The Fractional Brownian Motion Dataset for Evaluating Extreme Quantiles Forecasting Methods

Lyudmyla Kirichenko, Roman Lavrynenko and Nataliya Ryabova

Kharkiv National University of Radio Electronics, Nauky av., 14, Kharkiv, 61166, Ukraine

Abstract

Machine learning utilizes data for training. However, there are instances when the data is insufficient. To determine the degree of risk of extreme events, it is necessary to predict the values of extreme quantiles, which may occur once in a hundred years, having only 30 years of historical data. The data is clearly insufficient for conventional forecasting methods. The problem becomes even more complicated when the time series has fractal properties and contains long-term dependencies. Developing machine learning methods on real data for such a task often seems impossible, so we present a method for generating a dataset to obtain precise values of extreme quantiles for time series, which are realizations of fractional Brownian motion. A key feature of this data acquisition is the parallelization of the Hosking method, which is used for the simulation of a fractional Brownian motion.

Keywords

extreme quantile regression, fBm, Kaggle, probabilistic forecasting, time series

1. Introduction

Fractal time series is a class of time series characterized by the property of self-similarity, that is, the statistical properties of the series are preserved on different time scales. In recent decades, such time series have been found in many phenomena of the surrounding world, including weather data, financial data, biomedical data, etc. Forecasting fractal time series is of practical importance for decision making in various fields. For example, in economics and finance, fractal time series forecasting can help manage risk and make decisions about buying and selling stocks or other financial instruments.

Traditionally, point forecasting has been the primary approach in time series forecasting, where a single value is predicted as the most likely outcome. However, point forecasting does not capture the inherent uncertainty present in time series data, which can lead to unreliable and inaccurate predictions.

Probabilistic forecasting [30], on the other hand, provides a range of possible outcomes and their associated probabilities. Probabilistic forecasting allows decision-makers to understand the uncertainty in the forecast and make informed decisions based on the range of possible outcomes. Probabilistic forecasting can also capture important features of the underlying data distribution, such as seasonality, trend, and volatility.

To evaluate the performance of a probabilistic forecasting methods, it is necessary to compare the model's predicted probability distribution with the actual probability distribution of the ground truth data. One common way to do this is by calculating the quantiles of the predicted probability distribution and comparing them with the quantiles of the actual distribution.

Quantiles are simply points in the probability distribution that divide the data into groups. For example, the 50th percentile (also known as the median) is the value that divides the data into two equal groups, with 50% of the data above and 50% below this value. Other commonly used quantiles include

International Scientific Symposium «Intelligent Solutions» IntSol-2023, September 27–28, 2023, Kyiv-Uzhhorod, Ukraine
EMAIL: lyudmyla.kirichenko@nure.ua (L. Kirichenko); roman.lavrynenko.cpe@nure.ua (R. Lavrynenko); nataliya.ryabova@nure.ua (N.Ryabova)
ORCID: 0000-0002-2780-7993 (L. Kirichenko); 0000-0003-1969-1107 (R. Lavrynenko); 0000-0002-3608-6163 (N. Ryabova)

the 10th percentile, the 90th percentile, and the interquartile range (the difference between the 25th and 75th percentiles).

Although fractal time series have specific properties that can make forecasting difficult, in many cases it is possible to use the same forecasting methods as for conventional time series. However, the challenge arises in predicting extreme quantiles of the probability distribution, often referred to as the "tails" of the distribution. These extreme quantiles represent rare but critical events, such as catastrophic financial losses, extreme weather events, or catastrophic failures in infrastructure systems.

In fields such as hydrology and climate science [35], the concept of risk is frequently quantified using the T-year return level, symbolized as Q^T. This return level is a measure of the magnitude of an event that is expected, on average, to be exceeded once every T years. Consider Y as a certain variable, for which we record nY independent observations each year. For example, if we collect daily data, nY would be 365, representing the number of days in a year. The T-year return level, Q^T, is then computed as the quantile $Q(1 - 1/(nY * T))$.

Thus, the T-year return level is a probabilistic measure of the size of an event that is expected to be exceeded with a frequency of once every T years, based on historical data. It is a critical concept in risk assessment, particularly in understanding and preparing for extreme events. Predicting these quantiles accurately, especially in fractal time series, is essential in a number of fields to aid in risk management and policy planning. A significant challenge in extreme quantile prediction is the limited availability of data for estimation. For instance, predicting a 100-year return level becomes problematic when we only have training data from the previous 50 years. This scarcity of data makes the statistical estimation of extreme quantiles a challenging task, particularly in the context of probabilistic forecasting models. Consequently, developing robust methods for extreme quantile prediction under such data constraints is a crucial area of research.

Given the importance of forecasting extreme quantiles in fractal time series, we propose an approach for generating datasets of such time series specifically designed to evaluate extreme quantiles forecasting.

Contributions:

1. The method provides a way to efficiently compute multiple continuations of a single fractional Brownian motion (fBm) time series using the Hosking algorithm.
2. In result the dataset with ground truth extreme quantiles of possible continuations can be used for evaluating machine learning methods designed for probabilistic forecasting.

The code for generating a file of the dataset for a specific Hurst exponent can be found at the following link: https://www.kaggle.com/code/unfriendlyai/fbm-extreme-quantile-generator. Our fBm dataset is available at: https://www.kaggle.com/datasets/unfriendlyai/fbm-extreme-quantiles

2. Related Works

The increasing availability of data demands new processing and analysis methods for effective time series forecasting. Machine learning methods for time series forecasting are gaining significance, allowing for automated and faster prediction processes, as well as improved accuracy and quality of forecasts [4, 13, 14, 15]. Reviews have presented the main methods and approaches for time series forecasting using machine learning [6, 25, 31]. Although fractal time series have a wide application in scientific and technical fields, the application of machine learning in the field of fractal time series analysis has mainly affected classification methods [8, 16, 17, 20, 21] and clustering [18, 27], as well as methods for estimating the Hurst exponent by time realizations [3, 19, 27].

