

Comparative Analysis of Zero-Shot Capability of Time-Series Foundation Models in Short-Term Load Prediction

Lin, Nan; Yun, Dong; Xia, Weijie; Palensky, Peter; Vergara, Pedro P.

DOI

[10.1109/PowerTech59965.2025.11180251](https://doi.org/10.1109/PowerTech59965.2025.11180251)

Publication date

2025

Document Version

Final published version

Published in

2025 IEEE Kiel PowerTech, PowerTech 2025

Citation (APA)

Lin, N., Yun, D., Xia, W., Palensky, P., & Vergara, P. P. (2025). Comparative Analysis of Zero-Shot Capability of Time-Series Foundation Models in Short-Term Load Prediction. In *2025 IEEE Kiel PowerTech, PowerTech 2025* (2025 IEEE Kiel PowerTech, PowerTech 2025). IEEE.
<https://doi.org/10.1109/PowerTech59965.2025.11180251>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

**Green Open Access added to [TU Delft Institutional Repository](#)
as part of the Taverne amendment.**

More information about this copyright law amendment
can be found at <https://www.openaccess.nl>.

Otherwise as indicated in the copyright section:
the publisher is the copyright holder of this work and the
author uses the Dutch legislation to make this work public.

Comparative Analysis of Zero-Shot Capability of Time-Series Foundation Models in Short-Term Load Prediction

Nan Lin^{*1}, Dong Yun^{*2}, Weijie Xia^{*1}, Peter Palensky¹, Pedro P. Vergara¹

¹ Intelligent Electrical Power Grids (IEPG), Delft University of Technology, Delft, The Netherlands

² Department of Electrical and Photonics Engineering, Technical University of Denmark, Kgs., Lyngby, Denmark

Email: {W.Xia, N.Lin, P.P.VergaraBarrios, P.Palensky}@tudelft.nl, s232293@dtu.dk

Abstract—Short-term load prediction (STLP) is critical for modern power distribution system operations, particularly as demand and generation uncertainties grow with the integration of low-carbon technologies, such as electric vehicles and photovoltaics. In this study, we evaluate the zero-shot prediction capabilities of five Time-Series Foundation Models (TSFMs)—a new approach for STLP where models perform predictions without task-specific training—against two classical models, Gaussian Process (GP) and Support Vector Regression (SVR), which are trained on task-specific datasets. Our findings indicate that even without training, TSFMs like Chronos, TimesFM, and TimeGPT can surpass the performance of GP and SVR. This finding highlights the potential of TSFMs in STLP.

Index Terms—Short-term load prediction, Time-Series Foundation Model

I. INTRODUCTION

Short-term load prediction (STLP) is essential for the efficient operation of power distribution systems [1]. As modern power systems grow increasingly sophisticated, STLP serves as a fundamental tool for capturing dynamic fluctuations in load demand [2]. Consequently, it has significant applications such as demand response [3] and control optimization [4].

Before the prevalence of Deep Learning (DL), the STLP methods primarily relied on statistical models. For instance, in [5], an Auto-Regressive Moving Average (ARMA) model is introduced for STLP that incorporates considerations for non-Gaussian processes. The method utilizes bispectral analysis to assess the Gaussianity of load data. Results indicate that this approach achieves much lower error than a standard ARMA model. In [6], [7], the Decomposition technologies are used to decompose load data first, and then models like XGBoost and Support Vector Machine (SVM) are applied to model the trend and sub-series. These hybrid ways show better performance than a single model. Given the significant impact of weather factors on load, the study in [8] proposes an SVM model enhanced with a Grasshopper Optimization algorithm to address STLP under specific weather conditions. This method performs better than other hybrid approaches, such as GA-SVM. The study in [9] proposes an ensemble learning approach that uses models like regression trees and Random Forests as base learners. Like hybrid models, ensemble models consistently perform better than standalone models. The

study in [10] introduces a Support Vector Regression (SVR) model with a customized grid traverse algorithm for high-resolution (1-second) STLP. This approach clearly outperforms traditional models in high-resolution STLP, such as GA-SVM and ARIMA. In [11], [12], Gaussian Process (GP) models are applied for load prediction tasks. The integration of GP offers the advantage of reduced data requirements compared to standalone DL models.

However, with the growing availability of load data and DL advancements, DL-based STLP approaches have demonstrated superior performance compared to traditional statistical models in recent years. For instance, the study in [13] introduces a Graph Neural Network (NN) for STLP. Unlike models that consider only temporal information, the Graph NN approach incorporates spatial information, recognizing that the load patterns of households within a region exhibit similarities. The study in [14] proposes transforming load and weather data into an image format which allows Convolutional NN to be applied as forecasters. In [15], the author proposes a Long Short-Term Memory (LSTM) model with an Attention Mechanism, effectively capturing complex data correlations. This approach demonstrates better performance compared to a classical LSTM model. Despite the superior performance, one disadvantage of the above-mentioned models is they can only provide point prediction. The application of deep generative models in STLP tackles this issue. In [16], conditional Wasserstein GANs (cWGANs) are utilized for probabilistic load prediction, conditioned on weather data and historical load. This method outperforms traditional approaches, such as quantile regression. Additionally, studies in [17], [18] apply flow-based models for prediction, demonstrating that these models perform better than GANs and VAEs in probabilistic STLP.

