

# Deep Learning for Time Series Forecasting: Tutorial and Literature Survey

KONSTANTINOS BENIDIS\*, Amazon Research, Germany  
SYAMA SUNDAR RANGAPURAM, Amazon Research, Germany  
VALENTIN FLUNKERT, Amazon Research, Germany  
YUYANG WANG, Amazon Research, USA  
DANIELLE MADDIX, Amazon Research, USA  
CANER TURKMEN, Amazon Research, Germany  
JAN GASTHAUS, Amazon Research, Germany  
MICHAEL BOHLKE-SCHNEIDER, Amazon Research, Germany  
DAVID SALINAS, Amazon Research, Germany  
LORENZO STELLA, Amazon Research, Germany  
FRANÇOIS-XAVIER AUBET, Amazon Research, Germany  
LAURENT CALLOT, Amazon Research, Germany  
TIM JANUSCHOWSKI\*<sup>†</sup>, Zalando SE, Germany

Deep learning based forecasting methods have become the methods of choice in many applications of time series prediction or *forecasting* often outperforming other approaches. Consequently, over the last years, these methods are now ubiquitous in large-scale industrial forecasting applications and have consistently ranked among the best entries in forecasting competitions (e.g., M4 and M5). This practical success has further increased the academic interest to understand and improve deep forecasting methods. In this article we provide an introduction and overview of the field: We present important building blocks for deep forecasting in some depth; using these building blocks, we then survey the breadth of the recent deep forecasting literature.

\*Equal contribution.

<sup>†</sup>Work done while at AWS.

---

Authors' addresses: Konstantinos Benidis, kbenidis@amazon.com, Amazon Research, Charlottenstrasse 4, Berlin, Germany, 10969; Syama Sundar Rangapuram, rangapur@amazon.com, Amazon Research, Charlottenstrasse 4, Berlin, Germany, 10969; Valentin Flunkert, flunkert@amazon.com, Amazon Research, Charlottenstrasse 4, Berlin, Germany, 10969; Yuyang Wang, yuyawang@amazon.com, Amazon Research, 1900 University Ave., East Palo Alto, CA, USA, 94303; Danielle Maddix, dmmaddix@amazon.com, Amazon Research, 1900 University Ave., East Palo Alto, CA, USA, 94303; Caner Turkmen, atturkm@amazon.com, Amazon Research, Charlottenstrasse 4, Berlin, Germany, 10969; Jan Gasthaus, gasthaus@amazon.com, Amazon Research, Charlottenstrasse 4, Berlin, Germany, 10969; Michael Bohlke-Schneider, bohlkem@amazon.com, Amazon Research, Charlottenstrasse 4, Berlin, Germany, 10969; David Salinas, dsalina@amazon.com, Amazon Research, Charlottenstrasse 4, Berlin, Germany, 10969; Lorenzo Stella, stellalo@amazon.com, Amazon Research, Charlottenstrasse 4, Berlin, Germany, 10969; François-Xavier Aubet, aubetf@amazon.com, Amazon Research, Charlottenstrasse 4, Berlin, Germany, 10969; Laurent Callot, lcallot@amazon.com, Amazon Research, Charlottenstrasse 4, Berlin, Germany, 10969; Tim Januschowski, tim.januschowski@zalando.de, Zalando SE, Berlin, Germany.

---

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

© 2018 Association for Computing Machinery.

0004-5411/2018/8-ART111 \$15.00

<https://doi.org/10.1145/1122445.1122456>

50 CCS Concepts: • **General and reference** → **Surveys and overviews**; • **Computing methodologies**  
51 → **Probabilistic reasoning**; **Artificial intelligence**; • **Mathematics of computing** → **Probabilistic**  
52 **algorithms**.

53 Additional Key Words and Phrases: time series, forecasting, neural networks  
54

### 55 **ACM Reference Format:**

56 Konstantinos Benidis, Syama Sundar Rangapuram, Valentin Flunkert, Yuyang Wang, Danielle Maddix, Caner  
57 Turkmen, Jan Gasthaus, Michael Bohlke-Schneider, David Salinas, Lorenzo Stella, François-Xavier Aubet,  
58 Laurent Callot, and Tim Januschowski. 2018. Deep Learning for Time Series Forecasting: Tutorial and Literature  
59 Survey. *J. ACM* 37, 4, Article 111 (August 2018), 37 pages. <https://doi.org/10.1145/1122445.1122456>  
60

## 61 **1 INTRODUCTION**

62 Forecasting is the task of extrapolating time series into the future. It has many important appli-  
63 cations [54] such as forecasting the demand for items sold by retailers [15, 25, 41, 136, 156, 190],  
64 the flow of traffic [111, 118, 126], the demand and supply of energy [45, 117, 157, 170], or the  
65 covariance matrix, volatility and long-tail distributions in finance [12, 29, 30, 124, 197]. As such, it  
66 is a well-studied area (e.g., see [84] for an introduction) with its own dedicated research commu-  
67 nity. The machine learning, data science, systems, and operations research communities as well  
68 as application-specific research communities have also studied the problem intensively (e.g., see  
69 a series of recent tutorials [55–58]). In contrast to traditional forecasting applications, modern  
70 incarnations often exhibit large panels of related time series, all of which need to be forecasted  
71 simultaneously [91]. Although these problem characteristics make them amenable to deep learning  
72 or neural networks (NNs), as in many other domains over the course of history, NNs were not  
73 always a standard tool to tackle such problems. Indeed, their effectiveness has historically been  
74 regarded as mixed (e.g., [201]).  
75

76 The history of NNs starts in 1957 [152] and in 1964 for NNs in forecasting [83]. Since then, interest  
77 in NNs has oscillated, with upsurges in attention attributable to breakthroughs. The application of  
78 NNs in time series forecasting has followed the general popularity, typically with a lag of a few  
79 years. Examples of such breakthroughs include Rumelhart et al. [153, 154] that popularized the  
80 training of multilayer perceptrons (MLPs) using back-propagation. Significant advances were made  
81 subsequently such as the use of convolutional NNs (CNNs) [113], and Long Short Term Memory  
82 (LSTM) [81] cells that address the issue of recurrent NNs' (RNNs) training, just to name a few.  
83 Despite these advances, NNs remained hard to train and difficult to work with. Methods such as  
84 Support Vector Machines (SVMs) [26] and Random Forests [79] that were developed in the 1990s  
85 proved to be highly effective (LeCun et al. [115] found that SVMs were as good as the best designed  
86 NNs available at the time) and were supported by attractive theory. This shifted the interest of  
87 researchers away from NNs. Forecasting was no exception and results obtained with NNs were  
88 mostly mixed as reflected in a highly cited review [201]. The breakthrough that marked the dawn  
89 of the deep learning era came in 2006 when Hinton et al. [78] showed that it was possible to train  
90 NNs with a large number of layers (deep) if the weights are initialized appropriately. Accordingly,  
91 deep learning has had a sizable impact on forecasting [110] and NNs have long entered the canon  
92 of standard techniques for forecasting [84]. New models specifically designed for forecasting tasks  
93 have been proposed, taking advantage of deep learning to supercharge classical forecasting models  
94 or to develop entirely novel approaches. This recent burst of attention on *deep forecasting* models  
95 is the latest twist in a long and rich history.

96 Driven by the availability of (closed-source) large time series panels, the potential of deep  
97 forecasting models, i.e., forecasting models based on NNs, has been exploited primarily in applied  
98

99 industrial research divisions over the last years [64, 111, 156, 190].<sup>1</sup> With the overwhelming success  
 100 of deep forecasting methods in the M4 competition [169], this has convinced also formerly skeptical  
 101 academics [128, 129]. In the most recent M5 competition, deep forecasting methods were the second  
 102 and third placed solutions [130] although the competition was otherwise dominated by tree-based  
 103 forecasting methods such as LightGBM [99] and XGBoost [33], see e.g., [92]. Modern software  
 104 frameworks [1, 34, 143] have sped up the development of NN models and dedicated forecasting  
 105 packages available [4].