Datasets with modeled and real time series sets have been created for this, but in fact there are no datasets with modeled fractal time series that could be used to validate methods for extreme quantile prediction. At the moment, relatively few special methods have been developed for forecasting time series with fractal properties. Most of the existing methods are focused on predicting fractional Brownian motion (fBm) [7,28]. At the same time, the issue of forecasting extreme quantile of fractal time series, in particular fBm, remains open.

The development of probabilistic prediction is covered in reviews [30, 34]. There are two main approaches to modeling a probability distribution. In the first distribution shape is given beforehand (eg. Gaussian, exponential), and the model during training should determine 2-4 parameters of this distribution, depending on the input data [5]. For modeling more complex distributions more complex methods may be used like Generalized Additive Models for Shape, Scale, and Location (GAMLSS) [29] . The second type is when the model should approximate the conditional cumulative distribution function. This Maybe be achieved with direct quantiles or expectiles prediction [32] or using a novel Normalizing flow based approach [1, 33]. Existing models allow to obtain information about the distribution function in different ways [23].

The need for probabilistic forecasting is demonstrated in particular by the increasing number of Kaggle competitions that require predicting time series quantiles in the future [10, 11, 12, 22]. For example, in “M5 Forecasting – Uncertainty” competition the participants are asked to provide 28 days ahead point forecasts for all the series of the competition, as well as the corresponding median and 50%, 67%, 95%, and 99% prediction intervals. In the reference [35], modeling extreme quantile regression and risk assessment were explored, specifically with an application to forecasting flood risk. The study provided valuable insights into the practical usage of extreme quantile regression models for predicting rare and extreme events, such as floods.

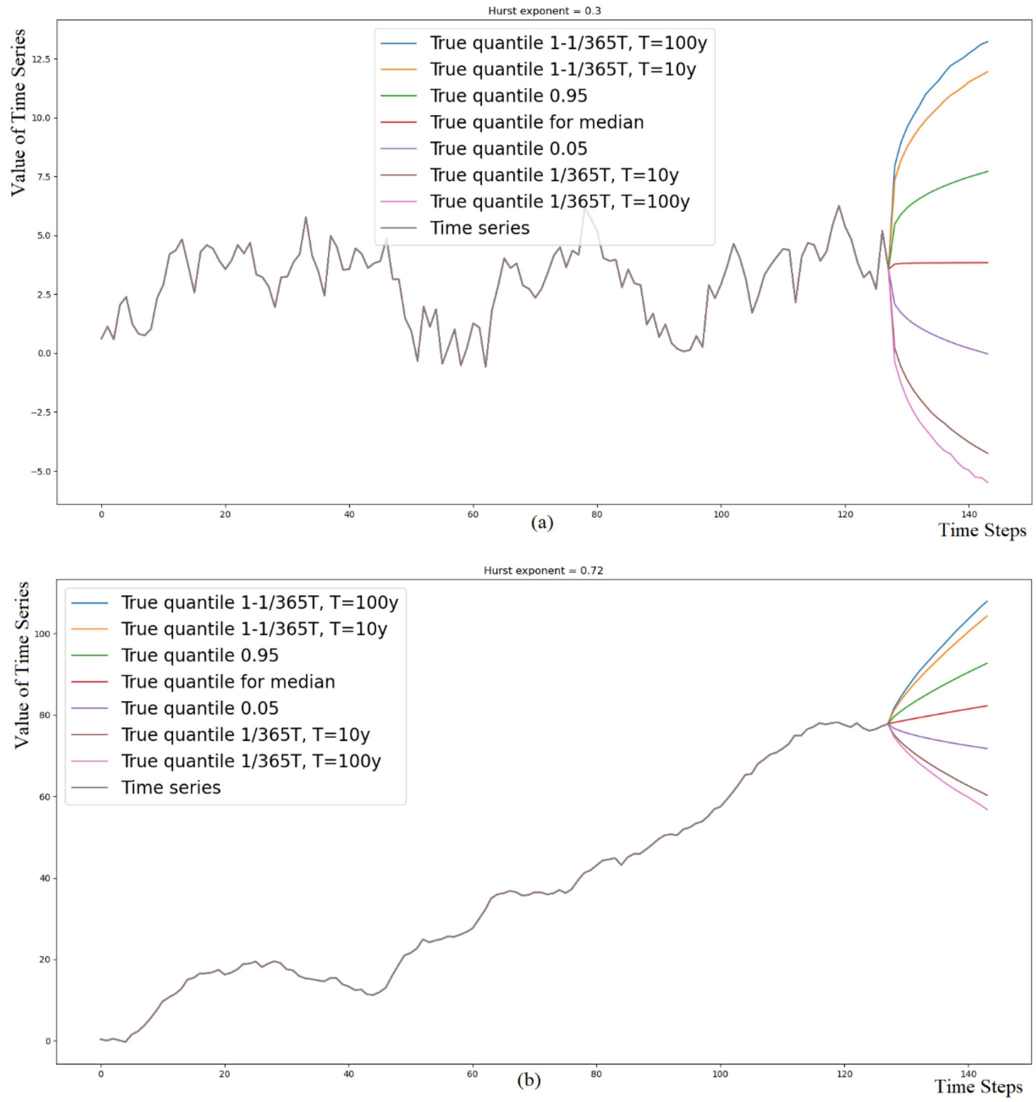

Figure 1: Visualization of example of proposed dataset. Common beginning of length 128 and quantiles of its continuation length 16 for antipersistent time series (a) and persistent time series (b)

3. Method of generating dataset for extreme quantiles problem

Hence, there is a clear need for creating specific datasets of fractal time series, with a particular emphasis on those designed for probabilistic forecasting of extreme quantiles. The goal of this study is to generate a dataset containing realizations of fractional Brownian motion (fbm) and the corresponding true quantiles of their possible continuations.

Fractional Brownian motion is an extension of classical Brownian motion, which is characterized by random walks of particles in space. The fBm has properties of self-similarity and scale invariance, which means that its structure and characteristics remain unchanged when the scale of observation changes. The increments $\Delta X(\tau) = X(t+\tau) - X(t)$ have a Gaussian distribution.