However, both statistical and DL-based models require training on specific datasets, which presents two challenges: (1) Despite increased data availability from metering devices, data for individual devices (e.g., households or substations) often remains limited, leading to potential data shortages during training [19]; (2) Training, particularly for DL models across various datasets, can be highly time-consuming. Time-Series Foundation Models (TSFM), characterized by the capability of zero-shot prediction, provide a potential solution to

* indicates equal contribution.

these challenges. The current prevalent TSFM models include Chronos [20], Moment [21], Lag-llama [22], TimesFM [23], and TimeGPT [24], a detailed introduction of these TSFM is provided in Sec. II. Most of these models are based on Transformer architectures and are pre-trained on large, diverse time series datasets from multiple domains for various prediction tasks. As such, they are designed to function as "general forecasters" similar to the Large Language Model (LLM). This allows them to be used without additional model training—particularly beneficial in data-scarce or zero-data scenarios.

In this paper, we assess the zero-shot (meaning no training process for the TSFMs) prediction performance of state-of-the-art TSFM for STLP, benchmarking them against the widely adopted classical models, SVR and GP. This paper aims to explore a potential new research direction for the STLP task. To our knowledge, this study represents the first exploration of TSFMs in the context of load prediction. The code and data of this research are accessible in the project repository¹.

II. TIMES-SERIES FOUNDATION MODELS

Foundation models are models that are designed to be trained on extremely diverse and large amounts of data, such that they can be applied to unseen data with little to no adaptation. In the popular LLM context, it would mean training a model on a very large language dataset in order that anyone can start using it immediately without further training. The principle is similar for time series. A TSFM is trained on a large amount of time series to predict the future steps given the history. With a capable enough model and adequate data, the model would learn to generalize to unseen data. That is why these models would be able to perform zero-shot (zero training) forecasts on load profiles.

Since TSFMs share similar design principles, this section primarily focuses on introducing and comparing the differences between Chronos and Lag-llama to provide a comprehensive understanding of TSFM architectures. For detailed information on model design, readers are encouraged to read the original papers [20]–[24].

The primary difference between time series and language series lies in their data types: while a language series consists of discrete tokens representing specific words, a time series comprises a sequence of real values. The time series tokenization approach is commonly adopted in TSFM to adapt the Transformer architecture for time series and bridge the format gap between these data types. For instance, in Lag-llama [22], the time series is segmented equally to create a sequence of tokens. In Chronos [20], tokenization is achieved through a tokenization function $q(\cdot)$ that maps real value x_i into a token $z_i = q(x_i)$, Fig. 1 illustrates these two tokenization approaches.

Once the time series is tokenized, a standard Transformer architecture can be applied to process the data. However, another

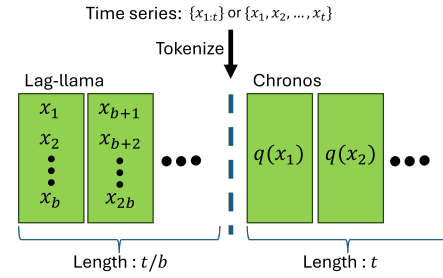


Fig. 1. Two tokenization approaches for time series. t is the length of the time series, and b is the length of the token. The Lag-llama's approach is to segment time series equally to create a sequence of tokens (assume t is divisible by b). The Chronos's approach is to map the real values to a token.




Fig. 2. Two output mechanisms. z_i is the token which is obtained by tokenization. $z_i = \{x_{i \times b + 1}, x_{i \times b + 2}, \dots, x_{i \times b + b}\}$ in Lag-llama, and $z_i = q(x_i)$ in Chronos.

challenge arises. Unlike LLMs, which output discrete tokens corresponding to specific words, the output of a TSFM is typically expected to be a sequence of continuous real values, reflecting the nature of the input time series data. To address this issue, Chronos adopts an approach similar to that of LLM, it outputs the probability distribution for each input token, and then samples input tokens based on these probabilities to obtain output tokens. The resulting tokens are subsequently detokenized using the detokenization function $d(\cdot)$ to generate predicted values $d(z)$. This method implicitly assumes that the time series remains stationary over time. In contrast, Lag-llama first predicts parameters of an exponential distribution rather than probabilities, then sampling values directly from this distribution as the prediction results. The output mechanism is shown in Fig. 2. Moreover, due to the output of these two models is essentially a distribution (or probability), we can sample multiple times from the distribution to obtain the probabilistic prediction results. Fig 3 illustrates a more detailed example of the TSFM model's prediction process using Chronos, including both tokenization and detokenization steps.

III. EXPERIMENTAL SETTING

A. Dataset

The dataset used in this study includes individual and aggregated electricity consumption data at 60-minute, 30-minute, and 15-minute resolutions from the UK [25], Germany [26], and the Netherlands [27]. Table I summarizes the characteristics of all datasets used in this research. The "Days" column in Table I indicates the total number of days for which load data is available.

