

Efficient Automated Deep Learning for Time Series Forecasting

Difan Deng^{✉1}, Florian Karl^{2,3},
Frank Hutter^{4,5}, Bernd Bischl^{2,3}, and Marius Lindauer¹

¹ Leibniz University Hannover, Hannover, Germany
`deng@tnt.uni-hannover.de`

² Ludwig-Maximilian University, Munich, Germany

³ Fraunhofer Institute for Integrated Circuits (IIS), Erlangen, Germany

⁴ University of Freiburg, Freiburg, Germany

⁵ Bosch Center for Artificial Intelligence, Renningen, Germany

Abstract. Recent years have witnessed tremendously improved efficiency of Automated Machine Learning (AutoML), especially Automated Deep Learning (AutoDL) systems, but recent work focuses on tabular, image, or NLP tasks. So far, little attention has been paid to general AutoDL frameworks for time series forecasting, despite the enormous success in applying different novel architectures to such tasks. In this paper, we propose an efficient approach for the joint optimization of neural architecture and hyperparameters of the entire data processing pipeline for time series forecasting. In contrast to common NAS search spaces, we designed a novel neural architecture search space covering various state-of-the-art architectures, allowing for an efficient macro-search over different DL approaches. To efficiently search in such a large configuration space, we use Bayesian optimization with multi-fidelity optimization. We empirically study several different budget types enabling efficient multi-fidelity optimization on different forecasting datasets. Furthermore, we compared our resulting system, dubbed **Auto-PyTorch-TS**, against several established baselines and show that it significantly outperforms all of them across several datasets.

Keywords: AutoML · Deep Learning · Time Series Forecasting · Neural Architecture Search

1 Introduction

Time series (TS) forecasting plays a key role in many business and industrial problems, because an accurate forecasting model is a crucial part of a data-driven decision-making system. Previous forecasting approaches mainly consider each individual time series as one task and create a local model [3, 7, 26]. In recent years, with growing dataset size and the ascent of Deep Learning (DL), research interests have shifted to global forecasting models that are able to learn information across all time series in a dataset collected from similar sources [20, 41].

Given the strong ability of DL models to learn complex feature representations from a large amount of data, there is a growing trend of applying new DL models to forecasting tasks [38, 46, 50, 57].

Automated machine learning (AutoML) addresses the need of choosing the architecture and its hyperparameters depending on the task at hand to achieve peak predictive performance. The former is formalized as neural architecture search (NAS) [14] and the latter as hyperparameter optimization (HPO) [17]. Several techniques from the fields of NAS and HPO have been successfully applied to tabular and image benchmarks [15, 18, 33, 62]. Recent works have also shown that jointly optimizing both problems provides superior models that better capture the underlying structure of the target task [61, 62].

Although the principle idea of applying AutoML to time series forecasting models is very natural, there are only few prior approaches addressing this [32, 37, 43, 52]. In fact, combining state-of-the-art AutoML methods, such as Bayesian Optimization with multi-fidelity optimization [16, 30, 34, 36], with state-of-the-art time series forecasting models leads to several challenges we address in this paper. First, recent approaches for NAS mainly cover cell search spaces, allowing only for a very limited design space, that does not support different macro designs [12, 60]. Our goal is to search over a large variety of different architectures covering state-of-the-art ideas. Second, evaluating DL models for time series forecasting is fairly expensive and a machine learning practitioner may not be able to afford many model evaluations. Multi-fidelity optimization, e.g. [36], was proposed to alleviate this problem by only allocating a fraction of the resources to evaluated configurations and promoting the most promising configurations to give them additional resources. Third, as a consequence of applying multi-fidelity optimization, we have to choose how different fidelities are defined, i.e. what kind of budget is used. Examples for such *budget types* are number of epochs, dataset size or time series length. Depending on the correlation between lower and highest fidelity, multi-fidelity optimization can boost the efficiency of AutoML greatly or even slow it down in the worst case. Since we are the first to consider multi-fidelity optimization for AutoML on time series forecasting, we studied the efficiency of different budget types across many datasets. Fourth, all of these need to be put together; to that effect, we propose a new open-source package for Automated Deep Learning (AutoDL) for time series forecasting, dubbed `Auto-PyTorch-TS`.⁶ Specifically, our contributions are as follows:

1. We propose the AutoDL framework `Auto-PyTorch-TS` that is able to jointly optimize the architecture and the corresponding hyperparameters for a given dataset for time series forecasting.
2. We present a unified architecture configuration space that contains several state-of-the-art forecasting architectures, allowing for a flexible and powerful macro-search.
3. We provide insights into the configuration space of `Auto-PyTorch-TS` by studying the most important design decisions and show that different architectures are reasonable for different datasets.

⁶ The code is available under <https://github.com/automl/Auto-PyTorch>.

4. We show that `Auto-PyTorch-TS` is able to outperform a set of well-known traditional statistical models and modern deep learning models with an average relative error reduction of 19% against the best baseline across many forecasting datasets.

2 Related Work

We start by discussing the most closely related work in DL for time series forecasting, AutoDL, and AutoML for time series forecasting.

2.1 Deep Learning based Forecasting

Early work on forecasting focused on building a local model for each individual series to predict future values, ignoring the correlation between different series. In contrast, global forecasting models are able to capture information of multiple time series in a dataset and use this at prediction time [31]. With growing dataset size and availability of multiple time series from similar sources, this becomes increasingly appealing over local models. We will in the following briefly introduce some popular forecasting DL models.

Simple feed-forward MLPs have been used for time series forecasting and extended to more complex models. For example, the N-BEATS framework [46] is composed of multiple stacks, each consisting of several blocks. This architectural choice aligns with the main principle of modern architecture design: Networks should be designed in a block-wise manner instead of layer-wise [63].

Additionally, RNNs [9, 23] were proposed to process sequential data and thus they are directly applicable to time series forecasting [22, 57]. A typical RNN-based model is the Seq2Seq network [9] that contains an RNN encoder and decoder. Wen et al. [57] further replaced the Seq2Seq’s RNN decoder with a multi-head MLP. Flunkert et al. [50] proposed DeepAR that wraps an RNN encoder as an auto-regressive model and uses it to iteratively generate new sample points based on sampled trajectories from the last time step.