Fractional Brownian motion (fBm) X is a type of self-similar stochastic process characterized by its Hurst exponent H ($0 < H < 1$) which determines the degree of long-range dependence of the process. A persistent time series is a series with a Hurst exponent greater than 0.5, indicating that future values are likely to exhibit a positive autocorrelation with past values. This means that when past values are higher (lower) than average, future values are also likely to be higher (lower) than average. On the other hand, an anti-persistent time series has a Hurst exponent less than 0.5, indicating that future values are likely to exhibit a negative autocorrelation with past values. This means that when past values are higher (lower) than average, future values are likely to be lower (higher) than average. When the Hurst exponent is equal to 0.5, this indicates a case of a random walk or white noise, where future values are independent of past values.

There are a number of exact methods to simulate fBm realizations [2]. The Hosking method is simple, popular and implemented in Python. The Hosking method involves simulating the fBm using the following steps:

1. Generate a sequence of independent and identically distributed (IID) random variables from a standard normal distribution with zero mean and variance one.
2. Compute the autocovariance function of the FBM using the formula above for a range of lags.
3. Use the autocovariance function to compute the Cholesky decomposition of the covariance matrix.
4. Multiply the IID random variables by the Cholesky factor to obtain a sequence of correlated random variables.
5. Compute the cumulative sum of the correlated random variables to obtain the simulated FBM.
6. Repeat steps 1-5 to obtain a sample of the FBM.

The Cholesky decomposition is a technique for decomposing a positive definite matrix into a product of a lower triangular matrix and its transpose. In the context of simulating FBM, the Cholesky factor is used to generate a sequence of correlated random variables from a sequence of IID random variables.

Also a very popular exact estimation method, particularly because of its fast speed is the Davies-Harte method. Davies-Harte method generates fBm by transforming a Gaussian white noise process with a discrete Fourier transform, applying a scaling factor based on the desired Hurst exponent, and then inverse transforming the noise process to produce the fBm. However, this method does not generate values iteratively and therefore cannot be used to continue time series with predefined values.

To verify that the examples generated for the dataset possess the required characteristics, we examine their Hurst exponent and the standard deviation of increments. The Whittle method [26] is a powerful tool for estimating the Hurst exponent of time series. The method has some drawbacks. In particular, it does not work well with non-Gaussian time series. However, in the case when it is known in advance that the time series are fBm, the method is one of the most accurate [9, 24, 26].

The problem addressed in this work is the need to evaluate extreme quantile forecasting methods for time series with long-range dependence, specifically those generated by fractional Brownian motion processes. To address this problem, we propose a method for creating a dataset of fBm time series and their ground truth continuations using the Hosking algorithm. Our method involves parallel calculation of M series, each with a common value of Hurst exponent, and the generation of a matrix of quantiles for each time step.

Our method for creating a single instance of the evaluation dataset for a specific Hurst exponent value, as illustrated in **Figure 1**, involves the following steps:

- Set the Hurst exponent value.
- Generate a matrix of normally distributed numbers of size M x N, where M is the number of continuations and $N = N_1 + N_2$. Here, N_1 represents the length of a common beginning of the time series, and N_2 represents the length of continuations.

- To obtain a common beginning of length N_1, make all M beginnings of the previous matrix equal. This is achieved by copying the first row into all other rows.
- Employ the Hosking method to iteratively calculate a set of M time series that are independent of each other. This method is selected due to its capacity for parallel calculation of the values of the succeeding time steps based on normally distributed random numbers using matrix operations.
- The output from the previous step is M time series, in which the first N_1 increments are identical and the following N_2 values are independent of each other but dependent on the N_1 initial identical increments.
- Calculate the target true quantiles for N_2 steps using M variants of continuation.
- The final result is stored as a covariate time series of length N_1 and a matrix of ground truth target quantiles of length N_2.

4. Experiment and results

As a result of the proposed method described above, the following fractal time series dataset was designed for the realizations of fractional Brownian motion. One can freely download or utilize the dataset on Kaggle platform by searching for the dataset with the name " fBm Extreme Quantiles" or by using the following link: https://www.kaggle.com/datasets/unfriendlyai/fbm-extreme-quantiles

The dataset was created with the following parameters:

- The set of Hurst exponent values comprises [0.3, 0.35, 0.45, 0.53, 0.6, 0.65, 0.72, 0.85, 0.9, 0.93]. This range is chosen to represent various types of time series behavior: antipersistent (values less than 0.5), nearly independent (around 0.5), and persistent (greater than 0.5) series. The diverse selection of Hurst values allows us to capture a broad range of potential dynamics in the data, thereby creating a more robust and comprehensive dataset for model evaluation.
- Number of records per Hurst exponent is 50. Each Hurst exponent value has 50 records in the evaluation dataset, sufficient for obtaining statistically significant experimental results when comparing different prediction methods.
- The length of the original time series N_1 is 128.
- The length of the continuations of the original series, for which quantiles are provided N_2 is 16.
- Number of continuation examples (M): 3,650,000 examples of continuations are used to calculate quantiles. This number was chosen for the convenience of determining the true quantile when an event occurs once every hundred years under daily observation [35]. The number of such outcomes for accurate computation is taken as 100 (365 days x 100 years x 100 events = 3,650,000).
- Set of true quantiles includes the median (0.5), usual quantiles (0.05 and 0.95), and quantiles corresponding to 100-year return levels T=100y (1/36500 and 1-1/36500) and 10-year return levels T=10y (1/3650 and 1-1/3650). This provides a comprehensive range of quantiles for analysis, from the most common to the most extreme (**Figure 1**).
- Time series increments are normalized (divided by STD of increments).
- Original time series are presented as cumulative sums of increments. Quantiles are calculated on their cumulative continuations.
- The training dataset consists of 10,000 examples of length 128 for covariates and values of the following 16 time steps for the target. These parameters are similar to those used in [36]. In this case, the data for calculating extreme quantiles of a hundred-year period is not enough, as 10,000 days are roughly three times fewer than a hundred years.