106 While the history of NNs for forecasting is rich, the focus of this article is on more recent  
 107 developments in NN for forecasting, roughly since the time that the term “deep learning” was coined.  
 108 As such, we do not attempt to give a complete historical overview and sacrifice comprehensiveness  
 109 for recency. The main objectives of this article are to educate on, review and popularize the recent  
 110 developments in forecasting driven by NNs for a general audience. Therefore, we place emphasis on  
 111 an educational aspect via a tutorial of deep forecasting in the first part (Section 2). In the second part,  
 112 Section 3, we provide an overview of the state-of-the-art of modern deep forecasting models. Our  
 113 exposition is driven by an attempt to identify the main building blocks of modern deep forecasting  
 114 models which hopefully enables the reader to digest the rapidly increasing literature more easily.  
 115 We do not attempt a taxonomy of all existing methods and our selection of the building blocks is  
 116 opinionated, motivated by our experience of innovating in this area with a strong focus on practical  
 117 applicability. Compared with other surveys [76, 120, 201], we provide a more comprehensive  
 118 overview with a particular focus on recent, advanced topics. Finally, in Section 4, we conclude and  
 119 speculate on potentially fruitful areas for future research.

## 120 121 2 DEEP FORECASTING: A TUTORIAL

122 In the following, we formalize the forecasting problem, summarize those advances in deep learning  
 123 that we deem as the most relevant for forecasting, expose important building blocks for NNs  
 124 and discuss archetypal models in detail. For general improvements that fueled the deep learning  
 125 renaissance, like weight initialization, optimization algorithms or general-purpose components such  
 126 as activation functions, we refer to standard textbooks like [70]. We are aware to be opinionated in  
 127 both the selection of topics as well as the style of exposition. We attempt to take a perspective akin  
 128 to a deep forecasting model builder who would compose a forecasting model out of several building  
 129 blocks such as NN architectures, input transformations and output representations. Although not  
 130 all models will fit perfectly into this exposition, it is our hope that this downside is outweighed by  
 131 the benefit of allowing the inclined reader to invent new models more easily.

### 132 133 2.1 Notation and Formalization of the Forecasting Problem

134 Matrices, vectors and scalars are denoted by uppercase bold, lowercase bold and lowercase normal  
 135 letters, i.e.,  $\mathbf{X}$ ,  $\mathbf{x}$  and  $x$ , respectively. Let  $\mathcal{Z} = \{z_{i,1:T_i}\}_{i=1}^N$  be a set of  $N$  univariate time series, where  
 136  $z_{i,1:T_i} = (z_{i,1}, \dots, z_{i,T_i})$ ,  $z_{i,t}$  is the value of the  $i$ -th time series at time  $t$  and  $\mathbf{Z}_{t_1:t_2}$  the values of all  
 137  $N$  time series at the time slice  $[t_1, t_2]$ . Typical examples for the domain of the time series values  
 138 include  $\mathbb{R}$ ,  $\mathbb{N}$ ,  $\mathbb{Z}$ ,  $[0, 1]$ . The set of time series is associated with a set of covariate vectors denoted by  
 139  $\mathcal{X} = \{\mathbf{X}_{i,1:T_i}\}_{i=1}^N$ , with  $\mathbf{x}_{i,t} \in \mathbb{R}^{d_x}$ . Note that each vector  $\mathbf{x}_{i,t}$  can include both time-varying or static  
 140 features. We denote by  $\alpha$  a general input in a model (that can be any combination of covariates and  
 141 lagged values of the target) and by  $\beta$  a general output. Since  $\alpha$  and  $\beta$  refer to a general case, we  
 142 always represent them with lowercase normal letters. We denote by  $\theta$  the parameters of a model  
 143

144  
 145 <sup>1</sup>Forecasting is an example of a sub-discipline in the machine learning community where the comparatively modest attention  
 146 it receives in published research is in stark contrast to a tremendous business impact.

(e.g., parameters of a distribution) and by  $\Phi$  the learnable free parameters of the underlying NN (e.g., the weights and biases).

In the most general form, the object of interest in forecasting is the conditional distribution

$$p(\mathbf{Z}_{t+1:t+h} | \mathbf{Z}_{1:t}, \mathbf{X}_{1:t+h}; \theta), \quad (1)$$

where  $\theta$  are the parameters of a (probabilistic) model. Eq. (1) is general in the sense that each  $\mathbf{z}_i \in \mathcal{Z}$  is multidimensional (the length of the time series),  $\mathcal{Z}$  is multivariate (the number of time series  $|\mathcal{Z}| = N > 1$ ) and the forecast is multi-step ( $h$  steps). Varying degrees of simplification of Eq. (1) are considered in the literature, for example by assuming factorizations of  $p$  and different ways of estimating  $\theta$ . In the following, we present the three archetypical models for addressing Eq. (1).

**Local univariate model:** A separate (*local*) model is trained independently for each of the  $N$  time series, modelling the predictive distribution

$$p(\mathbf{z}_{i,t+1:t+h} | \mathbf{z}_{i,1:t}, \mathbf{X}_{i,1:t+h}; \theta_i), \quad \theta_i = \Psi(\mathbf{z}_{i,1:t}, \mathbf{X}_{i,1:t+h}), \quad (2)$$

where  $\Psi$  is a generic function mapping input features to the parameters  $\theta_i$  of the probabilistic model that are local to the  $i$ -th time series. Note that one may use multidimensional covariates  $\mathbf{x}_{i,t}$  for each of the  $N$  models, but they are still solving a univariate problem, i.e., forecasting only one time series. The use of covariates common to all  $N$  models is possible but any pattern that is learned in one model is not used in another (unless provided explicitly which prohibits parallel training). Many classical approaches fall into this category and traditionally NNs were employed in this local fashion (e.g., [201]). Note that this approach is not suitable for cold start problems: i.e., forecasting a time series without historical values.

**Global univariate model:** A single, *global* model [90, 135] is trained using available data from all  $N$  time series. However, the model is still used to predict a univariate target. It does not produce joint forecasts of all time series but forecasts of any single time series at a time. This is also sometimes referred to as a cross-learning approach, e.g., [161]. In a more general form, global univariate models specialize Eq. (1) to

$$p(\mathbf{z}_{i,t+1:t+h} | \mathbf{Z}_{1:t}, \mathbf{X}_{1:t+h}; \theta_i), \quad \theta_i = \Psi(\mathbf{z}_{i,1:t}, \mathbf{X}_{i,1:t+h}, \Phi), \quad (3)$$

where  $\Phi$  are shared parameters among all  $N$  time series.

In this article,  $\Psi$  in global models is usually a NN and  $\mathbf{X}_i$  include item-specific features to allow the model to distinguish between the time series. Although the parameters  $\theta_i$  of the probabilistic model for each time series are different, they are still predicted using shared parameters (or weights)  $\Phi$  in  $\Psi$ . This allows for efficient learning since the model pools information from all time series and in particular improves inference for shorter time series compared to local univariate models. Such a model is expected to learn some advanced features (“embeddings”) exploiting information across time series. Once these advanced features are learned via  $\Psi$ , the global model is then used to forecast each time series *independently*. That is, although during training the model sees all the related time series together, the prediction is done by looking at each time series individually. Note that the embeddings learned in the global model are useful beyond the  $N$  time series used in the training. This addresses the cold start problem in the sense that the global model can be used to provide forecasts for time series without historical values. Global models are also referred to as cross-learning or panel models in econometrics and statistics and have been the subject of considerable study, e.g., via dynamic factor models [66].

**Multivariate model:** Here, a single model is learned for all  $N$  time series using all available data, directly predicting the *multivariate* target:

$$p(\mathbf{Z}_{t+1:t+h} | \mathbf{Z}_{1:t}, \mathbf{X}_{1:t+h}; \theta), \quad \theta = \Psi(\mathbf{Z}_{1:t}, \mathbf{X}_{1:t+h}, \Phi). \quad (4)$$

Table 1. Summary of deep forecasting models based on forecast and model type. For one-step and multi-step forecasting models  $h = 1$  and  $h > 1$ , respectively.

Forecast type	Model type	Formulation
Point	Local univariate	$\hat{z}_{i,t+1:t+h} = \Psi(z_{i,1:t}, X_{i,1:t+h})$
	Global univariate	$\hat{z}_{i,t+1:t+h} = \Psi(z_{i,1:t}, X_{i,1:t+h}, \Phi)$
	Multivariate	$\hat{Z}_{t+1:t+h} = \Psi(Z_{1:t}, X_{1:t+h}, \Phi)$
Probabilistic	Local univariate	$P(z_{i,t+1:t+h}   z_{i,1:t}, X_{i,1:t+h}; \theta_i), \quad \theta_i = \Psi(z_{i,1:t}, X_{i,1:t+h})$
	Global univariate	$P(z_{i,t+1:t+h}   Z_{1:t}, X_{1:t+h}; \theta_i), \quad \theta_i = \Psi(z_{i,1:t}, X_{i,1:t+h}, \Phi)$
	Multivariate	$P(Z_{t+1:t+h}   Z_{1:t}, X_{1:t+h}; \theta), \quad \theta = \Psi(Z_{1:t}, X_{1:t+h}, \Phi)$

Note that the model also learns the dependency structure among the time series. Technically speaking, Eq. (4) is a global multivariate model and a further distinction from local multivariate models, such as VARMA [125], is possible.

