Multi-Step State Forecasting with Conformal Prediction for High-Rate Dynamic Systems

YANG KANG CHUA, ARMAN RAZMARASHOOLI, MOHAMMAD MUNDIWALA, DANIEL A. SALAZAR MARTINEZ, METRID OKUMU, CHAO HU and SIMON LAFLAMME

ABSTRACT

High-rate systems, such as hypersonic vehicles and impact mitigation mechanisms, exhibit rapid, nonlinear, and rapid time-varying dynamics. These systems are often subjected to extreme accelerations (e.g., > 100 g_n) over very short durations (e.g., < 100 ms), making them highly susceptible to uncertainties, non-stationarities, and external disturbances. In such environments, accurate and reliable state forecasting is essential for real-time decision-making and control. This paper presents a novel framework for high-rate multi-step forecasting that integrates recurrent neural networks (RNNs), topological data analysis (TDA), and conformal prediction for uncertainty quantification. RNNs are used to model temporal dependencies in high-dimensional system dynamics, while TDA-derived features capture the underlying topological structure of the system state. To provide robust and interpretable uncertainty estimates, we employ conformal prediction techniques—particularly inductive and block-based conformal methods—to construct statistically valid prediction intervals. These methods adapt to the non-exchangeable nature of time-series data, enabling coverage guarantees under temporal dependence. The effectiveness of the proposed framework is evaluated on experimental data from the *Dynamic Reproduction of Projectiles in Ballistic Environments for Advanced Research* (DROPBEAR) testbed. Results demonstrate that the model produces accurate multi-step forecasts with well-calibrated uncertainty bounds, which can be used to support high-rate system monitoring and decision-making.

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
Arman Razmarashooli, Ph.D student, Email: shooli@iastate.edu. Department of Civil, Construction, and Environmental Engineering, Iowa State University, Ames, IA, 50011, USA
Mohammad Mundiwala, Ph.D student, Email: mohammad.mundiwala@uconn.edu. School of Mechanical, Aerospace, and Manufacturing Engineering, University of Connecticut, Storrs, CT 06269, USA Daniel A. Salazar Martinez, Ph.D student, Email: daniel94@iastate.edu. Department of Civil, Construction, and Environmental Engineering, Iowa State University, Ames, IA, 50011, USA
Metrid Okumu, Ph.D student, Email: okumu@iastate.edu. Department of Civil, Construction, and Environmental Engineering

INTRODUCTION

High-rate dynamic systems, such as hypersonic vehicles, automobile crash safety mechanisms, and structural health monitoring, are characterized by rapid and extreme changes occurring within milliseconds [1]. These systems often experience accelerations exceeding 100 g_{n} in less than 100 ms, where g_{n} represents gravitational acceleration, leading to significant challenges in real-time state forecasting due to large uncertainties, nonstationarities, and external disturbances. Accurate and timely monitoring of such systems is critical for ensuring structural integrity and enabling prompt control interventions. Failure to do so can result in severe consequences, such as undetected control surface deviations in hypersonic vehicles or delayed responses in blast mitigation systems [2]. However, the presence of unmodeled dynamics, abrupt configuration changes, and severe measurement noise necessitates forecasting frameworks capable of capturing complex system behaviors while also quantifying uncertainty in real time.

Several approaches have been proposed for high-rate dynamics estimation, generally categorized into physics-based and data-driven methods. Physics-based methods use first-principles models such as differential equations and finite element models to describe and predict system behavior. For instance, Yan et al. [3] proposed a real-time adaptive estimator based on reduced-order physical models and concurrent learning to track high-rate system dynamics. Similarly, Downey et al. [4] developed a millisecond-scale FEM updating method to monitor structural changes during high-rate events using real-time physics-based simulations. In contrast, data-driven techniques, such as recurrent neural networks (RNNs), have shown promise in modeling nonlinear, high-dimensional time-series data directly from measurements. For example, recent work has demonstrated the use of RNN ensembles with multi-rate sampling and long short-term memory (LSTM) units for sub-millisecond prediction in dynamic environments [5]. More recent advances have incorporated topological data analysis (TDA) to extract multi-scale geometric features from system trajectories, leading to improved forecasting performance in nonstationary conditions [6, 7].

