Conformal Prediction for Time-series Forecasting with Change Points

Conformal prediction has been explored as a general and efficient way to provide uncertainty quantification for time series. However, current methods struggle to handle time series data with change points — sudden shifts in the underlying data-generating process. In this paper, we propose a novel Conformal Prediction for Time-series with Change points (CPTC) algorithm, addressing this gap by integrating a model to predict the underlying state with online conformal prediction to model uncertainties in non-stationary time series. We prove CPTC's validity and improved adaptivity in the time series setting under minimum assumptions, and demonstrate CPTC's practical effectiveness on 6 synthetic and real-world datasets, showing improved validity and adaptivity compared to state-of-the-art baselines.

This paper builds upon the practical observation that in real-world scenarios, distribution shifts in time-series are often *predictable*. Take for example the task of forecasting electricity demands: we know the underlying dynamics differs between day and night, and during weekdays and weekends. CPTC explores the case of online conformal prediction when we can *anticipate* distribution shift. The algorithm is designed such that it greatly improves coverage compared to purely reactive methods, while still maintaining the ability to adapt to unknown shifts (and does not need to re-train a model every time step). There have not been other algorithms that explore the same problem; CPTC provides practitioners with a simple and sound solution, when they have the ability (such as via a state space model) to predict shifts.

Our contributions are:

- A new algorithm **C**onformal **P**rediction for **T**ime series with **C**hange points (CPTC) that utilizes state transition predictions to improve uncertainty quantification for time-series forecasts. Leveraging properties of a SDS model, CPTC consolidates multiple future forecasts to adaptively adjust its prediction intervals when underlying dynamics shift.
- We prove that CPTC achieves asymptotic valid coverage, *without any assumptions* on the time series data generation process, or state transition model accuracy. When predicted state transitions align well with data distribution shifts, CPTC can *anticipate* uncertainty and adapt faster.
- We show strong empirical results on 3 synthetic and 3 real-world datasets. Compared to online conformal prediction baselines, CPTC achieves more robust coverage with comparable prediction intervals sharpness (example in Figure 1); compared to residual regression baselines, CPTC is computationally much more efficient and can scale to long series.

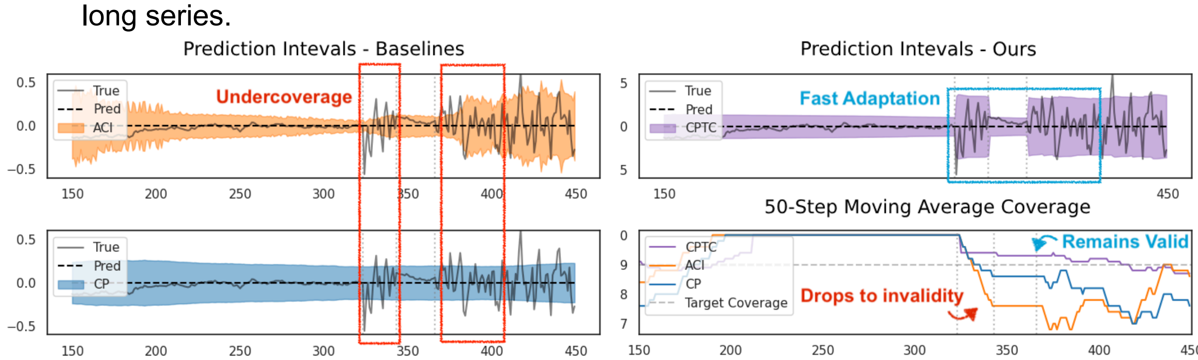

Figure 1. Comparison of the prediction intervals obtained by our algorithm CPTC (purple) against online conformal prediction baselines on synthetic data. The vertical dashed line marks the distribution shifts; ideal behavior is consistent coverage at the horizontal dashed line in the final panel. The bottom right panel shows that CPTC achieves fast adaptation and remains valid when change points occur.