



HAL
open science

AutoML Meets Time Series Regression Design and Analysis of the AutoSeries Challenge

Zhen Xu, Wei-Wei Tu, Isabelle Guyon

► **To cite this version:**

Zhen Xu, Wei-Wei Tu, Isabelle Guyon. AutoML Meets Time Series Regression Design and Analysis of the AutoSeries Challenge. ECML PKDD 2021, 2021, Barcelona (virtual), Spain. hal-03299144v1

HAL Id: hal-03299144

<https://hal.science/hal-03299144v1>

Submitted on 26 Jul 2021 (v1), last revised 6 Jan 2022 (v2)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

AutoML Meets Time Series Regression

Design and Analysis of the AutoSeries Challenge

Zhen Xu

4Paradigm, China

XUZHEN@4PARADIGM.COM

Wei-Wei Tu

4Paradigm, China

TUWEIWEI@4PARADIGM.COM

Isabelle Guyon

LISN CNRS/INRIA, France

University Paris-Saclay, France

ChaLearn, USA

GUYON@CHALEARN.ORG

Abstract

Analyzing better time series with limited human effort is of interest to academia and industry. Driven by business scenarios, we organized the first Automated Time Series Regression challenge (AutoSeries) for the WSDM Cup 2020. We present its design, analysis, and post-hoc experiments. The code submission requirement precluded participants from any manual intervention, testing automated machine learning capabilities of solutions, across many datasets, under hardware and time limitations. We prepared 10 datasets from diverse application domains (sales, power consumption, air quality, traffic, and parking), featuring missing data, mixed continuous and categorical variables, and various sampling rates. Each dataset was split into a training and a test sequence (which was streamed, allowing models to continuously adapt). The setting of “time series regression”, differs from classical forecasting in that covariates at the present time are known. Great strides were made by participants to tackle this AutoSeries problem, as demonstrated by the jump in performance from the sample submission, and post-hoc comparisons with AutoGluon. Simple yet effective methods were used, based on feature engineering, LightGBM, and random search hyper-parameter tuning, addressing all aspects of the challenge. Our post-hoc analyses revealed that providing additional time did not yield significant improvements. The winners’ code was open-sourced¹.

1. Introduction

Machine Learning (ML) has made remarkable progress in the past few years in time series-related tasks, including time series classification, time series clustering, time series regression, and time series forecasting Wang et al. (2017); Lim and Zohren (2020). To foster research in time series analysis, several competitions have been organized, since the onset of machine learning. These include the Santa Fe competition², the Sven Crone competitions³,

-
1. <https://www.4paradigm.com/competition/autoseries2020>
 2. <https://archive.physionet.org/physiobank/database/santa-fe/>
 3. <http://www.neural-forecasting-competition.com/>

several Kaggle competitions including M5 Forecasting⁴, Web Traffic Time Series Forecasting⁵, to name a few. While time series forecasting remains a very challenging problem for ML, successes have been reported on problems of time series regression and classification in practical applications Tan et al. (2020); Wang et al. (2017).

Despite these advances, switching domain, or even analysing a new dataset from the same domain, still requires considerable human engineering effort. To address this problem, recent research has been directed to Automated Machine Learning (AutoML) frameworks Hutter et al. (2018); Yao et al. (2018), whose charter is to reduce human intervention in the process of rolling out machine learning solutions to specific tasks. AutoML approaches include designing (or meta-learning) generic reusable pipelines and/or learning machine architectures, fulfilling specific task requirements, and designing optimization methods devoid of (hyper-)parameter choices. To stimulate research in this area, we launched with our collaborators a series of challenges exploring various application settings⁶, whose latest editions include the Automated Graph Representation Learning (AutoGraph) challenge at the KDD Cup AutoML track⁷, Automated Weakly Supervised Learning (AutoWeakly) challenge at ACML 2019⁸, Automated Computer Vision (AutoCVLiu et al. (2020)) challenges at IJCNN 2019 and ECML PKDD 2019, etc.

This paper presents the design and results of the Automated Time Series Regression (AutoSeries) challenge, one of the competitions of the WSDM Cup 2020 (Web Search and Data Mining conference) that we co-organized, in collaboration with 4Paradigm and ChaLearn.

This challenge addresses “time series regression” tasks Hyndman and Athanasopoulos (2021). In contrast with “strict” forecasting problems in which *forecast* variable(s) \mathbf{y}_t should be predicted from **past** values **only** (often \mathbf{y} values alone), **time series regression** seeks to predict \mathbf{y}_t using **past** $\{t - t_{min}, \dots, t - 1\}$ AND **present** t values of one (or several) “covariate” *feature* time series $\{\mathbf{x}_t\}$ ⁹. Typical scenarios in which \mathbf{x}_t is known at the time of predicting \mathbf{y}_t include cases in which \mathbf{x}_t values are scheduled in advance or hypothesized for decision making purposes. Examples include: *scheduled events* like upcoming sales promotions, *recurring events* like holidays, or *forecasts* obtained by external accurate simulators, like weather forecasts. This challenge addresses in particular **multivariate** time series regression problems, in which \mathbf{x}_t is a feature vector or a matrix of *covariate* information, and \mathbf{y}_t is a vector. The domains considered include air quality, sales, parking, and city traffic forecasting. Data are feature-based and represented in a “tabular” manner. The challenge was run with **code submission** and the participants were evaluated on the Codalab challenge platform, without any human intervention, on five datasets in the feedback phase and five different datasets in the final “private” phased (with full blind testing of a single submission).

While future AutoSeries competitions might address other difficulties, this particular competition focused on the following 10 questions:

4. <https://www.kaggle.com/c/m5-forecasting-accuracy>

5. <https://www.kaggle.com/c/web-traffic-time-series-forecasting>

6. <http://automl.chalearn.org>, <http://autodl.chalearn.org>

7. <https://www.automl.ai/competitions/3>

8. <https://autodl.lri.fr/competitions/64>

9. In some application domains (not considered in this paper), even **future** $\{t + 1, \dots, t + t_{max}\}$ values of the covariates may be considered. An example would be “simultaneous translation” with a small lag.

- Q1: **Beyond autoregression: Time series regression.** Do participants exploit covariates/features $\{\mathbf{x}_t\}$ to predict y_t , as opposed to only past y ?
- Q2: **Explainability.** Do participants make an effort to provide an explainable model, e.g. by identifying the most predictive features in $\{\mathbf{x}_t\}$?
- Q3: **Multivariate/multiple time series.** Do participants exploit the joint distribution/relationship of various time series in a dataset?
- Q4: **Diversity of sampling rates.** Can methods developed handle different sampling rates (hourly, daily, etc.)?
- Q5: **Heterogeneous series length.** Can methods developed handle truncated series either at the beginning or the end?
- Q6: **Missing data.** Can methods developed handle (heavily) missing data?
- Q7: **Data streaming.** Do models update themselves according to newly acquired streaming test data (to be explained in Sec 2.2)?
- Q8: **Joint model and HP selection.** Can models select automatically learning machines and hyper-parameters?
- Q9: **Transfer/Meta learning.** Are solutions provided generic and applicable to new domains or at least new datasets of the same domain?
- Q10: **Hardware constraints.** Are computational/memory limitations observed?