This paper focuses on advancing real-time forecasting of dynamic systems using a data-driven framework that integrates TDA features with RNNs for multi-step prediction. We emphasize the role of conformal prediction in generating statistically valid uncertainty intervals for each forecasted step. First, we explore several conformal prediction methods that are potentially useful for high-rate state estimation. Second, we investigate how calibration data can be constructed to support dynamic updates. Third, we draw conclusions from the evaluation of different conformal prediction methods and provide a discussion on adapting these methods to high-rate problems.

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

CONFORMAL PREDICTION

Conformal prediction has emerged as a popular, distribution-free method for uncertainty quantification, offering finite-sample coverage guarantees under minimal assumptions, notably exchangeability of the data. As introduced in [8], it uses conformity scores to assess how unusual a new observation is relative to past data, which enables the construction of prediction intervals with a specified coverage level (e.g., 95%). A widely used and efficient variant, inductive conformal prediction (ICP), splits the data into a training set for model fitting and a calibration set for computing conformity scores [9]. This setup avoids retraining for each test point, making the method more scalable for practical use.

However, both approaches still rely on the assumption that the data are exchangeable, making them not directly applicable to time series data, where temporal dependence violates these assumptions. In the context of time series forecasting, this presents a fundamental challenge for conformal prediction. When data arrive sequentially, as in a single time series, methods like adaptive conformal inference (ACI) and its extensions adapt conformal prediction to update prediction intervals one step at a time while aiming for asymptotic coverage [10, 11]. Other techniques, such as ensemble batch prediction intervals (EnbPI), further address this setting by leveraging ensemble methods to improve robustness [12]. Nevertheless, these methods typically offer only asymptotic or marginal coverage guarantees and are not designed for joint multi-step prediction.

In contrast, multi-horizon forecasting produces forecasting intervals for multiple time steps ahead at once. This is typical for i.i.d. time series data, allowing methods to exploit instance independence. Conformal forecasting RNN (CF-RNN) was among the earliest methods to address this by using a Bonferroni correction to ensure the family-wise error rate [13]. This approach is simple but yields excessively conservative intervals with strong temporal dependence. To overcome this, Sun and Yu suggested the application of Copulas as a dependence structure model, with an extra calibration dataset [14]. Lindemann et al. applied the CF-RNN concept to motion planning through the combination of prediction and planning [15]. These examples highlight a growing interest in tackling the multi-horizon forecasting problem, but they also underscore ongoing challenges in maintaining coverage guarantees and managing computational complexity. In this work, we adopt a recently proposed method, ConForME, which addresses several challenges, such as coverage guarantees and data costs, which is presented in detail in the methodology.

DROPBEAR EXPERIMENTAL TESTBED

The DROPBEAR experimental test bed was developed to study structural responses under high-rate dynamic conditions [16]. The current setup includes a $51 \times 6.66 \times 501$ mm cantilever beam outfitted with two accelerometers and two strain gauges to record time series vibration signals. A movable roller support system introduces controlled variations in boundary conditions by following specific motion profiles. High-resolution analog-to-digital converters (14-bit for displacement and 24-bit for acceleration) enable precise capture of the beam's dynamic behavior. This platform provides a rich dataset