The code for generating a file of the dataset for a specific Hurst exponent can be found at the following link: https://www.kaggle.com/code/unfriendlyai/fbm-extreme-quantile-generator

The obtained dataset was tested using three machine learning methods known for their efficiency and speed, often used in Kaggle competitions, that can predict specific quantiles: LightGBM [36], CatBoost, and Statsmodels QuantReg [37]. The results of the three models for one time series are shown in Figures 2, 3. The 0.05, 0.50, and 0.95 quantiles were satisfactorily predicted by all three methods. The extreme quantiles were predicted unsatisfactorily for persistent fBm time series. For antipersistent fBm time series with H=0.30, LightGBM and CatBoost show more satisfactory results for extreme quantiles. However, the prediction from Statsmodels QuantReg is unsatisfactory.

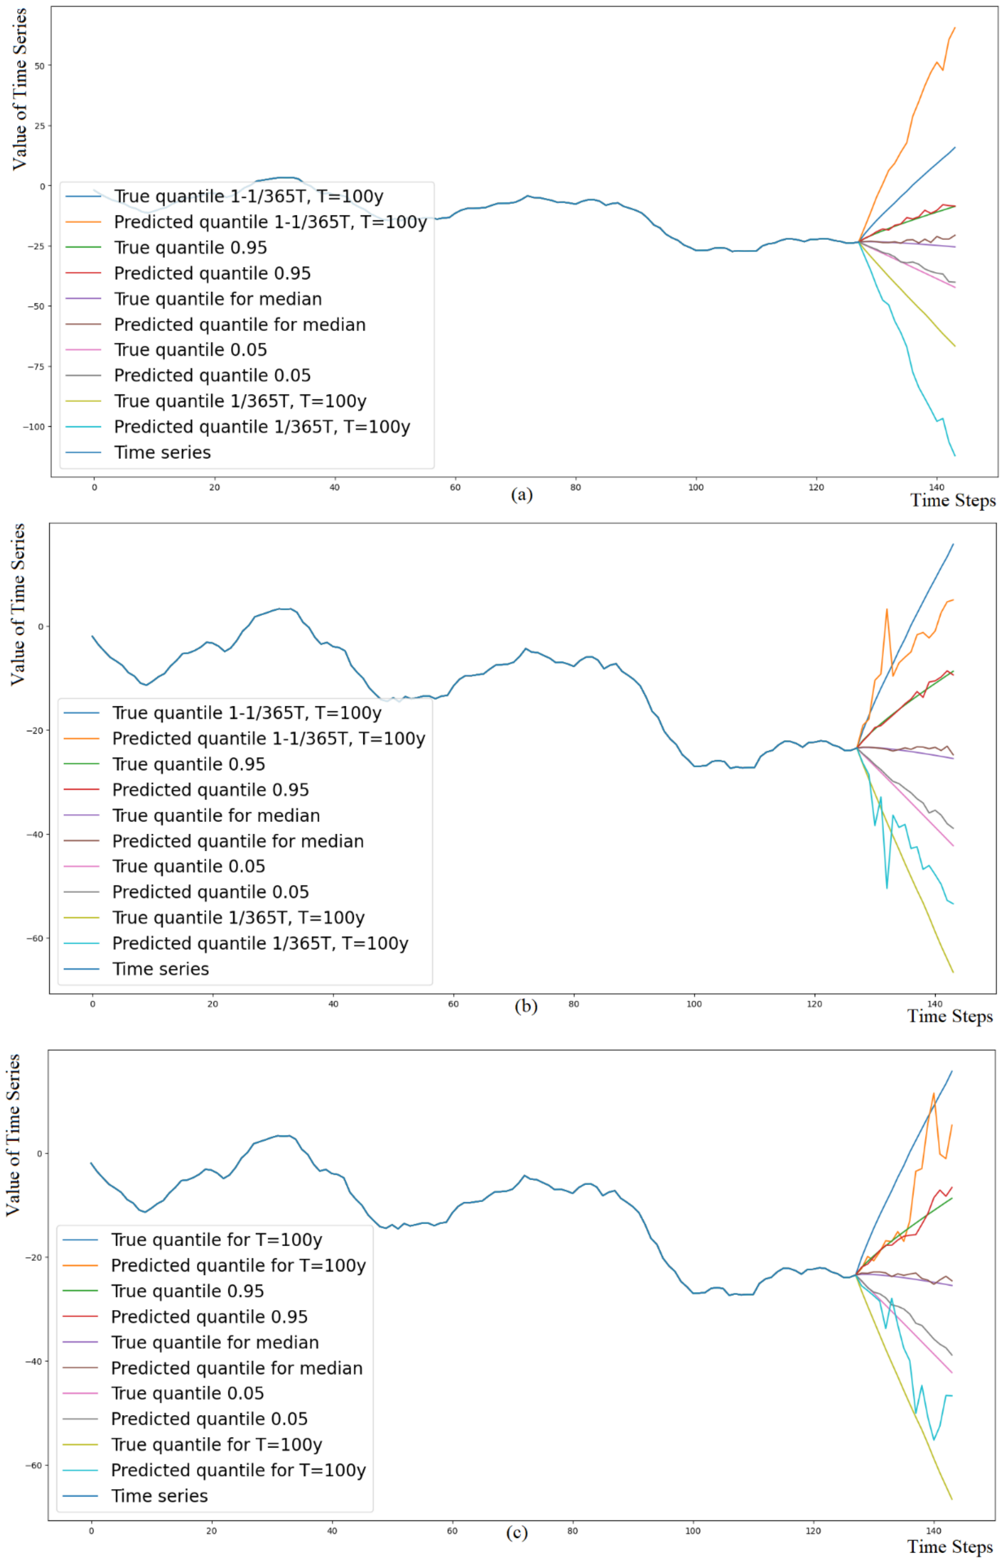

Figure 2: LightGBM (a), CatBoost (b) and Statsmodels QuantReg (c) results for fBm time series H=0.93. The 0.05, 0.50, and 0.95 quantiles were satisfactorily predicted by all three methods. However, the prediction of quantiles corresponding to an event occurring once in a hundred years proved to be extremely unsatisfactory.