2. Challenge Setting

2.1 Phases

The AutoSeries challenge had three phases: a **Feedback Phase**, a **Check Phase** and a **Private Phase**. In the Feedback Phase, five “feedback datasets” were provided to evaluate participants’ AutoML models. The participants could read error messages in log files made available to them (e.g. if their model failed due to missing values) and obtain performance and ranking feedback on a leaderboard. When the Feedback Phase finished, five new “private datasets” were used in the Check Phase and the Private Phase. The Check Phase was a brief transition phase in which the participants submitted their models to the platform to verify whether the model ran properly. No performance information or log files were returned to them. Using a Check Phase is a particular feature of this challenge, to avoid disqualifying participants on the sole ground that their models timed out, used an excessive amount of memory, or raised another exception possible to correct without specific feedback on performance. Finally in the Private Phase, the participants submitted blindly their debugged models, to be evaluated by the same five datasets as in Check Phase.

As previously indicated, in addition to the five feedback datasets and five private datasets, two public datasets were provided for offline practice.

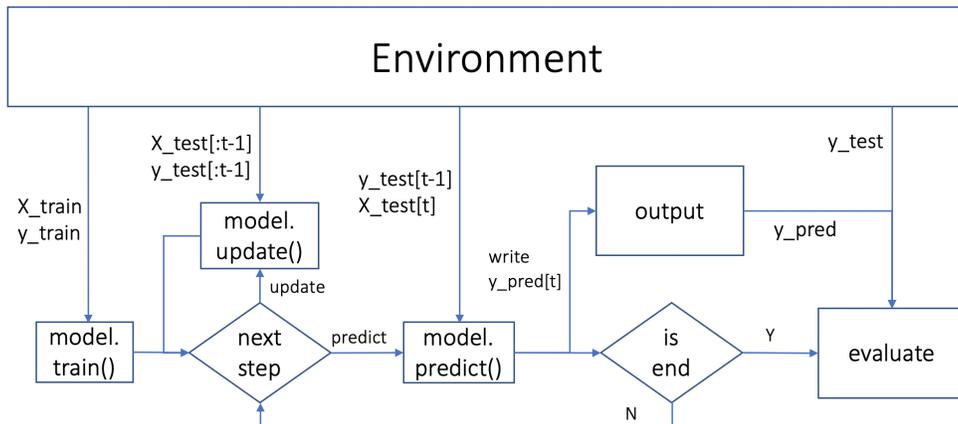


Figure 1: **Challenge protocol.** `train`, `update`, and `predict` methods must be provided by participants. Such methods are under control of timers, omitted in the figure.

2.2 Protocol

The AutoSeries challenge was designed based on real business scenarios, emphasizing **automated machine learning** (AutoML) and **data streaming**. First, as in other AutoML challenges, algorithms were evaluated on various datasets entirely hidden to the participants, **without any human intervention**. In other time series challenges, such as Kaggle’s Web Traffic Time Series Forecasting¹⁰), participants downloaded and explored past training data, and manually tuned features or models. The AutoSeries challenge forced the participants to design **generic methods**, instead of developing *ad hoc* solutions. Secondly, test data were streamed such that at each time point t , historical information of past time steps $\mathbf{x}_{train}[:t-1]$, $\mathbf{y}_{train}[:t-1]$ and features of time t , $X_{test}[t]$ were available for predicting \mathbf{y}_t . In addition to the usual `train` and `predict` methods, the participants had to prepare a method `update`, together with a strategy to update their model at an appropriate frequency, once fully trained on the training data. Updating too frequently might lead to run out of time; updating not frequently enough could result in missing recent useful information and performance degradation. The protocol is illustrated in Figure 1.

2.3 Datasets

The datasets from the Feedback Phase and final Private Phase are listed in Table 1. We purposely chose datasets from various domains, having a diversity of types of variables (continuous/categorical), number of series, noise level, amount of missing values, and sampling frequency (hourly, daily, monthly), and level of nonstationarity. Still, we eased the difficulty by including in each of the two phases datasets having some resemblance.

Two types of tabular formats are commonly used: the “wide format” and “long format”¹¹. The wide format facilitates visualization and direct use with machine learning packages. It consists in one time record per line, with feature values (or series) in columns.

10. <https://www.kaggle.com/c/web-traffic-time-series-forecasting>

11. <https://doc.dataiku.com/dss/latest/time-series/data-formatting.html>

However, for large number of features and/or missing values, the long format is preferred. In that format, a minimum of 3 columns are provided: (1) date and time (referred to as “Main_Timestamp”), (2) feature identifier (referred to as “ID_Key”), (3) feature value. Pivoting is an operation, which allows converting the wide format into the long format and vice-versa. From the long format, given one value of ID_Key (or a set of ID_Keys), a particular time series is obtained by ordering the feature values by Main_Timestamp. In this challenge, since we address a time series regression problem, we add a fourth column (4) “Label/Regression_Value” providing the target value, which must always be provided. A data sample is found in Table 2 and data visualizations in Figure 2.

Table 1: **Statistics of all 10 datasets.**

Sampling “Period” is indicated in (M)inutes, (H)ours, (D)ays. “Row” and “Col” are the total number of lines and columns, in the long format. Columns includes: Timestamp, (multiple) Id_Keys, (multiple) Features, and Target. “KeyNum” is the number of Id_Keys (called Id_Key combination, e.g. in a sales problem Product_Id and Store_Id.) “FeatNum” indicates the number of features for each Id_Key combination (e.g. for a given Id_Key corresponding to a product in a given store, features include price, and promotion.) “ContNum” is the number of continuous features and “CatNum” is the number of categorical features; CatNum + ContNum = FeatNum. “IdNum” means the number of unique Id_Key combinations. One can verify that Col = 1 (timestamp) + KeyNum + FeatNum + 1 (target). “Budget” is the time in seconds that we allow participants’ models to run.

Dataset	Domain	Period	Row	Col	KeyNum	FeatNum	ContNum	CatNum	IdNum	Budget
fph1	Power	M	39470	29	1	26	26	0	2	1300
fph2	AirQuality	H	716857	10	2	6	5	1	21	2000
fph3	Stock	D	1773	65	0	63	63	0	1	500
fph4	Sales	D	3147827	23	2	19	10	9	8904	3500
fph5	Sales	D	2290008	23	2	19	10	9	5209	2000
pph1	Traffic	H	40575	9	0	7	4	3	1	1600
pph2	AirQuality	H	721707	10	2	6	5	1	21	2000
pph3	Sales	D	2598365	23	2	19	10	9	6403	3500
pph4	Sales	D	2518172	23	2	19	10	9	6395	2000
pph5	Parking	M	35501	4	1	1	1	0	30	350

Table 2: **Sample data for dataset fph2.** A1 = timestamp. A2, A3, A4, A5, A7 = continuous features. A6 = categorical feature (hashed). A8, A9 = Id columns (hashed). Hashing is used for privacy. A10 = target.

A1	A2	A3	A4	A5	A6	A7	A8	A9	A10
2013-03-01 00:00:00	-2.3	1020.8	-19.7	0.0	-457...578	0.5	657...216	-731...089	13.0
2013-03-01 01:00:00	-2.5	1021.3	-19.0	0.0	511...667	0.7	657...216	-731...089	6.0
2013-03-01 02:00:00	-3.0	1021.3	-19.9	0.0	511...667	0.2	657...216	-731...089	22.0
...
2017-02-28 19:00:00	10.3	1014.2	-12.4	0.0	495...822	1.8	784...375	156...398	27.0
2017-02-28 20:00:00	9.8	1014.5	-9.9	0.0	-286...752	1.5	784...375	156...398	47.0
2017-02-28 21:00:00	9.1	1014.6	-12.7	0.0	-213...128	1.7	784...375	156...398	18.0