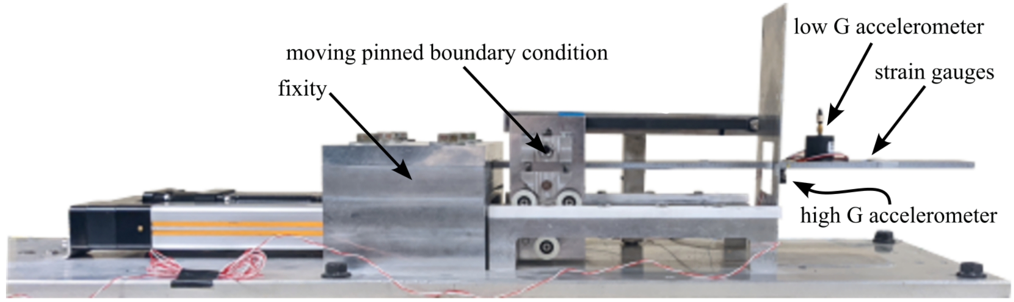

Figure 1. The DROPBEAR testbed configured with a 51 mm $\times$ 6.66 mm $\times$ 501 mm aluminum cantilever beam, instrumented with two accelerometers and two strain gauges for dynamic response measurement [18].

for developing and benchmarking advanced forecasting and state estimation techniques tailored to high-rate systems. One of the more recent publicly available datasets, the DROPBEAR Dataset-2 [17], includes square wave, sinusoidal, and impulse motion profiles designed to induce distinct structural responses. This dataset was later extended into DROPBEAR Dataset-8, which incorporated a broader range of motion patterns, enhancing its suitability for machine learning applications [18]. In this work, we use the most recent version, Dataset-8.

METHODOLOGY

To construct a forecasting pipeline for high-rate dynamic systems, we extend the standard RNN framework for multi-step prediction. RNNs are well-suited for time-series modeling due to their ability to capture temporal dependencies via a recurrent hidden state mechanism [19]. In our configuration, the model is tasked with predicting multiple future steps of the system state based on a window of past observations that include topological features derived from beam acceleration vibrations. At each time step t, the hidden state $h_t \in \mathbb{R}^N$ evolves according to the recurrence relation:

$$h_t = \sigma(Wx_t + Uh_{t-1} + b), \tag{1}$$

where $x_t \in \mathbb{R}^M$ is the input feature vector comprising topological descriptors computed from recent acceleration signals; $W \in \mathbb{R}^{N\times M}$ and $U \in \mathbb{R}^{N\times N}$ are learnable weight matrices; $b \in \mathbb{R}^N$ is a bias term and σ denotes a non-linear activation function such as $\tanh$ or ReLU. We construct training samples using a sliding window of length L and define a prediction horizon of H steps into the future to enable multi-step forecasting. Each training pair consists of an input sequence of L consecutive time steps and an output sequence of H future values, which is expressed mathematically as

$$(\mathbf{X}_t, \mathbf{Y}_{t+1:t+H}) = \big(\{x_{t-L+1}, \ldots, x_t\}, \{y_{t+1}, \ldots, y_{t+H}\}\big), \tag{2}$$

where $\mathbf{X}_t \in \mathbb{R}^{L\times M}$ is the input feature matrix and $\mathbf{Y}_{t+1:t+H} \in \mathbb{R}^{H\times K}$ is the target output sequence. The RNN processes the input sequence $\mathbf{X}_t$ sequentially, updating its hidden state over time. The output from the final time step, h_t, is then passed through a fully