¹Code and data are available in the following repositories:

- 1) Personal Repository
- 2) TU Delft Project Repository (need to be given)

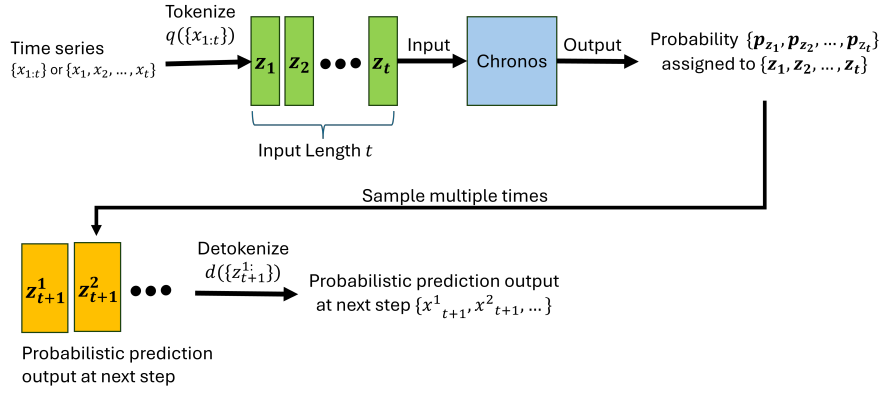


Fig. 3. The prediction process of Chronos. $\{z_{t+1}^i\}$ is sampled from $\{z_{1:t}\}$ based on the probability $\{p_{z_i}; i = 1, 2, \dots, t\}$, where $\sum_i^t p_{z_i} = 1$.

TABLE I
DATASETS USED FOR THE MODEL COMPARISON.

ID	Country	Resolution	Days	Date Range
<i>Individual Load Data</i>				
NL-I-60	NL	60-minutes	15,600	01.2013-12.2013
GE-I-60	GE	60-minutes	1,200	12.2014-05.2019
GE-I-30	GE	30-minutes	1,200	12.2014-05.2019
GE-I-15	GE	15-minutes	1,200	12.2014-05.2019
UK-I-30	UK	30-minutes	7,200	01.2013-12.2013
UK-I-60	UK	60-minutes	7,200	01.2013-12.2013
<i>Aggregated Load Data</i>				
NL-A-60	NL	60-minutes	600	01.2013-12.2013
GE-A-60	GE	60-minutes	200	12.2014-05.2019
GE-A-30	GE	30-minutes	200	12.2014-05.2019
GE-A-15	GE	15-minutes	200	12.2014-05.2019
UK-A-30	UK	30-minutes	2,400	01.2013-12.2013
UK-A-60	UK	60-minutes	2,400	01.2013-12.2013

Days: The total number of days for which load data is available.

TABLE II
COMPARISON OF MODEL FEATURES.

Model	Type	Probabilistic	Scale
Chronos-tiny	TFSM	Yes	8M
Chronos-small	TFSM	Yes	46M
Moment	TFSM	No	37.9M
TimeGPT	TFSM	No	/
TimeFM	TFSM	No	200M
Lag-llama	TFSM	YES	2.45M
SVR	Statistical Model	No	/
GP	Statistical Model	YES	/

Probabilistic: Whether the model supports probabilistic prediction.

Scale: The number of parameters of the model, "M" means millions. The scale of TimeGPT is not publicly released.

B. Experimental Design

Five TFSM models—Chronos, Moment, Lag-llama, TimesFM, and TimeGPT—are compared with SVR and GP as classical approaches for STLP. Table II summarizes the key features of all models used. For Chronos, multiple pre-trained model versions are available [20], and we selected the two versions shown in Table II, which differ only in the number of model parameters. The electricity consumption data is divided into training and test sets. The TFSMs are supplied with three days of historical load data in the test sets, resulting in 72×1 , 72×2 , and 72×4 data points for 60-minute, 30-minute, and 15-minute resolutions, respectively. These models then

forecast the load for the following day, producing 24×1 , 24×2 , and 24×4 data points. For the SVR and GP models, training is conducted on the training set (60%), with testing on the test sets (40%), allowing for direct comparison with the TFSMs.

C. Evaluation

For evaluation metrics, we calculate the Mean Absolute Error (MAE) and Root Mean Squared Error (RMSE) between the predicted average values (or the single predicted value for models that do not support probabilistic predictions) and the actual values. These metrics are expressed below

$$\text{MAE} = \frac{1}{T} \sum_{t=1}^T |x_t - \hat{x}_t|, \quad (1)$$

$$\text{RMSE} = \sqrt{\frac{1}{T} \sum_{t=1}^T (x_t - \hat{x}_t)^2}, \quad (2)$$

where x_t and \hat{x}_t are the real load value and predicted (average) load value at time index t . T is the prediction length. For models that support probabilistic predictions, we additionally compute the quantile losses at the 10%, 50%, and 90% levels. The metric for computing the quantile loss is expressed as

$$\text{QL}_\gamma = \frac{1}{T} \sum_{i=1}^T (\mathbb{I}_{\hat{x}_t \geq x_t} (1 - \gamma) |z_t - \hat{x}_t| + \mathbb{I}_{\hat{x}_t < x_t} \gamma |x_t - \hat{x}_t|), \quad (3)$$

where QL_γ is the quantile loss at quantile level $\gamma \in \{0.1, 0.5, 0.9\}$, $\mathbb{I}(\cdot)$ is Indicator function.