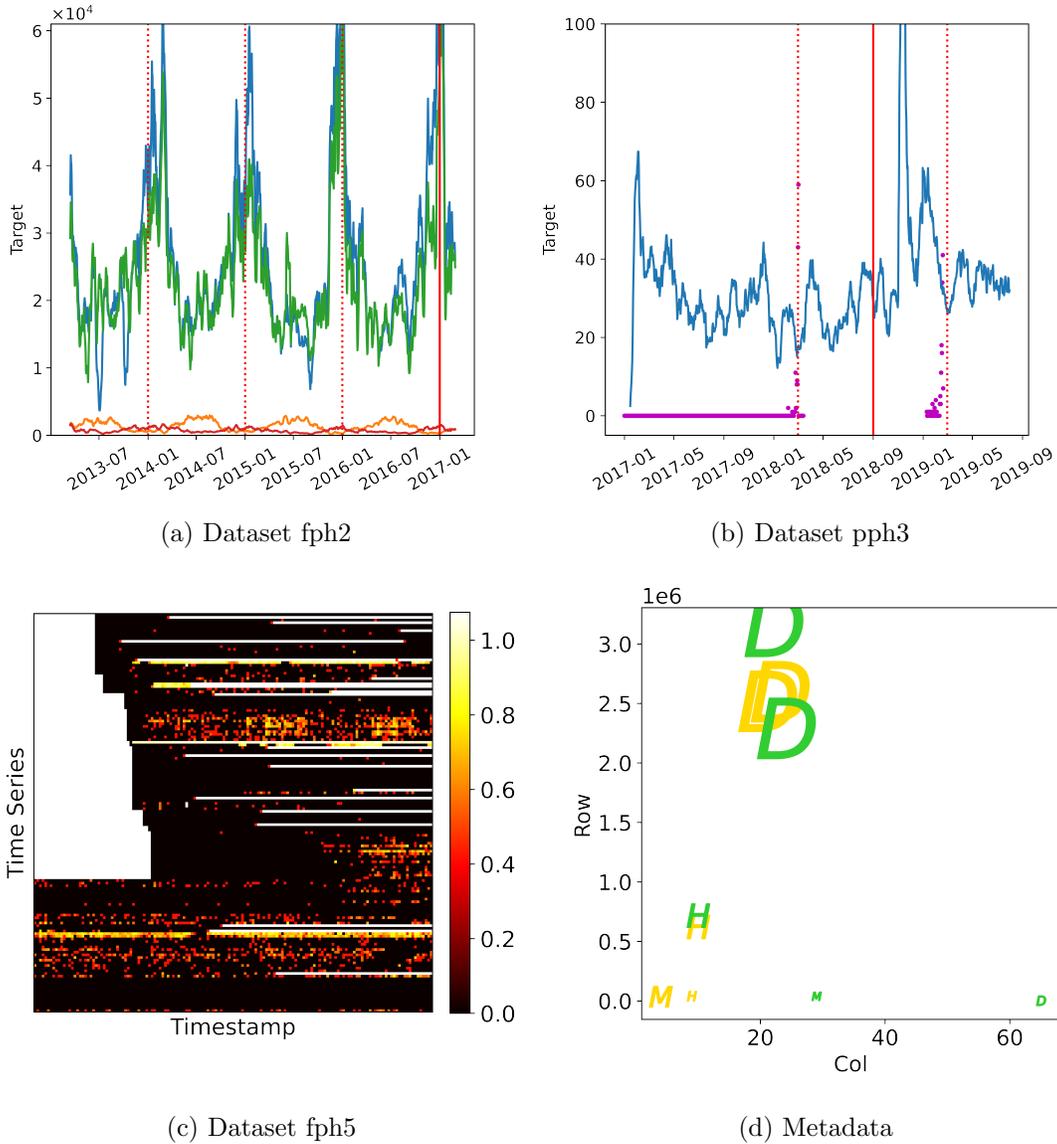


Figure 2: **Dataset visualization.** (a) **Sample visualization of dataset fph2.** Four time series (after smoothing and resampling). Training data is available until end of 2016 (red vertical solid line). Yearly periodicity is indicated by dashed vertical lines to highlight seasonalities. One notices large differences in series amplitudes and patterns of seasonality. (b) **Sample visualization of dataset pph3.** Two time series (after smoothing). Training data is available until end of 2018-08 (red vertical solid line). The purple time series suffers from missing values (certain items have zero sales most of the time). A clear trend exists in blue time series, unlike the example shown in (a). (c) **Heatmap visualization of dataset fph5.** White means missing value. Black means zero target value (sales). Several common issues can be observed: (1) many items don't sell most of the time; (2) presence of many missing values; (3) time series vary in lengths and are not aligned; (4) different time series have totally different scales. (d) **All dataset metadata visualization.** X axis is the number of columns. Y axis is the number of rows. The symbol letter shape represents the time period: Monthly, Daily, or Hourly. The symbol color represents the phase: green for “feedback” and orange for “private”. The symbol size represents the number of lines in the dataset.

2.4 Metrics

The metric used to judge the participants is the RMSE. For each datasets, the participant’s submissions are run in the same environment, and ranked according to the RMSE for each dataset. Then, an overall ranking is obtained from the average dataset rank, in a given phase. In post challenge analyses, we also used two other metrics: SMAPE and Correlation (CORR). The formulas are provided below. y means ground truth target. \hat{y} is the prediction. \bar{y} is the mean. N is total number of unique Id combinations (IdNum in Table 1) and T is number of timestamps. For evaluation, these metrics are run on the test sequences only.

$$\text{RMSE} = \sqrt{\frac{1}{NT} \sum_{n=1}^N \sum_{t=1}^T (y_{nt} - \hat{y}_{nt})^2} \tag{1}$$

$$\text{SMAPE} = \frac{1}{NT} \sum_{n=1}^N \sum_{t=1}^T \frac{|y_{nt} - \hat{y}_{nt}|}{(|y_{nt}| + |\hat{y}_{nt}| + \epsilon)/2} \tag{2}$$

$$\text{CORR} = \frac{\sum_{n=1}^N \sum_{t=1}^T (y_{nt} - \bar{y})(\hat{y}_{nt} - \bar{\hat{y}})}{\sqrt{\sum_{n=1}^N \sum_{t=1}^T (y_{nt} - \bar{y})^2} \sqrt{\sum_{n=1}^N \sum_{t=1}^T (\hat{y}_{nt} - \bar{\hat{y}})^2}} \tag{3}$$

2.5 Platform, Hardware and Limitations

The AutoSeries challenge is hosted on CodaLab¹², an open sourced challenge platform. We provide 4-core 30GB memory CPU and no GPU is available. Participants may submit at most 5 times per day. A docker is provided¹³ for executing submissions and for offline development. Participants can also install external packages if necessary.

2.6 Baseline

To help participants get started, we provided a baseline method, which is simple but contains necessary modules in the processing pipeline. Many participants’ submissions were derived from this baseline. In what follows, we decompose solutions (baseline and winning methods) into three modules: **feature engineering** (including time processing, numerical features, categorical features), **model training** (including models used, hyperparameter tuning, ensembling) and **update strategy** (including when and how to update models with the steaming test data). For the baseline, such modules include:

- **Feature engineering.** Multiple calendar features are extracted from the time stamp: *year*, *month*, *day*, *weekday*, and *hour*. Categorical variables (or strings) are hashed to unique integers. No preprocessing is applied to numerical features.
- **Model training.** A single LightGBM Ke et al. (2017) model is used. A LightGBM regressor is instantiated by predetermined hyperparameters and there is no hyperparameter tuning.

12. <https://autodl.lri.fr/>

13. <https://hub.docker.com/r/vergilgxw/autotable>

- **Update strategy.** Since the test data comes in a streaming way, we need an update strategy to incorporate new test data and adjust our model. However, due to time limit on update procedure, we can't update too frequently. The update strategy used in baseline is simple. We split all test timestamps by 5 segments and for every segment, we retrain the lightGBM with old training data and new segment of test data.