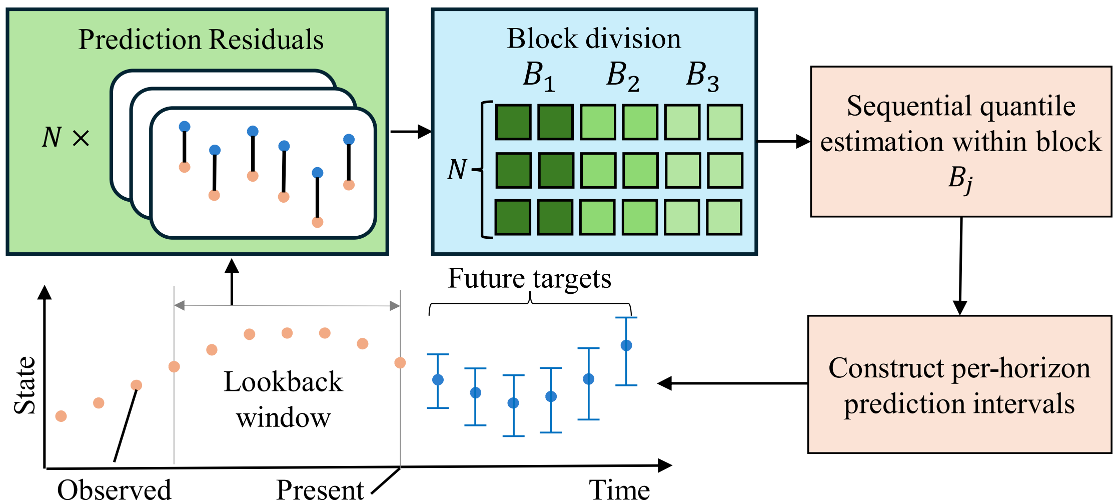

Figure 2. Overview of the online ConForME framework for probabilistic state forecasting.

connected linear layer to produce a multi-step forecast:

$$\hat{\mathbf{Y}}_{t+1:t+H} = Vh_t + c, \tag{3}$$

where $V \in \mathbb{R}^{H\times N}$ is the weight matrix of the linear layer and $c \in \mathbb{R}^H$ is the bias vector. This formulation allows the model to generate all H future steps in a single forward pass. In this paper, we set $L = 50$ samples (0.40 s of history) and $H = 25$ forecast steps (0.20 s into the future). We trained our model on the standard movement profile and evaluated it on the random movement profile from DROPBEAR Dataset-8.

Dynamic conformal prediction

To quantify uncertainty in our multi-step forecasting framework, we adopt ConForME (**Con**formal **For**ecast for **M**ulti-horizon pr**E**diction), a block-based conformal prediction method that constructs valid prediction intervals without relying on strong distributional assumptions [20]. ConForME partitions the prediction horizon into k non-overlapping blocks $B_1, \ldots, B_k$, each assigned an error budget α_j, such that the total error satisfies $\sum_{j=1}^{k} \alpha_j \leq \alpha$, where α is the desired overall error rate.

Prediction intervals within each block are constructed sequentially. For example, consider the forecast steps $t_1, t_2, \ldots, t_n \in B_j$, where n is the size of block B_j. For the first step t_1, the full calibration set is used to construct the conformal interval. For each subsequent step $t_2, t_3, \ldots$, the calibration set is iteratively filtered to retain only those sequences for which all previous intervals in the block correctly covered the true values. This conditioning ensures that each interval is valid given the correctness of the earlier ones, thereby preserving the joint blockwise coverage guarantee. The total miscoverage across block B_j is thus bounded by

$$1 - \prod_{t\in B_j} (1 - \alpha_t) \leq \alpha_j. \tag{4}$$

This sequential filtering strategy allows flexible allocation of uncertainty across forecast steps and yields tighter intervals than uniform-error strategies such as Bonferroni correction. Unlike Bonferroni, which allocates error uniformly and often leads to overly conservative intervals, ConForME adapts to the relative difficulty of each step, resulting in more informative uncertainty estimates.

To extend ConForME for real-time forecasting, we introduce an online calibration mechanism. Figure 2 illustrates this process, showing how recent prediction residuals are collected and grouped into blocks for horizon-specific conformal quantile computation. At each time t, predictions are made using input features $\{x_{t-L+1}, \ldots, x_t\}$, for which the corresponding target y_t is already observed. This enables dynamic updating of the calibration set. For each forecast horizon $h = 1, \ldots, H$, we maintain a rolling calibration window of N recent predictions, each associated with its forecasted and observed values at horizon h. This setup ensures that residuals used for conformal calibration at each horizon are based on sequences for which the model has already issued horizon-h predictions, supporting horizon-specific and adaptively updated prediction intervals.