*Remarks.* Note that in Eq. (1) and in the following model-specific cases we have chosen the multi-step ahead predictive distribution. We can always obtain a multi-step predictive distribution via a rolling one-step predictive distribution. In our discussion so far, we presented *probabilistic forecast* models that learn the entire distribution of the future values. However, it may be desirable to model specific values such as the mean, median or some other quantile, instead of the whole probability distribution. These are called *point-forecast* models and the optimal choice of the summary statistics to turn a probabilistic forecast into a point forecast depends on the metric used to judge the quality of the point forecast [104]. More concretely, a point-forecast global univariate model learns a quantity  $\hat{z}_{i,t+1:t+h} = \Psi(z_{i,1:t}, X_{i,1:t+h}, \Phi)$ , where  $\hat{z}_{i,t+1:t+h}$  is some point estimate of the future values of the time series. Table 1 summarizes the various modelling options based on the forecast and model types.

## 2.2 Neural Network Architectures

NNs are compositions of *differentiable* functions formed from simple building blocks to learn an approximation of some unknown function from data. An NN is commonly represented as a directed acyclic graph consisting of nodes and edges. The edges between the nodes contain weights (also called parameters) that are learned from the data. The basic unit of every NN is a neuron (illustrated in Fig. 1a), consisting of an input, an affine transformation with learnable weights and (optionally) a nonlinear activation function. Different types of NNs arrange these components in different ways. We refer to other reviews [120] for more details on the main architectures. Here, we only offer a high-level summary for completeness, focusing instead on forecasting specific ingredients for NNs such as input processing and loss functions.

**2.2.1 Multilayer perceptron.** In multilayer perceptrons (MLPs) or synonymously feedforward NNs, layers of neurons are stacked on top of each other to learn more complex nonlinear representations of the data. An MLP consists of an input and an output layer, while the intermediate layers are called *hidden*. The nodes in each layer of the network are fully connected to all the nodes in the previous layer. The output of the last hidden layer can be seen as some nonlinear feature representation (also called an *embedding*) obtained from the inputs of the network. The output layer then learns a mapping from these nonlinear features to the actual target. Learning with MLPs, and more generally with NNs, can be thought of as the process of learning a nonlinear feature map of the

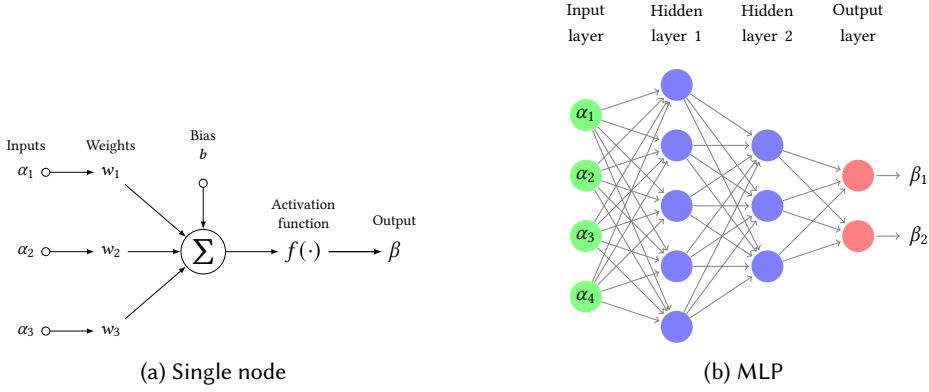


Fig. 1. (a) Structure of a single node or neuron. An affine transformation is applied to the input followed by an activation function, i.e.,  $\beta = f(\sum \alpha_i w_i + b)$ . The weights and bias parameters are learned during training. (b) Illustration of the MLP structure. Each circle in the hidden and output layers is a node, i.e., it applies an affine transformation followed by a nonlinear activation to the set of its inputs.

inputs and the relationship between this feature map and the actual target. Figure 1b illustrates the structure of an MLP with two hidden layers. Modern incarnations of the MLP have added important details to alleviate problems like vanishing gradients [80]. For example, ResNet [75], contains direct connections between hidden layers  $\ell - 1$  and  $\ell + 1$ , skipping over the hidden layer  $\ell$ .

One of the main limitations of MLPs is that they do not exploit the structure often present in the data in applications such as computer vision, natural language processing and forecasting. Moreover, the number of inputs and outputs is fixed making them inapplicable to problems with varying input and output sizes as in forecasting. Next, we discuss more complex architectures that overcome these limitations, for which MLPs are often used as the basic building blocks.

**2.2.2 Convolutional neural networks.** Convolutional neural networks (CNNs) [112] are a special class of NNs that are designed for applications where inputs have a known ordinal structure such as images and time series [70]. CNNs are *locally* connected NNs that use *convolutional* layers to exploit the structure present in the input data by applying a convolution function to smaller neighborhoods of the input data. Convolution here refers to the process of computing moving weighted sums by sliding the so-called *filter* or *kernel* over different parts of the input data. The size of the filter as well as how the filter is slid across the input are part of the hyperparameters of the model. A nonlinear activation, typically ReLU [68], is then applied to the output of the convolution operation.

In addition to convolutional layers, CNNs also typically use a *pooling* layer to reduce the size of the feature representation as well as to make the features extracted from the convolutional layer more robust. For example, a commonly used max-pooling layer, which is applied to the output of convolutional layers, extracts the maximum value of the features in a given neighborhood. Similarly to the convolution operation, the pooling operation is applied to smaller neighborhoods by sliding the corresponding filter over the input. A pooling layer, however, does not have any learnable weights and hence both the convolution and the pooling layer are counted as one layer in CNNs.

Of particular importance for forecasting are the so-called *causal* convolutions, defined as

$$h_j = \sum_{d \in \mathcal{D}} w_d \alpha_{j-d},$$




Fig. 2. (a) Structure of a CNN consisting of a stack of three causal convolution layer. The input layer (green) is non-dilated and the other two are dilated. (b) Structure of an unrolled RNN. At each timestep  $t$  the network receives an external input  $\alpha_t$  and the output of the hidden units from the previous time step  $h_{t-1}$ . The hidden units all share the same weights. The internal state of the network is updated to  $h_t$  that is going to play the role of the previous state in the next timestep  $t + 1$ . Finally, the network outputs  $\beta_t$  which is a function of  $\alpha_t$  and  $h_t$ .

where  $h_j$  is the output of a hidden node,  $\alpha$  denotes the input,  $\mathcal{D} = \{1, \dots, n\}$  for some  $n$ ,  $|\mathcal{D}|$  is the *width* of the causal convolution (or also called the receptive field) and  $\mathbf{w}$  are the learnable parameters. In other words, causal convolutions are weighted moving averages which only take inputs into account which are before  $j$  hence the reference to causality in its name. A variation are *dilated* causal convolutions where we vary the index set  $\mathcal{D}$ , e.g., such that it does not necessarily contain consecutive values, but only every  $k$ -th value. Typically, these dilated causal convolutions are stacked on top of each other where the output of one layer of dilated causal convolutions is the input of another layer of causal convolution and the dilation grows by the depth of the NN. Figure 2a illustrates the general structure of a CNN with dilated causal convolutions.

**2.2.3 Recurrent neural networks.** Recurrent neural networks (RNNs) are NNs specifically designed to handle sequential data that arise in applications such as time series, natural language processing and speech recognition. The core idea consists of connecting *recurrently* the NNs' hidden units back to themselves with a time delay [94, 95]. Since hidden units learn some kind of feature representations of the raw input, feeding them back to themselves can be interpreted as providing the network with a dynamic memory. One crucial detail here is that the same network is used for all timesteps, i.e., the weights of the network are shared across timesteps. This weight-sharing idea is similar to that in CNNs where the same filter is used across different parts of the input. This allows the RNNs to handle sequences of varying length during training and, more importantly, generalize to sequence lengths not seen during training. Figure 2b illustrates the general structure of an (unrolled) RNN.

Although RNNs have been widely used in practice, training them is difficult given that they are typically applied to long sequences of data. A common issue while training very deep NNs by gradient-based methods using back-propagation is that of vanishing or exploding gradients [142] which renders learning challenging. Hochreiter and Schmidhuber [81] proposed Long short-term memory networks (LSTM) to address this problem. Similar to Resnet, via the skip-connections, LSTMs (and a simplified version Gated recurrent units (GRU) [36]) always offer a path where the gradient does not vanish or explode.