IV. RESULTS AND DISCUSSION

Tables III and IV present the experimental results for aggregated and individual datasets, respectively. Chronos consistently achieves superior performance in probabilistic prediction across both dataset types, attaining the lowest quantile losses at the 10%, 50%, and 90% levels. Overall, Chronos-small exhibits a slight performance advantage over Chronos-tiny, likely due to its larger model size as shown in Table II. Notably, despite training, GP does not outperform Chronos

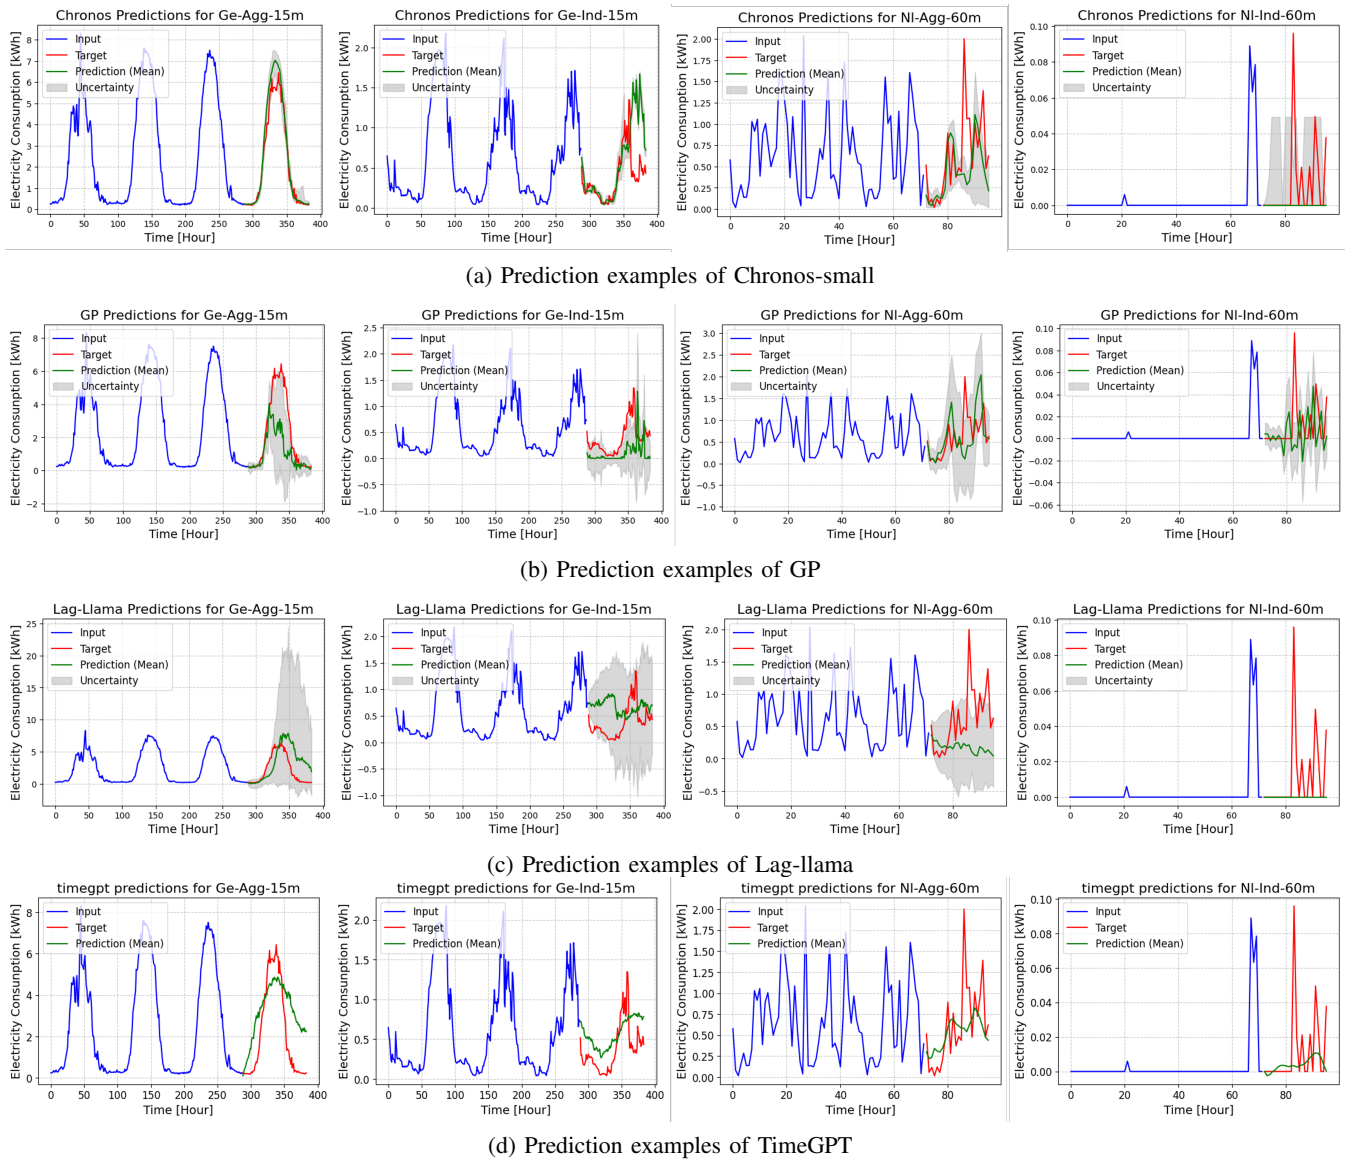


Fig. 4. Prediction examples of GP, TimeGPT, Lag-llama, and Chronos.

in any experiment except NL-I-60. Even compared with Lag-llama, a relatively weaker TSFM, GP fails to demonstrate clear advantages. For instance, Lag-llama outperforms GP in the NL-A-60, UK-A-60, and UK-I-30 examples.

Fig. 4 (a-d) provides prediction examples of prediction results from Chronos, GP, Lag-llama, and TimeGPT. As illustrated in Fig. 4 (b), GP frequently predicts negative values, which theoretically do not exist in the dataset (as load values are always positive). This issue, which contributes to higher quantile losses, arises from GP's reliance on the assumption that time series follow a multivariate Gaussian distribution—a simplification inadequate for modeling load pattern volatility, especially in individual datasets. In contrast, Chronos, as shown in Fig. 4 (a), avoids these inaccuracies. Additionally, Chronos more accurately captures peak load values in the Ge-A-15 and Ge-I-15 examples than GP. However, we observed that Chronos, along with nearly all models, fails to accurately

predict the load for the highly volatile IN-I-60 example shown in Fig. 4.

Lag-llama, as another TSFM model, performs worse than Chronos in terms of probabilistic prediction, as indicated in Tables III and IV. Fig. 4 (c) shows that, similar to GP, Lag-llama generates negative values with even larger variation (reflected by broad uncertainty intervals) and diverges notably in the NL-A-60 example. Lag-llama's weaker performance can be attributed to (1) its smaller scale, evident in the NL-A-60 example in Fig. 4 (c), where Lag-llama fails to capture the load data pattern, and (2) its reliance on a t-distribution assumption, which inevitably generates negative values similar to GP, as seen in the GE-A-15 and GE-I-15 examples in Fig. 4 (c). This limitation stems from the long-tailed nature of the t-distribution, which may be inadequate for modeling load data. In contrast, Chronos's good performance could first be attributed to its tokenization methods. In Chronos's prediction

