TDA-Informed Feedforward Recurrent Neural Networks for High-Rate State Estimation Application

---

ARMAN RAZMARASHOOLI, DANIEL SALAZAR MARTINEZ,
YANG KANG CHUA, METRID OKUMU, SIMON LAFLAMME,
CHAO HU and PAUL T. SCHRADER

ABSTRACT

Real-time state estimation is essential for conducting efficient structural health assessment and enabling feedback control strategies in high-rate systems. High-rate systems are dynamic systems that experience extreme acceleration ($> 100\ g_n$ where g_n is the gravitational constant near the earth's surface) within very short periods (< 1 millisecond), commonly observed in applications such as hypersonic systems and impact mitigation mechanisms. These engineering systems require control and feedback methodologies capable of operating within the sub-millisecond ranges, posing significant challenges for traditional prediction methods due to the underlying nonlinear and nonstationary behavior. This paper introduces and investigates a deep learning algorithm that combines topological data analysis (TDA) features with machine learning to enhance state estimation of high-rate systems. The algorithm consists of an ensemble of recurrent neural networks (RNNs) constructed with a parallel arrangement of long short-term memory cells. The ensemble incorporates predictions from multiple RNNs trained on varying window sizes and delays, allowing the system to capture both short- and long-term data dependencies. A Feedforward Neural Network is employed to predict and analyze the contribution of each RNN's output to an overall state estimation using TDA-derived features such as the maximum persistence of the first dimensional persistence homology group, H_1. The proposed approach is validated using data from the dynamic reproduction of projectiles in ballistic environments for advanced research (DROPBEAR) testbed. Results demonstrate the promise of the algorithm at predicting time series and system states over longer prediction horizons.

Arman Razmarashooli, Ph.D student, Email: shooli@iastate.edu. Department of Civil, Construction, and Environmental Engineering, Iowa State University, Ames, IA, 50011, USA
Daniel Salazar Martinez, Ph.D student, Email: daniel94@iastate.edu. Department of Civil, Construction, and Environmental Engineering, Iowa State University, Ames, IA, 50011, USA
Yang Kang Chua, Ph.D student, Email:_yang kang.chua@uconn.edu. School of Mechanical, Aerospace, and Manufacturing Engineering, University of Connecticut, Storrs, CT 06269, USA
Metrid Okumu, Ph.D student, Email: okumu@iastate.edu. Department of Civil, Construction, and Environmental Engineering

INTRODUCTION

High-rate dynamic systems are systems subject to extreme dynamic events, typically involving accelerations greater than 100 g_n and of durations shorter than 100 milliseconds (ms) [1]. Examples include active blast mitigation systems, advanced weaponry, hypersonic vehicles, and automotive airbag deployment mechanisms [1, 2]. Of particular interest is high-rate structural health monitoring (HRSHM), which focuses on the real-time state estimation of such systems to enable feedback mechanisms that have the potential to improve operational safety and performance [3]. However, HRSHM is challenging due to three primary factors: (1) large uncertainties in external excitations, (2) nonstationary behavior with high levels of disturbance, and (3) unmodeled dynamics resulting from abrupt changes in system configuration. These characteristics complicate the accurate estimation of system states, especially under the sub-millisecond response time often required for effective feedback [2].

Applicable methods are generally categorized into physics-based observers, data-driven methodologies, and hybrid approaches. Physics-based observers including Kalman filters, sliding mode observers, and model reference adaptive systems (MRAS) rely on detailed mathematical models of the system to perform state estimation [4, 5]. Although these methods provide robust solutions, they are computationally expensive, have slower convergence, and face difficulty in modeling the nonlinearities that frequently exist in high-rate dynamics [6].