**2.2.4 Transformer.** A more recent architecture is based on the *attention* mechanism which has received increased interest in other sequence learning tasks [37, 38, 116, 184] for its ability to improve on long sequence prediction tasks over RNNs. One natural way to address this issue is to

learn more than one feature representation (contrary to RNNs), e.g., one for each time step of the input sequence and decide which of these representations are useful to predict the current element of the target sequence. Bahdanau et al. [10] suggest using a weighted sum of the representations where the weights are jointly learned along with the feature representation learning and the prediction. Note that at each time step in the prediction, one needs to learn a separate set of weights for the representations. This is essentially training the predictor to learn to which parts of the input sequence it should pay *attention* to produce a prediction. This attention mechanism has been shown to be instrumental for the state of the art in speech recognition and machine translations tasks [37, 38]. Inspired by the success of attention models, Vaswani et al. [184] developed the so-called *Transformer* model and showed that attention alone is sufficient, thus making the training amenable for parallelization and large number of parameters [28, 44]. In the literature, the term Transformer can refer to both the specific model and to the overall architecture as well.

### 2.3 Input Transformations

The careful handling of the input (parameters  $\alpha_t$  in Fig. 1 and 2) is a practically important ingredient for deep learning models in general and deep forecasting models in particular. Deep forecasting models are most commonly deployed as so-called global models (see Section 2.1), which means that the weights of the NN are trained across the panel of time series. Hence, it is important that the scale of the input is comparable. Standard techniques such as mean-variance scaling carry over to the forecasting setting. In practice, it is important to avoid leakage of future values in normalization schemes, so that mean and variance are taken over past windows (similar to causal convolutions).

Traditionally, the forecasting literature has used transformations such as the Box-Cox, i.e.,

$$h = \frac{z^\lambda - 1}{\lambda}, \quad (5)$$

where  $z$  is the input of the transformation,  $h$  is the output and  $\lambda$  is a free parameter. Box-Cox is a popular heuristic to have the input data more closely resemble the Gaussian distribution. A Box-Cox transformation can be readily integrated into an NN, with the free parameter  $\lambda$  optimized as part of the training process jointly with the other parameters of the network. More sophisticated approaches based on probability integral transformation (PIT) or Copulas are similarly possible, see e.g., [88] (and references therein) for a recent example.

A further standard technique is the discretization of input into categorical values or *bins*, for example by choosing the number and borders of bins such that each bins contains equal mass, see e.g., [145] for an example in forecasting.

We note that any input transformation must be reversed also to obtain values in the actual domain of interest. It is a choice for the modeller where/when to apply this reversal. Two extreme choices are to have transformation of the input and output fully outside the NN or have the input transformations as part of the NN and hence be subjected to learning.

### 2.4 Output Models and Loss Functions

Similar to the input, the output ( $\beta_t$  in Fig. 1 and 2) deserve a special discussion. Closely related is the question on the choice of loss function which we use to train a NN. The simplest form of an output is a single value, also referred to as a *point* forecast. For this case, the output  $\hat{z}_{i,t}$  is the best (w.r.t. the chosen loss function) estimate for the true value  $z_{i,t}$ . Standard regression loss functions (like  $\ell_p$  losses with their regularized modifications) can be used or more sophistication accuracy metrics specifically geared towards forecasting such as the MASE, sMAPE or others [85].

As remarked in Section 2.1, a point estimate  $\hat{z}_{i,t}$  can be seen as a particular realization from a probabilistic estimate of  $p(z_{i,t})$ . Depending on the accuracy metric used in forecasting, a different






Fig. 3. For a Gaussian distribution, its density function  $f$  is on the left-hand panel, the corresponding cumulative density function  $F$  (the primitive integral of  $f$ ) in the central panel and the quantile function  $F^{-1}$  on the right-hand panel.

realization may be appropriate [104]. So, even for obtaining point forecasts, *probabilistic* forecasts are important. More importantly, forecasts are often used in downstream optimization problem where some form of *expected* cost is to be minimized and for this, an estimate of the entire probability distribution is required. The probability distribution can be represented equivalently by its probability density function (PDF), the cumulative density function (CDF) or its inverse, the quantile function. Fig. 3 contains a visualization of the different representations for the Gaussian distribution. Across the deep forecasting landscape, most approaches (e.g., [64, 146, 148, 151, 156]), have chosen the PDF and quantile function to represent  $p(z_{i,t})$  and we will discuss general recipes next. Since the CDF has typically not been chosen to represent  $p(z_{i,t})$ , we do not discuss it further.

**2.4.1 PDF.** Arguably the most common way to represent a probability distribution in forecasting is via its PDF. The literature contains examples of using the standard parametric distribution families to represent probabilistic forecasts. For example, the output layer of an NN may produce the mean and variance parameter of a Gaussian distribution. So, the parameter  $\beta_t$  in Fig. 1 and 2 is a two-dimensional vector corresponding to  $\mu_t$  and  $\sigma_t$  of a Gaussian distribution. We typically achieve  $\sigma_t \geq 0$  by mapping the corresponding parameter through a softplus function. For the loss function, a natural choice is the negative log-likelihood (NLL) since a PDF allows to readily compute the likelihood of a point under it.

Beyond Gaussian likelihood, a number of differentiable parametric distributions have been used in the literature depending on the nature of the forecasting problem, e.g., the student-t distribution or the Tweedie distribution for continuous data, the negative binomial distribution for count data and more flexible approaches via mixtures of Gaussian. Although forecasting is most commonly done for domains of numerical values (i.e., we assume  $z_{i,t}$  to be in  $\mathbb{R}$  or  $\mathbb{N}$ ), other distributions such as the multinomial have also been employed successfully in forecasting even though they have no notion of the order on the domain [145]. The deployment of a multinomial distribution requires a binning of the input values (see Section 2.3). An alternative approach is to cut the output space in bins and treat each of them as a uniform distribution, while modelling the tails with a parametric distribution [50], this results in a piecewise linear CDF.

**2.4.2 Quantile function.** Another representation of  $p(z_{i,t})$  is via the quantile function which has a particular importance for forecasting. Often, a particular quantile is of practical interest. For example, in a simplified supply chain scenario for inventory control, there is a direct correspondence between the chosen quantile and a safety stock level in the newsvendor problem [54].




Fig. 4. An illustration how a quantile function parametrized by linear splines (left panel) corresponds to a piece-wise linear CDF (middle) which in turn corresponds to a piece-wise constant PDF as assumed in an adaptive binning strategy (right panel).

So naturally, estimating the quantiles directly via quantile regression approaches [103] is a common choice in forecasting either via choosing a single quantile (in a point-forecasting approach) or multiple quantiles simultaneously [51, 190]. Essentially, this discretizes the quantile function and estimates specific points only. A common choice for the loss function is the quantile loss or pinball loss. For the  $q$ -th quantile and  $F^{-1}$  the quantile function, the quantile loss is defined as

$$QS_q(\hat{F}_{i,t}^{-1}(q), z_{i,t}) := 2 \left( \mathbb{1}_{\{z_{i,t} \leq \hat{F}_{i,t}^{-1}(q)\}} - q \right) \left( \hat{F}_{i,t}^{-1}(q) - z_{i,t} \right), \quad (6)$$

where  $\mathbb{1}_{\{cond\}}$  is the indicator function that is equal to 1 if *cond* is true and 0 otherwise. The output of the NN is  $\hat{F}_{i,t}^{-1}(q)$ , i.e., the estimated value of the  $q$ -th quantile. For  $q = 0.5$  this reduces to the median of the forecast distribution and is a common choice of point forecasts.

As an alternative to a quantile regression approach, we can make a parametric assumption on the quantile function and estimate it directly. The main requirements for modelling a quantile function are that its domain should be constrained to  $[0, 1]$  and the function should be monotonically increasing. This can be achieved easily via linear splines for example, so the output of the NN's last layers are the corresponding free parameters. For the loss function, a rich theory around the continuous ranked probability score (CRPS) exists [69, 132] and CRPS can be used as a loss function directly. CRPS can be defined [108] to summarize all possible quantile losses as

$$CRPS(\hat{F}_{i,t}, z_{i,t}) := \int_0^1 QS_q(\hat{F}_{i,t}^{-1}(q), z_{i,t}) dq. \quad (7)$$

Multivariate extensions such as the energy score [69] exist.