Table 3: Answers to the 10 challenge question. All of them are tackled to certain extent. Orange checkmark means the solution is trivial, though answers the question.

Question	Answered?	Comment
Q1 Beyond autoregression	✓	Features $\{x_t\}$ are leveraged
Q2 Explainability	✓	LightGBM outputs feature importance
Q3 Multivariate/multiple time series	✓	All training data is used to fit
Q4 Diversity of sampling rates	✓	Multiple calendar features are extracted
Q5 Heterogeneous series length	✓	Long format data facilitates the issue
Q6 Missing data	✓	Missing data is imputed by mean value
Q7 Data streaming	✓	Models are retrained every few steps
Q8 Joint model and HP selection	✓	Randomized grid search is applied
Q9 Transfer/Meta Learning	✓	Metadata (size, IdNum) is considered
Q10 Hardware constraints	✓	Model training time is recorded

2.7 Results

The AutoSeries challenge lasted one month and a half. We received over 700 submissions and more than 40 teams from both Academia (University of Washington, Nanjing University, etc.) and Industry (Oura, DeepBlue Technology, etc.), coming from various countries including China, United States, Singapore, Japan, Russia, Finland, etc. In the Feedback Phase¹⁴, the top five participants are: **rekcahd**, **DeepBlueAI**, **DenisVorotyntsev**, **DeepWisdom**, **Kon** while in the Private Phase, the top five participants are: **DenisVorotyntsev**, **DeepBlueAI**, **DeepWisdom**, **rekcahd**, **bingo**. It can be seen that team **rekcahd** seems to overfit on the Feedback Phase (additional experiments are provided in Sec 3.2). All winners use LightGBM Ke et al. (2017) which is boosting ensemble of decision trees dominating most tabular challenges. Only 1st winner and 2nd winner implements hyperparameter tuning module which is really a key to successful generalisation in AutoSeries. We briefly summarize the solutions and provide a detailed account in Appendix.

- **Feature engineering.** Calendar features e.g *year*, *month*, *day* were extracted from timestamp. Lag/shift and diff features added to original numerical features. Categorical features were encoded in various ways to integers.
- **Model training.** Only linear regression models and LightGBM were used. Most participants used default or fixed hyperparameters. Only the first winner made use

14. <https://autodl.lri.fr/competitions/149#results>

of HPO. The second winner optimized only the learning rate. LightGBM provides built-in feature importance/selection. Model ensembling was obtained by weighting models based on their performance in the previous round.

- **Update strategy.** All participants updated their models. The update period was either hard coded, computed as a fixed fraction of the time budget, or re-estimating on-the-fly, given remaining time.

We verified (in Table 3) that the challenge successfully answered the ten questions we wanted addressed (see Section 1).

3. Post Challenge Experiments

This section presents systematic experiments, which consolidate some of our findings and extend them. We are particularly interested in verifying the generalisation ability of winning solutions on a larger number of tasks, and comparing them with open-sourced AutoSeries solutions. We also revisit some of our challenge design choices to provide guidelines for future challenges, including time budget limitations, and choice and number of datasets.

3.1 Reproducibility

First, we reproduce the solutions of the top four participants and the baseline methods, on the 10 datasets of the challenge (from both phases). In the AutoSeries challenge, we only used the RMSE (Eq 1) for evaluation. For a more thorough comparison, we also include the SMAPE (Eq 2) here for calculating the relative error (which is particularly useful when the ground truth target is small, e.g. in the case of sales). The results are shown in the Table 4a, Table 4b, Table 8, Figure 3a and Figure 3b. We ran each method on each dataset for 10 times with different random seed. For simplicity, we use *1st DV*, *2nd DB*, *3rd DW*, *4th Rek* to denote solutions from top 4 winners.

We can observe that clear improvements have been made by the top winners, compared to the baseline, and both RMSE and SMAPE are significantly reduced. From Figure 3a we can further visualize that, while sometimes the winners’ solutions are close in RMSE, their SMAPE are totally different, which implies the necessity of using multiple metrics for evaluation.

3.2 Overfitting and generalisation

Based on our reproduced results, we analyse potential overfitting which is visualized in Figure 3c. Among each run (based on a different random seed), we rank solutions on feedback phase datasets and private phase datasets separately. Rankings are based on RMSE as in AutoSeries challenge. After 10 runs, we plot the mean and std of the ranking as a region. This shows that 4th Rek overfits to feedback datasets since it performs very well in feedback phase but poorly in private phase. But it is also interesting to visualize that 1st DV has a good generalisation: although it is not the best in feedback phase, it achieves great results in private phase. Including hyperparameter search may have provided the winner with a key advantage.

Table 4: **Post-challenge runs.** We repeated 10 times all runs on all datasets for the top ranking submissions of the private phase. Each run is based on a different random seed. Error bar are indicated (one standard deviation) unless no variance was observed (algorithm with no stochastic component, such as 4th Rek).

Dataset	Phase	Baseline	1st DV	2nd DB	3rd DW	4th Rek
fph1	Feedback	100±10	40.7±0.2	40.0±0.1	40.1±0.2	40.7
fph2	Feedback	18000±2000	237±2	244±1	243.9±0.3	230.7
fph3	Feedback	3000	600±20	53.04±0.01	52.4	108.4
fph4	Feedback	6.9±0.3	3.66±0.02	2.760±0.007	NA	2.632
fph5	Feedback	8.6±0.7	5.76±0.01	5.760±0.005	5.780±0.003	5.589
pph1	Private	400±6	200±2	223.5±0.7	200±10	420.8
pph2	Private	17000±4000	240±2	253.7±0.5	260±2	246.7
pph3	Private	9±3	6.20±0.02	6.330±0.007	6.40±0.03	6.56
pph4	Private	12±2	4.0±0.2	3.80±0.04	3.700± 0.004	3.32
pph5	Private	300±30	50±1	100	60±20	167.8

(a) RMSE comparison.

Dataset	Phase	Baseline	1st DV	2nd DB	3rd DW	4th Rek
fph1	Feedback	140±20	100.0±0.5	104.00±0.07	104.00±0.04	40.77
fph2	Feedback	140±10	30±1	40.0±0.3	38.8±0.1	33.9
fph3	Feedback	38.49	5.0±0.1	0.770±0.001	0.700±0.001	1.674
fph4	Feedback	190±1	191.0±0.1	191.0±0.1	NA	186.2
fph5	Feedback	170±1	173.5±0.1	174.1± 0.1	172.9±0.1	170.6
pph1	Private	12.7±0.6	6.1±0.2	6.49±0.03	6.4±0.8	12.14
pph2	Private	140±10	24.0±0.5	35.0±0.3	31.0±0.1	32.75
pph3	Private	180±3	180.0±0.7	181.0±0.1	180±0.1	174.7
pph4	Private	170±3	170±1	167.8±0.1	170.0±0.2	164
pph5	Private	40±2	6.0±0.5	9	8±3	30.61

(b) SMAPE comparison.

Table 5: Supported features comparison between various open-source packages and the AutoSeries winning solution (also open-sourced).