Data-driven techniques, particularly those involving machine learning models such as recurrent neural networks (RNNs) and long short-term memory (LSTM) networks, provide an acceptable alternative as they do not need to explicit model the dynamics [3]. RNNs have shown effective in analyzing time-series data by modeling and forecasting nonlinear and nonstationary signals [7]. Additionally, ensemble methodologies, which involve training multiple LSTM networks each using a distinct set of inputs, have improved the ability to capture complex multi-temporal dynamics [8]. However, these ensemble approaches still face significant challenges. Firstly, they typically depend on attention mechanisms to aggregate model predictions, which increases computational demands and can reduce prediction accuracy for multi-step forecasts [3]. Secondly, preparing labeled data necessary for training these models is expensive and challenging due to the complexity of experimental setups, low repeatability of high-rate events, and potential risks of system damage [3]. Nelson et al. [5] proposed a hybrid model known as the neural state estimator, which combines LSTM-based forecasting with an MRAS for subsequent state estimation. While this method effectively shifts computational de-

---

ronmental Engineering, Iowa State University, Ames, IA, 50011, USA

Simon Laflamme, Professor, Email: laflamme@iastate.edu. Department of Civil, Construction, and Environmental Engineering, Iowa State University, Ames, IA, 50011, USA

Chao Hu, Associate Professor, Email: chao.hu@uconn.edu. School of Mechanical, Aerospace, and Manufacturing Engineering, University of Connecticut, Storrs, CT 06269, USA

Paul T. Schrader, Research Mathematician, Email: paul.schrader.1@afrl.af.mil. Air Force Research Laboratory, Information Directorate, Information Exploitation Branch, NY, USA

mands forward in time, thus reducing or eliminating deadline overshoots, it remains computationally expensive due to the MRAS component.

Recently, Topological Data Analysis (TDA) has gained attention as a powerful tool for capturing intrinsic data characteristics of nonlinear time series by evaluating their embeddings into point clouds [9, 10]. TDA uses concepts from algebraic topology to explore data as geometric and topological objects, detecting topological invariants such as connected components, boundaries of 1-dimensional holes or loops, boundary hypersurfaces of higher dimensions holes or voids, within datasets geometricaly/topologically expressed as *simplicial complexes* [11]. TDA has shown promise in the analysis of complex and chaotic dynamics, for example demonstrating improved time series classification capabilities [12]. In prior work, we have studied the use of TDA features as inputs to an ensemble of RNNs, and demonstrated that TDA features could be utilized as dynamic descriptors [3, 9]. Here, we build on our prior work and investigate if TDA features could be used as a feed-forward mechanism to an ensemble of RNNs, with the objective to reduce computation time.

Specifically, our approach employs multiple parallel LSTM networks, each trained independently on data constructed with a distinct delay vector representing a unique temporal scale. Unlike existing approaches that rely on attention mechanisms, we introduce a simpler and more computationally efficient Feedforward Neural Network (FNN). This FNN dynamically calculates the weights to optimally combine outputs from the parallel LSTMs based on real-time features derived from TDA. These TDA features are obtained using persistent homology by measuring the maximum persistence of topological features, including the 0^{th} dimensional persistent homology(H_0) that encodes connected component characteristics and 1^{st} dimensional persistent homology(H_1) that encodes the 1-D holes all from the given dataset's associated topology preserving simplicial complex filtration. This information is obtained from two sliding windows moving over the input signal. Each LSTM network within the ensemble is exclusively trained on synthetic harmonic signals generated at different frequencies. By using these synthetic signals instead of experimental data, we reduce the dependency on labeled datasets, which are typically difficult and expensive to collect in high-rate environments [3]. We conduct numerical experiments using a synthetic dataset and experimental data obtained from the DROPBEAR (Dynamic Reproduction of Projectiles in Ballistic Environments for Advanced Research) testbed.

The rest of the paper is organized as follows. The next section describes the algorithm developed for real-time multi-step ahead prediction and also the DROPBEAR testbed. The subsequent section presents and discusses the results obtained from synthetic and experimental datasets. The last section reports conclusions and recommendations.