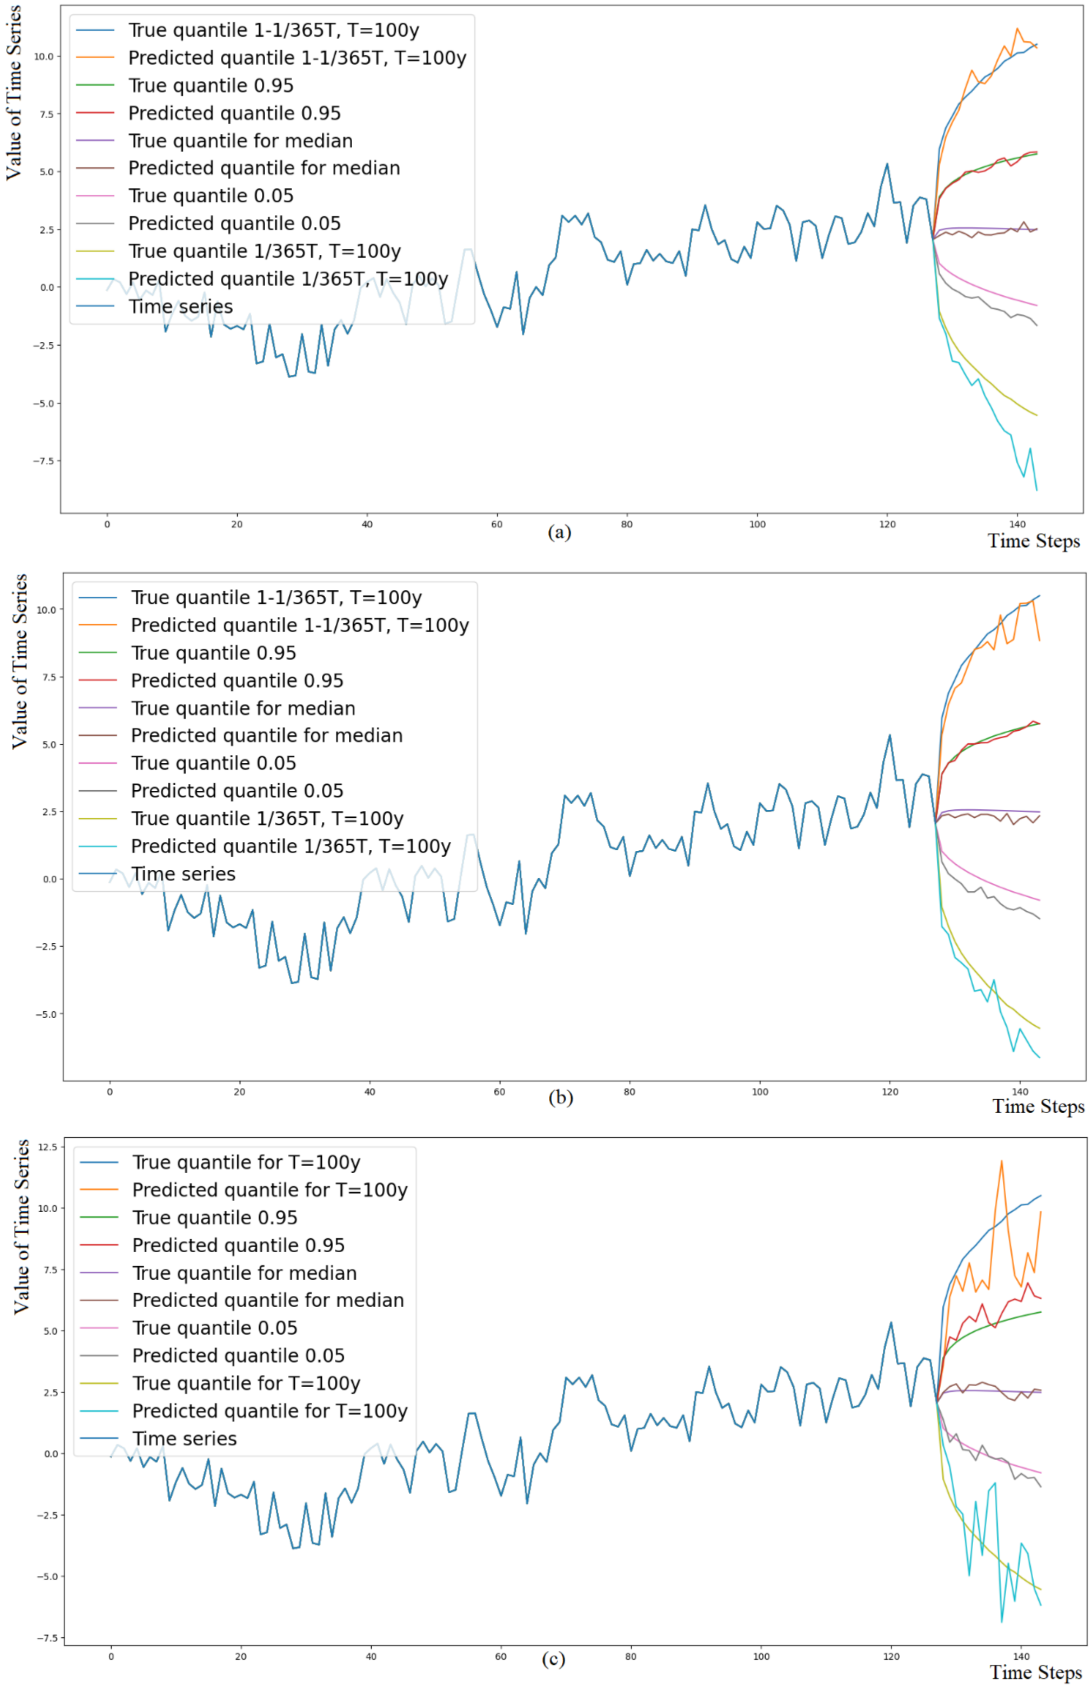

Figure 3: For antipersistent fBm time series with H=0.30, LightGBM (a) and CatBoost (b) show more satisfactory results for extreme quantiles. However, the prediction from Statsmodels QuantReg (c) is unsatisfactory.