Solutions	FeatureEngineering	ModelTraining	StreamingUpdate	TimeManagement
FeaturetoolsKanter and Veeramachani (2015)	Tabular	✗	✗	✗
tsfresh ¹⁵	Temporal	✗	✗	✗
ProphetTaylor and Letham (2017)	✗	✓	✗	✗
GluonTSAlexandrov et al. (2020)	Temporal	✓	✗	✗
AutoKerasJin et al. (2019)	✗	✓	✗	✓
AutoGluonErickson et al. (2020)	Tabular	✓	✗	✓
Google AutoTable ¹⁶	Tabular	✓	✓	✓
AutoSeries	Temporal	✓	✓	✓

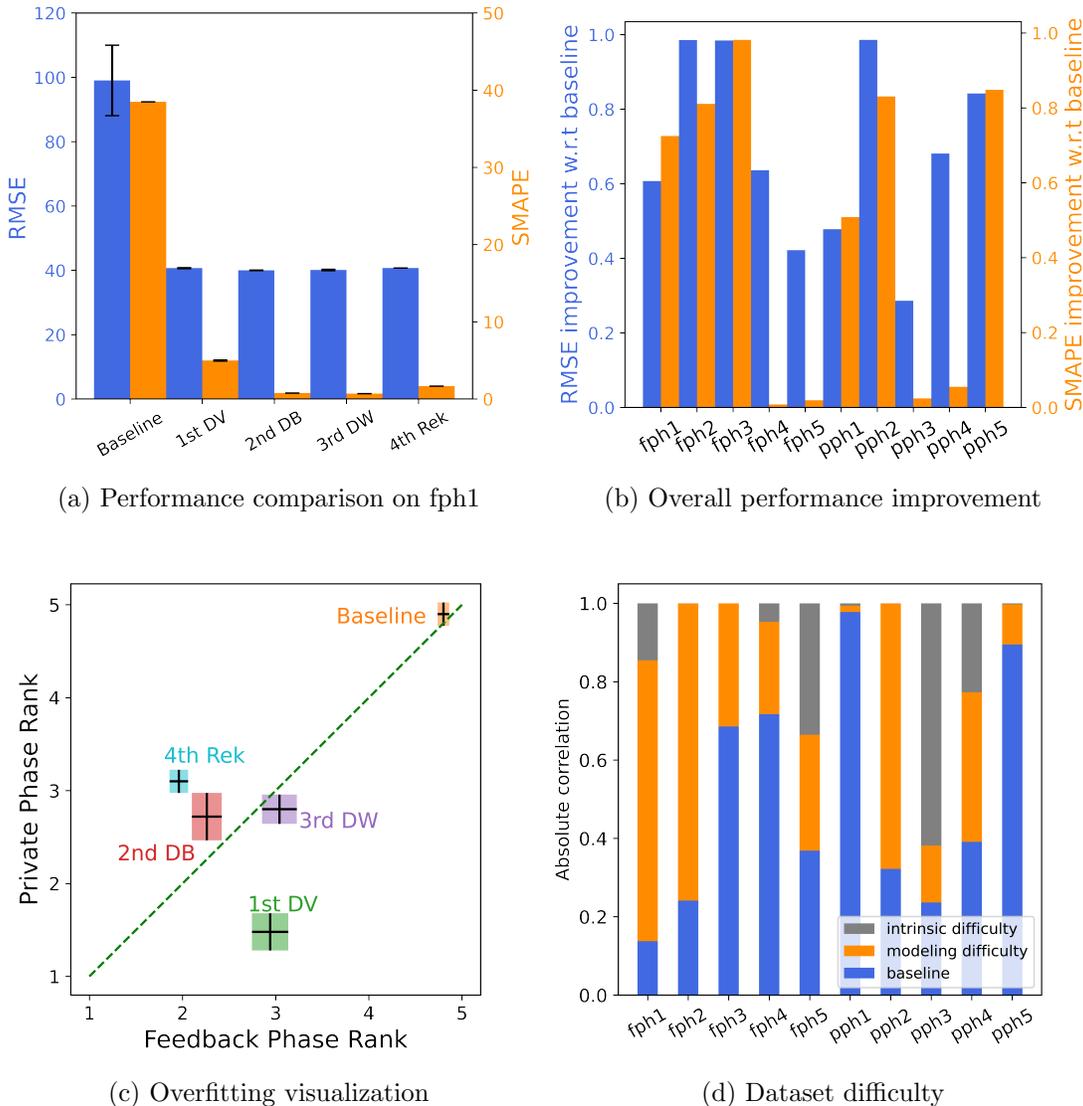


Figure 3: **Post challenge experiments.** (a) **Performance comparison on dataset fph1.** We compare RMSE and SMAPE of all solutions on dataset fph1. Performances in RMSE are significantly better than the baseline for all winning teams, but all winners perform similarly. In contrast the SMAPE metric differentiates the winners, which focuses more on relative error. (b) **Performance improvement on all datasets.** Both RMSE and SMAPE errors from best methods are compared to the baseline performance. The improvement ratio is calculated by (baseline score - best score) / baseline score. (c) **Did the participants overfit the feed-back phase tasks?** Rankings are based on 10 runs: for each run, we rank separately in order to validate the stability of methods. Regions show the mean and std of rankings over multiple runs. Methods in the upper triangle are believed to overfit, e.g. 4th winner’s solution. Methods in the lower triangle are believed to generalize well, e.g. 1st winner’s solution. (d) **How well did we choose the 10 datasets?** As explained in Sec 3.5, we use absolute correlation as a bounded metric for calculating a notion of intrinsic difficulty (gray bar) and modeling difficulty (orange bar) of the 10 datasets. Datasets with high modeling difficulty and low intrinsic difficulty are better choices in a benchmark.

3.3 Comparison to open source AutoML solutions

In this section, we turn our attention to comparing AutoSeries with similar open-source solutions. However, to the best of our knowledge, there is no publicly available AutoML framework dedicated to time series data. Current features (categorized by three modules of solutions as in Sec) of open source packages, which can be used to tackle the problems of the challenge with some engineering effort, are summarized in Table 5.

Packages like Featuretools, tsfresh focus on (tabular, temporal) feature engineering; they do not provide trainable models and should be used in conjunction with another package. Prophet and GluonTS are known for rapid prototyping with time series, but they are not AutoML packages (in the sense that they do not come with automated model selection and hyper-parameter selection). AutoKeras is an package focusing more on image and text, with KerasTuner¹⁷ for neural architecture search. Google AutoTable meets most of our requirements, but is not open sourced, and is not dedicated to time series. Moreover, Google AutoTable costs around 19 dollars per hour in order to train on 92 computing instances at the same time, which is far more than our challenge settings.

At last, we selected AutoGluon for comparison, as being closest to our use case. AutoGluon provides end-to-end automated pipelines to handle tabular data without any human intervention (e.g. hyperparameter tuning, data preprocessing). AutoGluon includes many more candidate models and fancier ensemble methods than the winning solutions, but its feature engineering is not dedicated to multivariate time series. For example, it doesn't distinguish time series Id combinations to summarize statistics of one particular time series. We ran AutoGluon on all 10 datasets with default parameters except for using RMSE as evaluation metric and best_quality as presets parameter. The results are summarized in Table 6 column AutoGluon. Not surprisingly, vanilla AutoGluon can only beat the baseline, and it is significantly worse than the winning solutions. We further compile AutoGluon with 1st winner's time series feature engineering and update the models the same way as in baseline. The results are in Table 6 column FE+AutoGluon. AutoGluon can now indeed achieve comparable results with best winner and sometimes even better, which strongly implies the importance of time series feature engineering. Note that we didn't limit strictly AutoGluon's running time as in our challenge. In general, AutoGluon takes 10 times more time than the winning solution and it still can't output a valid performance on the four datasets in a reasonable time. For the six AutoGluon's feasible datasets, we further visualize in Figure 4 by algorithm groups. AutoGluon contains mainly three algorithm groups: Neural Network (MXNet, FastAI), Ensemble Trees (LightGBM, Catboost, XGBoost) and K-Nearest Neighbors. We first plot on the left the average RMSE for Neural Networks models and ensemble tree models each (we omit KNN methods since they are usually the worst). Note that among the 6 datasets, 3 datasets don't use Neural Network for final ensemble (so their RMSE are set to be a large number for visualization). On 2 datasets (bottom left corner), however, Neural Networks can be competitive. This encourages us to explore in the future the effectiveness of deep models on time series which evolve quickly these days. On the right, we average the training/inference time per algorithm group and find that KNN can be used for very fast prediction if needed. Neural Networks take significantly more time. Points above the dotted line mean that no NN models or KNN models are chosen for

17. <https://keras-team.github.io/keras-tuner/>

Table 6: **Comparison with AutoGluon.** NA means a missing value: AutoGluon did not terminate within a reasonable time.