METHODOLOGY

This section describes the methodology used to design, train, and evaluate a physics-informed state estimation algorithm using an ensemble of Long Short-Term Memory (LSTM) cells guided by TDA features. The goal is to predict fast-changing dynamics under nonstationary conditions using minimal experimental training data. The architecture for prediction is shown in Figure 1. At each time step, the output from all LSTM

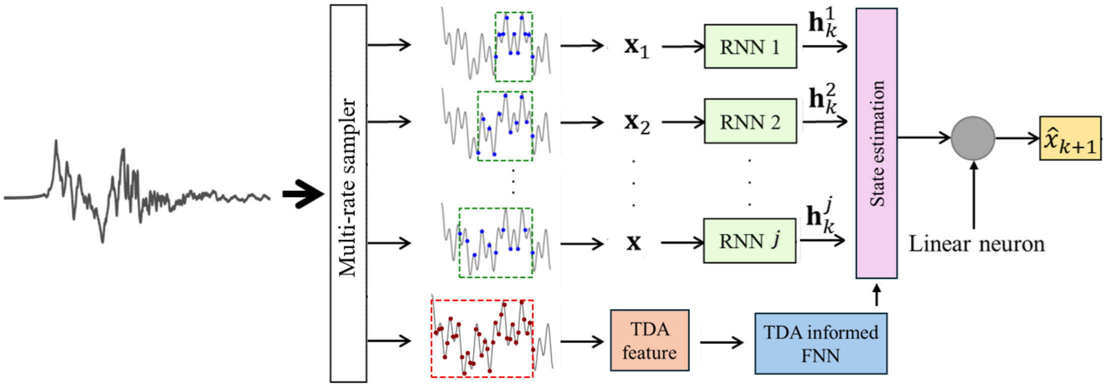

Figure 1. Schematics of the high-rate time series estimation algorithm.

networks is combined to form a final prediction $\hat{x}_{k+m}$ for an m-*step ahead* forecast. In previous work, this combination was handled by an attention layer followed by a linear neuron [3]. In this study, we replace the attention layer with the FNN architecture, which is trained using real-time topological features (maximum persistence of H_1).

The algorithm comprises three major steps:

1. Extract TDA features such as maximum persistence and one-dimensional homology groups (H_1).

2. Feed the extracted TDA features into the FNN, which outputs a set of participation weights $\alpha_i \in \mathbb{R}$ corresponding to each LSTM. The FNN, been pre-trained using windowed signal inputs, with the corresponding maximum persistence of H_1 is used as the target output. The optimum window size and time delay are selected based on the previous research $d = 2$ and $\tau < 0.25/f_{max}$ [9].

3. Combine the weighted outputs of each LSTM using a linear neuron to produce an m-step ahead prediction.

There are two methods to predict m steps ahead. The first one is to repeat m times recursively, each time feeding the previous output back as input. As a caveat, this procedure might result in wrong predictions due to accumulative/propagating errors, especially in the absence of real training labeled data. The second method involves directly predicting m step ahead without computing intermediate predictions. This second method is used in this study. The FNN enables adaptive weighting across different time resolutions, guided by real-time TDA features, and eliminates the need for manually labeled training data leveraging synthetic harmonic signals for initial training. This structure improves accuracy, reduces computational cost, and allows generalization to unseen dynamic behaviors.

To handle dynamics that evolve at different temporal scales, we use a multi-resolution windowing strategy to build delay vectors. Each LSTM is assigned a specific delay vector characterized by a different time delay τ and dimension d. The choice of τ directly affects the resolution of dynamics captured: smaller delays are better for high-frequency signals, while larger delays capture low-frequency signals [13]. Each LSTM receives its own delay vector, and all are synchronized to produce predictions from the same time step. This setting helps each sub-model specialize in different temporal structures.