In contrast, CNNs can extract local, spatially-invariant relationships. Similarly, time series data may have time-invariant relationships, which makes CNN-based models suitable for time series tasks, e.g. WaveNet [6, 45] and Temporal Convolution Networks (TCN) [4]. Similar to RNNs, CNNs could also be wrapped by an auto-regressive model to recursively forecast future targets [6, 45].

Last but not least, attention mechanisms and transformers have shown superior performance over RNNs on natural language processing tasks [56] and over CNNs on computer vision tasks [13]. Transformers and RNNs can also be combined; e.g. Lim et al. [38] proposed temporal fusion transformers (TFT) that stack a transformer layer on top of an RNN to combine the best of two worlds.

2.2 Automated Deep Learning (AutoDL)

State-of-the-art AutoML approaches include Bayesian Optimization (BO) [18], Evolutionary Algorithms (EA) [44], reinforcement learning [63] or ensembles [15].

Most of them consider AutoML system as a black-box optimization problem that aims at finding the most promising machine learning models and their optimal corresponding hyperparameters. Neural Architecture Search (NAS), on the other hand, only contains one search space: its architecture. NAS aims at finding the optimal architecture for the given task with a fixed set of hyperparameters. Similar to the traditional approach, the architecture could be optimized with BO [33, 62], EA [49] or Reinforcement Learning [63] among others, but there also exist many NAS-specific speedup techniques, such as one-shot models [59] and zero-cost proxies [1]. In this work we follow the state-of-the-art approach from Auto-PyTorch [62] and search for both the optimal architecture and its hyperparameters with BO.

Training a deep neural network requires lots of computational resources. Multi-fidelity optimization [16, 30, 36] is a common approach to accelerate AutoML and AutoDL. It prevents the optimizer from investing too many resources on the poorly performing configurations and allows for spending more on the most promising ones. However, the correlation between different fidelities might be weak [60] for DL models, in which case the result on a lower fidelity will provide little information for those on higher fidelities. Thus, it is an open question how to properly select the budget type for a given target task, and researchers often revert to application-specific decisions.

2.3 AutoML for Time Series Forecasting

While automatic forecasting has been of interest in the research community in the past [28], dedicated AutoML approaches for time series forecasting problems have only been explored recently [21, 32, 35, 42, 51]. Optimization methods such as random search [55], genetic algorithms [10], monte carlo tree search and algorithms akin to multi-fidelity optimization [51] have been used among others. Paldino et al. [47] showed that AutoML frameworks not intended for time series forecasting originally - in combination with feature engineering - were not able to significantly outperform simple forecasting strategies; a similar approach is presented in [10]. As part of a review of AutoML for forecasting pipelines, Meisenbacher et al. [42] concluded that there is a need for optimizing the entire pipeline as existing works tend to only focus on certain parts. We took all of these into account by proposing `Auto-PyTorch-TS` as a framework that is specifically designed to optimize over a flexible and powerful configuration space of forecasting pipelines.

3 AutoPyTorch Forecasting

For designing an AutoML system, we need to consider the following components: optimization targets, configuration space and optimization algorithm. The high-level workflow of our `Auto-PyTorch-TS` framework is shown in Figure 1; in many ways it functions similar to existing state-of-the-art AutoML frameworks [17, 62]. To better be able to explain unique design choice for time series forecasting, we

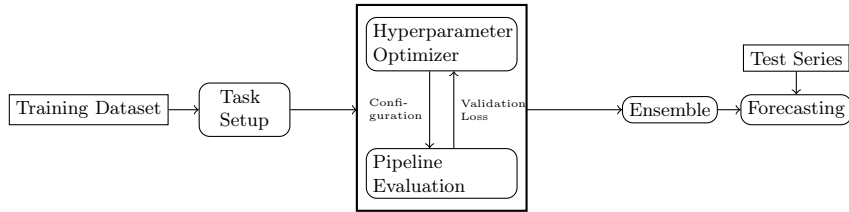


Fig. 1: An overview of Auto-PyTorch-TS. Given a dataset, Auto-PyTorch-TS automatically prepares the data to fit the requirement of a forecasting pipeline. The AutoML optimizer will then use the selected budget type to search for desirable neural architectures and hyperparameters from the pipeline configuration space. Finally, we create an ensemble out of the most promising pipelines to do the final forecasting on the test sets.

first present a formal statement of the forecasting problem and discuss challenges in evaluating forecasting pipelines before describing the components in detail.

3.1 Problem Definition

A multi-series forecasting task is defined as follows: given a dataset that contains N series: $\mathcal{D} = \{\mathcal{D}_i\}_{i=1}^N$ and \mathcal{D}_i represents one series in the dataset: $\mathcal{D}_i = \{\mathbf{y}_{i,1:T_i}, \mathbf{x}_{i,1:T_i}^{(p)}, \mathbf{x}_{i,T_i+1:T_i+H}^{(f)}\}^T$, where T is the number of time steps until forecasting starts; H is the forecasting horizon that the model is required to predict; $\mathbf{y}_{1:T}$, $\mathbf{x}_{1:T}^{(p)}$ and $\mathbf{x}_{T_i+1:T_i+H}^{(f)}$ are the sets of observed past targets, past features and known future features values, respectively. The task of time series forecasting is to predict the possible future values with a model trained on \mathcal{D} :

$$\hat{\mathbf{y}}_{T+1:T+H} = f(\mathbf{y}_{1:T}, \mathbf{x}_{1:T+H}; \boldsymbol{\theta}) \quad (1)$$

where $\mathbf{x}_{1:T+H} := [\mathbf{x}_{1:T}^{(p)}, \mathbf{x}_{T+1:T+H}^{(f)}]$, $\boldsymbol{\theta}$ are the model parameters that are optimized with training losses \mathcal{L}_{train} , and $\hat{\mathbf{y}}_{T+1:T+H}$ are the predicted future target values. Depending on the model type, $\hat{\mathbf{y}}_{T+1:T+H}$ can be distributions [50] or scalar values [46]. Finally, the forecasting quality is measured by the discrepancy between the predicted targets $\hat{\mathbf{y}}_{T+1:T+H}$ and the ground truth future targets $\mathbf{y}_{T+1:T+H}$ according to a defined loss function \mathcal{L} . The most commonly applied metrics include mean absolute scaled error (MASE), mean absolute percentage error (MAPE), symmetric mean absolute percentage error (sMAPE) and mean absolute error (MAE) [19, 29, 46].