RESULTS AND DISCUSSION

To better understand how dynamic ConForME operated, we analyzed its behavior at a specific time step. We targeted a coverage level of 90% ($\alpha = 0.1$). For ConForME, we set the block size to $k = 1$, treating the entire prediction horizon as a single block, and distributed the total error budget evenly across all forecast steps using an error rate of α/H at each time step. Figure 3 illustrates an example of the ConForME process. At time $t = 7.058$ s, the method first collects the most recent N calibration samples that predicted y_t at each forecast horizon h, as shown in Figure 3(a). It then computes the conformity scores for each respective horizon, illustrated in Figure 3(b). These scores are used to derive empirical distributions, from which quantile thresholds are extracted, as seen in Figure 3(c). Finally, the resulting quantiles are used to construct the prediction intervals at time $t = 7.058$ s, visualized in Figure 3(d). This process is repeated for each horizon, with a filtration step that includes only calibration sequences where all previous intervals in the block successfully covered their corresponding true values. It can be observed that the conformity score increases at higher forecast horizons. This is because the model tends to produce larger errors as the prediction extends further into the future, reflecting growing uncertainty over time.

We compared our proposed extended ConForME method with other baseline methods. Figure 4 plots the mean prediction interval width against empirical coverage for each method, evaluated at varying calibration set sizes (expressed as a percentage of the training data). The results show that our extended ConForME consistently achieved or slightly exceeded the target coverage level of 90% in all 10 trials, even at lower calibration sizes (e.g. 10%), with an average empirical coverage of 93%. In contrast, the standard ICP method frequently underperformed, achieving coverage as low as 49% at the 10% calibration size, indicating its sensitivity to distributional shifts in the data. The dynamic ICP improved upon this, reaching up to average empirical coverage of 76%, but still fell short of the target in all trials. The dynamic ICP improved upon this, reaching up to an average empirical coverage of 76%, but still fell short of the target in all trials.