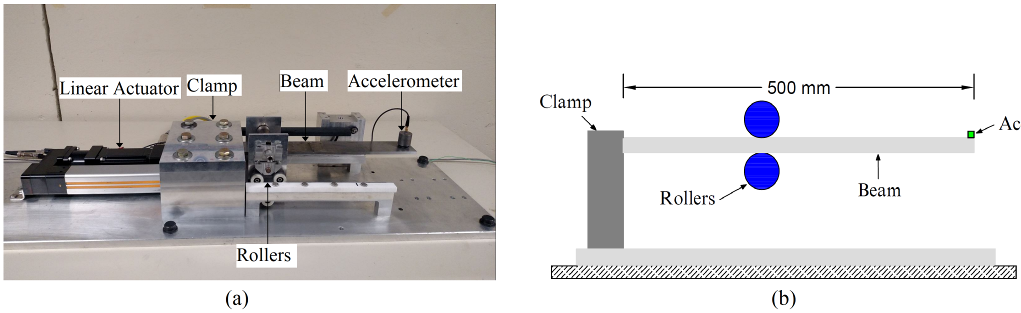

Figure 2. The DROPBEAR testbed.

TDA features are extracted in real-time using a sliding window approach. At each step, a local segment of the signal $\boldsymbol{x}_w(t) = [x(t-1), x(t-2), \cdots, x(t-w)]$ is embedded into a delay vector using different τ and d parameters as the LSTM ensemble. From the resulting point cloud, topological features are used to measure the shape of the underlying trajectory and track changes in dynamic regimes. The choice of window size w aims at capturing at least one complete dynamic cycle. Based on previous studies, the optimal τ for a single harmonic f is $0.25/f$ resulting in a circular point cloud; smaller values stretch the cloud into ellipses, while larger ones result in folding the state space. A fixed embedding dimension of $d = 2$ is suggested for a single harmonic function.

Realistic performance of the algorithm is evaluated on the Dynamic Reproduction of Projectiles in Ballistic Environments for Advanced Research (DROPBEAR) testbed. DROPBEAR is a cantilever beam with a moving boundary condition, designed to simulate high-rate dynamics (see Figure 2). The beam is clamped at one end and connected to a movable pin at the other, which introduces sudden changes in stiffness. An impact hammer is used to create repeated impact excitations. An accelerometer is placed 300 mm from the fixed end to measure vibrations. The beam response can be simplified as a harmonic signal assuming the first mode dominates as shown in Equation 1 [14]. We selected DROPBEAR case for two reasons: (1) the experimental data are publicly available [5], and (2) the DROPBEAR datasets provide a rigorous benchmark for evaluating high-rate estimation algorithms [5].

$$x(t) = \cos(2\pi f t) \tag{1}$$

To address the challenge of data scarcity in high-rate regimes, transfer learning is implemented using synthetic harmonic signals with varying frequencies. Each LSTM in the ensemble is independently pre-trained on a fixed-frequency signal of the form $x(t) = \cos(2\pi f t)$, where f is specific to that LSTM. The training process relies on delay vectors constructed from time-delayed signal samples and eliminates the need for labeled experimental data. During inference, the system does not perform feedback-based error correction. Instead, the maximum persistence of H_1 is extracted in real time from the test signal and mapped to participation weights α_i via the FNN. These weights determine the contribution of each LSTM to the final prediction, which is computed using a linear neuron that aggregates their outputs. The ensemble thus adapts to the geometric structure of the incoming signal and predicts future steps one at a time, using only the current estimated state.

TABLE I. Hyper-parameters used for training in synthetic datasets.

	Config 1	Config 2	Config 3	TDA
Time delay (s)	0.500	0.250	0.167	0.030
Dimension	2	2	2	0.103
Epochs	3	3	3	25
Learning rate	0.01	0.01	0.01	0.01

Two performance metrics are used to assess the prediction accuracy. Metric J_1 is the root mean squared error (RMSE) between the true location of the moving boundary and the estimated location across the entire signal duration. Metric $J_{2,i}$ measures the percentage of time the prediction error exceeds a threshold i, where i is taken as 5%, 10%, and 20%. Metric $J_{3,i}$ measures the percentage of time the prediction error exceeds a threshold i when the cart position is constant. These thresholds represent application-relevant tolerances for structural response tracking under high-rate conditions.