3.2 Evaluating Forecasting Pipelines

We split each sequence into three parts to obtain: a training set $\mathcal{D}_{train} = \{\mathbf{y}_{1:T-H}, \mathbf{x}_{1:T+H}\}$, a validation set $\mathcal{D}_{val} = \{\mathbf{y}_{T-H+1:T}, \mathbf{x}_{T-H+1:T}\}$ and a test

⁷ For the sake of brevity, we omit the sequence index i in the following part of this paper unless stated otherwise.

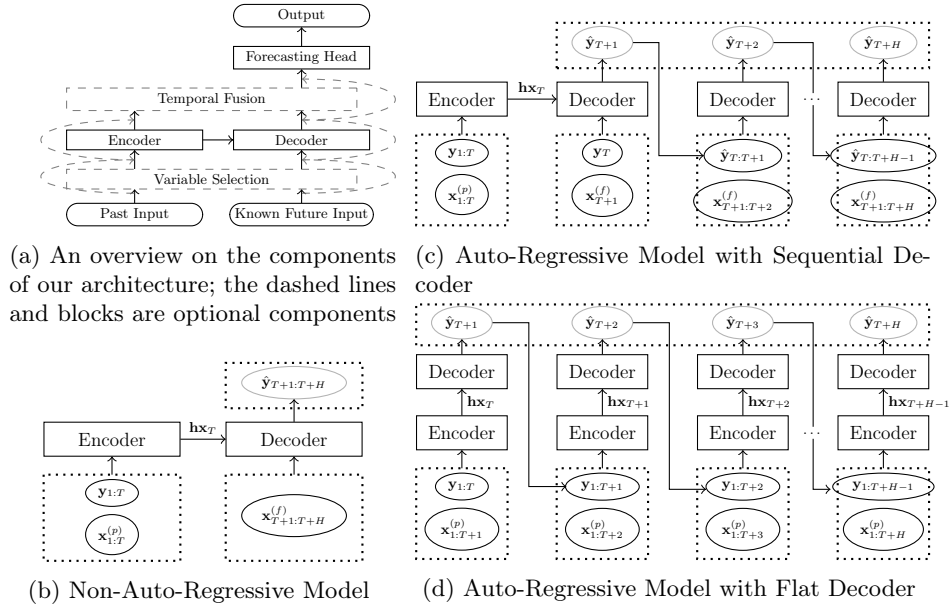


Fig. 2: Overview of the architectures that can be built by our framework. (a) shows the main components of our architecture space. (b)-(d) are specific instances of (a) and its data flow given different architecture properties.

set $\mathcal{D}_{\text{test}} = \{\mathbf{y}_{T+1:T+H}, \mathbf{x}_{T+1:T+H}\}$, i.e., the tails of each sequences are reserved as \mathcal{D}_{val} . At each iteration, our AutoML optimizer suggests a new hyperparameter and architecture configuration λ , trains it on $\mathcal{D}_{\text{train}}$ and evaluates it on \mathcal{D}_{val} .

Both in AutoML frameworks [18, 62] and in forecasting frameworks [46], ensembling of models is a common approach. We combine these two worlds in `Auto-PyTorch-TS` by using ensemble selection [8] to construct a weighted ensemble that is composed of the best k forecasting models from the previously evaluated configurations $\mathcal{D}_{\text{hist}}$. Finally, we retrain all ensemble members on $\mathcal{D}_{\text{val}} \cup \mathcal{D}_{\text{train}}$ before evaluating on $\mathcal{D}_{\text{test}}$.

3.3 Forecasting Pipeline Configuration Space

Existing DL packages for time series forecasting [2, 5] follow the typical structure of traditional machine learning libraries: models are built individually with their own hyperparameters. Similar to other established AutoML tools [15, 18, 44], we designed the configuration space of `Auto-PyTorch-TS` as a combined algorithm selection and hyperparameter (CASH) problem [53], i.e., the optimizer first selects the most promising algorithms and then optimizes for their optimal hyperparameter configurations, with a hierarchy of design decisions. Deep neural networks, however, are built with stacked blocks [63] that can be disentangled to fit different requirements [58]. For instance, Seq2Seq [9], MQ-RNN [57] and

Table 1: An overview of the possible combinations and design decisions of the models that exists in our configuration space. Only the TFT Network contains the optional components presented in Figure 2a.

	Encoder	Decoder	auto-regressive	Architecture Class
Flat Encoder	MLP	MLP	No	Feed Forward Network
	N-BEATS	N-BEATS	No	N-BEATS [46]
Seq. Encoder	RNN/Transformer	RNN/Transformer	Yes	Seq2Seq [9]
			No	TFT [38]
	MLP	Yes	DeepAR [50]	
		No	MQ-RNN [57]	
	TCN	MLP	Yes	DeepAR [50]/WaveNet [45]
			No	MQ-CNN [57]

DeepAR [50] all contain an RNN as their encoders. These models naturally share common aspects and cannot be simply treated as completely different models. To fully utilize the relationships of different models, we propose a configuration space that includes all the possible components in a forecasting network.

As shown in Figure 2a, most existing forecasting architectures can be decomposed into 3 parts: *encoder*, *decoder* and *forecasting heads*: the encoder receives the past target values and embeds them into the latent space. The latent embedding, together with the known future features (if applicable), are fed to the decoder network; the output of the decoder network is finally passed to the forecasting head to generate a sequence of scalar values or distributions, depending on the type of forecasting head. Additionally, the *variable selection*, *temporal fusion* and *skip connection layers* introduced by TFT [38] can be seamlessly integrated into our networks and are treated as optional components.