Interestingly, a popular discretization strategy, adaptive binning, used with multinomial distributions corresponds to quantile functions parametrized by piece-wise linear splines, see Fig. 4.

**2.4.3 Further approaches.** The recent deep learning literature contains more advanced examples for density estimation, most prominently via Generalized Adversarial Networks (GANs). We discuss them in Section 3.8 and discuss *normalizing flows* here which have arguably resonated more strongly in forecasting. Normalizing flows are invertible NNs that transform a simple distribution to a more complex output distribution. Invertibility guarantees the conservation of probability mass and allows the evaluation of the associated density function everywhere. The key observation is that the probability density of an observation  $z_{i,t}$  can be computed using the change of variables formula:

$$p(z_{i,t}) = p_{y_{i,t}}(f^{-1}(z_{i,t})) |\det[\text{Jac} f_{i,t}^{-1} z_{i,t}]|, \quad (8)$$




Fig. 5. Canonical versus sequence-to-sequence models.

where the first term  $p_{y_{i,t}}(f^{-1}(z_{i,t}))$  is the (in general simple) density of a variable  $y_{i,t}$ , and the second is the absolute value of the determinant of the Jacobian of  $f_{i,t}^{-1}$ , evaluated at  $z_{i,t}$ .

The invertible function  $f$  is typically parametrized by an NN. A particular instantiation is the Box-Cox transformation, Eq. (5). The field of normalizing flows (e.g., [46, 101, 137]) studies invertible NNs that typically transform isotropic Gaussians to more complex data distributions. The choice of a particular instantiation of  $f$  can facilitate the computation of the likelihood of a given point when the NLL is amenable as a loss function. Alternatively, generating samples may be computationally more viable for other instantiations (this is typically the cases with generative adversarial networks as well). In this case, the NLL can be replaced by other loss functions such as CRPS.

A number of extensions are possible. For example, more complex models for  $p(z_{i,t})$  are possible such as Hidden Markov Models or Linear Dynamical Systems. NNs can output the free parameters of these models but then need to be combined with the learning and inference schemes associated with these models, such as Kalman Filtering/Smoothing in the case of Linear Dynamical System [42, 146] or the Forward/Backward Algorithm in the case of Hidden Markov Models [5]. Another avenue is to relax constraints on the representation of  $p(z_{i,t})$  to obtain closely related objects with more favorable computational properties. For example, *energy based models* (EBMs) approximate the unnormalized log-probability [77, 114]. EBMs perform well in learning high dimensional distributions at the cost of being difficult to train [174] and have been employed in forecasting [150].

## 2.5 Archetypical Architectures

With all key components in place, in this section we present in more details popular forecasting architectures. In particular we focus on the widely-used RNN-based architecture that takes as input its previous hidden state, the currently available information and produces an one-step ahead estimate of the target time series. There are subtle details on how to handle a multi-step unrolled model during training (e.g., [109]), which we will skip over. We further examine the sequence-to-sequence (seq2seq) modelling approach where the model takes an *encoding sequence* as input and maps it to a *decoding sequence* (of predetermined length) on which the loss is computed against the actual values  $z$  during training. A typical instance in the training set in this approach consists of the target and covariate values up to a certain point in time  $t$  as the encoding sequence and the outputs of the NN are a predetermined number of target values after time  $t$ . Figure 5 contrasts both approaches. In the following we present two popular deep forecasting models, DeepAR and MQRNN/MQCNN, in some details to illustrate the core concepts. They represent the one-step-ahead RNN-based and seq2seq approach, respectively.

2.5.1 *DeepAR*. Among the first of the modern deep forecasting models is DeepAR [156], a global univariate model (see Table 1) that consists of an RNN backbone (typically an LSTM).<sup>2</sup> The input of the model is a combination of lagged target values and relevant covariates. The output is either a point forecast with a standard loss function or, in the basic variant, a probabilistic forecast via the parameters of a PDF (e.g.,  $\mu$  and  $\sigma$  of a Gaussian distribution), where the loss function is then the NLL. The output modelling of DeepAR has been the subject of follow-up work, e.g., Jeon and Seong [93] propose a Tweedie loss, Mukherjee et al. [136] propose a mixture of Gaussians as the distribution and domain specific feature processing blocks. Figure 6 summarizes the architecture. The dotted arrows in the picture correspond to drawing a sample that can be used as alternative input (as a lagged target) during training (even though  $\alpha_t$  may be available or in the case where an  $\alpha_t$  is missing) and during prediction to obtain multi-step ahead forecasts.

It is also possible to change the output of DeepAR to model the quantile function and use CRPS, Eq. (7), as the loss function [64, 72, 96]. While this in general computationally challenging, special cases are amendable for practical computation. For example, we can assume a parametrization of the quantile function by linear isotonic regression splines:

$$s(q; \gamma, b, d) = \gamma + \sum_{\ell=0}^L b_{\ell} (q - d_{\ell})_+ \quad (9)$$

where  $q \in [0, 1]$  is the quantile level,  $\gamma \in \mathbb{R}$  is the intercept term,  $b \in \mathbb{R}^{L+1}$  are weights describing the slopes of the function pieces,  $d \in \mathbb{R}^{L+1}$  is a vector of knot positions,  $L$  the number of pieces of the spline and  $(x)_+ = \max(x, 0)$  is the ReLU function. In order for  $s(\cdot)$  to represent a quantile function we need to guarantee its monotonicity and restrict its domain to  $[0, 1]$ . Both of these constraints can readily be achieved using standard NN tooling using a reparametrization of Eq. (9), while CRPS can be solved in closed form for linear splines (see [64]).

*Bag of tricks*. While the general setup of DeepAR is straightforward, a number of algorithmic optimizations turn it into a robust, general-purpose forecasting model. The handling of missing values via sample replacement from the probability distribution is one such example. Another one is oversampling of “important” training examples during training, where importance typically corresponds to time series with larger absolute values. Adding lagged values further help improve predictive accuracy. Lags can be chosen heuristically based on the frequency of the time series. For example, in a time series with daily frequency, a lag of 7 days often helps. Similarly, covariates corresponding to calendar events (e.g., indicator variables for weekends or holidays) can help further.

2.5.2 *MQRNN/MQCNN*. As an example for another type of deep forecasting model, we discuss the multi-horizon quantile recurrent forecaster (MQRNN) [190] next which was conceived concurrently to DeepAR. Contrary to DeepAR, it is most naturally deployed as a discriminative, seq2seq model in a quantile regression setting. For each time point  $t$  in the forecast horizon, MQRNN outputs a




Fig. 6. DeepAR: The model outputs parameters of a previously chosen family of distributions. Samples from this distribution can be fed back into the model during prediction (dotted lines) or in case of  $\alpha_t$  being missing.

<sup>2</sup>Hewamalage et al. [76] provide an overview specifically targeted at RNNs for forecasting.

589 chosen number of estimates for corresponding quantiles and the loss function in MQRNN is Eq. (6),  
590 i.e., the pinball loss summed over all quantiles and time points.

591 While MQRNN can use multiple configurations, often a CNN-based architecture is chosen in  
592 practice in the encoder (MQCNN) for computational efficiency reasons over RNN-based methods  
593 and two MLPs in the decoder. The first MLP captures all inputs during the forecast horizon and the  
594 context provided by the encoder. A second, local MLP applies only to specific horizons for which it  
595 uses the corresponding available input and the output of the MLP. A further innovation provided  
596 by MQCNN is the training scheme via the so-called *forking sequences* where the model forecasts  
597 by placing a series of decoders with shared parameters at each timestep in the encoder. Thus, the  
598 model can structurally forecast at each timestep, while the optimization process is stabilized by  
599 updating the gradients from the sequences together. An additional component of MQRNN is a local  
600 MLP component that aims to model spikes and events specifically.

### 602 3 LITERATURE REVIEW

603 In the prior section, we provided an in-depth introduction to selected, basic topics. Building on  
604 these topics, we survey the literature on modern deep forecasting models more broadly in this  
605 section. Given the breadth of the literature available, our selection is necessarily subjective.

606 We proceed as follows. In Section 3.1 we present probabilistic forecasting models, both one-step  
607 and multi-step. Similarly, in Section 3.2 we summarize point forecast models. We remark that, after  
608 Section 2, we have recipes at hand to turn an one-step ahead forecasting model into a multi-step  
609 forecasting model and a point forecasting model into a probabilistic model. We discuss hybrids  
610 of deep learning with state space models in Section 3.3, multivariate forecasting in Section 3.4,  
611 physics-based model in Section 3.5, global-local models in Section 3.6, models for intermittent time  
612 series in Section 3.7 and generative adversarial networks for forecasting in Section 3.8. We close  
613 this section with an overview of the large number of available models in Section 3.9 where we also  
614 provide guidelines on where to start the journey with deep forecasting models.