RESULTS AND DISCUSSION

This section evaluates the prediction performance of the proposed TDA-informed ensemble of LSTM networks across synthetic and experimental datasets designed to replicate high-rate dynamic systems. The algorithm is validated using two types of datasets: (1) a synthetic chirp signal with frequency varying from 1 Hz to 3 Hz, and (2) experimental acceleration measurements from the DROPBEAR testbed. The synthetic chirp signal is taken as

$$x(t) = \cos\big(2\pi f(t)t\big) \tag{2}$$

where $f(t)$ varies between 1 and 3 Hz with a 1000 Hz sampling rate. The excitation is plotted in Figure 3 (gray solid line). In this excitation, the frequency remains constant at 1 Hz for the first 2 seconds, then increases to 3 Hz over the next 2 seconds. It stays at 3 Hz for an additional 2 seconds before returning to 1 Hz, where it remains constant for the final 2 seconds. The size of the moving window and time delay to embedded data for each LSTM and TDA is reported in Table I, and Figure 3 plots results for TDA-informed and attention-based ensemble networks for 100 steps ahead (100 milliseconds). Results from the Figure show that the attention-based ensemble is incapable of predicting the acceleration magnitude that far ahead, while the TDA-informed ensemble can, except during transition regions mainly attributable to the time delay. An inspection of results reported in Table II reveals that under similar training conditions, TDA-informed ensemble achieved significantly better accuracy in the steady state region due to its awareness of the system's state, predicting acceleration magnitudes within 5% accuracy 70% of the time. The attention-based ensemble is not capable of predicting acceleration that far ahead. These results from 100-step ahead predictions demonstrate the robustness of the TDA-informed ensemble.

Figure 4 shows how metrics change across different prediction steps, and the accuracy is presented using metrics J_1 and $J_{2,i}$ for various thresholds. The degradation in performance as the prediction step increases is expected and visible across the attention-based ensemble model because of less information about the future compared to TDA-

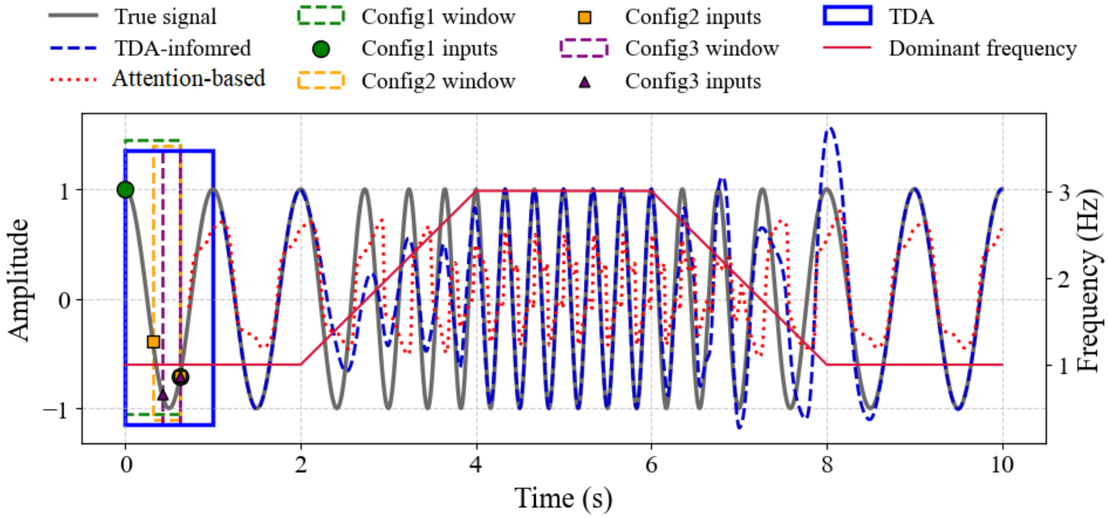

Figure 3. TDA-informed and attention-based ensemble predictions for synthetic datasets, 100 steps (100 milliseconds) ahead.