Table 1 lists all possible choices of encoders, decoders, and their corresponding architectures in our configuration space. Specifically, we define two types of network components: sequential encoder (Seq. Encoder) and flat encoder (Flat Encoder). The former (e.g., RNN, Transformer and TCN) directly processes sequential data and output a new sequence; the latter (e.g., MLP and N-BEATS) needs to flatten the sequential data into a 2D matrix to fuse the information from different time steps. Through this configuration space, `Auto-PyTorch-TS` is able to encompass the “convex hull” of several state-of-the-art global forecasting models and tune them.

As shown in Figure 2, given the properties of encoders, decoders, and models themselves, we construct three types of architectures that forecast the future targets in different ways. Non-Auto-Regressive models (Figure 2b), including MLP, MQ-RNN, MQ-CNN, N-BEATS and TFT, forecast the multi-horizontal predictions within one single step. In contrast, Auto-Regressive models do only one-step forecasting within each forward pass. The generated forecasting values are then iteratively fed to the network to forecast the value at the next time step. All the auto-regressive models are trained with teacher forcing [22]. Only

sequential networks could serve as an encoder in auto-regressive models, however, we could select both sequential and flat decoders for auto-regressive models. Sequential decoders are capable of independently receiving the newly generated predictions. We consider this class of architectures as a Seq2Seq [9] model: we first feed the past input values to the encoder to generate its output \mathbf{hx} and then pass \mathbf{hx} to the decoder, as shown in Figure 2c. Having acquired \mathbf{hx} , the decoder then generates a sequence of predictions with the generated predictions and known future values by itself. Finally, Auto-Regressive Models with flat decoders are classified as the family of DeepAR models [50]. As the decoder could not collect more information as the number of generated samples increases, we need to feed the generated samples back to the encoder, as shown in Figure 2d.

Besides its architectures, hyperparameters also play an important role on the performance of a deep neural network [61], for the details of other hyperparameters in our configuration space, we refer to the Appendix.

3.4 Hyperparameter Optimization

We optimize the loss on the validation set $\mathcal{L}_{\mathcal{D}_{\text{val}}}$ with BO [17]. It is known for its sample efficiency, making it a good approach for expensive black-box optimization tasks, such as AutoDL for expensive global forecasting DL models. Specifically, we optimize the hyperparameters with SMAC [25]⁸ that constructs a random forest to model the loss distribution over the configuration space.

Similar to other AutoML tools [18, 62] for supervised classification, we utilize multi-fidelity optimization to achieve better any-time performance. Multi-fidelity optimizers start with the lowest budget and gradually assign higher budgets to well-performing configurations. Thereby, the choice of what budget type to use is essential for the efficiency of a multi-fidelity optimizer. The most popular choices of budget type in DL tasks are the number of epochs and dataset size. For time series forecasting, we propose the following four different types of budget:

- Number of Epochs (*#Epochs*)
- Series Resolution (*Resolution*)
- Number of Series (*#Series*)
- Number of Samples in each Series (*#SMPs per Ser.*)

A higher *Resolution* indicates an extended sample interval. The sample interval is computed by the inverse of the fidelity value, e.g., a resolution fidelity of 0.1 indicates for each series we take every tenth point: we shrink the size of the sliding window accordingly to ensure that the lower fidelity optimizer does not receive more information than the higher fidelity optimizer. *#Series* means that we only sample a fraction of sequences to train our model. Finally, *#SMPs per Ser.* indicates that we decrease the expected value of the number of samples within each sequence; see Section 3.2 for sample-generation method. Next to these multi-fidelity variants, we also consider vanilla Bayesian optimization (*Vanilla BO*) using the maximum of all these fidelities.

⁸ We used SMAC3 [39] from <https://github.com/automl/SMAC3>

3.5 Proxy-Evaluation on Many Time Series

All trained models must query every series to evaluate \mathcal{L}_{val} . However, the number of series could be quite large. Additionally, many forecasting models (e.g., DeepAR) are cheap to be trained but expensive during inference time. As a result, rather than training time, inference time is more likely to become a bottleneck to optimize the hyperparameters on a large dataset (for instance, with 10k series or more), where configuration with lower fidelities would no longer provide the desirable speed-up when using the full validation set. Thereby, we consider a different evaluation strategy on large datasets (with more than 1k series) and lower budgets: we ask the model to only evaluate a fraction of the validation set (we call this fraction “proxy validation set”) while the other series are predicted by a dummy forecaster (which simply repeats the last target value in the training series, i.e., \mathbf{y}_T , H times). The size of the proxy validation set is proportional to the budget allocated to the configuration: maximal budget indicates that the model needs to evaluate the entire validation set. We set the minimal number of series in the proxy set to be 1k to ensure that it contains enough information from the validation set. The proxy validation set is generated with a grid to ensure that all the configurations under the same fidelity are evaluated on the same proxy set.

4 Experiments

We evaluate **Auto-PyTorch-TS** on the established benchmarks of the Monash Time Series Forecasting Repository [20]⁹. This repository contains various datasets that come from different domains, which allows us to assess the robustness of our framework against different data distributions. Additionally, it records the performance of several models, including local models [3, 7, 11, 26, 27], global traditional machine learning models [48, 54], and global DL models [2, 6, 46, 50, 56] on \mathcal{D}_{test} , see [20] for details. For evaluating **Auto-PyTorch-TS**, we will follow the exact same protocol and dataset splits. We focus our comparison of **Auto-PyTorch-TS** against two types of baselines: (i) the overall single best baseline from [20], assuming a user would have the required expert knowledge and (ii) the best dataset-specific baseline. We note that the latter is a very strong baseline and a priori it is not known which baseline would be best for a given dataset; thus we call it the *theoretical oracle baseline*. Since the Monash Time Series Forecasting Repository does not record the standard deviation of each method, we reran those baselines on our cluster for 5 times. Compared to the repository, our configuration space includes one more strong class of algorithms, TFT [38], which we added to our set of baselines to ensure a fair and even harder comparison.