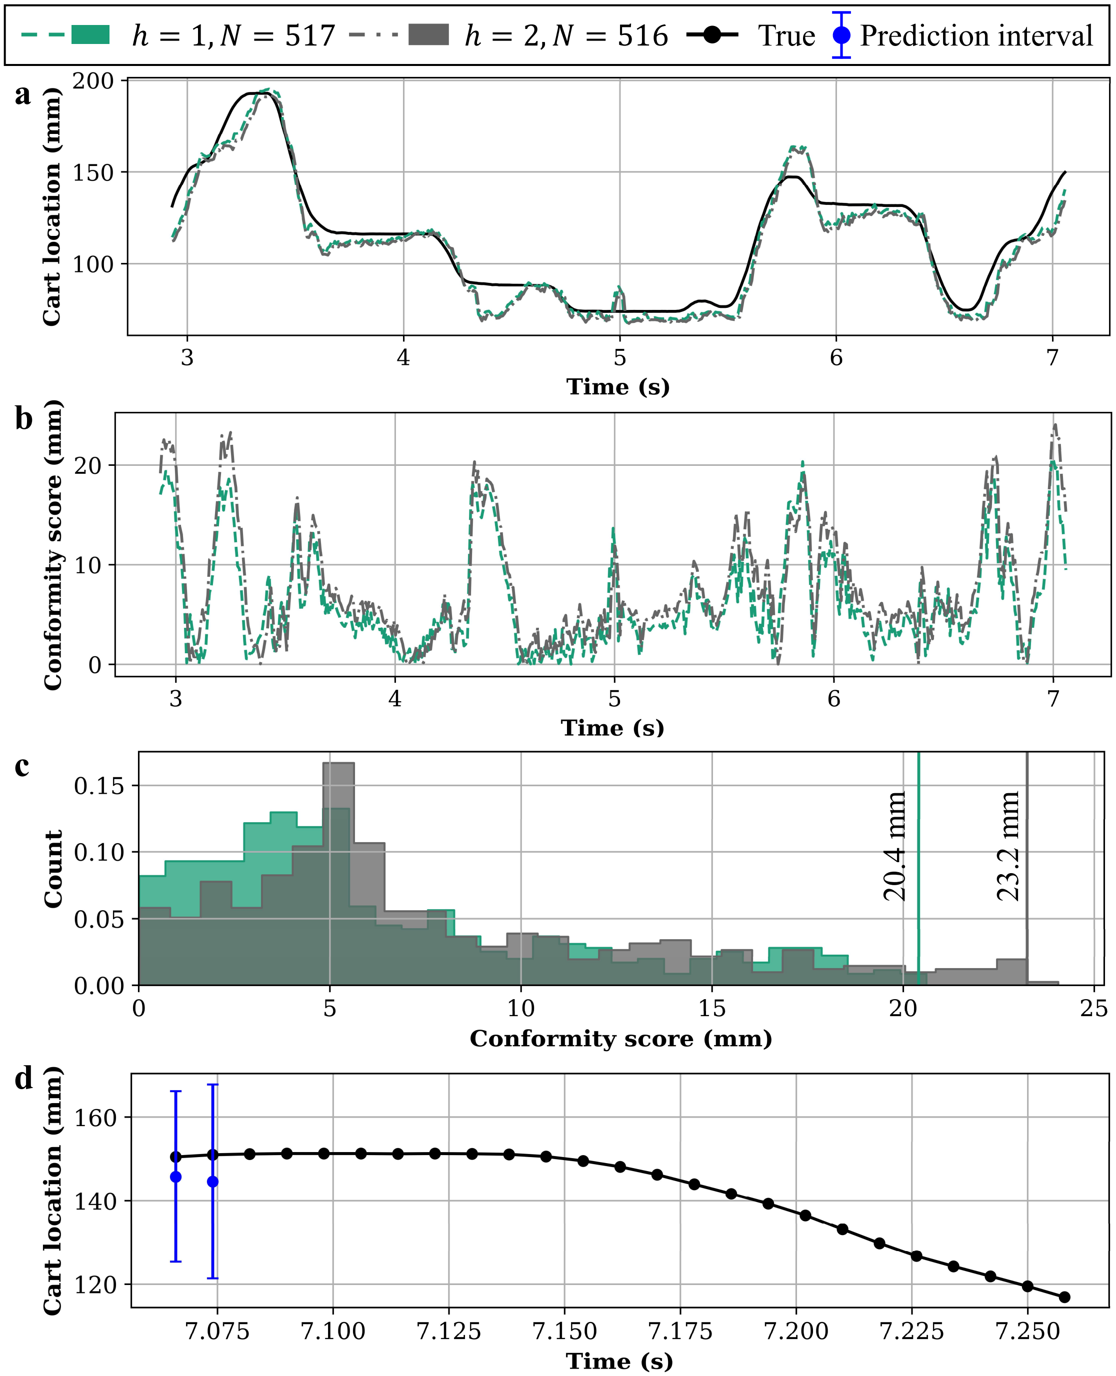

Figure 3. Overview of the dynamic ConForME process. (a) The most recent N predicted values for each forecast horizon h are compared to the corresponding true values to compute conformity scores. (b) Conformity scores are collected for each horizon. (c) Quantile thresholds (solid vertical lines) are extracted from the empirical distributions. (d) Final prediction intervals are constructed. The process is repeated for $h = 2$ with fewer calibration samples due to the filtering constraint.

While ConForME demonstrated reliability in terms of coverage, this came at the cost of wider prediction intervals. On average, ConForME produced intervals of approximately 93.42 mm, compared to 60.72 mm for dynamic ICP and 30.37 mm for standard ICP. For context, the beam length in the DROPBEAR setup is 501 mm, so the intervals represent roughly 18.64% of the total length. This suggests a trade-off between achieving reliable coverage and maintaining narrow intervals.