Dataset	Phase	Baseline		1st DV		AutoGluon		FE + AutoGluon	
		RMSE	SMAPE	RMSE	SMAPE	RMSE	SMAPE	RMSE	SMAPE
fph1	Feedback	99.04	142.59	40.69	102.19	90.19	26.45	40.57	105.31
fph2	Feedback	17563	142.64	236.6	26.63	14978	59.94	263.74	25.51
fph3	Feedback	3337	38.49	623.32	4.99	6365	116.14	3159	31.08
fph4	Feedback	6.91	187.58	3.66	190.94	NA	NA	NA	NA
fph5	Feedback	8.63	174.45	5.76	173.54	NA	NA	NA	NA
pph1	Private	422.37	12.65	218.83	6.11	2770.70	9.46	212.68	5.85
pph2	Private	16851	139.31	242.41	23.46	15028	57.04	269.85	22.98
pph3	Private	8.78	178.45	6.21	177.08	NA	NA	NA	NA
pph4	Private	11.54	174.94	3.74	168.4	NA	NA	NA	NA
pph5	Private	309.33	39.2	50.37	5.91	949.4	20.52	65.22	6.65

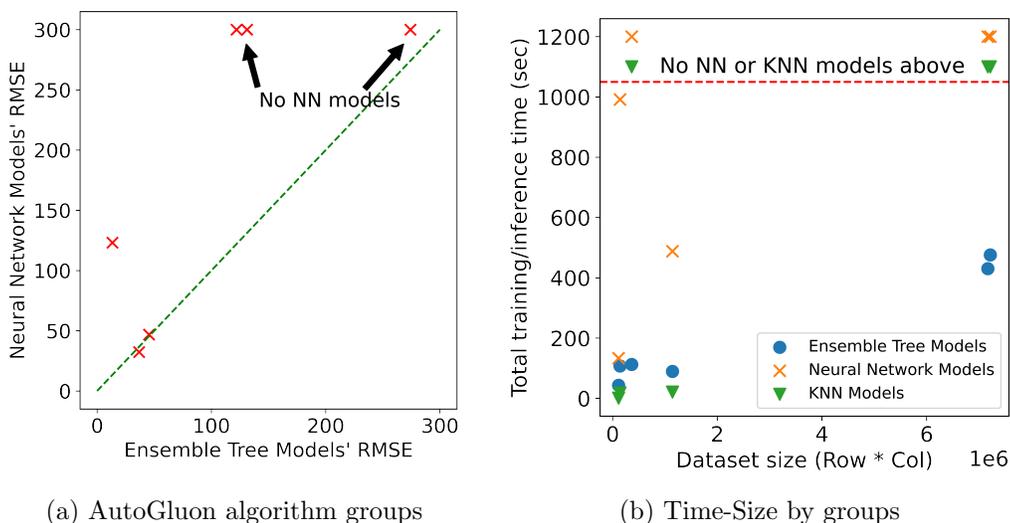


Figure 4: **AutoGluon experiments.** (a) **Average performance for two algorithm groups.** Here we compare the average RMSE of Neural Network models and Ensemble Tree models. Among the six feasible datasets for AutoGluon with time series feature engineering, 3 of them don't choose Neural Networks as ensemble candidates. Ensemble trees have always significantly better performances. On 2 datasets, Neural Networks are quite competitive. (b) **Average time costs of candidate models.** When dataset is large, only ensemble tree models are chosen. When dataset is medium, KNN is fastest, followed by tree models. Neural Networks take significantly more time.

this dataset (either due to performance or time cost). Only the tree-based methods provide solutions across the range of dataset sizes.

3.4 Impact of time budget

In the AutoSeries challenge, time management is an important aspect. Different time budgets are allowed for different datasets (as shown in Table 1). Ideally, AutoSeries solutions should take into account the allowed time budget and adapt all modules in the pipeline (i.e. different feature engineering, model training and updating strategy based on different allowed time budgets). We double the time budget and compare the performance in Appendix. In general, no obviously stable improvement can be observed. We also try to half the time budget and most solutions can't even produce valid predictions meaning that no single model training is finished. This could be because that we set the default budgets too tight but it also shows from another perspective that participants' solutions overfit to the challenge design (default time budget).

3.5 Dataset Difficulty

After a challenge finishes, another important issue for the organizers is to validate the choice of datasets. This is particularly interesting for AutoML challenges since the point is to generalize to a wide variety of tasks. Inspired by difficulty measurements in Liu et al. (2020), we want to define intrinsic difficulty and modeling difficulty. By intrinsic difficulty we mean the irreducible error. As a surrogate to the intrinsic difficulty, we use the error of the best model. By modeling difficulty, we mean the range or spread of performances of candidate models. To separate well competition participants, we want to choose datasets of low intrinsic difficulty and high modeling difficulty. In Liu et al. (2020), a notion of intrinsic difficulty and modeling difficulty is introduced for classification problems. Here we adapt such ideas and choose another bounded metric, the correlation (CORR) (Eq 3). In fact, correlation has been used in many time series papers as a metric Lai et al. (2018); Wang et al. (2019). We calculate the absolute correlation between the prediction sequence and ground truth test sequence. We define **Intrinsic difficulty** as 1 minus the best solution's absolute correlation score; and **Modeling difficulty** as the difference between the best solution's absolute correlation score and the provided baseline score.

These difficulty measures are visualized in Figure 3d. It is obvious that both intrinsic difficulty and modeling difficulty differ from datasets to datasets. A posteriori, we can observe that some datasets like pph1 and pph5 are too easy, while pph3 is too difficult. In general, feedback datasets are of higher quality than private datasets, which is unfortunate. However, it is also possible that participants overfit the feedback datasets and thus, by using the best performing methods to estimate the intrinsic difficulty, we obtain a biased estimation.

4. Conclusion and future work

In this challenge, we introduce an AutoML setting with streaming test data, aiming at pushing forward research on Automated Time Series, and also having an impact on industry. Since there were no open sourced AutoML solutions dedicated to time series prior to our challenge, we believe the open sourced AutoSeries solutions fill this gap and provide a useful tool to researchers and practitioners. AutoSeries solutions don't need a GPU which facilitates their adoption.

The solutions of the winners are based on lightGBM. They addressed all challenge questions, demonstrating the feasibility of automating time series regression on datasets of the type considered. Significant improvements were made compared to the provided baseline. Our generalisation and overfitting experiments show that hyperparameter search is key to generalize. Still, some of the questions were addressed in a rather trivial way and deserve further attention. Explainability boils down to the feature importance delivered by lightGBM. In future challenge designs, we might want to quantitatively evaluate this aspect. Missing data were trivially imputed with the mean value. Hyper-parameters were not thoroughly optimized by most participants, and simple random search was used (if at all). Our experiments with the AutoGluon package demonstrate that much can be done in this direction to further improve results. Additionally, no sophisticated method of transfer learning or meta-learning was used. Knowledge transfer was limited to the choice of features and hyper-parameters performed on the feedback phase datasets. New challenge designs could include testing meta-learning capabilities on the platform, by letting the participant’s code meta-train on the platform, e.g. not resetting the model instances when presented with each new dataset.

