{\varphi,k}.$$

5.3 Final NS-SDN output equation

The emission is:

$$\hat{y}_n = B_n + \sum_{k=1}^K A_{n,k} \sin(\theta_{n,k} + \varphi_{n,k})$$

Optionally, for numerical symmetry one may use the trigonometric identity:

$$\sin(\theta + \varphi) = \sin(\theta) \cos(\varphi) + \cos(\theta) \sin(\varphi),$$

so:

$$\hat{y}_n = B_n + \sum_{k=1}^K A_{n,k} \left[\sin(\theta_{n,k}) \cos(\varphi_{n,k}) + \cos(\theta_{n,k}) \sin(\varphi_{n,k}) \right].$$

6 Interpretation and positioning

The resulting architecture is conceptually related to FM synthesis in signal processing (Chowning, 1973) and also closely related to instantaneous frequency representations (Huang et al., 1998) and neural oscillator architectures (Chen et al., 2018). Conversely, NS-SDN differs in two critical ways:

1. **Nonstationary envelopes and frequencies are state-driven:** A_n and ω_n evolve with h_n , allowing time-varying cycles rather than fixed Fourier bases.
2. **Econometric state-space structure:** The model explicitly separates latent dynamics h_n (transition equation) from measurement construction $\hat{y}_n$ (spectral emission), paralleling classical trend-cycle state-space decomposition (Harvey, 1989; Durbin & Koopman, 2012) while allowing nonlinear, data-driven evolution.

For these reasons, the architecture is naturally described as a Non-Stationary Spectral Decomposition Network (NS-SDN).

7 Econometric Application

7.1 Targeted use case

The proposed NS-SDN architecture is designed for macroeconomic and financial time series that exhibit secular trend, evolving cycle durations and volatility (Hamilton, 1989; Lubik & Matthes, 2015), where the amplitude, frequency, and phase of cycles may evolve over time due to regime changes, policy shocks, or structural breaks.

Canonical examples include: - real GDP and output gap measures - inflation and inflation expectations - interest rates and yield spreads - unemployment and labor market indicators - commodity prices and real exchange rates

These series are known to exhibit: 1. persistent low-frequency trends, 2. medium-frequency business-cycle dynamics, 3. episodic nonstationarity following shocks (e.g. monetary tightening, supply disruptions).

Traditional Linear constant-parameter models may fail to capture evolving macroeconomic dynamics (Lubik & Matthes, 2015).

7.2 Mapping NS-SDN to econometric structure

NS-SDN can be interpreted as a nonlinear, nonstationary unobserved-components model, where: - the latent state h_n plays the role of an unobserved economic state, - the trend component B_n corresponds to a stochastic trend, - the sinusoidal components correspond to stochastic cycles with time-varying amplitude and frequency, - the observation equation is explicitly additive and interpretable.

Unlike classical unobserved components models or Kalman-filter-based stochastic cycle models, NS-SDN: - does not impose fixed frequencies, - does not assume linear Gaussian dynamics, - allows regime-dependent evolution of spectral structure.

This makes NS-SDN particularly suitable for environments where cyclical behavior itself is nonstationary, such as post-crisis macroeconomic regimes.

7.3 Forecasting task formulation

Let $\{y_1, \dots, y_T\}$ denote an observed time series. The forecasting task is defined as predicting future values:

$$\hat{y}_{T+1}, \hat{y}_{T+2}, \dots, \hat{y}_{T+H}$$

given information available up to time T .

NS-SDN is trained in a recursive one-step-ahead forecasting framework, where: - the latent state h_n is updated sequentially, - predictions are generated via the spectral emission equation, - multi-step forecasts are generated recursively, consistent with standard practice in neural autoregressive forecasting (Salinas et al., 2020).

This setup mirrors the forecasting protocols used for VARs, state-space models, and neural sequence models, enabling fair comparison.

8 Evaluation and Forecasting Efficacy

8.1 Benchmark models

To evaluate the forecasting performance of NS-SDN, it will be compared against a set of widely used econometric and machine learning benchmarks, chosen to span increasing levels of model flexibility:

Linear econometric models - AR(p) - ARIMA(p,d,q) - VAR (for multivariate extensions)

State-space and decomposition models - local-level / local-trend models - unobserved components models with stochastic cycles - Kalman-filter-based trend-cycle decompositions

Neural network baselines - feedforward neural networks (MLP) - recurrent neural networks (LSTM / GRU)

These benchmarks are well understood, commonly used in applied econometrics, and provide a meaningful reference for assessing gains from nonstationary spectral modeling.

8.2 Training and evaluation protocol

To ensure comparability across models, all methods will be evaluated under a rolling-origin out-of-sample forecasting design standard in forecast comparison (West, 1996): 1. Split the data into an initial estimation window and an evaluation window. 2. Estimate each model using data up to time t . 3. Generate forecasts for horizons $h = 1, 4, 8, 12$ (depending on data frequency). 4. Roll the estimation window forward and repeat.