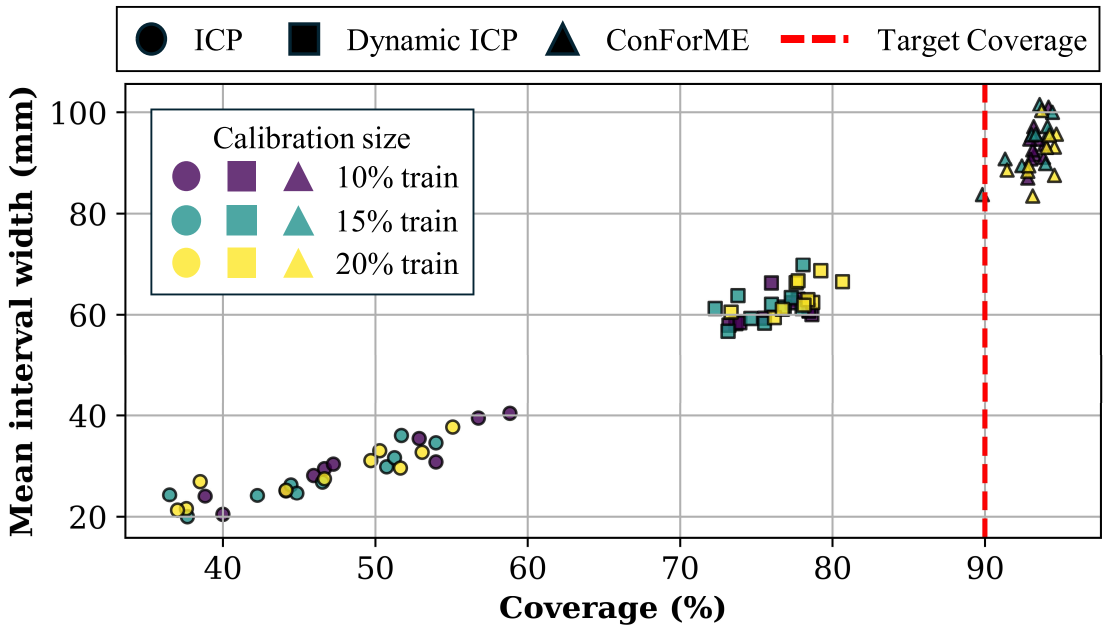

Figure 4. Performance of different conformal prediction methods across various calibration set sizes, expressed as a percentage of the training data. Multiple icons of the same type represent multiple trials at the same calibration size.

Discussion

The objective of this study was to introduce conformal prediction as a method for uncertainty quantification in high-rate state estimation forecasting. The primary focus was on comparing multiple conformal prediction methods to determine which best achieves the target empirical coverage. Our evaluation revealed that the standard ICP method performed poorly in the presence of distribution shifts within the dataset, leading to the lowest empirical coverage. Notably, the model was trained on a standard movement profile but tested on a different movement profile, contributing to the observed distribution shift. To handle the distribution shift, we extended ICP into a dynamic version with a sliding calibration window; this improved coverage, though it still fell short of the target. This may be attributed to its underlying assumption of independence across forecast horizons, which neglects the temporal correlations in prediction errors. In contrast, our proposed extension of ConForME consistently achieved the desired coverage level. However, despite this success, there remains room for improvement—the primary limitation being the width of the prediction intervals. Future work includes exploring more advanced conformal prediction techniques that may reduce interval width without compromising coverage. Additionally, part of the wide intervals may be due to forecasting errors from the base model itself, indicating that improving model accuracy could further enhance uncertainty quantification.

CONCLUDING REMARKS

This paper presented a preliminary study on the use of conformal prediction as a method for uncertainty quantification in high-rate state estimation forecasting. The problem of interest was developing a conformal prediction method capable of achieving uncertainty quantification. Several conformal prediction methods were explored, and their ability to achieve uncertainty quantification was assessed. A promising conformal prediction method, namely extended ConForME, was selected. To address distribution shifts in the dataset, a dynamically updating calibration dataset was proposed and applied to laboratory data obtained from the DROPBEAR testbed. The results showed that the extended ConForME method was able to achieve the target coverage. However, further investigations and refinement are required to improve the width of the prediction intervals. Future work will focus on optimizing error rate assignments to reduce the interval width by improving the baseline model to achieve better performance, before incorporating the conformal prediction methods.