#### 616 3.1 Probabilistic Forecast Models

617  
618 3.1.1 *One-step forecast.* The DeepAR model presented in Sec. 2.5.1, is an example of one-step  
619 canonical forecasting model. In its base variant, DeepAR is a global univariate model which learns  
620 a univariate distribution; we discuss multivariate extensions in Sec. 3.4. DeepAR can be equipped  
621 with outputs representing a parametrized PDF including Gaussian Mixture Distributions Mukherjee  
622 et al. [136] or quantile functions Gasthaus et al. [64].

623 Rasul et al. [150] propose TimeGrad which, like DeepAR, is an RNN model using LSTM or GRU  
624 cells for which samples are drawn from the data distribution at each time step, with the difference  
625 that in TimeGrad the RNN conditions a diffusion probabilistic model [172] which allows the model  
626 to easily scale to multivariate time series and accurately use the dependencies between dimensions.  
627 Replacing the RNN-backbone of DeepAR with dilated causal convolutions has been proposed as  
628 both point and probabilistic forecasting models [4, 20, 183].

629  
630 3.1.2 *Multi-step forecast.* Contrary to the some of the models in Section 3.1.1 which produce  
631 one-step ahead forecasts, multi-step forecasts can be obtained directly with a seq2seq architecture.  
632 In Section 2.5.2, we reviewed the MQRNN/MQCNN architecture [190] as a seq2seq architecture for  
633 probabilistic forecasting. The main advantage of seq2seq over one-step ahead forecast models is that  
634 the decoder architecture can be chosen to output all future target values at once. This removes the  
635 need to unroll over the forecast horizon which can lead to *error accumulation* since early forecast  
636

638 errors propagate through the forecast horizon. Thus, the decoder of seq2seq forecasting models is  
639 typically an MLP while other architectures are also used for the encoder [138, 190].

640 Wen and Torkkola [189] extended the MQCNN model with a generative quantile copula. This  
641 model learns the conditional quantile function that maps the quantile index, which is a uniform  
642 random variable conditioned on the covariates, to the target. During training, the model draws the  
643 quantile index from a uniform distribution. This turns MQCNN into a generative, marginal quantile  
644 model. The authors combine this approach with a Gaussian copula to draw correlated marginal  
645 quantile index random values. They show that the Gaussian copula component improves the  
646 forecast at the distribution tails. Chen et al. [35] proposed DeepTCN, another seq2seq model where  
647 the encoder is the dilated causal convolution with residual blocks, and the decoder is simply an  
648 MLP with residual connections. Structure-wise, DeepTCN is almost the same as the basic structure  
649 of MQCNN [190], i.e., without the local MLP component that aims to model spikes and events.

650 Park et al. [141] propose the incremental quantile functions (IQF), a flexible and efficient  
651 distribution-free quantile estimation framework that resolves quantile crossing with a simple  
652 NN layer. A seq2seq encoder-decoder structure is used although the method can be readily applied  
653 to recurrent models with one-step ahead forecasts [156]. IQF is trained using the CRPS loss (Eq. (7))  
654 similar to [64].

655 A combination of recurrent and encoder-decoder structures has also been explored. In [205], the  
656 authors use an LSTM with Monte Carlo dropout as both the encoder and decoder. However, unlike  
657 other models that directly use RNNs to generate forecasts, the learned embedding at the end of the  
658 decoding step is fed into an MLP prediction network and is combined with other external features  
659 to generate the forecast. Along a similar line, Laptev et al. [111] employ an LSTM as a feature  
660 extractor (LSTM autoencoder), and use the extracted features, combined with external inputs to  
661 generate the forecasts with another LSTM.

662 Van Den Oord et al. [183] introduced the WaveNet architecture, a generative model for speech  
663 synthesis, which uses dilated causal convolutions to learn the long range dependencies important  
664 for audio signals. Since this architecture is based on convolutions, training is very efficient on GPUs  
665 – prediction is still sequential and further changes are necessary for fast inference. Adaptations of  
666 WaveNet for forecasting are available [4].

667

## 668 3.2 Point Forecast Models

669 Point forecast models do not model the probability distribution of the future values of a time series  
670 but rather output directly a point forecast that typically corresponds to a summary statistic of the  
671 predictive distribution. We have discussed generic recipes on how to turn a point forecasting model  
672 into a probabilistic forecasting model in Section 2.4 and the literature contains further examples  
673 (see e.g., [74, 175] for recent complementary approaches).

674

675 *3.2.1 One-step forecast.* A considerable amount of attention of the community is dedicated to  
676 one-step forecasting. LSTNet [107] is a model using a combination of a CNN and an RNN. Targeting  
677 multivariate time series, LSTNet uses a convolution network (without pooling) to extract short-term  
678 temporal patterns as well as correlations among variables. The output of the convolution network  
679 is fed into a recurrent layer and a temporal attention layer which, combined with the autoregressive  
680 component, generates the final point forecast. While LSTNet uses a standard point forecast loss  
681 function, it can readily be turned into a probabilistic forecast model using the components described  
682 in Sec. 2, e.g., by modifying LSTNet to output the parameters of a probability distribution and using  
683 NLL as a loss function. Qiu et al. [144] proposed an ensemble of deep belief networks for forecasting.  
684 The outputs of all the networks is concatenated and fed into a support vector regression model  
685 (SVR) that gives the final prediction. The NNs and the SVR are not trained jointly though. Hsu [82]

686

687 proposed an augmented LSTM model which combines autoencoders with LSTM cells. The input  
688 observations are first encoded to latent variables, which is equivalent to feature extraction, and are  
689 fed into the LSTM cells. The decoder is an MLP which maps the LSTM output into the predicted  
690 values. For point forecast multivariate forecasting, Yoo and Kang [198] proposed time-invariant  
691 attention to learn the dependencies between the dimensions of the time series and use them with a  
692 convolution architecture to model the time series.

693 Building upon the success of CNNs in other application domains, Borovykh et al. [24] proposed  
694 an adjustment to WaveNet [183] that makes it applicable to conditional forecasting. They evaluated  
695 their model on various datasets with mixed results, concluding that it can serve as a strong baseline  
696 and that various improvements could be made. In a similar vein, inspired by the Transformer  
697 architecture [184] Song et al. [173] proposed an adjustment that makes the architecture applicable  
698 to time series. Their method is applied to both regression and classification tasks.

700 **3.2.2 Multi-step forecast.** N-BEATS [138] is an NN architecture purpose-built for the forecasting  
701 task that relies on a deep, residual stack of MLP layers to obtain point forecasts. The basic building  
702 block in this architecture is a forked MLP stack that takes the block input and feeds the intermediate  
703 representation into separate MLPs to learn the parameters of the context (the authors call it backcast)  
704 and forecast time series models. The residual architecture removes the part of the context signal  
705 it can explain well before passing to the next block and adds up the forecasts. The learned time  
706 series model can have free parameters or be constrained to follow a particular, functional form.  
707 Constraining the model to trend and seasonality functional forms does not have a big impact on  
708 the error and generates models whose stacks are interpretable since the trend and seasonality  
709 components of the model can be separated and analyzed. N-BEATS has also been interpreted as a  
710 meta-learning model [139], where the repeated application of residual blocks can be seen as an  
711 inner optimization loop. N-BEATS generalizes better than other architectures when trained on a  
712 source dataset (e.g., M4-monthly) and applied to a different target datasets (e.g., M3-monthly).

713 Lv et al. [126] propose a *stacked autoencoder* (SAE) architecture to learn features from spatio-  
714 temporal traffic flow data. On top of the autoencoder, a logistic regression layer is used to output  
715 predictions of the traffic flow at all locations in a future time window. The resulting architec-  
716 ture is trained layer-wise in a greedy manner. The experimental results show that the method  
717 significantly improves over other shallow architectures, suggesting that the SAE is capable of  
718 extracting latent features regarding the spatio-temporal correlations of the data. In the same con-  
719 text of spatio-temporal forecasting and under the seq2seq framework, Li et al. [118] proposed the  
720 Diffusion Convolutional Recurrent NN (DCRNN). Diffusion convolution is employed to capture the  
721 dependencies on the spatial domain, while an RNN is utilized to model the temporal dependencies.  
722 Finally, Asadi and Regan [6] proposed a framework where the time series are decomposed in  
723 an initial preprocessing step to separately feed short-term, long-term, and spatial patterns into  
724 different components of a NN. Neighbouring time series are clustered based on their similarity of  
725 the residuals as there can be meaningful short-term patterns for spatial time series. Then, in a CNN  
726 based architecture, each kernel of a multi-kernel convolution layer is applied to a cluster of time  
727 series to extract short-term features in neighbouring areas. The output of the convolution layer is  
728 concatenated by trends and is followed by a convolution-LSTM layer to capture long-term patterns  
729 in larger regional areas.