TABLE II. Performance comparison for synthetic datasets, 100 steps (100 milliseconds) ahead.

Model	J_1 (-)	$J_{2,5}$ (%)	$J_{2,10}$ (%)	$J_{2,20}$ (%)	$J_{3,5}$ (%)	$J_{3,10}$ (%)	$J_{3,20}$ (%)
Ensemble	0.58	97.75	95.86	91.63	99.05	98.20	96.20
TDA-informed	0.32	50.61	46.21	40.78	31.88	29.14	23.94

informed ensemble. TDA-informed ensemble maintains a relatively stable accuracy due to embedded topological knowledge of the system.

The investigation is continued on experimental datasets taken from the DROPBEAR testbed. The dataset corresponds to Dataset-6, where the experiment does not involve an impact hammer. The system's dominant frequency varies between 28.8 Hz and 54.4 Hz. The ensemble LSTM part consists of two configurations: config 1 trained at 28.8 Hz and config 2 trained at 54.4 Hz, representing the minimum and maximum dominant frequencies of the system. A fully connected FNN is employed to map topological features extracted from the signal to the true cart location. The network consists of 3 hidden layers, each with 512 ReLU-activated neurons. The model is trained using the Adam optimizer with mean squared error as the loss function and parameters used for training are summarized in Table III.

Figure 5 shows the true and predicted cart locations and the TDA-informed and attention-based ensemble for DROPBEAR dataset for 2500 steps (100 milliseconds) ahead. Metrics summarized in Table IV illustrate the challenges faced by attention-based ensemble prediction due to a lack of labeled data, resulting in substantially higher errors. The significant improvement in the metric J_3 indicates that the TDA-informed ensemble is better at maintaining prediction stability under steady state conditions. While performance through these metrics appears generally low even for the TDA-informed ensemble, a visual inspection of Figure 5 reveals that the cart can be tracked, and that

