Model Card for TACTiS

Intended use

TACTiS is a flexible model designed for probabilistic prediction of multivariate time series: forecasting, interpolation, or arbitrary combination of these patterns. TACTiS is intended for experts in stochastic optimization and strategic forecasting in the absence of specialized models.

Model type

Transformer-based multivariate time-series prediction model with an attentional copula and normalizing flows on the decoder to model inverse marginal CDFs.

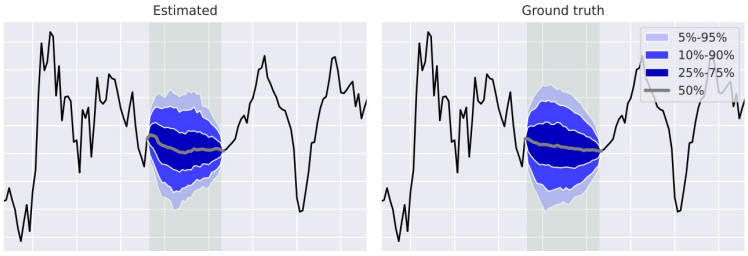

Comparison of TACTiS interpolation with ground truth.

Metrics

CRPS, CRPS-sum, Energy score: for all of them, lower is better. Note that they all tend to overlook variable dependencies and thus their interpretation must be completed by further inspection of learned variable dependencies.

Model Performances

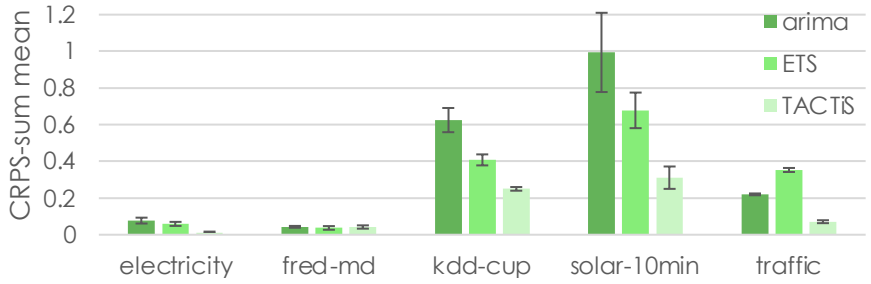

We report CRPS-sum means. Similar patterns can be observed for the other metrics.

Factors and limitations

Salient factors affecting model performance depend on the application domain in which TACTiS is used. The model does not mitigate factors. TACTiS assumes continuous data and is not intended for domain transfer. Appendix E.1 of the paper analyzes some bad forecasting examples.

Ethical considerations

TACTiS can be used in many applications, with some relying on sensitive data. According to the application domain, significant errors in the model can have a larger magnitude of harm though TACTiS is not recommended for high-stakes applications that require increased robustness, particularly long-tail events.

Training TACTiS requires significant computing resources and took up to 3 days for the above datasets, which have environmental and financial impacts.

Model details

Released: June 2022
Resources: **Paper** **Repo** **BibTeX**
License: Apache 2.0
Contact: Queries can be addressed on our repo.

Evaluation data

Electricity contains 321 series with hourly measures and daily & weekly seasonalities. Its prediction length is 24.

fred-md contains 107 series with monthly measures and a yearly seasonality. Its prediction length is 12.

kdd-cup contains 270 series with hourly measures and daily & weekly seasonalities. Its prediction length is 48.

solar-10min contains 137 series with measurements every 10 minutes and a daily seasonality. Its prediction length is 72.

traffic contains 862 series with hourly measures and daily & weekly seasonalities. Its prediction length is 24.

Inspired by industry practices, back-testing splits start in January for fred-md and on Monday for the other datasets. Missing data is padded with 0.

Training data

The training was carried out on the above datasets. Series alignment relies on **gluonTS**. Missing data is padded with the series mean.

Additional information

Usage that require further testing include zero-shot and long-tail event predictions and a domain transfer.