TABLE III
EXPERIENTIAL RESULTS OF AGGREGATED LOAD DATA

Model	Q 10%	Q 50%	Q 90%	MAE	RMSE
<i>Experiment ID : GE-A-15</i>					
Moment	/	/	/	1.3186	1.5784
GP	0.1969	0.4796	0.2676	0.9592	2.2039
SVR	/	/	/	0.8406	1.6161
Chronos-small	0.1014	0.2258	0.1540	0.4517	0.7040
Chronos-tiny	<u>0.0968</u>	<u>0.2171</u>	<u>0.1503</u>	<u>0.4343</u>	<u>0.6863</u>
Lag-llama	0.2968	0.5885	0.5249	1.1771	1.6455
TimeGPT	/	/	/	1.0419	1.2109
TimesFM	/	/	/	1.1747	1.6302
<i>Experiment ID : GE-A-30</i>					
Moment	/	/	/	2.6974	3.1888
GP	<u>0.3876</u>	0.9202	0.5104	1.8404	2.8721
SVR	/	/	/	1.6922	2.5012
Chronos-small	0.4834	<u>0.6457</u>	0.4023	<u>1.2914</u>	1.9006
Chronos-tiny	0.4719	0.6524	<u>0.3895</u>	1.3048	1.9307
Lag-llama	0.4999	1.0307	0.7717	2.0614	2.8865
TimeGPT	/	/	/	2.4966	2.9227
TimesFM	/	/	/	1.3126	<u>1.8802</u>
<i>Experiment ID : GE-A-60</i>					
Moment	/	/	/	5.3237	6.3017
GP	0.7755	1.8548	1.0545	3.7095	5.7497
SVR	/	/	/	3.4079	5.4665
Chronos-small	0.4835	0.8018	0.3894	1.6037	2.4774
Chronos-tiny	<u>0.4616</u>	<u>0.7859</u>	<u>0.3859</u>	<u>1.5718</u>	<u>2.4190</u>
Lag-llama	0.6847	1.9988	1.6902	3.9977	5.4701
TimeGPT	/	/	/	2.0881	2.9514
TimesFM	/	/	/	2.1064	3.1958
<i>Experiment ID : UK-A-30</i>					
Moment	/	/	/	0.2262	0.2947
GP	0.2602	0.1176	0.2971	0.2353	0.4246
SVR	/	/	/	0.2591	0.3133
Chronos-small	0.0217	<u>0.0691</u>	0.0868	<u>0.1383</u>	0.2392
Chronos-tiny	<u>0.0213</u>	0.0692	<u>0.0862</u>	<u>0.1383</u>	0.2382
Lag-llama	0.0452	0.1192	0.1351	0.2384	0.3276
TimeGPT	/	/	/	0.1456	<u>0.2066</u>
TimesFM	/	/	/	0.1560	0.2317
<i>Experiment ID : UK-A-60</i>					
Moment	/	/	/	0.2299	0.3002
GP	0.0448	0.0852	<u>0.0536</u>	0.1703	0.2386
SVR	/	/	/	0.1801	0.2190
Chronos-small	0.0222	0.0717	0.0826	<u>0.1433</u>	0.2324
Chronos-tiny	<u>0.0221</u>	0.0720	0.0828	0.1441	0.2407
Lag-llama	0.0390	0.1133	0.1231	0.2266	0.3163
TimeGPT	/	/	/	0.1504	0.2202
TimesFM	/	/	/	0.1449	<u>0.2188</u>
<i>Experiment ID : NL-A-60</i>					
Moment	/	/	/	2.7385	3.1118
GP	0.3064	0.6320	<u>0.3115</u>	1.2639	1.5064
SVR	/	/	/	1.3312	1.6444
Chronos-small	0.2578	0.5330	0.3974	1.0660	1.3988
Chronos-tiny	<u>0.2471</u>	<u>0.5290</u>	0.3980	<u>1.0580</u>	1.3936
Lag-llama	0.8456	1.8532	1.7357	3.7064	4.1879
TimeGPT	/	/	/	1.0879	<u>1.3697</u>
TimesFM	/	/	/	1.1175	1.4481

