

# Summary

**Thesis Title: Efficacy of Detecting and Characterizing Anomalies Using Cortical Algorithms in Time Series Data**

**Candidate: Fatemeh Ahmadvand**

This thesis focuses on the challenge of detecting anomalies in univariate time-series data and equally emphasizes identifying anomalies alongside describing their severity, duration, and nature. The task is increasingly relevant in areas like finance, industrial maintenance, healthcare, and cybersecurity, where time-series signals guide critical decisions. Compared to static datasets, time-series data arrives sequentially, with each new observation influenced by preceding points, making small deviations potentially indicative of serious issues such as failures or fraud. This project, therefore, examines and contrasts two distinct methods grounded in different theoretical frameworks— **Hierarchical Temporal Memory (HTM)** which is inspired by cortical learning algorithms and excels at online sequence prediction, against **Sparse LSTM AutoEncoder for Anomaly Detection in Time-Series (SLADiT)**, which imposes sparsity constraints on a deep autoencoder architecture to reconstruct normal patterns and flag deviations.

Time-series data introduces unique challenges because observations arrive in sequence and depend on prior values, requiring specialized methods that handle temporal correlations, as well as seasonality, trends, and abrupt changes. Traditional statistical techniques such as ARIMA-based forecasting or moving averages can struggle with sudden shifts or complex temporal patterns. Deep learning methods can capture non-linear relationships but often need extensive labeled data and high computational resources—prohibitive when anomalies are rare, labels are scarce, or rapid detection is critical. Consequently, this thesis explores HTM and SLADiT, which both offer distinct ways to learn from time-series data. A major challenge, however, is model tuning—HTM, for instance, has around 21 hyperparameters, and slight adjustments can significantly affect anomaly detection performance. HTM draws on cortical structures for continuous, unsupervised learning, while SLADiT leverages recurrent networks and sparsity constraints to build compact latent representations for reconstruction-based anomaly detection.

Hierarchical Temporal Memory (HTM) emulates how the neocortex processes sensory information, featuring two main components. The first, the Spatial Pooler (SP), transforms raw input into Sparse Distributed Representations (SDRs) by activating only those “mini-columns” that receive sufficient matching input bits, thereby keeping a small fraction of columns active at any given time. The second, the Temporal Memory (TM), learns sequential transitions in data by forming dendritic segments that predict the next active set of neurons (or mini-columns). If the observed input differs from that prediction, the anomaly score increases. HTM naturally adapts to non-stationary environments because it continuously updates its learned representations in real time, making it particularly promising for streaming or industrial monitoring tasks where data patterns may shift. However, in univariate setups—where there is only a single variable evolving over time—the spatial correlations that HTM typically exploits can be less pronounced, sometimes diminishing its advantage. However, in univariate contexts, HTM’s reliance on spatial correlations can offer limited benefits, and it does not inherently learn seasonal or periodic structures. Experiments on UCR datasets showed that HTM struggles when data lacks time-based encoding, prompting the use of Fast Fourier Transform (FFT) to introduce time-based features. By concatenating synthetic sine and cosine dimensions, cyclical behaviors were captured even in datasets not explicitly time-stamped, preserving temporal structure and enhancing HTM’s ability to detect subtle or non-temporal anomalies.

Sparse LSTM AutoEncoder for Anomaly Detection in Time-Series (SLADiT) adds a sparsity constraint to the classic autoencoder design while leveraging Long Short-Term Memory (LSTM) layers to capture temporal dependencies. Its pipeline begins with an encoder composed of multiple stacked LSTM layers that compress a sequence of observations into a lower-dimensional latent vector. The latent space then imposes a sparsity constraint, often enforced through Kullback-Leibler divergence, which ensures that only a subset of latent neurons remain active and pushes the model to learn the most salient features of normal behavior. Finally, a decoder reconstructs the original time-series from this latent representation, and the reconstruction error signals how much an input deviates from the learned normal patterns, indicating an anomaly if the error exceeds a certain threshold. This approach generally excels at capturing subtle point anomalies in univariate time-series, aided by sparsity that reduces noise and overfitting. However, SLADiT often requires periodic retraining when patterns shift significantly, making continuous deployment more challenging in real-time applications.

This thesis employs the UCR Time-Series Anomaly Archive—widely used for benchmarking univariate anomaly detection—and the NYC Taxi dataset from the Numenta Anomaly Benchmark (NAB) for near real-time testing. Each dataset is split into training (mostly normal data) and testing (with known anomalies). Both HTM and SLADiT undergo hyperparameter tuning to maximize performance.

For HTM, parameters include column count, cells per column, and permanence thresholds; for SLADiT, layer depth, hidden dimensions, and sparsity factors are refined. Evaluation relies on precision, recall, the F1-score (which balances false positives and negatives), and the Area Under the ROC Curve (AUC), measuring the trade-off between correctly identifying anomalies and avoiding false alarms.

When balancing accuracy and real-time adaptability, HTM’s online learning provides an edge for streaming applications, as it dispenses with a separate retraining phase. However, fine-tuning HTM is highly sensitive, and in purely univariate contexts its performance can sometimes lag behind SLADiT. Meanwhile, SLADiT offers robust reconstruction-based detection of point anomalies, often achieving high recall and lower false-positive rates—but its lack of built-in continuous adaptation can hamper real-time responsiveness. Both methods benefit from added data dimensionality: HTM draws on its Spatial Pooler more effectively with multiple correlated signals, while SLADiT can leverage deeper representations in both univariate and multivariate cases. In practical terms, SLADiT’s minibatch gradient-based training often runs efficiently on GPUs but requires retraining. HTM can operate continuously on CPUs but grows computationally heavier if configured with many columns and cells. Through systematic experiments on representative univariate time-series datasets, this thesis demonstrates that both HTM and SLADiT can excel at detecting anomalies, but their suitability depends on the application context. HTM’s biologically inspired, continuous learning suits scenarios where data patterns evolve unpredictably, provided its numerous hyperparameters are properly tuned. SLADiT’s reconstruction-based framework, enhanced by latent sparsity, yields strong detection performance in more stable conditions. The choice between these approaches depends on factors like tolerance for false alarms, adaptability to shifting processes, computational resources, and the nature of anomalies.

Future research could investigate hybrid or ensemble strategies that leverage the complementary strengths of HTM and SLADiT in real-time contexts, further enhancing the reliability and efficiency of anomaly detection in high-stakes environments. Transitioning from a multivariate, time-aware HTM model to a univariate implementation introduced significant challenges, as the model’s effectiveness became heavily dependent on dataset characteristics. Low-noise datasets with clear periodic structures allowed for reasonable detection performance, whereas abrupt anomalies and non-stationary trends proved more problematic. The removal of contextual cues, particularly those derived from timestamps or external features, significantly reduced HTM’s ability to infer meaningful predictions. Future improvements should focus on reintroducing time-based encodings, refining hyperparameters, and exploring hybrid methods that integrate HTM’s continuous learning with more structured encoding strategies, thereby enhancing generalizability and adaptability for real-world univariate time-series anomaly detection.