730 Bandara et al. [14] addressed the problem of predicting a set of disparate time series, which may  
731 not be well captured by a single global model. For this reason, the authors propose to cluster the  
732 time series according to a vector of features extracted using the technique from [87] and the Snob  
733 clustering algorithm [186]. Only then, an RNN is trained per cluster, after having decomposed the  
734 series into trend, seasonality and residual components. The RNN is followed by an affine neural  
735

736 layer to project the cell outputs to the dimension of the intended forecast horizon. This approach is  
 737 applied to publicly available datasets from time series competitions, and appears to consistently  
 738 improve against learning a single global model. In subsequent work, Bandara et al. [13] continued  
 739 to mix heuristics, in this instance seasonality decomposition techniques, known from classical  
 740 forecasting methods with standard NN techniques. Their aim is to improve on scenarios with  
 741 multiple seasonalities such as inter and intra daily. The findings are that for panels of somewhat  
 742 unrelated time series, such decomposition techniques help global models whereas for panels of  
 743 related or homogeneous time series this may be harmful. The authors do not attempt to integrate  
 744 these steps into the NN architecture itself, which would allow for end-to-end learning.

745 Cinar et al. [39] proposed a content attention mechanism that seats on top of any seq2seq RNN.  
 746 The idea is to select a combination of the hidden states from the history and combine them using a  
 747 pseudo-period vector of weights to the predicted output step.

748 Li et al. [116] introduce two modifications to the Transformer architecture to improve its per-  
 749 formance for forecasting. First, they include causal convolutions in the attention to make the key  
 750 and query context dependent, which makes the model more sensitive to local contexts. Second,  
 751 they introduce a sparse attention, meaning the model cannot attend to all points in the history,  
 752 but only to selected points. Through exponentially increasing distances between these points,  
 753 the memory complexity can be reduced from quadratic to  $O(T(\log T)^2)$ , where  $T$  is the sequence  
 754 length, which is important for long sequences that occur frequently in forecasting. Other archi-  
 755 tectural improvements to the Transformer model have also been used more recently to improve  
 756 accuracy and computational complexity in forecasting applications. For example, Lim et al. [119]  
 757 introduce the Temporal Fusion Transformer (TFT), which incorporates novel model components  
 758 for embedding static covariates, performing “variable selection”, and gating components that skip  
 759 over irrelevant parts of the context. The TFT is trained to predict forecast quantiles, and promotes  
 760 forecast interpretability by modifying self-attention and learning input variable importance. Eise-  
 761 nach et al. [51] propose MQ-Transformer, a Transformer architecture that employs novel attention  
 762 mechanisms in the encoder and decoder separately, and consider learning positional embeddings  
 763 from event indicators. The authors discuss the improvements not only on forecast accuracy, but also  
 764 on excess forecast volatility where their model improves over the state of the art. Finally, Zhou et al.  
 765 [204] recently proposed the Informer, a computationally efficient Transformer architecture, that  
 766 specifically targets applications with long forecast horizons. The Informer introduces a *ProbSparse*  
 767 attention layer and a distilling mechanism to reduce both the time complexity and memory usage  
 768 of learning to  $O(T \log T)$ , while improving forecast performance over deep forecasting benchmark.

### 770 3.3 Deep State Space Models

771 In contrast to pure deep learning methods for time series forecasting introduced in Section 2,  
 772 Rangapuram et al. [146] propose to combine classical state space models (SSM) [49, 86] with  
 773 deep learning. The main motivation is to bridge the gap between SSMs that provide a principled  
 774 framework for incorporating structural assumptions but fail to learn patterns across a collection of  
 775 time series, and NNs that are capable of extracting higher order features but results in models that are  
 776 hard to interpret. Their method parametrizes a linear Gaussian SSM using an RNN. The parameters  
 777 of the RNN are learned jointly from a dataset of raw time series and associated covariates. Instead of  
 778 learning the SSM parameters  $\theta_{i,1:T_i}$  for the  $i$ -th time series individually or locally in the terminology  
 779 of Section 2.1, the model is global and learns a shared mapping from the covariates associated  
 780 with each target time series to the parameters of a linear SSM. This mapping  $\theta_{i,t} = f(\mathbf{x}_{i,1:t}; \Phi)$ , for  
 781  $i = 1, \dots, N$  and  $t = 1, \dots, T_i + h$ , is implemented by an RNN with weights  $\Phi$  which are shared across  
 782 different time series as well as different time steps. Note that  $f$  depends on the entire covariate  
 783 time series up to time  $t$  as well as the set of shared parameters  $\Phi$ . Since each individual time series  
 784



785  $i$  is modelled using an SSM with parameters  $\Theta_i$ , assumptions such as temporal smoothness in  
 786 the forecasts are easily enforced. The shared model parameters  $\Phi$  are learned by maximizing the  
 787 likelihood given the observations  $\mathcal{Z} = \{z_{i,1:T_i}\}_{i=1}^N$ . The likelihood terms for each time series reduce  
 788 to the standard likelihood computation under the linear-Gaussian SSM, which can be carried out  
 789 efficiently via Kalman filtering [16]. Once the parameters  $\Phi$  are learned, it is straightforward to  
 790 obtain the forecast distribution via the SSM parameters  $\theta_{i,T_i+1:T_i+h}$ .

791 There are two major limitations of the method proposed in Rangapuram et al. [146]: first, the  
 792 observations are assumed to follow a Gaussian distribution and second, the underlying latent  
 793 process that generates observations is assumed to evolve linearly. de Bézenac et al. [42] address  
 794 the first limitation via *Normalizing Kalman Filters* (NKF) by augmenting SSMs with normalizing  
 795 flows [46, 101, 137] thereby giving them the flexibility to model non-Gaussian, multimodal data.  
 796 Their main idea is to map the non-Gaussian observations  $\{z_{i,1:T_i}\}$  to more Gaussian-like observations  
 797 via a sequence of *learnable*, nonlinear transformations (e.g., a normalizing flow) so that the method  
 798 in [146] can then be applied on the transformed observations. While being more flexible, their  
 799 method still retains attractive properties of linear Gaussian SSMs, namely, tractability of exact  
 800 inference and likelihood computation, efficient sampling, and robustness to noise.

801 In a concurrent work to [42], Kurle et al. [106] improve the method in [146] by addressing both  
 802 limitations. In particular, to model nonlinear latent dynamics, they propose a recurrent switching  
 803 Gaussian SSM, which uses additional latent variables to switch between different linear dynamics.  
 804 Moreover, to handle non-Gaussian observations, they propose a nonlinear emission model via a  
 805 decoder-type NN [61]. Although the exact inference is no longer tractable with these improvements,  
 806 they show that the approximate inference and likelihood estimation can be Rao-Blackwellised; i.e.,  
 807 the inference for the Gaussian latent states can be done exactly while the inference for the switch  
 808 variables needs to be approximated.

809 Finally, Ansari et al. [5] propose to extend [146] via incorporating switching dynamics. The  
 810 recurrent explicit duration switching dynamical system (RED-SDS) is a flexible model that is  
 811 capable of identifying both state- and time-dependent switching dynamics of a time series. State-  
 812 dependent switching is enabled by a recurrent state-to-switch connection and an explicit duration  
 813 count variable is used to improve the time-dependent switching behavior. A hybrid algorithm that  
 814 approximates the posterior of the continuous states via an inference network and performs exact  
 815 inference for the discrete switches and counts provides efficient inference. The method is able to  
 816 infer meaningful switching patterns from the data and extrapolate the learned patterns into the  
 817 forecast horizon.

### 818 3.4 Multivariate Forecasting

820 The models presented up to this point are mainly global univariate models, i.e., they are trained on  
 821 all time series but they are still used to predict a univariate target. When dealing with multivariate  
 822 time series, one should be able to exploit the dependency structure between the different time  
 823 series in the panel in a generalization of Eq. (3) to Eq. (4).