Q $h\%$: Meaning the quantile loss at quantile level h .

process, as shown in Fig. 3 Chronos's prediction is essentially randomly sampled tokens from the historical tokens, this will ensure that the predicted load values are non-negative (as there are no negative historical load values). However, this tokenization strategy might affect Chronos's ability to predict unseen peak values.

In terms of point prediction, GP and SVR—classical models that require training—do not show advantages over TSFMs, consistently yielding relatively higher MAE and RMSE values.

TABLE IV
EXPERIENTIAL RESULTS OF INDIVIDUAL LOAD DATA

Model	Q 10%	Q 50%	Q 90%	MAE	RMSE
<i>Experiment ID : GE-I-15</i>					
Moment	/	/	/	0.2761	0.3513
GP	0.0475	0.1216	0.0831	0.2432	0.3498
SVR	/	/	/	0.2415	0.3463
Chronos-small	0.0319	0.0716	0.0687	0.1433	<u>0.2202</u>
Chronos-tiny	<u>0.0289</u>	<u>0.0700</u>	<u>0.0679</u>	<u>0.1399</u>	0.2231
Lag-llama	0.0526	0.1425	0.1290	0.2849	0.3730
TimeGPT	/	/	/	0.1845	0.2313
TimesFM	/	/	/	0.1779	0.2622
<i>Experiment ID : GE-I-30</i>					
Moment	/	/	/	0.5183	0.6636
GP	0.0905	0.2574	0.2173	0.5148	0.7506
SVR	/	/	/	0.4267	0.6020
Chronos-small	0.0550	<u>0.1410</u>	<u>0.1366</u>	<u>0.2821</u>	0.4593
Chronos-tiny	<u>0.0497</u>	0.1422	0.1375	0.2845	0.4636
Lag-llama	0.0818	0.2362	0.2318	0.4723	0.6644
TimeGPT	/	/	/	0.4108	0.5106
TimesFM	/	/	/	0.2947	<u>0.4177</u>
<i>Experiment ID : GE-I-60</i>					
Moment	/	/	/	1.0095	1.2799
GP	0.1845	0.5346	0.4401	1.0692	1.2763
SVR	/	/	/	0.8028	1.1808
Chronos-small	0.1122	0.2232	<u>0.1636</u>	0.4464	0.7146
Chronos-tiny	<u>0.0945</u>	<u>0.2189</u>	0.1656	<u>0.4378</u>	0.7034
Lag-llama	0.1504	0.5067	0.5046	1.0135	1.3523
TimeGPT	/	/	/	0.4807	0.6693
TimesFM	/	/	/	0.4428	<u>0.6633</u>
<i>Experiment ID : UK-I-30</i>					
Moment	/	/	/	0.1847	0.2428
GP	0.1774	0.1644	0.2123	0.3287	1.0757
SVR	/	/	/	0.2262	0.2989
Chronos-small	0.0196	<u>0.0568</u>	0.0662	<u>0.1136</u>	<u>0.1748</u>
Chronos-tiny	0.0190	0.0569	<u>0.0658</u>	0.1138	0.1935
Lag-llama	0.0362	0.0899	0.0956	0.1798	0.2501
TimeGPT	/	/	/	0.1273	0.1799
TimesFM	/	/	/	0.1194	0.1795
<i>Experiment ID : UK-I-60</i>					
Moment	/	/	/	0.2685	0.3419
GP	0.0625	0.1322	0.1076	0.2645	0.3680
SVR	/	/	/	0.2634	0.3368
Chronos-small	<u>0.0240</u>	0.0746	0.0860	0.1492	<u>0.2257</u>
Chronos-tiny	<u>0.0240</u>	<u>0.0745</u>	<u>0.0858</u>	<u>0.1490</u>	<u>0.2469</u>
Lag-llama	0.0496	0.1320	0.1422	0.2640	0.3547
TimeGPT	/	/	/	0.1588	0.2313
TimesFM	/	/	/	0.1505	0.2263
<i>Experiment ID : NL-I-60</i>					
Moment	/	/	/	0.2641	0.3237
GP	0.0456	0.1022	<u>0.0662</u>	0.2044	0.2869
SVR	/	/	/	0.2239	0.2675
Chronos-small	<u>0.0219</u>	<u>0.0710</u>	0.0726	0.1420	0.2305
Chronos-tiny	<u>0.0219</u>	<u>0.0710</u>	0.0725	0.1421	0.2306
Lag-llama	0.0421	0.1173	0.1392	0.2345	0.3227
TimeGPT	/	/	/	0.1416	<u>0.1950</u>
TimesFM	/	/	/	<u>0.1397</u>	0.2001