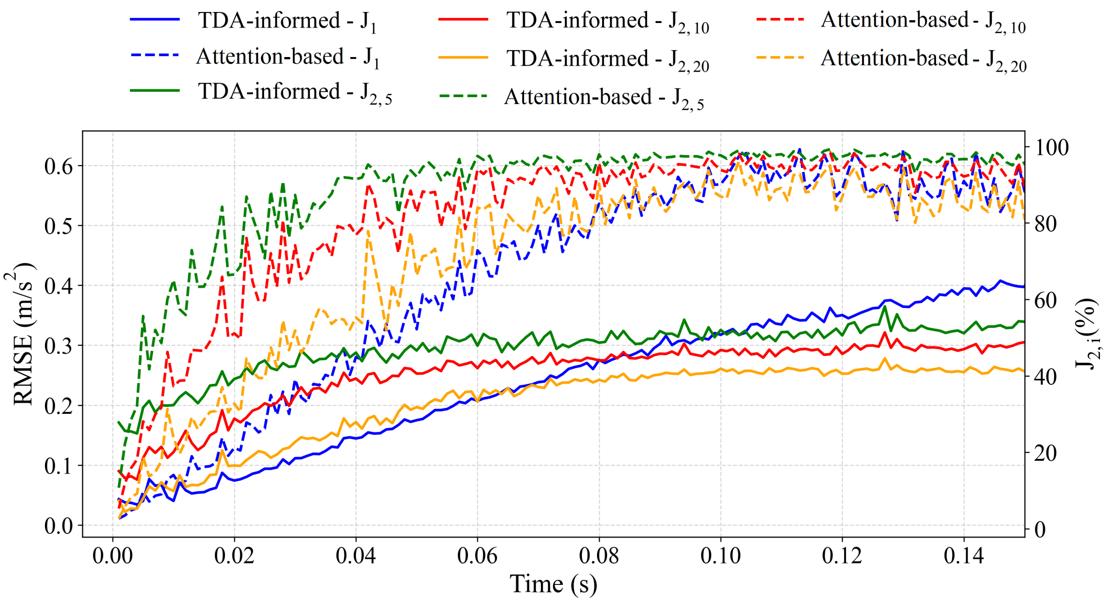

Figure 4. Accuracy comparison for synthetic datasets over prediction horizons (1 step = 1 millisecond).

TABLE III. Hyper-parameters used for training in DROPBEAR.

	Config 1	Config 2	TDA
Time delay (s)	0.01024	0.00512	0.0012
Dimension	2	2	0.0612
Epochs	3	3	25
Learning rate	0.01	0.01	0.01

performance metrics are highly influence by high chattering in the predicted cart location. Further investigations on training and physics-informed predictions could significantly enhance algorithmic performance; this is left to future work.

CONCLUSION

This paper evaluated the performance of an FNN architecture to conduct step-head predictions of nonstationary time series based on TDA features. The TDA-informed FNN is used to estimate the system state and compute the participation weight of each

TABLE IV. Performance comparison for DROPBEAR, 2500 steps (100 milliseconds) ahead.

Model	J_1 (m/s^2)	$J_{2,5}$ (%)	$J_{2,10}$ (%)	$J_{2,20}$ (%)	$J_{3,5}$ (%)	$J_{3,10}$ (%)	$J_{3,20}$ (%)
Attention-based ensemble	4.32	98.68	95.89	90.65	98.84	98.02	95.84
TDA-informed ensemble	4.36	96.07	92.34	85.33	88.78	78.32	58.32

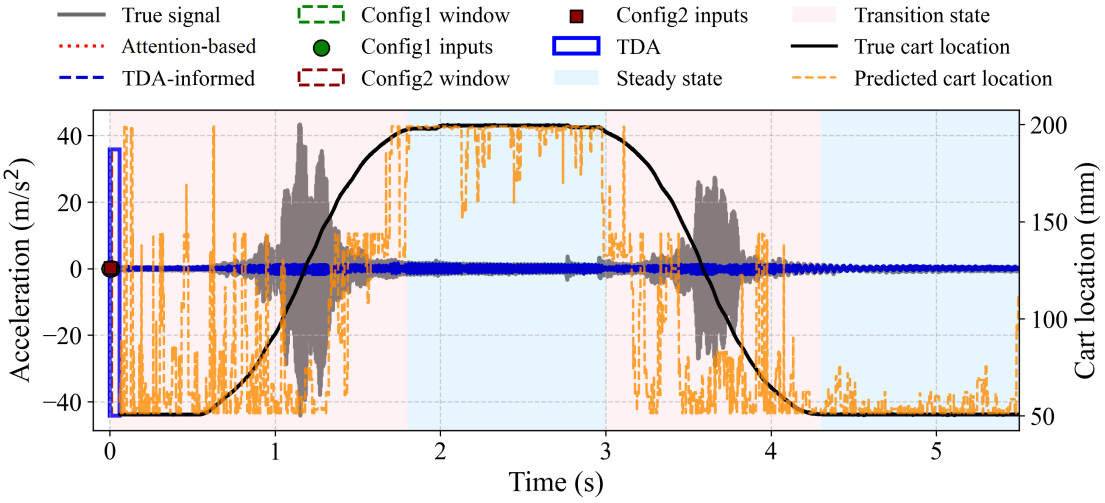

Figure 5. TDA-informed and attention-based ensemble prediction for DROPBEAR, 2500 steps (100 milliseconds) ahead.