Other self criticisms of our design include that some datasets in the private phase may have been too easy or too difficult. Additionally, the RMSE alone could not separate well solutions, while a combination of metrics might be more revealing. Lastly, GPUs were not provided. On one hand this forced the participants to deliver practical rapid solutions; on the other hand, this precluded them from exploring neural time series models, which are rapidly progressing in this field.

Finally, winning solutions overfitted to the provided time budgets (no improvement with more time and fail with less time). An incentive to encourage participants to deliver “any-time-learning” solutions as opposed to “fixed-time-learning” solutions is to use the area under the learning curve as metric, as we did in other challenges. We will consider this for future designs.

References

- Alexander Alexandrov, Konstantinos Benidis, Michael Bohlke-Schneider, Valentin Flunkert, Jan Gasthaus, Tim Januschowski, Danielle C. Maddix, Syama Sundar Rangapuram, David Salinas, Jasper Schulz, Lorenzo Stella, Ali Caner Türkmen, and Yuyang Wang. Gluonts: Probabilistic and neural time series modeling in python. *J. Mach. Learn. Res.*, 2020.
- Nick Erickson, Jonas Mueller, Alexander Shirkov, Hang Zhang, Pedro Larroy, Mu Li, and Alexander Smola. Autogluon-tabular: Robust and accurate automl for structured data. 2020.
- Frank Hutter, Lars Kotthoff, and Joaquin Vanschoren, editors. *Automated Machine Learning: Methods, Systems, Challenges*. Springer, 2018. <http://automl.org/book>.
- Rob J. Hyndman and George Athanasopoulos, editors. *Forecasting: principles and practice*. OTexts, 2021. OTexts.com/fpp3. Accessed on 2021/03/25.

- Haifeng Jin, Qingquan Song, and Xia Hu. Auto-keras: An efficient neural architecture search system. In *KDD*, 2019.
- James Max Kanter and Kalyan Veeramachaneni. Deep feature synthesis: Towards automating data science endeavors. In *IEEE International Conference on Data Science and Advanced Analytics, DSAA*, 2015.
- Guolin Ke, Qi Meng, Thomas Finley, Taifeng Wang, Wei Chen, Weidong Ma, Qiwei Ye, and Tie-Yan Liu. Lightgbm: A highly efficient gradient boosting decision tree. In *Advances in Neural Information Processing Systems*, 2017.
- Guokun Lai, Wei-Cheng Chang, Yiming Yang, and Hanxiao Liu. Modeling long- and short-term temporal patterns with deep neural networks. In *SIGIR*, 2018.
- Bryan Lim and Stefan Zohren. Time series forecasting with deep learning: A survey. 2020.
- Zhengying Liu, Zhen Xu, Sergio Escalera, Isabelle Guyon, Júlio C. S. Jacques Júnior, Meysam Madadi, Adrien Pavao, Sébastien Treguer, and Wei-Wei Tu. Towards automated computer vision: analysis of the autocv challenges 2019. *Pattern Recognit. Lett.*, 2020.
- Liudmila Ostroumova Prokhorenkova, Gleb Gusev, Aleksandr Vorobev, Anna Veronika Dorogush, and Andrey Gulin. Catboost: unbiased boosting with categorical features. In *Advances in Neural Information Processing Systems*, 2018.
- Chang Wei Tan, Christoph Bergmeir, Francois Petitjean, and Geoffrey I. Webb. Time series extrinsic regression. 2020.
- Sean J. Taylor and Benjamin Letham. Forecasting at scale. *PeerJ Prepr.*, 2017.
- Lijing Wang, Jiangzhuo Chen, and Madhav Marathe. DEFSI: deep learning based epidemic forecasting with synthetic information. In *AAAI*, 2019.
- Zhiguang Wang, Weizhong Yan, and Tim Oates. Time series classification from scratch with deep neural networks: A strong baseline. In *International Joint Conference on Neural Networks*, 2017.
- Quanming Yao, Mengshuo Wang, Yuqiang Chen, Wenyuan Dai, Yu-Feng Li, Wei-Wei Tu, Qiang Yang, and Yang Yu. Taking human out of learning applications: A survey on automated machine learning. 2018.

Appendix A. Detailed descriptions of winning methods

A.1 First place: DenisVorotyntsev

The 1st winning is from team DenisVorotyntsev. Their code is open sourced on GitHub¹⁸.

Feature engineering. First, a small LightGBM model is fit on training data and the top 3 most important numerical features are extracted. Then pairwise arithmetic operations are conducted on these top candidates to generate extra numerical features. Afterwards, a large number of lag features are generated. To deal with multivariate time series which contain one or more ID columns for indexing time series, a *batch_id* column is created by concatenating all ID columns (as strings). This *batch_id* column will be used by *groupby* in further processing steps.

For generating lag/shift features on target column, data is grouped by *batch_id* and for each group (i.e. one particular time series dataframe), lagged/shifted target and differences with respect to lagged/shifted target are recorded as additional features. Lags are by default a list of small number e.g. 1,2,3,5,7. Same lag feature generation is performed on numerical features as well, with fewer lags e.g. 1,2,3.

Categorical columns are concatenated with corresponding time series *batch_id*. They are encoded by CatBoost Prokhorenkova et al. (2018). The Target is linearly transformed to have minimum equal to 1. A possible target difference operation is optional and is added to the hyperparameter search space. Calendar features are extracted as in the baseline method.

Model training. LightGBM is the only used model. For the first time training, three steps are performed: base model fit, feature selection, hyperparameter search. For a base model, a lightGBM model is fit on all features including generated feature columns. Then a feature selection step is performed to remove least important features. They search in order 5 different splits *0.2*, *0.5*, *0.75*, *0.05*, *0.1* of most important features and fit on the selected ones. The best performing ratio is recorded. Lastly, a relatively large hyperparameter search space is defined and in total 2880 configurations of hyperparameters are randomly searched under training time limit. These hyperparameters include *num_leaves*, *min_child_samples*, *subsample_freq*, *colsample_bytree*, *subsample*, *lambda_l2*.

Update strategy. The core of update strategy is to determine the update frequency. They first calculate the affordable training time from the training time budget and the first training time cost, with a coefficient buffer. Then the update frequency is simply calculated with the number of test steps divided by affordable training rounds. After the update frequency is determined, the update function is called per certain steps, which calls the training module with the training data and newly streamed data.

18. <https://github.com/DenisVorotyntsev/AutoSeries>

A.2 Second place: DeepBlueAI

The second winning solution is from the team DeepBlueAI. Their code is open sourced on GitHub¹⁹.

Feature engineering. A mean imputer is used for missing data. Calendar features including *year, month, day, weekday, hour, minute* are extracted. In the case of multivariate time series, ID columns are merged into a unique identifier. Categorical features are encoded by pandas Categorical. For numerical features, many more features are generated after grouping by unique ID: mean, max, min, lag, division, etc. No feature selection is applied.

Model training. Two models are used: linear regression and LightGBM. For linear regression, they distinguish data with few features (less than 20) and data with rich features. In the former case, a simple fit is used. When there are lots of features, they select features based on a F-score of regression. For LightGBM, most hyperparameters are predefined except for learning rate. They choose learning rate by fitting different LightGBM models with different hyperparameters according to dataset size. What’s particular in the training and predict part is that they maintain two coefficients for ensembling linear model and LightGBM model. These coefficients are searched by comparing to ground truth.

Update strategy. The update strategy is the same as in the baseline method.

A.3 Third place: DeepWisdom

The third winning solution is from the team DeepWisdom. Their code is open sourced on GitHub²⁰.

