Forecasting Power Grid Frequency Trajectories with Structured State Space Models

Sebastian Pütz
Karlsruhe Institute of Technology
Karlsruhe, Germany
sebastian.puetz@kit.edu

Benjamin Schäfer
Karlsruhe Institute of Technology
Karlsruhe, Germany
benjamin.schaefer@kit.edu

ABSTRACT

Improving our ability to model, predict, and understand power system dynamics is becoming increasingly important as we face the challenges of transitioning to a carbon-neutral energy system. The power grid frequency is central to power system control as it is the primary observable for balancing generation and demand on short time scales. By facilitating frequency control actions, accurate prediction of grid frequency can improve system stability. In recent years, promising new deep learning techniques for time series forecasting tasks have emerged. Here, we explore the application of structured state space models (S4) to high-resolution power system frequency time series. S4 models have previously demonstrated good performance for long-term dependence tasks, but how useful are they for high-resolution energy time series?

KEYWORDS

time series forecasting, power grid, deep learning, state-space model

ACM Reference Format:
Sebastian Pütz and Benjamin Schäfer. 2023. Forecasting Power Grid Frequency Trajectories with Structured State Space Models. In *The 14th ACM International Conference on Future Energy Systems (e-Energy '23 Companion), June 20–23, 2023, Orlando, FL, USA.* ACM, New York, NY, USA, 2 pages. https://doi.org/10.1145/3599733.3604769

1 INTRODUCTION

A power grid is not in itself a storage facility, but only transmits energy. The stable operation of a power grid therefore crucially depends on power generation and demand being balanced at all times. The preservation of this supply and demand balance is reflected in the deviation from the nominal frequency of the power system, which makes it the central metric for assessing system stability and determining control measures on short time scales.

The grid frequency dynamics exhibits a complex mixture of deterministic (e.g. market-driven) and stochastic (e.g. fluctuations of renewable generation) characteristics and is impacted by various external factors. The ability to model and anticipate the grid frequency supports the planning of control measures and thereby improves grid stability.

Recent work on forecasting the power grid frequency investigated its predictability via Weighted Nearest Neighbour (WNN) searches for the profiles with the closest resemblance [3, 5]. While these approaches showed promising results in matching past and future patterns, we aim to learn the frequency dynamics itself and thereby freely evolve trajectories in time, potentially allowing interpretation of the learned model or transfer-learning.

In the past decade, various powerful deep learning techniques for time series forecasting have emerged. However, if or how grid frequency forecasting benefits from these techniques largely remains an open question. In this work we explore the use of structured state space models introduced by Gu et al. [2] to forecast power grid frequency. These models have two distinctive properties: First, they are very good at processing information over large time distances and thus learning long-term dependencies, and second, they inherently treat the time series as sampled from an underlying continuous-time process.

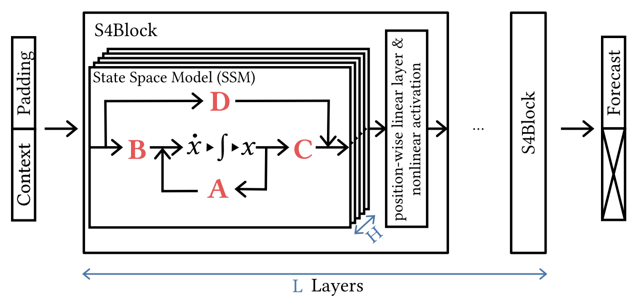

Figure 1: Structure of the S4 model. The model receives an input sequence zero-padded with the forecast horizon which is propagated through L layers. Each layer consists of H stacked independent linear State Space Models (SSMs). The H output sequences of these SSMs are then mixed with a position-wise linear layer with nonlinear activation functions. Finally, only the second part is used as the forecast while the first part is discarded.

2 DATA

We have used power grid frequency time series recorded on the Balearic Islands at a resolution of one second. Based on previous work, the Balearic grid displayed very regular and deterministic patterns, making it a good first candidate to explore the feasibility of S4 forecasting. These frequency time series were recorded by Electrical Data Recorders (EDRs) developed at KIT and are publicly available, see [6] for details and links to data repositories.

3 METHODS

In this work, we utilize the S4 model (see Fig. 1). The S4 model is a sequence-to-sequence model that is built from H parallel and L

consecutive blocks of linear State Space Models (SSMs) defined by:

$$\dot{x}(t) = Ax(t) + Bu(t) \tag{1}$$

$$y(t) = Cx(t) + Du(t), \tag{2}$$

where the discretized counterparts of the state matrix A as well as the other SSM parameters (B, C & D) are learned during training. The parameters of the SSMs are initialized to map onto a basis of orthogonal polynomials and the state matrix is structured to enable fast computation of the SSMs via convolution kernels. In each block, the H output sequences of the SSMs are mixed via a position-wise linear layer with nonlinear activations.

For our forecasting model, we input our context (initial time series) concatenated with a sequence of zeros set to the length of our forecasting window. In the end, we only consider the last entries of the output sequence to calculate the loss.

We aim to achieve forecasting horizons on the order of one hour (several 1000 time steps), since one hour marks a central time scale for market activity and regular patterns, while simultaneously the grid needs to be balanced for every single second. A challenge for such long horizons is to combine very short-term and longer-term dynamics in a single model. Our contribution is to explore how well S4 models can handle this challenge.