When the Hurst exponent is close to 0.5, the prediction values do not depend on the previous values of the series and these predictions are not interesting. When the Hurst exponent deviates from 0.5, the predictions for the 0.05, 0.50, and 0.95 quantiles in all cases depended on the previous values of the series, and all three models captured these dependencies.

However, for extreme quantiles in this case, it was noted that the predictions often did not depend (or almost did not depend) on the previous values. Whether for a persistent descending series or an ascending one, the predictions were quite symmetrical relative to the last value of the series, unlike the adequately predicted median (Figure 4).

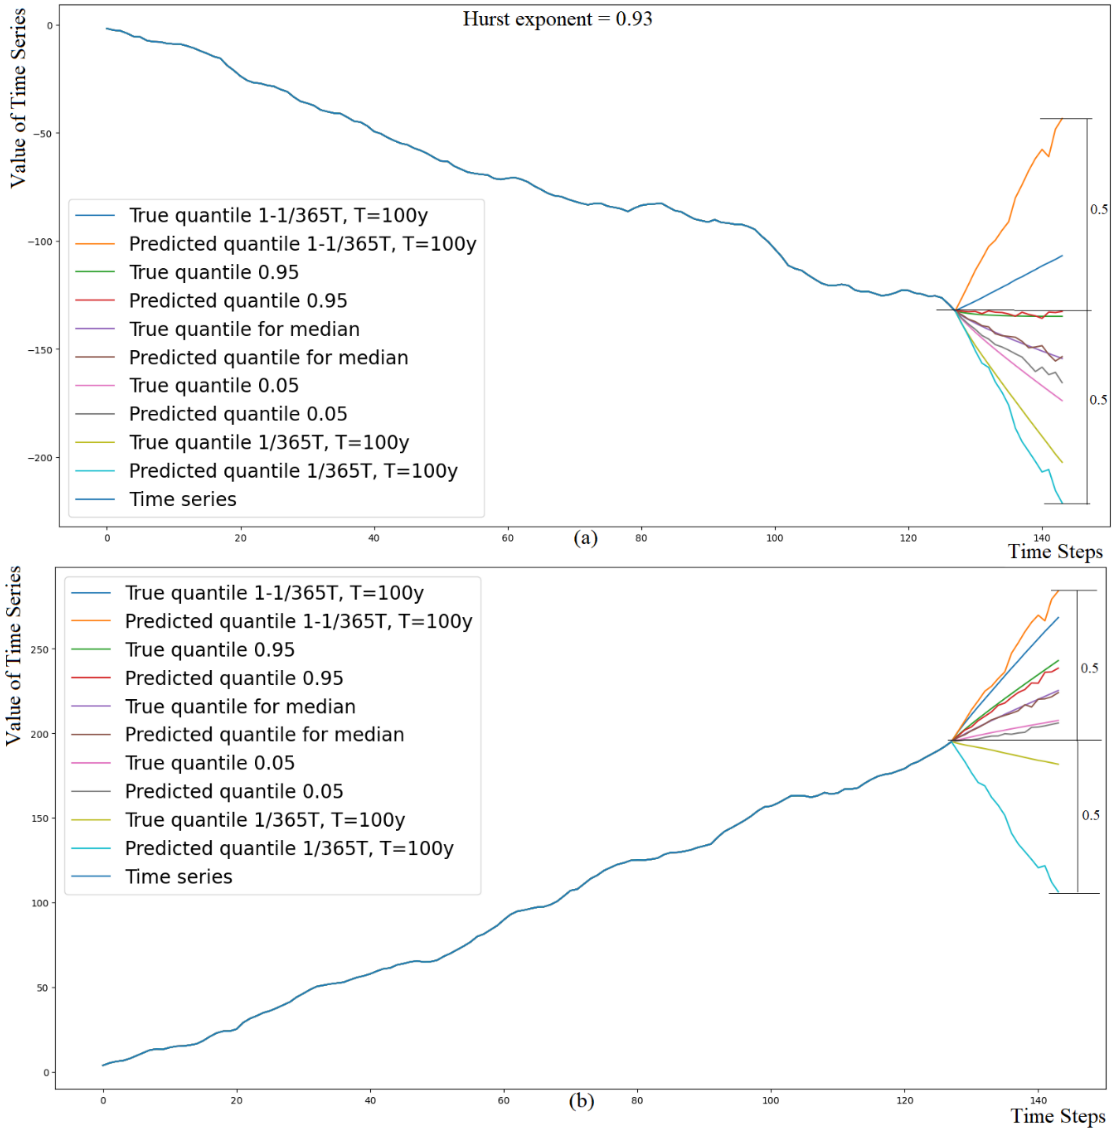

Figure 4: The predictions of extreme quantiles by LightGBM did not depend on the previous values of the time series. Regardless of whether the series was persistently descending (a) or ascending (b), the predictions were quite symmetrical relative to the last value of the series.

To determine correctness of examples in dataset and limits for calculating Hurst exponent of predicted time series, the values of the Hurst exponent for each time series of length 128 were determined using the Whittle algorithm. The results are shown in Table 1.

To compare the correctness, we generated the same number of time series for each value of the Hurst exponent with the Davies-Harte method. The estimate of the Hurst exponent for the time series of the dataset and the fBm generated by the Davies-Harte method has similar scatter values.

The results of comparison with the Davies-Harte method show that the time series by the Hosking's method with parallel calculation of M the continuations of one row are executed correctly.