We set up our task following the method described in Section 3.2: HPO is only executed on $\mathcal{D}_{train/val}$ while H is given by the original repository. As described in Section 3.2, we create an ensemble with size 20 that collects multiple

⁹ <https://forecastingdata.org/>

models during the course of optimization. When the search finishes, we refit the ensemble to the union of $\mathcal{D}_{train/val}$ and evaluate the refitted model on \mathcal{D}_{test} . Both \mathcal{L}_{val} and \mathcal{L}_{test} are measured with the mean value of MASE [29] across all the series in the dataset. To leverage available expert knowledge, **Auto-PyTorch-TS** runs an initial design with the default configurations of each model in Table 1. Please note that this initial design will be evaluated on the smallest available fidelity. All multi-fidelity variants of **Auto-PyTorch-TS** start with the cheapest fidelity of 1/9, use then 1/3 and end with the highest fidelity (1.0). The runs of **Auto-PyTorch-TS** are repeated 5 times with different random seeds.

We ran all the datasets on a cluster node equipped with 8 Intel Xeon Gold 6254@ 3.10GHz CPU cores and one NVIDIA GTX 2080TI GPU equipped with PyTorch 1.10 and Cuda 11.6. The hyperparameters were optimized with SMAC3 v1.0.1 for 10 hours, and then we refit the ensemble on $\mathcal{D}_{train/val}$ and evaluate it on the test set. All the jobs were finished within 12 hours.

4.1 Time Series Forecasting

Table 2 shows how different variants of **Auto-PyTorch-TS** perform against the two types of baselines across multiple datasets. Even using the theoretical oracle baseline for comparison, **Auto-PyTorch-TS** is able to outperform it on 18 out of 24 datasets. On the other 6 datasets, it achieved nearly the same performance as the baselines. On average, we were able to reduce the MASE by up to 5% against the oracle and by up to 19% against the single best baseline, establishing a new robust state-of-the-art overall.

Surprisingly, the forecasting-specific budget types did not perform significantly better than the number of epochs (the common budget type in classification). Nevertheless, the optimal choice of budget type varies across datasets, which aligns with our intuition that on a given dataset the correlation between lower and higher fidelities may be stronger for certain budget types than for other types. If we were to construct a theoretically optimal budget-type selector, which utilizes the best-performing budget type for a given dataset, we would reduce the relative error by 2% over the single best (i.e., # SMPs per Ser.).

4.2 Hyperparameter Importance

Although HPO is often considered as a black-box optimization problem [17], it is important to shed light on the importance of different hyperparameters to provide insights into the design choice of DL models and to indicate how to design the next generation of AutoDL systems.

Here we evaluate the importance of the hyperparameters with a global analysis based on fANOVA [24], which measures the importance of hyperparameters by the variance caused by changing one single hyperparameter while marginalizing over the effect of all other hyperparameters. Results on individual datasets can be found in appendix.

Table 2: We compare variants of **Auto-PyTorch-TS** against the single best baseline (TBATS) and a theoretically optimal oracle of choosing the correct baseline for each dataset wrt mean MASE errors on the test sets. We show the mean and standard deviation for each dataset. The best results are highlighted in bold-face. We computed the relative improvement wrt the Oracle Baseline on each dataset and used the geometric average for aggregation over the datasets.

	Auto-PyTorch-TS				Vanilla BO	Best dataset-specific baseline	Overall single best baseline
	# Epochs	Resolution	# Series	# SMPs per Ser.			
M3 Yearly	2.73(0.10)	2.66(0.05)	2.76(0.09)	2.64 (0.09)	2.68(0.08)	2.77(0.00)	3.13(0.00)
M3 Quarterly	1.08 (0.01)	1.10(0.01)	1.10(0.01)	1.09(0.02)	1.12(0.03)	1.12(0.00)	1.26(0.00)
M3 Monthly	0.85 (0.01)	0.89(0.02)	0.86(0.01)	0.87(0.04)	0.86(0.02)	0.86(0.00)	0.86(0.00)
M3 Other	1.90(0.07)	1.82(0.03)	1.98(0.13)	1.92(0.05)	1.95(0.15)	1.81 (0.00)	1.85(0.00)
M4 Quarterly	1.15(0.01)	1.13(0.01)	1.13 (0.01)	1.15(0.01)	1.15(0.02)	1.16(0.00)	1.19(0.00)
M4 Monthly	0.93(0.02)	0.93(0.02)	0.93(0.02)	0.93 (0.02)	0.96(0.02)	0.95(0.00)	1.05(0.00)
M4 Weekly	0.44(0.01)	0.45(0.02)	0.43 (0.02)	0.44(0.02)	0.45(0.01)	0.48(0.00)	0.50(0.00)
M4 Daily	1.14(0.01)	1.18(0.07)	1.16(0.06)	1.14(0.04)	1.38(0.41)	1.13 (0.02)	1.16(0.00)
M4 Hourly	0.86(0.12)	0.95(0.11)	0.78 (0.07)	0.85(0.07)	0.85(0.06)	1.66(0.00)	2.66(0.00)
M4 Yearly	3.05 (0.03)	3.08(0.04)	3.05(0.01)	3.09(0.04)	3.10(0.02)	3.38(0.00)	3.44(0.00)
Tourism Quarterly	1.61(0.03)	1.57(0.05)	1.59(0.05)	1.59(0.02)	1.55(0.03)	1.50 (0.01)	1.83(0.00)
Tourism Monthly	1.42 (0.03)	1.44(0.03)	1.45(0.04)	1.47(0.02)	1.42(0.02)	1.44(0.02)	1.75(0.00)
Dominick	0.51(0.04)	0.49(0.00)	0.49 (0.01)	0.49(0.01)	0.49(0.01)	0.51(0.00)	0.72(0.00)
Kdd Cup	1.20(0.02)	1.18(0.02)	1.18(0.03)	1.18(0.03)	1.20(0.03)	1.17 (0.01)	1.39(0.00)
Weather	0.63(0.08)	0.58(0.04)	0.59(0.02)	0.59(0.06)	0.57 (0.00)	0.64(0.01)	0.69(0.00)
NN5 Daily	0.79(0.01)	0.80(0.01)	0.81(0.04)	0.78 (0.01)	0.79(0.01)	0.86(0.00)	0.86(0.00)
NN5 Weekly	0.76(0.01)	0.76(0.03)	0.76(0.01)	0.77(0.01)	0.76 (0.01)	0.77(0.01)	0.87(0.00)
Hospital	0.76(0.01)	0.76(0.00)	0.76(0.00)	0.75 (0.01)	0.75(0.01)	0.76(0.00)	0.77(0.00)
Traffic Weekly	1.04(0.07)	1.10(0.03)	1.04(0.05)	1.08(0.09)	1.03(0.07)	0.99 (0.03)	1.15(0.00)
Electricity Weekly	0.78(0.04)	1.06(0.13)	0.80(0.04)	0.74 (0.07)	0.85(0.11)	0.76(0.01)	0.79(0.00)
Electricity Hourly	1.52(0.05)	1.54(0.00)	1.58(0.08)	1.54(0.06)	1.51 (0.05)	1.60(0.02)	3.69(0.00)
Kaggle Web Traffic Weekly	0.56(0.01)	0.56(0.01)	0.55 (0.00)	0.57(0.01)	0.59(0.01)	0.61(0.02)	0.62(0.00)
Covid Deaths	5.11(1.60)	4.54(0.05)	4.43 (0.13)	4.58(0.30)	4.53(0.24)	5.16(0.04)	5.72(0.00)
Temperature Rain	0.76(0.05)	0.75(0.01)	0.73(0.02)	0.73(0.03)	0.71(0.04)	0.71 (0.03)	1.23(0.00)
∅ Rel. Impr Best	0.82	0.83	0.81	0.81	0.82	0.86	1.0
∅ Rel. Impr Oracle	0.96	0.97	0.95	0.95	0.96	1.0	1.17