Feature engineering. Four types of features are generated sequentially: KeysCross, TimeData, Delta and LagFeat. KeysCross applies in the case of multiple ID_Key for indexing time series, for example in the retail setting, we may have *shop_id* and *item_id* together to index the time series for certain item in certain shop. KeysCross converts string ID columns to integers first and then for each ID columns c , existing ID encoding is multiplied by $c.max$ and is added by $c.values$. TimeData is the calendar feature. It extracts *year, month, day, hour, weekday* from the timestamp column. Both Delta and LagFeat deal with the target column. They are basically diff and lag/shift features. Delta calculates first order difference, second order difference and numerical exchange ratio of the target. LagFeat calculates *mean, std, max, min* for a lag window of size *3,7,14,30*. For linear regression models used in later steps (as mentioned in next paragraph), one hot encoding of categorical variables and a mean imputer for numerical variables are used.

Model training. Three models are used for training: lasso regression, ridge regression and LightGBM. For all these models, hyperparameters are predefined manually and there is no hyperparameter search. During the prediction step, all trained models will produce outputs and a weighted linear combination is used for ensembling these models. The weights are

19. <https://github.com/DeepBlueAI/AutoSeries>

20. <https://github.com/DeepWisdom/AutoSeries2019>

inversely proportional to a model’s previous performance of RMSE.

Update strategy. The update strategy is the same as that of the 1st Place DenisVorotyntsev.

A.4 Fourth place: Rekcahd

The fourth place participant is from team rekcahd. Their code isn’t open sourced since this team is not among top three, and thus it was not a requirement for them to publish their solution. This solution relies a lot on the provided baseline and it achieved first place in the Feedback Phase.

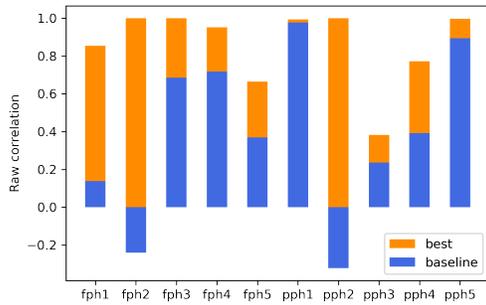
Feature engineering. Missing values are filled by feature means. Calendar features are extracted as in baseline Sec 2.6. The type adaptation module deals with categorical features particularly. In each group of a categorical feature (i.e. a sub dataframe where this particular category of same value. Note that all time series exist in the group), this categorical value (as long as there are not too few instances in the group), is mapped to a linear combination of target mean of this group and target mean of the whole training data.

For multivariate time series datasets, they add extra features. Concretely, two shifted target features are added (shifted once and twice); if target is always positive, square root of the target is added; difference with respect to shifted once values are calculated for numerical features; another feature for each categorical feature indicating whether it is same as previous time’s value is added. However, these shift, difference and category change indicator features are not calculated on univariate time series datasets, which explains to some extent the higher error of this solution on two univariate time series (fph3, pph1).

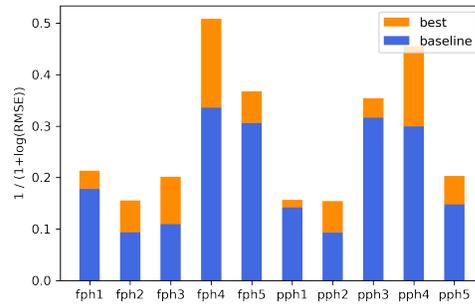
Model training. A linear regression model is fit on training data and is served as starting score for a LightGBM model. All hyperparameters of LightGBM are predefined except for *num_boost_round*, which is determined by the best iteration after fitting another LightGBM model.

Update strategy. The update strategy is the same as in baseline.

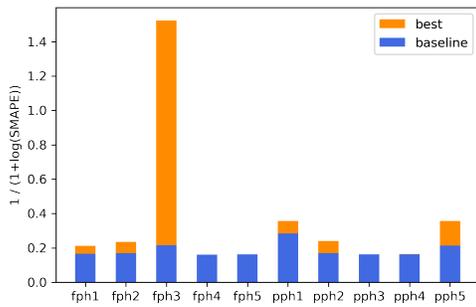
Appendix B. Dataset difficulty based on other metrics



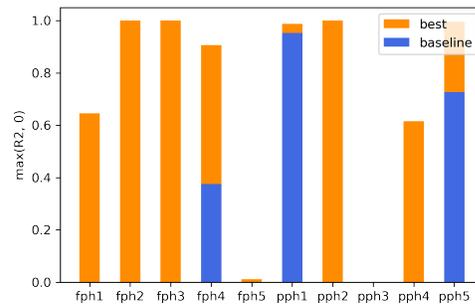
(a) Raw corr difficulty measure of 10 datasets.



(b) RMSE difficulty measure of 10 datasets.



(c) SMAPE difficulty measure of 10 datasets.



(d) R2 difficulty measure of 10 datasets.

Figure 5: Dataset difficulty based on other metrics.

Appendix C. Running time comparison

Table 7: **Performance improvement with double time budget.** A positive number (in bold) means performances have been improved (in percentage) given more time.

Dataset	Phase	Baseline		1st DV		2nd DB		3rd DW		4th Rek	
		RMSE	SMAPE	RMSE	SMAPE	RMSE	SMAPE	RMSE	SMAPE	RMSE	SMAPE
fph1	Feedback	23.50	18.24	-0.02	-1.47	0.05	0.00	-0.18	0.00	0.00	0.00
fph2	Feedback	26.70	3.84	-1.57	1.76	-0.90	-1.61	78.53	98.17	0.00	0.00
fph3	Feedback	0.00	0.00	0.65	0.50	-0.04	0.00	0.00	0.00	0.00	0.00
fph4	Feedback	2.59	-1.39	0.35	0.05	0.04	0.00	NA	NA	0.00	0.00
fph5	Feedback	-15.63	1.43	-0.38	0.06	-0.05	0.06	0.00	0.00	0.00	0.00
pph1	Private	-1.11	-3.65	1.17	7.97	0.00	0.03	-14.54	-43.05	0.00	0.00
pph2	Private	47.32	18.08	1.35	-4.05	0.35	0.57	0.61	0.92	0.00	0.00
pph3	Private	-13.18	-1.72	1.23	0.11	-0.27	-0.17	NA	NA	59.90	-6.58
pph4	Private	-119.24	5.67	-1.23	0.24	-0.21	-0.06	NA	NA	0.00	0.00
pph5	Private	-6.98	-13.37	2.97	-16.32	0.00	0.00	37.44	35.26	0.00	0.00

Table 8: **Duration comparison.** Mean and Std calculated on 10 runs. (Unit: sec). Each run is the same setting except for a different random seed. The method consuming most allowed time is bolded.

Dataset	Phase	Budget	Baseline	1st DV	2nd DB	3rd DW	4th Rek
fph1	Feedback	1300	70±10	700±100	600±100	600±50	260±60
fph2	Feedback	2000	150±20	1200±100	550±80	1300±70	200±30
fph3	Feedback	500	16±2	150±10	70±10	160±20	14±3
fph4	Feedback	3500	600±60	1200±40	1000±80	NA	800±70
fph5	Feedback	2000	200±20	900±30	900±80	1100±50	470±40
pph1	Private	1600	100±20	900±150	900±200	890±70	150±30
pph2	Private	2000	150±15	1100±100	600±80	1400±50	170±30
pph3	Private	3500	300±30	1300±40	900±80	2000±60	550±50
pph4	Private	2000	450±50	860±60	870±100	1000±80	500±50
pph5	Private	350	15±2	200±8	140±20	170±6	26±4