Table 1
Hurst exponent statistics (mean and standard deviation) determined with Whittle algorithm

Target H value	H (Dataset) mean	H (Davies-Harte) mean	H (Dataset) std	H (Davies-Harte) std
0.30	0.299	0.295	0.062	0.071
0.35	0.348	0.355	0.053	0.051
0.45	0.440	0.447	0.048	0.051
0.53	0.526	0.525	0.054	0.073
0.60	0.582	0.589	0.057	0.064
0.65	0.651	0.644	0.057	0.056
0.72	0.716	0.697	0.063	0.074
0.85	0.822	0.824	0.072	0.060
0.90	0.882	0.859	0.055	0.042
0.93	0.901	0.875	0.047	0.052

The Python code of experiments is available for review on Kaggle platform at https://www.kaggle.com/datasets/unfriendlyai/fbm-extreme-quantiles/code

5. Discussions

In this study, we successfully employed parallel computation using the Hosking method to obtain the true values of extreme quantiles for fBm (fractional Brownian motion) time series. This was achieved by substituting normally distributed random numbers at the begining of the series with identical numbers for all numerous variants of the series. By using 100 times more series variants than the number of days in 100 years, we were able to prepare a dataset of time series and their true quantiles with sufficient accuracy for the evaluation of machine learning methods predicting extreme quantiles.

Although we could generate an unlimited number of training examples, inspired by [35], we limited the training dataset to the same size, specifically 7000 training examples, and an additional 3000 for validation used for hyperparameter selection (in our case, early stopping points). This corresponds to approximately 19 years of daily observations. At the same time, the task was to predict an event with a frequency of once in 100 years of daily observations (quantile 1/36500).

Given such a limited volume of training data, conventional machine learning models were unable to predict extreme quantiles. However, there were no issues with predicting the median and 0.95 quantile.

We also noted that the presence of long-term dependencies in the time series, due to the fact that the time series is a realization of fractional Brownian motion, was necessary. In the absence of long-term dependencies (Hurst exponent close to 0.5), it was impossible to compare the effectiveness of prediction methods. The prediction of extreme quantiles of persistent time series compared to antipersistent ones proved to be a significant challenge and deserves special attention.

Some models predicted extreme quantiles independently of the previous values (increments) of the time series. This was noticeable by the symmetrical distribution of quantiles 1/36500 and 1-1/36500 relative to the last value of the time series for strongly persistent series (consistently increasing or decreasing). Although in this case it sometimes seems that one of the extreme quantiles is predicted, this is refuted by the symmetrical opposite quantile, which clearly does not depend on the input data. This feature also needs to be taken into account when calculating the effectiveness of methods.

6. Conclusions

This study presents a novel method for generating a dataset to evaluate the prediction of extreme quantiles in fractional Brownian motion time series. Despite the challenges posed by limited training data and long-term dependencies, our approach provides a foundation for further research into refining existing prediction methods and exploring new machine learning approaches for this task.

These findings highlight the challenges and potential avenues for improving the prediction of extreme quantiles in fractal time series data, a task of significant relevance in risk assessment and other fields. It's important to note that while we have proposed a method for creating a dataset for evaluation, we have not proposed the prediction methods themselves. Further research is needed to refine these methods, such as those proposed in [35], and to explore other machine learning approaches for this task. This work lays the groundwork for such future investigations.

7. References

[1] M. Arpogaus, M. Voss, B. Sick, M. Nigge-Uricher, and O. Dürr, 'Probabilistic Short-Term LowVoltage Load Forecasting using Bernstein-Polynomial Normalizing Flows', in ICML 2021, Workshop Tackling Climate Change with Machine Learning, 2021.

[2] O. Banna, Fractional Brownian Motion - Approximations and Projections. London, England: ISTE, 2019. doi: 10.1002/9781119476771.

[3] S. Bo, F. Schmidt, R. Eichhorn, and G. Volpe, 'Measurement of anomalous diffusion using recurrent neural networks', Phys. Rev. E., vol. 100, no. 1–1, p. 010102, Jul. 2019. doi: 10.1103/PhysRevE.100.010102.

[4] G. Bontempi, S. Ben Taieb, and Y.-A. Le Borgne, 'Machine learning strategies for time series forecasting', in Business Intelligence, Berlin, Heidelberg: Springer Berlin Heidelberg, 2013, pp. 62–77. doi: 10.1007/978-3-642-36318-4_3.

[5] T. Duan, A. Avati, D. Y. Ding, S. Basu, A. Ng, and A. Schuler, NGBoost: Natural Gradient Boosting for Probabilistic Prediction. International Conference on Machine Learning. 2019.

[6] Y. P. Faniband, I. Ishak, and S. M. Sait, 'A review of open source software tools for Time Series Analysis', 2022. doi: 10.48550/arXiv.2203.05195.

[7] M. Garcin, 'Forecasting with fractional Brownian motion: a financial perspective', 2021. doi: 10.48550/arXiv.2105.09140.

[8] N. Granik et al., 'Single-particle diffusion characterization by deep learning', Biophys. J., vol. 117, no. 2, pp. 185–192, Jul. 2019. doi: 10.1016/j.bpj.2019.06.015.

[9] A. H. Hamza and M. Y. Hmood, "Comparison of Hurst exponent estimation methods", JEAS, vol. 27, no. 128, pp. 167–183, Jun. 2021. doi:10.33095/jeas.v27i128.2162.

[10] Addison Howard, Jay Evan Reid, Michael Lopez, Will Cukierski. (2019). NFL Big Data Bowl. Kaggle. https://kaggle.com/competitions/nfl-big-data-bowl-2020.

[11] Addison Howard, inversion, Spyros Makridakis, Vangelis. (2020). M5 Forecasting – Uncertainty. Kaggle. https://kaggle.com/competitions/m5-forecasting-uncertainty

[12] Addison Howard, inversion. (2020). COVID19 Global Forecasting (Week 5). Kaggle. https://kaggle.com/competitions/covid19-global-forecasting-week-5

[13] X. Jin, X. Yu, X. Wang, Y. Bai, T. Su, and J. Kong, 'Prediction for time series with CNN and LSTM', in Proceedings of the 11th International Conference on Modelling, Identification and Control (ICMIC2019), Singapore: Springer Singapore, 2020, pp. 631–641. doi: 10.1007/978-981-15-0474-7_59.

[14] S. Khlamov and V. Savanevych, 'Big astronomical datasets and discovery of new celestial bodies in the solar system in automated mode by the CoLiTec software', in Knowledge Discovery in Big Data from Astronomy and Earth Observation, Elsevier, 2020, pp. 331–345. doi: 10.1016/B978-0-12-819154-5.00030-8.