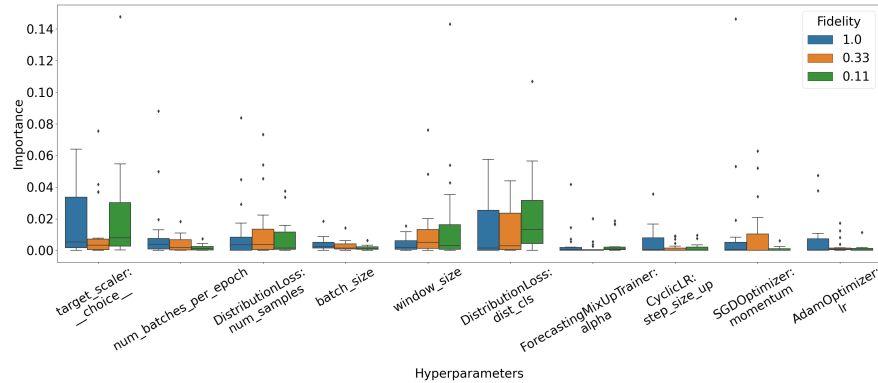


Fig. 3: Hyperparameter importance with fANOVA across all datasets of Table 2

For each of the 10 most important hyperparameters in our configuration space (of more than 200 dimensions), Figure 3 shows a box plot of the importance across our datasets. The most important hyperparameters are closely associated with the training procedure: 3 of them control the optimizer of the neural network and its learning rate. Additionally, 4 hyperparameters (*window_size*, *num_batches_per_epoch*, *batch_size*, *target_scaler*) contribute to the sampler and data preprocessing, showing the importance of the data fed to the network. Finally, the fact that two hyperparameters controlling the data distribution are amongst the most important ones indicates that identifying the correct potential data distribution might be beneficial to the performance of the model.

4.3 Ablation Study

In Section 3.5, we propose to partially evaluate the validation set on larger datasets to further accelerate the optimization process. To study the efficiency gain of this approach, we compare evaluation on the full validation set vs the proxy-evaluation on parts of the validation set. We ran this ablation study on the largest dataset, namely "Dominick" (115 704 series).

Figure 4 shows the results. It takes much less time for our optimizer (blue) to finish the first configuration evaluations on the lowest fidelity, improving effi-

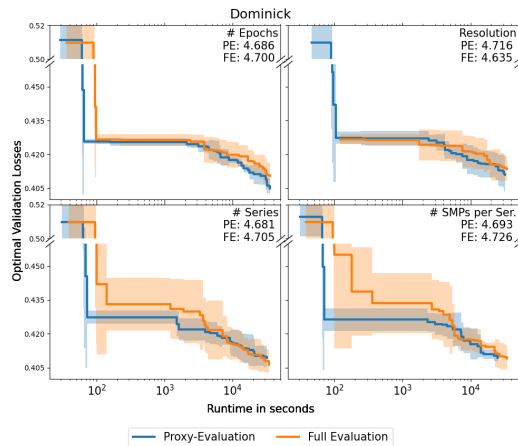


Fig. 4: Validation losses over time with different multi-fidelity approaches. We compute the area under the curve (AUC) of our approach (PE) and naive multi-fidelity optimizer (FE) and list them in the figures.

ciency early on and showing the need of efficient validation and not only training. We note that the final performance does not change substantially between the different methods. Overall, `Auto-PyTorch-TS` achieves the best any-time performance. We note that `Auto-PyTorch-TS` has not converged after 10h and will most likely achieve even better performance if provided with more compute resources. The results on the other datasets show a similar trend and can be found in the appendix.

5 Conclusion and Future Work

In this work, we introduced `Auto-PyTorch-TS`, an AutoDL framework for the joint optimization of architecture and hyperparameters of DL models for time series forecasting tasks. To this end, we propose a new flexible configuration space encompassing several state-of-the-art forecasting DL models by identifying key concepts in different model classes and combining them into a single framework.

Given the flexibility of our configuration space, new developers can easily adapt their architectures to our framework under the assumption that they can be formulated as an encoder-decoder-head architecture. Despite recent advances and competitive results, DL methods have until now not been considered the undisputed best approach in time series forecasting tasks: Traditional machine learning approaches and statistical methods have remained quite competitive [20, 40]. By conducting a large benchmark, we demonstrated, that our proposed `Auto-PyTorch-TS` framework is able to outperform current state-of-the-art methods on a variety of forecasting datasets from different domains and even improves over a theoretically optimal oracle comprised of the best possible baseline model for each dataset.

While we were able to show superior performance over existing methods, our results suggest, that a combination of DL approaches with traditional machine learning and statistical methods could further improve performance. The optimal setup for such a framework and how to best utilize these model classes side by side poses an interesting direction for further research. Our framework makes use of BO and utilizes multi-fidelity optimization in order to alleviate the costs incurred by the expensive training of DL models. Our experiments empirically demonstrate, that the choice of budget type can have an influence on the quality of the optimization and ultimately performance.