We compare our results to very simple baselines such as the 50 Hz constant predictor and the average daily frequency profile. Moreover, we compare our model to the WNN predictor from [3, 5], using a Root Mean Squared Error (RMSE) to assess performance.

4 PRELIMINARY RESULTS

The trained S4 model outperforms the simple baseline models at the beginning of the forecast horizon (see Fig. 2). After about 10 minutes (600 timesteps), the gap diminishes, but even after more than 1000 steps, it remains on par with the baseline models.

These initial results are already encouraging to pursue this approach further. We will now try to improve the performance, especially with regard to the long-term dependencies, through a more thorough search for suitable hyperparameters and changes in the model structure and training procedure.

5 DISCUSSION AND FUTURE WORK

In this contribution, we have motivated the use of S4 models to forecast highly resolved time series, such as power grid frequency trajectories. Our very preliminary findings indicate that long-term dependencies of frequency dynamics may indeed be learned by the S4 model and existing benchmarks might be outperformed on certain time horizons. These first results will be complemented in several aspects in the future.

Since the power system is a non-autonomous dynamical system influenced by various external factors, a natural next step is to extend the frequency prediction models to include those factors for which data is available. Useful features include scheduled generation as well as load, weather or renewable generation forecasts. An interesting extension of our approach could take the intrinsic stochastic variability [1] into account when generating predictions.

Most deep learning models are inherently black boxes, i.e. they provide no insight into their mappings of inputs to outputs. In recent years, however, eXplainable Artificial Intelligence (XAI) has gained traction in the energy informatics community [4]. XAI is a rapidly evolving field that is producing various methods that enable knowledge discovery from AI models. A major focus of future work will be to attempt to extract scientific insights from frequency forecasting models using methods from XAI.

Furthermore, we plan to compare performance with other deep neural network architectures, assess the capability of transfer or zero-shot learning of frequency dynamics of individual grids, and apply our learning to other highly resolved power system data.

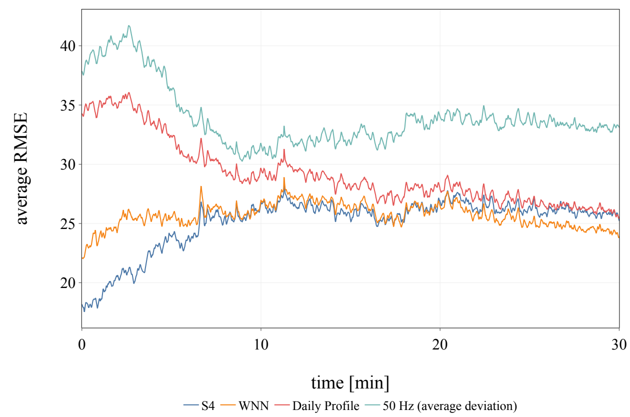

Figure 2: Preliminary results of the S4 model against benchmarks. The S4 model outperforms all baselines at the beginning of the forecast horizon and still keeps up with the baselines after more than 1000 timesteps.

ACKNOWLEDGMENTS

We gratefully acknowledge funding from the Helmholtz Association Initiative and Networking Fund through Helmholtz AI, the HAICORE@KIT partition and grant no. VH-NG-1727.

REFERENCES

[1] Leonardo Rydin Gorjão, Mehrnaz Anvari, Holger Kantz, Christian Beck, Dirk Witthaut, Marc Timme, and Benjamin Schäfer. 2020. Data-Driven Model of the Power-Grid Frequency Dynamics. 8 (2020), 43082–43097. https://doi.org/10.1109/ACCESS.2020.2967834

[2] Albert Gu, Karan Goel, and Christopher Ré. 2022. Efficiently Modeling Long Sequences with Structured State Spaces. In *The International Conference on Learning Representations (ICLR)*. https://openreview.net/forum?id=uYLFoz1vlAC

[3] Johannes Kruse, Benjamin Schäfer, and Dirk Witthaut. 2020. Predictability of Power Grid Frequency. 8 (2020), 149435–149446. https://doi.org/10.1109/ACCESS.2020.3016477

[4] R. Machlev, L. Heistrene, M. Perl, K. Y. Levy, J. Belikov, S. Mannor, and Y. Levron. 2022. Explainable Artificial Intelligence (XAI) Techniques for Energy and Power Systems: Review, Challenges and Opportunities. 9 (2022), 100169. https://doi.org/10.1016/j.egyai.2022.100169

[5] Thorbjørn Lund Onsaker, Heidi S. Nygård, Damiá Gomila, Pere Colet, Ralf Mikut, Richard Jumar, Heiko Maass, Uwe Kühnapfel, Veit Hagenmeyer, and Benjamin Schäfer. 2023. Predicting the Power Grid Frequency of European Islands. 4, 1 (2023), 015012. https://doi.org/10.1088/2632-072X/acbd7f

[6] Leonardo Rydin Gorjão, Richard Jumar, Heiko Maass, Veit Hagenmeyer, G Cigdem Yalcin, Johannes Kruse, Marc Timme, Christian Beck, Dirk Witthaut, and Benjamin Schäfer. 2020. Open database analysis of scaling and spatio-temporal properties of power grid frequencies. *Nature communications* 11, 1 (2020), 6362.