[15] L. Kirichenko, T. Radivilova, and I. Zinkevich, 'Forecasting weakly correlated time series in tasks of electronic commerce', in 2017 12th International Scientific and Technical Conference on Computer Sciences and Information Technologies (CSIT), Lviv, 2017. doi: 10.1109/STC-CSIT.2017.8098793.

[16] L. Kirichenko, B. Vitalii, and T. Radivilova, 'Machine learning classification of multifractional Brownian motion realizations (2020) CEUR Workshop Proceedings', pp. 980–989, 2608.

[17] L. Kirichenko, P. Zinchenko, and T. Radivilova, 'Classification of time realizations using machine learning recognition of recurrence plots', in Advances in Intelligent Systems and Computing, Cham: Springer International Publishing, 2021, pp. 687–696. doi: 10.1007/978-3-030-54215-3_44.

[18] L. Kirichenko, O. Pichugina, and H. Zinchenko, 'Clustering time series of complex dynamics by features', CEUR Workshop Proceedings, vol. 3132, pp. 83–93, 2022.

[19] L. Kirichenko, K. Pavlenko, and D. Khatsko, ‘Wavelet-based estimation of Hurst exponent using neural network’, in 2022 IEEE 17th International Conference on Computer Sciences and Information Technologies (CSIT), Lviv, Ukraine, 2022. doi: 10.1109/CSIT56902.2022.10000906.
[20] P. Kowalek, H. Loch-Olszewska, and J. Szwabiński, ‘Classification of diffusion modes in single-particle tracking data: Feature-based versus deep-learning approach’, Phys. Rev. E., vol. 100, no. 3–1, p. 032410, Sep. 2019. doi: 10.1103/PhysRevE.100.032410.
[21] X. Li, J. Yu, L. Xu, and G. Zhang, ‘Time series classification with deep neural networks based on Hurst exponent analysis’, in Neural Information Processing, Cham: Springer International Publishing, 2017, pp. 194–204. doi: 10.1007/978-3-319-70087-8_21.
[22] S. Makridakis et al., ‘The M5 uncertainty competition: Results, findings and conclusions’, Int. J. Forecast., vol. 38, no. 4, pp. 1365–1385, Oct. 2022. doi: 10.1016/j.ijforecast.2021.10.009.
[23] A. März and T. Kneib, ‘Distributional Gradient Boosting Machines’, 2022. doi: 10.48550/arXiv.2204.00778.
[24] G. Millán, R. Osorio-Comparán and G. Lefranc, "Preliminaries on the Accurate Estimation of the Hurst Exponent Using Time Series," 2021 IEEE International Conference on Automation/XXIV Congress of the Chilean Association of Automatic Control (ICA-ACCA), Valparaíso, Chile, 2021, pp. 1-8, doi: 10.1109/ICAACCA51523.2021.9465274.
[25] B. Ramadevi and K. Bingi, ‘Chaotic time series forecasting approaches using machine learning techniques: A review’, Symmetry (Basel), vol. 14, no. 5, p. 955, May 2022. doi: 10.3390/sym14050955.
[26] H. L. Shang, ‘A comparison of Hurst exponent estimators in long-range dependent curve time series’, J. Time Ser. Econom., vol. 12, no. 1, Jun. 2020. doi: 10.1515/jtse-2019-0009.
[27] G. S. Sidhu, A. Ibrahim Ali Metwaly, A. Tiwari, and R. Bhattacharyya, ‘Short term trading models using Hurst exponent and machine learning’, SSRN Electron. J., 2021. doi: 10.2139/ssrn.3824032.
[28] W. Song, M. Li, Y. Li, C. Cattani, and C.-H. Chi, ‘Fractional Brownian motion: Difference iterative forecasting models’, Chaos Solitons Fractals, vol. 123, pp. 347–355, Jun. 2019. doi: 10.1016/j.chaos.2019.04.021.
[29] D. M. Stasinopoulos and R. A. Rigby, “Generalized Additive Models for Location Scale and Shape (GAMLSS) in R”, J. Stat. Soft., vol. 23, no. 7, pp. 1–46, Dec. 2007. doi:10.18637/jss.v023.i07.
[30] H. Tyralis and G. Papacharalampous, ‘A review of probabilistic forecasting and prediction with machine learning’, 2022. doi: 10.48550/arXiv.2209.08307.
[31] J. F. Torres, D. Hadjout, A. Sebaa, F. Martínez-Álvarez, and A. Troncoso, ‘Deep learning for time series forecasting: A survey’, Big Data, vol. 9, no. 1, pp. 3–21, Feb. 2021. doi: 10.1089/big.2020.0159.
[32] C. Wan, J. Lin, J. Wang, Y. Song, and Z. Y. Dong, ‘Direct quantile regression for nonparametric probabilistic forecasting of wind power generation’, IEEE Trans. Power Syst., vol. 32, no. 4, pp. 2767–2778, Jul. 2017. doi: 10.1109/TPWRS.2016.2625101.
[33] P. Wielopolski and M. Zięba, ‘TreeFlow: Going beyond tree-based Gaussian probabilistic regression’, 2022. doi: 10.48550/arXiv.2206.04140.
[34] X. Zhou, H. Liu, F. Pourpanah, T. Zeng, and X. Wang, ‘A survey on epistemic (model) uncertainty in supervised learning: Recent advances and applications’, Neurocomputing, vol. 489, pp. 449–465, Jun. 2022. doi: 10.1016/j.neucom.2021.10.119.
[35] O. C. Pasche and S. Engelke, ‘Neural networks for extreme quantile regression with an application to forecasting of flood risk’, 2022. doi: 10.48550/arXiv.2208.07590.
[36] G. Ke et al., ‘LightGBM: A Highly Efficient Gradient Boosting Decision Tree’, in Advances in Neural Information Processing Systems, 2017, vol. 30.
[37] R. Koenker and K. F. Hallock, ‘Quantile regression’, J. Econ. Perspect., vol. 15, no. 4, pp. 143–156, Nov. 2001. https://doi.org/10.1257/jep.15.4.143.