To the best of our knowledge there is currently no research concerning the choice of fidelity when utilizing multi-fidelity optimization for architecture search and HPO of DL models; not only for time series forecasting, but other tasks as well. This provides a great opportunity for future research and could further improve current state-of-the-art methods already utilizing multi-fidelity optimization. Additionally, we used our extensive experiments to examine the importance

of hyperparameters in our configuration space and were able to identify some of the critical choices for the configuration of DL architectures for time series forecasting. Finally, in contrast to previous AutoML systems, to the best of our knowledge, time series forecasting is the first task, where not only efficient training is important but also efficient validation. Although we showed empirical evidence for the problem and took a first step in the direction of efficient validation, it remains an open challenge for future work. Auto-PyTorch-TS can automatically optimize the hyperparameter configuration for a given task and can be viewed as a benchmark tool that isolates the influence of hyperparameter configurations of the model. This makes our framework an asset to the research community as it enables researchers to conveniently compare their methods to existing DL models.

Acknowledgements

Difan Deng and Marius Lindauer acknowledge financial support by the Federal Ministry for Economic Affairs and Energy of Germany in the project CoyPu under Grant No. 01MK21007L. Bernd Bischl acknowledges funding by the German Federal Ministry of Education and Research (BMBF) under Grant No. 01IS18036A. Florian Karl acknowledges support by the Bavarian Ministry of Economic Affairs, Regional Development and Energy through the Center for Analytics – Data – Applications (ADACenter) within the framework of BAY-ERN DIGITAL II (20-3410-2-9-8). Frank Hutter acknowledges support by European Research Council (ERC) Consolidator Grant “Deep Learning 2.0” (grant

no. 101045765).  Funded by
the European Union

References

1. Abdelfattah, M.S., Mehrotra, A., Dudziak, Ł., Lane, N.D.: Zero-cost proxies for lightweight NAS. In: ICLR (2021)
2. Alexandrov, A., et al.: Gluonts: Probabilistic and neural time series modeling in python. *Journal of Machine Learning Research* **21** (2020)
3. Assimakopoulos, V., Nikolopoulos, K.: The theta model: a decomposition approach to forecasting. *International journal of forecasting* **16**(4) (2000)
4. Bai, S., Kolter, J.Z., Koltun, V.: An empirical evaluation of generic convolutional and recurrent networks for sequence modeling. arXiv:1803.01271 (2018)
5. Beitner, J.: PyTorch Forecasting: Time series forecasting with PyTorch (2020)
6. Borovykh, A., Bohte, S., Oosterlee, C.W.: Conditional time series forecasting with convolutional neural networks. arXiv:1703.04691 (2017)
7. Box, G.E., Jenkins, G.M., Reinsel, G.C., Ljung, G.M.: *Time series analysis: forecasting and control* (2015)
8. Caruana, R., Niculescu-Mizil, A., Crew, G., Ksikes, A.: Ensemble selection from libraries of models. In: ICML. (2004)
9. Cho, K., et al.: Learning phrase representations using RNN encoder-decoder for statistical machine translation (2014)

10. Dahl, S.M.J.: TSPO: an autoML approach to time series forecasting. Ph.D. thesis, Universidade NOVA de Lisboa (2020)
11. De Livera, A.M., Hyndman, R.J., Snyder, R.D.: Forecasting time series with complex seasonal patterns using exponential smoothing. *Journal of the American statistical association* **106**(496) (2011)
12. Dong, X., Yang, Y.: NAS-Bench-201: Extending the scope of reproducible neural architecture search. In: *ICLR* (2020)
13. Dosovitskiy, A., et al: An image is worth 16x16 words: Transformers for image recognition at scale. In: *ICLR* (2021)
14. Elsken, T., Metzen, J.H., Hutter, F.: Neural architecture search. In: *Automatic Machine Learning: Methods, Systems, Challenges* (2019)
15. Erickson, N., et al.: Autogluon-tabular: Robust and accurate automl for structured data. [arXiv:2003.06505](https://arxiv.org/abs/2003.06505) (2020)
16. Falkner, S., Klein, A., Hutter, F.: BOHB: robust and efficient hyperparameter optimization at scale. In: *ICML* (2018)
17. Feurer, M., Hutter, F.: Hyperparameter optimization. In: *Automatic Machine Learning: Methods, Systems, Challenges* (2019)
18. Feurer, M., Klein, A., Eggensperger, K., Springenberg, J.T., Blum, M., Hutter, F.: Efficient and robust automated machine learning. In: *NeurIPS* (2015)
19. Flores, B.E.: A pragmatic view of accuracy measurement in forecasting. *Omega* (1986)
20. Godahewa, R., Bergmeir, C., Webb, G.I., Hyndman, R.J., Montero-Manso, P.: Monash time series forecasting archive. In: *NeurIPS Track on Datasets and Benchmarks* (2021)
21. Halvari, T., Nurminen, J.K., Mikkonen, T.: Robustness of automl for time series forecasting in sensor networks. In: *IFIP Networking Conference* (2021)
22. Hewamalage, H., Bergmeir, C., Bandara, K.: Recurrent neural networks for time series forecasting: Current status and future directions. *International Journal of Forecasting* (2021)
23. Hochreiter, S., Schmidhuber, J.: Long short-term memory. *Neural Comput.* (1997)
24. Hutter, F., Hoos, H., Leyton-Brown, K.: An efficient approach for assessing hyperparameter importance. In: *ICML* (2014)
25. Hutter, F., Hoos, H.H., Leyton-Brown, K.: Sequential model-based optimization for general algorithm configuration. In: *Learning and Intelligent Optimization* (2011)
26. Hyndman, R., Koehler, A.B., Ord, J.K., Snyder, R.D.: Forecasting with exponential smoothing: the state space approach (2008)
27. Hyndman, R.J., Athanasopoulos, G.: *Forecasting: principles and practice* (2021)
28. Hyndman, R.J., Khandakar, Y.: Automatic time series forecasting: the forecast package for r. *Journal of statistical software* **27** (2008)
29. Hyndman, R.J., Koehler, A.B.: Another look at measures of forecast accuracy. *International Journal of Forecasting* (2006)
30. Jamieson, K.G., Talwalkar, A.: Non-stochastic best arm identification and hyperparameter optimization. In: *AISTA* (2016)
31. Januschowski, T., et al.: Criteria for classifying forecasting methods. *International Journal of Forecasting* **36**(1) (2020)
32. Javeri, I.Y., Toutiaee, M., Arpinar, I.B., Miller, J.A., Miller, T.W.: Improving neural networks for time-series forecasting using data augmentation and automl. In: *BigDataService* (2021)
33. Jin, H., Song, Q., Hu, X.: Auto-keras: An efficient neural architecture search system. In: *SIGKDD* (2019)