This procedure mimics real-time forecasting and avoids look-ahead bias.

8.3 Forecast accuracy metrics

Forecast performance will be assessed using standard loss functions: - Mean Squared Error (MSE) - Root Mean Squared Error (RMSE) - Mean Absolute Error (MAE)

For each horizon h , loss is computed as: $\text{MSE}(h) = \frac{1}{N} \sum_{i=1}^N (y_{t_i+h} - \hat{y}_{t_i+h})^2$.

Results will be reported both per horizon and aggregated across horizons, enabling comparison of short-term versus medium-term forecasting ability.

8.4 Statistical comparison of forecasting performance

To formally assess whether NS-SDN delivers statistically significant improvements, forecast errors will be compared using standard forecast comparison tests, such as: - Diebold–Mariano tests for equal predictive accuracy (Diebold & Mariano, 1995) - horizon-specific loss differentials - cumulative squared forecast error plots

This ensures that any observed improvements are not due to sampling variability alone.

8.5 Diagnostic analysis and interpretability

Beyond point forecast accuracy, NS-SDN allows inspection of learned internal components: - time-varying amplitudes $A_{n,k}$, - evolving frequencies $\omega_{n,k}$, - inferred latent state trajectories h_n .

These diagnostics will be analyzed to determine whether the model captures economically meaningful patterns, such as: - changes in business-cycle duration, - post-shock damping or amplification, - regime-dependent cyclical behavior.

Unlike black-box neural models, NS-SDN provides interpretable component structure (Oreshkin et al., 2020; Lim et al., 2021).

8.6 Hypotheses

The empirical analysis will test the following hypotheses: 1. **Forecasting hypothesis:** NS-SDN achieves lower out-of-sample forecast error than traditional linear econometric models for medium-horizon forecasts. 2. **Nonstationarity hypothesis:** Gains from NS-SDN are largest during periods of structural change or regime transition. 3. **Interpretability hypothesis:** Learned spectral components correspond to economically plausible cyclical dynamics rather than noise-fitting.

9 References

Sitzmann, V., Martel, J., Bergman, A., Lindell, D., & Wetzstein, G. (2020). Implicit neural representations with periodic activation functions. *Advances in Neural Information Processing Systems* (NeurIPS)

Huang, N. E., Shen, Z., Long, S., Wu, M., Shih, H., Zheng, Q., Yen, N., Tung, C., & Liu, H. (1998). The empirical mode decomposition and the Hilbert spectrum for nonlinear and non-stationary time series analysis. *Proceedings of the Royal Society A*, 454(1971), 903–995.

Lubik, T. A., & Matthes, C. (2015). Time-varying parameter vector autoregressions: Specification, estimation, and an application. *Federal Reserve Bank of Richmond Economic Quarterly*, 101(4), 323–352.

Harvey, A. C. (1989). *Forecasting, Structural Time Series Models and the Kalman Filter*. Cambridge University Press.

Beveridge, S., & Nelson, C. R. (1981). A new approach to decomposition of economic time series into permanent and transitory components. *Journal of Monetary Economics*, 7(2), 151–174.

Chen, T. Q., Rubanova, Y., Bettencourt, J., & Duvenaud, D. (2018). Neural ordinary differential equations. *NeurIPS*.

Chowning, J. M. (1973). The synthesis of complex audio spectra by means of frequency modulation. *Journal of the Audio Engineering Society*, 21(7), 526–534.

Durbin, J., & Koopman, S. J. (2012). *Time Series Analysis by State Space Methods* (2nd ed.). Oxford University Press.

Salinas, D., Flunkert, V., Gasthaus, J., & Januschowski, T. (2020). DeepAR: Probabilistic forecasting with autoregressive recurrent networks. *International Journal of Forecasting*, 36(3), 1181–1191.

Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. *Neural Computation*, 9(8), 1735–1780.

Cho, K., et al. (2014). Learning phrase representations using RNN encoder–decoder for statistical machine translation. *EMNLP*.

West, K. D. (1996). Asymptotic inference about predictive ability. *Econometrica*, 64(5), 1067–1084.

Diebold, F. X., & Mariano, R. S. (1995). Comparing predictive accuracy. *Journal of Business & Economic Statistics*, 13(3), 253–263

Oreshkin, B. N., Carpow, D., Chapados, N., & Bengio, Y. (2020). N-BEATS: Neural basis expansion analysis for interpretable time series forecasting. *ICLR*.

Lim, B., Arik, S. Ö., Loeff, N., & Pfister, T. (2021). Temporal Fusion Transformers for interpretable multi-horizon time series forecasting. *International Journal of Forecasting*, 37(4), 1748–1764.