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 awards numbers CCF-1937460 and 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.

REFERENCES

1. Dodson, J., A. Downey, S. Laflamme, M. D. Todd, A. G. Moura, Y. Wang, Z. Mao, P. Avitabile, and E. Blasch. 2022. "High-rate structural health monitoring and prognostics: An overview," in *Data Science in Engineering, Volume 9: Proceedings of the 39th IMAC, A Conference and Exposition on Structural Dynamics 2021*, Springer, pp. 213–217.
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. Yan, J., S. Laflamme, J. Hong, and J. Dodson. 2021. "Online parameter estimation under non-persistent excitations for high-rate dynamic systems," *Mechanical Systems and Signal Processing*, 161:107960.
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.
5. 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.
6. 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.

7. Razmarashooli, A., D. A. S. Martinez, Y. K. Chua, S. Laflamme, and C. Hu. 2024. "Real-time state estimation using recurrent neural network and topological data analysis," in *Non-destructive Characterization and Monitoring of Advanced Materials, Aerospace, Civil Infrastructure, and Transportation XVIII*, SPIE, vol. 12950, pp. 106–114.
8. Shafer, G. and V. Vovk. 2008. "A tutorial on conformal prediction." *Journal of Machine Learning Research*, 9(3).
9. Papadopoulos, H., K. Proedrou, V. Vovk, and A. Gammerman. 2002. "Inductive confidence machines for regression," in *Machine learning: ECML 2002: 13th European conference on machine learning Helsinki, Finland, August 19–23, 2002 proceedings 13*, Springer, pp. 345–356.
10. Gibbs, I. and E. Candes. 2021. "Adaptive conformal inference under distribution shift," *Advances in Neural Information Processing Systems*, 34:1660–1672.
11. Feldman, S., L. Ringel, S. Bates, and Y. Romano. 2022. "Achieving risk control in online learning settings," *arXiv preprint arXiv:2205.09095*.
12. Xu, C. and Y. Xie. 2023. "Conformal prediction for time series," *IEEE transactions on pattern analysis and machine intelligence*, 45(10):11575–11587.
13. Stankeviciute, K., A. M Alaa, and M. van der Schaar. 2021. "Conformal time-series forecasting," *Advances in neural information processing systems*, 34:6216–6228.
14. Sun, S. H. and R. Yu. 2023. "Copula conformal prediction for multi-step time series prediction," in *The Twelfth International Conference on Learning Representations*.
15. Lindemann, L., M. Cleaveland, G. Shim, and G. J. Pappas. 2023. "Safe planning in dynamic environments using conformal prediction," *IEEE Robotics and Automation Letters*, 8(8):5116–5123.
16. Joyce, B., J. Dodson, S. Laflamme, and J. Hong. 2018. "An Experimental Test Bed for Developing High-Rate Structural Health Monitoring Methods," *Shock and Vibration*, 2018:1–10, ISSN 1875-9203, doi:10.1155/2018/3827463.
17. Downey, A., J. Hong, J. Dodson, M. Carroll, and J. Scheppegrell. 2020, "Dataset-2-DROPBEAR-Acceleration-vs-Roller-Displacement," .
18. Vereen, A., A. Downey, J. Dodson, and A. G. Moura. 2023, "Dataset-8-DROPBEAR-Acceleration-vs-Roller-Displacement," .
19. Jordan, M. I. 1997. "Serial order: A parallel distributed processing approach," in *Advances in psychology*, Elsevier, vol. 121, pp. 471–495.
20. Lopes, A. G., E. Goubault, S. Putot, and L. Pautet. 2024. "ConForME: Multi-horizon conformal time series forecasting," *Proceedings of Machine Learning Research*, 230:1–21.