Q $h\%$: Meaning the quantile loss at quantile level h .

Chronos continues to demonstrate the strongest performance in point prediction. Additionally, TimesFM and TimesGPT exhibit competitive results, frequently achieving the lowest error rates, as seen in several scenarios. In contrast, Moment and Lag-llama generally show weaker performance in point prediction. Moment employs tokenization methods similar to those used in Lag-llama [21]. Given that Moment has a large scale among the TSFMs discussed, it is expected to deliver strong performance. However, its comparatively

weaker results may be attributed to the fact that, unlike other TSFMs specifically designed for predictive tasks, Moment is intended as a foundational model capable of supporting diverse tasks, such as classification and detection.

In summary, Chronos demonstrates competitive performance in probabilistic prediction compared to classical models like GP, even without training. By contrast, Lag-llama does not show clear advantages over GP. For point prediction, Chronos maintains superior accuracy, with TimeGPT also performing well. Moment appears to predict only average values across experiments, while SVR, GP, and Lag-llama generally show weaker performance than Chronos and TimeGPT.

V. CONCLUSION

In this paper, we examine the zero-shot performance of five TSFMs including Chronos, Moment, TimesFM, Lag-llama, and TimeGPT comparing them to two classical models, GP and SVR. The experimental results indicate that, under the same conditions, pre-trained TSFMs such as TimeGPT and Chronos can outperform classical models that are trained on the dataset. We observe, that TSFMs, especially Chronos show a strong performance in both probabilistic and point prediction.

We consider TSFMs to be a promising approach for future STLP tasks, as once trained, TSFMs require little to no additional data for specific training, effectively addressing data scarcity for data privacy or data collection issues without compromising much performance.

However, current TSFMs cannot incorporate external conditions, which is their primary limitation compared to other models. As shown in [8], external conditions and information can effectively improve prediction accuracy. Therefore, a key direction for future research is to integrate external conditions, such as weather information, into the forecasting mechanism.

ACKNOWLEDGMENTS

This publication is part of the project ALIGN4Energy (with project number NWA.1389.20.251) of the research programme NWA ORC 2020 which is (partly) financed by the Dutch Research Council (NWO). This research utilized the Dutch National e-Infrastructure with support from the SURF Cooperative (grant number: EINF-5398).