34. Klein, A., Tiao, L., Lienart, T., Archambeau, C., Seeger, M.: Model-based asynchronous hyperparameter and neural architecture search (2020)
35. Kurian, J.J., Dix, M., Amihai, I., Ceusters, G., Prabhune, A.: BOAT: A bayesian optimization automl time-series framework for industrial applications. In: Big-DataService (2021)
36. Li, L., Jamieson, K.G., DeSalvo, G., Rostamizadeh, A., Talwalkar, A.: Hyperband: A novel bandit-based approach to hyperparameter optimization. *J. Mach. Learn. Res.* (2017)
37. Li, T., Zhang, J., Bao, K., Liang, Y., Li, Y., Zheng, Y.: Autost: Efficient neural architecture search for spatio-temporal prediction. In: SIGKDD (2020)
38. Lim, B., Arik, S.Ö., Loeff, N., Pfister, T.: Temporal fusion transformers for interpretable multi-horizon time series forecasting. *International Journal of Forecasting* (2021)
39. Lindauer, M., et al.: SMAC3: A versatile bayesian optimization package for hyperparameter optimization. *Journal of Machine Learning Research* (2022)
40. Makridakis, S., Spiliotis, E., Assimakopoulos, V.: The m4 competition: Results, findings, conclusion and way forward. *International Journal of Forecasting* (2018)
41. Makridakis, S., Spiliotis, E., Assimakopoulos, V.: The m4 competition: 100,000 time series and 61 forecasting methods. *International Journal of Forecasting* (2020)
42. Meisenbacher, S., et al: Review of automated time series forecasting pipelines. *arXiv: 2202.01712* (2022)
43. Montero-Manso, P., Athanasopoulos, G., Hyndman, R.J., Talagala, T.S.: Fforma: Feature-based forecast model averaging. *International Journal of Forecasting* **36**(1) (2020)
44. Olson, R.S., Bartley, N., Urbanowicz, R.J., Moore, J.H.: Evaluation of a tree-based pipeline optimization tool for automating data science. In: GECCO (2016)
45. van den Oord, A. et al: Wavenet: A generative model for raw audio. In: ISCA Speech Synthesis Workshop (2016)
46. Oreshkin, B.N., Carpov, D., Chapados, N., Bengio, Y.: N-BEATS: neural basis expansion analysis for interpretable time series forecasting. In: ICLR (2020)
47. Paldino, G.M., De Stefani, J., De Caro, F., Bontempi, G.: Does automl outperform naive forecasting? In: *Engineering Proceedings*. vol. 5 (2021)
48. Prokhorenkova, L., Gusev, G., Vorobev, A., Dorogush, A.V., Gulin, A.: Catboost: Unbiased boosting with categorical features. In: *NeurIPS* (2018)
49. Real, E., Aggarwal, A., Huang, Y., Le, Q.V.: Regularized evolution for image classifier architecture search. In: *AAAI* (2019)
50. Salinas, D., Flunkert, V., Gasthaus, J., Januschowski, T.: Deepar: Probabilistic forecasting with autoregressive recurrent networks. *International Journal of Forecasting* **36**(3) (2020)
51. Shah, S.Y., et al.: AutoAI-ts: Autoai for time series forecasting. In: *SIGMOD* (2021)
52. Talagala, T.S., Hyndman, R.J., Athanasopoulos, G., et al.: Meta-learning how to forecast time series. *Monash Econometrics and Business Statistics Working Papers* **6** (2018)
53. Thornton, C., Hutter, F., Hoos, H.H., Leyton-Brown, K.: Auto-WEKA: combined selection and hyperparameter optimization of classification algorithms. In: *SIGKDD* (2013)
54. Trapero, J.R., Kourentzes, N., Fildes, R.: On the identification of sales forecasting models in the presence of promotions. *Journal of the Operational Research Society* (2015)
55. van Kuppevelt, D., Meijer, C., Huber, F., van der Ploeg, A., Georgievskaya, S., van Hees, V.: Mcfly: Automated deep learning on time series. *SoftwareX* **12** (2020)

56. Vaswani, A., et al.: Attention is all you need. In: NeurIPS (2017)
57. Wen, R., Torkkola, K., Narayanaswamy, B., Madeka, D.: A multi-horizon quantile recurrent forecaster. 31st Conference on NeurIPS, Time Series Workshop (2017)
58. Wu, B., et al.: FBNetV5: Neural architecture search for multiple tasks in one run. arXiv:2111.10007 (2021)
59. Xiao, Y., Qiu, Y., Li, X.: A survey on one-shot neural architecture search. In: IOP Conference Series: Materials Science and Engineering. vol. 750. IOP Publishing (2020)
60. Ying, C., Klein, A., Christiansen, E., Real, E., Murphy, K., Hutter, F.: NAS-bench-101: Towards reproducible neural architecture search. In: ICML (2019)
61. Zela, A., Klein, A., Falkner, S., Hutter, F.: Towards automated deep learning: Efficient joint neural architecture and hyperparameter search. In: ICML 2018 AutoML Workshop (2018)
62. Zimmer, L., Lindauer, M., Hutter, F.: Auto-pytorch tabular: Multi-fidelity meta-learning for efficient and robust autodl. IEEE Transactions on Pattern Analysis and Machine Intelligence (2021)
63. Zoph, B., Vasudevan, V., Shlens, J., Le, Q.V.: Learning transferable architectures for scalable image recognition. In: CVPR (2018)