LSTM to perform multi-step-ahead prediction. Validation was performed using both synthetic and experimental datasets from the DROPBEAR testbed. Results for the synthetic datasets showed that the attention-based ensemble works better for small prediction horizons, but that the TDA-informed ensemble is more effective at predicting over extended prediction horizons.

ACKNOWLEDGMENT

The authors would like to acknowledge the financial support from the Defense Established Program to Stimulate Competitive Research (DEPSCoR) award number FA9550-22-1-0303,the Air Force Office of Scientific Research (AFOSR) award number FA9550-23-1-0033, and the National Science Foundation award number CCF-2234919. Any opinions, findings, conclusions, or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the sponsors/organizations, the United States Air Force, or the US government.

REFERENCES

1. Dodson, J., A. Downey, S. Laflamme, M. D. Todd, A. G. Moura, Y. Wang, Z. Mao, P. Avitabile, and E. Blasch. 2021. "High-Rate Structural Health Monitoring and Prognostics: An Overview," in *Data Science in Engineering, Volume 9*, Springer International Publishing, pp. 213–217, doi:10.1007/978-3-030-76004-5_23.
2. Hong, J., S. Laflamme, J. Dodson, and B. Joyce. 2018. "Introduction to state estimation of high-rate system dynamics," *Sensors*, 18(1):217.
3. Barzegar, V., S. Laflamme, C. Hu, and J. Dodson. 2022. "Ensemble of recurrent neural networks with long short-term memory cells for high-rate structural health monitoring," *Mechanical Systems and Signal Processing*, 164:108201.

4. Downey, A., J. Hong, J. Dodson, M. Carroll, and J. Scheppegrell. 2020. “Millisecond model updating for structures experiencing unmodeled high-rate dynamic events,” *Mechanical Systems and Signal Processing*, 138:106551, doi:10.1016/j.ymssp.2019.106551.
5. Nelson, M., V. Barzegar, S. Laflamme, C. Hu, A. R. Downey, J. D. Bakos, A. Thelen, and J. Dodson. 2023. “Multi-step ahead state estimation with hybrid algorithm for high-rate dynamic systems,” *Mechanical Systems and Signal Processing*, 182:109536.
6. Hong, J., S. Laflamme, L. Cao, J. Dodson, and B. Joyce. 2020. “Variable input observer for nonstationary high-rate dynamic systems,” *Neural computing and applications*, 32:5015–5026.
7. Han, Z., J. Zhao, H. Leung, K. F. Ma, and W. Wang. 2019. “A review of deep learning models for time series prediction,” *IEEE Sensors Journal*, 21(6):7833–7848.
8. Liu, H., R. Yang, and Z. Duan. 2020. “Wind speed forecasting using a new multi-factor fusion and multi-resolution ensemble model with real-time decomposition and adaptive error correction,” *Energy Conversion and Management*, 217:112995.
9. Razmarashooli, A., Y. K. Chua, V. Barzegar, D. Salazar, S. Laflamme, C. Hu, A. R. Downey, J. Dodson, and P. T. Schrader. 2025. “Real-time state estimation of nonstationary systems through dominant fundamental frequency using topological data analysis features,” *Mechanical Systems and Signal Processing*, 224:112048.
10. Khasawneh, F. A. and E. Munch. 2016. “Chatter detection in turning using persistent homology,” *Mechanical Systems and Signal Processing*, 70:527–541.
11. Carlsson, G. and M. Vejdemo-Johansson. 2021. *Topological data analysis with applications*, Cambridge University Press, doi:10.1017/9781108975704.011.
12. Umeda, Y. 2017. “Time series classification via topological data analysis,” *Information and Media Technologies*, 12:228–239.
13. Barzegar, V., S. Laflamme, C. Hu, and J. Dodson. 2021. “Multi-time resolution ensemble lstms for enhanced feature extraction in high-rate time series,” *Sensors*, 21(6):1954.
14. Kausel, E. 2017. *Advanced Structural Dynamics*, Cambridge University Press, doi: 10.1017/9781316761403.