REFERENCES

- [1] L. Zhang, J. Wen, Y. Li, J. Chen, Y. Ye, Y. Fu, and W. Livingood, "A review of machine learning in building load prediction," *Applied Energy*, vol. 285, p. 116452, 2021.
- [2] G. Gong, X. An, N. K. Mahato, S. Sun, S. Chen, and Y. Wen, "Research on short-term load prediction based on seq2seq model," *Energies*, vol. 12, no. 16, p. 3199, 2019.
- [3] L. Wen, K. Zhou, and S. Yang, "Load demand forecasting of residential buildings using a deep learning model," *Electric Power Systems Research*, vol. 179, p. 106073, 2020.
- [4] D. J. Trudnowski, W. L. McReynolds, and J. M. Johnson, "Real-time very short-term load prediction for power-system automatic generation control," *IEEE Trans. Control Systems Technology*, vol. 9, no. 2, pp. 254–260, 2001.
- [5] S.-J. Huang and K.-R. Shih, "Short-term load forecasting via arma model identification including non-gaussian process considerations," *IEEE Trans. power systems*, vol. 18, no. 2, pp. 673–679, 2003.
- [6] Y. Wang, S. Sun, X. Chen, X. Zeng, Y. Kong, J. Chen, Y. Guo, and T. Wang, "Short-term load forecasting of industrial customers based on svm and xgboost," *Int. J of Electrical Power & Energy Systems*, vol. 129, p. 106830, 2021.
- [7] X. Zhang, J. Wang, and K. Zhang, "Short-term electric load forecasting based on singular spectrum analysis and support vector machine optimized by cuckoo search algorithm," *Electric Power Systems Research*, vol. 146, pp. 270–285, 2017.
- [8] M. Barman, N. D. Choudhury, and S. Sutradhar, "A regional hybrid go-svm model based on similar day approach for short-term load forecasting in assam, india," *Energy*, vol. 145, pp. 710–720, 2018.
- [9] F. Divina, A. Gilson, F. Gómez-Vela, M. García Torres, and J. F. Torres, "Stacking ensemble learning for short-term electricity consumption forecasting," *Energies*, vol. 11, no. 4, p. 949, 2018.
- [10] H. Jiang, Y. Zhang, E. Muljadi, J. J. Zhang, and D. W. Gao, "A short-term and high-resolution distribution system load forecasting approach using support vector regression with hybrid parameters optimization," *IEEE Trans. Smart Grid*, vol. 9, no. 4, pp. 3341–3350, 2016.
- [11] D. Cao, J. Zhao, W. Hu, Y. Zhang, Q. Liao, Z. Chen, and F. Blaabjerg, "Robust deep gaussian process-based probabilistic electrical load forecasting against anomalous events," *IEEE Trans. Industrial Informatics*, vol. 18, no. 2, pp. 1142–1153, 2021.
- [12] P. Zhao, D. Cao, Y. Wang, Z. Chen, and W. Hu, "Gaussian process-aided transfer learning for probabilistic load forecasting against anomalous events," *IEEE Trans. power systems*, vol. 38, no. 3, pp. 2962–2965, 2023.
- [13] W. Lin, D. Wu, and B. Boulet, "Spatial-temporal residential short-term load forecasting via graph neural networks," *IEEE Trans. Smart Grid*, vol. 12, no. 6, pp. 5373–5384, 2021.
- [14] H. J. Sadaei, P. C. d. L. e Silva, F. G. Guimaraes, and M. H. Lee, "Short-term load forecasting by using a combined method of convolutional neural networks and fuzzy time series," *Energy*, vol. 175, pp. 365–377, 2019.
- [15] Z. Fazlipour, E. Mashhour, and M. Joorabian, "A deep model for short-term load forecasting applying a stacked autoencoder based on lstm supported by a multi-stage attention mechanism," *Applied Energy*, vol. 327, p. 120063, 2022.
- [16] Y. Wang, G. Hug, Z. Liu, and N. Zhang, "Modeling load forecast uncertainty using generative adversarial networks," *Electric Power Systems Research*, vol. 189, p. 106732, 2020.
- [17] W. Xia, C. Wang, P. Palensky, and P. P. Vergara, "A flow-based model for conditional and probabilistic electricity consumption profile generation and prediction," *arXiv preprint arXiv:2405.02180*, 2024.
- [18] L. Lin, C. Chen, B. Wei, H. Li, J. Shi, J. Zhang, and N. Huang, "Residential electricity load scenario prediction based on transferable flow generation model," *Journal of Electrical Engineering & Technology*, vol. 18, no. 1, pp. 99–109, 2023.
- [19] Y. Song, T. Wang, P. Cai, S. K. Mondal, and J. P. Sahoo, "A comprehensive survey of few-shot learning: Evolution, applications, challenges, and opportunities," *ACM Computing Surveys*, vol. 55, no. 13s, pp. 1–40, 2023.
- [20] A. F. Ansari, L. Stella, C. Turkmen, X. Zhang, P. Mercado, H. Shen, O. Shchur, S. S. Rangapuram, S. P. Arango, S. Kapoor *et al.*, "Chronos: Learning the language of time series," *arXiv preprint arXiv:2403.07815*, 2024.
- [21] M. Goswami, K. Szafer, A. Choudhry, Y. Cai, S. Li, and A. Dubrawski, "Moment: A family of open time-series foundation models," *arXiv preprint arXiv:2402.03885*, 2024.
- [22] K. Rasul, A. Ashok, A. R. Williams, H. Ghonia, R. Bhagwatkar, A. Khorasani, M. J. D. Bayazi, G. Adamopoulos, R. Riachi, N. Hassen *et al.*, "Lag-llama: Towards foundation models for probabilistic time series forecasting," *Preprint*, 2024.
- [23] A. Das, W. Kong, R. Sen, and Y. Zhou, "A decoder-only foundation model for time-series forecasting," *arXiv preprint arXiv:2310.10688*, 2023.
- [24] A. Garza, C. Challu, and M. Mergenthaler-Canseco, "Timegpt-1," *arXiv preprint arXiv:2310.03589*, 2023.
- [25] "Smart meter data - london households," <https://data.london.gov.uk/dataset/smartmeter-energy-use-data-in-london-households>.
- [26] "Household data - open power system data," https://data.open-power-system-data.org/household_data/2020-04-15.
- [27] "Liander open data," <https://www.liander.nl/partners/datadiensten/open-data/data>.