824 Toubeau et al. [179] and Salinas et al. [155] combined RNN-based models with copulas to  
 825 model multivariate distributions. The model in [179] uses a nonparametric copula to capture  
 826 the multivariate dependence structure. In contrast, the work in [155] uses a Gaussian copula  
 827 process approach. Salinas et al. [155] use a low-rank covariance matrix approximation to scale to  
 828 thousands of dimensions. Additionally, the model implements a non-parametric transformation of  
 829 the marginals to deal with varying scales in the dimensions and non-Gaussian data. More recently,  
 830 Rasul et al. [151] proposed to represent the data distribution with a type of normalizing flows called  
 831 Masked Autoregressive Flows [140] while using either an RNN or a Transformer [184] to model  
 832 the multivariate temporal dynamics of time series. Normalizing flows were also used to bring deep  
 833

834 SSMs [146] to a flexible, multivariate scenario [42]. Rasul et al. [150] propose TimeGrad which,  
835 like DeepAR, is an RNN model for which samples are drawn from the data distribution at each  
836 time step, with the difference that in TimeGrad the RNN conditions a diffusion probabilistic model  
837 [172] which allows the model to easily scale to multivariate time series and accurately use the  
838 dependencies between dimensions.

839 A recent application of global multivariate models is for hierarchical forecasting problems [7,  
840 17, 176, 191]. Typically, in such problems an aggregation structure is defined (e.g., via a product  
841 hierarchy) and a trade-off between forecast accuracy and forecast coherency with respect to the  
842 aggregation structure must be managed. Here, forecast coherency or consistency means that the  
843 forecasts conforms to the aggregation structure, so that aggregated forecasts are the same as  
844 forecasts of aggregated time series. This aggregation structure is typically encoded via linear  
845 constraints where the aggregation structure is captured in a matrix  $S$ . Rangapuram et al. [148]  
846 propose to use a multivariate model such as [155] and enforce consistency of forecast samples via  
847 incorporation of a projection of the samples with  $S$  into the learning problem. Dedicated work  
848 exists for aggregation along the time dimension [8, 147, 178].

849 In some multivariate forecasting settings the different dimensions are tied together by some  
850 interpretable connections other than a hierarchy and this can be modelled as part of the input  
851 layer rather than the output as discussed so far. One can for example think of forecasting the traffic  
852 network of a city where the traffic at each of the location in the city is mostly influenced by the  
853 traffic at the neighboring locations, like in PEMS-BAY and METR-LA [118]. Graph Neural Networks  
854 (GNN) have been used in this forecasting setting [43, 63, 102, 163, 193, 206] where, in addition to  
855 the forecasting task, the challenge is to best use the graph information that is provided or even  
856 learn the graph if none is available. The methods that propose to learn the graph do so by looking  
857 for the graph that allows to produce the most accurate forecasts. An embedding is learned for  
858 each dimension, and similarity scores are computed between every two dimension using these  
859 embeddings from which the adjacency matrix is obtained, either by taking the  $K$ -top edges [43, 193]  
860 or sampling from them [102, 163]. As of now, two main strategies have been proposed to learn the  
861 node embeddings, either simply by gradient descent [43, 193] or by taking representation from the  
862 time series [102, 163], with the latter approach to seemingly yielding better results. While these  
863 methods were all presented as point forecasting method, one could obtain probabilistic forecasts  
864 by training these models to parametrize a predictive distribution as explained in Section 2.4.

865  
866

### 3.5 Physics-based Models

867 In *physics-based* models, deep forecasting methods have been proposed that model the underlying  
868 dynamics in sophisticated ways. Chen et al. [32] proposed the Neural ODE (NODE) model, where  
869 an ordinary differential equation (ODE) is solved forward in time, and the adjoint equation is solved  
870 backwards in time using backpropagation. One limitation of the Neural ODE model is that the  
871 unknown parameters  $\theta$  are assumed to be constant in time. Other limitations such as computational  
872 complexity have been addressed in follow-up work, e.g., [18]. Vialard et al. [185] extends the NODE  
873 model to allow the parameters  $\theta(t)$  to be time-varying by introducing a shooting formulation. In  
874 the shooting formulation, the optimal  $\theta$  is determined by minimizing a regularized loss function.  
875 Vialard et al. [185] also shows that a residual network (ResNet) can be expressed as the Forward  
876 Euler discretization of an ODE with time step  $\Delta t = 1$ . Wang et al. [187] compares successful time  
877 series deep sequence models, such as [146, 156] to NODE and other hybrid deep learning models to  
878 model COVID-19 dynamics, as well as the population dynamics using the Lotka-Volterra equations.  
879 Through their benchmarking study, the authors show that distribution shifts can pose problems  
880 for deep sequence models on these tasks, and propose a hybrid model AutoODE to model the  
881 underlying dynamics.

882

### 3.6 Global-local

With *local models*, the free parameters of the model are learned individually for each series in a collection, see Section 2.1. Classical local time series models such as SSMs, ARIMA, and exponential smoothing (ETS) [84] excel at modelling the complex dynamics of individual time series given a sufficiently long history. Other local models include Gaussian SSMs, which are computationally efficient, e.g., via a Kalman filter, and Gaussian Processes (GPs) [27, 67, 149, 159]. These methods provide uncertainty estimates, which are critical for optimal downstream decision making. Since these methods are local, they learn one model per time series and cannot effectively extract information across multiple time series. These methods are unable to address cold-start problems where there is a need to generate predictions for a time series with little or no observed history.

Conversely, recall that in *global models*, their free parameters are learned jointly on every series in a collection of time series. NNs have proven particularly well suited as global models [64, 111, 146, 156, 190]. Global methods can extract patterns from collections of irregular time series even when these patterns would not be distinguishable using a single series.

Global-local models have been proposed to combine the advantages of both global and local models. Examples include mixed effect models [40], which consist of two kinds of effects: fixed (global) effects that describe the whole population, and random (local) effects that capture the idiosyncratic of individuals or subgroups. A similar mixed approach is used in hierarchical Bayesian [65] methods, which combine global and local models to jointly model a population of related statistical problems. In an early example of hierarchical Bayesian models, [31] combined global and local features for intermittent demand forecasting in retail planning. In [3, 123], other combined global and local models are detailed.

A recent global-local family of models, Deep Factors [188] provide an alternative way to combine the expressive power of NNs with the data efficiency and uncertainty estimation abilities of classical probabilistic local models. Each time series, or its latent function for non-Gaussian data, is represented as the weighted sum of a global time series and a local model. The global part is given by a linear combination of a set of deep dynamic factors, where the loading is temporally determined by attentions. The local model is stochastic. Typical choices include white noise processes, linear dynamical systems, GPs [127] or RNNs. The stochastic local component allows for the uncertainty to propagate forward in time, while the global NN model is capable of extracting complex nonlinear patterns across multiple time series. The global-local structure extracts complex nonlinear patterns globally while capturing individual random effects for each time series locally.

The Deep Global Local Forecaster (DeepGLO) [162] is a method that “thinks globally and acts locally” to forecast collections of up to millions of time series. It crucially relies on a type of temporal convolution (a so-called leveled network), that can be trained across a large amount time series with different scales without the need for normalization or rescaling. DeepGLO is a hybrid model that uses a global matrix factorization model [200] regularized by a temporal deep leveled network and a local temporal deep level network to capture patterns specific to each time series. Each time series is represented by a linear combination of  $k$  basis time series, where  $k \ll N$ , with  $N$  the total number of time series. The global and local models are combined through data-driven attention for each time series.

A further example in the global-local model class is the ES-RNN model proposed by Smyl [169] that has recently attracted attention by winning the M4 competition [128] by a large margin on both evaluation settings. In the ES-RNN model, locally estimated level and trend components are multiplicatively combined with an RNN model. Apart from its global-local nature, it also integrates aspects of different model classes into a a single model similar to Deep State Space models (Section 3.3). In particular, the  $h$ -step ahead prediction  $\hat{\mathbf{z}}_{i,t+1:t+h} = l_{i,t} \cdot \mathbf{s}_{i,t+1:t+h} \cdot \exp(\text{RNN}(\mathbf{x}_{i,t}))$

932 consists of a level  $l_{i,t}$  and a seasonal component  $s_{i,t}$  obtained through local exponential smoothing,  
 933 and the output of a global RNN model  $\text{RNN}(\mathbf{x}_{i,t})$ , where  $\mathbf{x}_{i,t}$  is a vector of preprocessed data  
 934 extracted from deseasonalized and normalized time series  $\mathbf{x}_{i,t} = \log(\mathbf{z}_{i,t-K:t}/(s_{i,t-K:t}l_{i,t}))$  cut in a  
 935 window of length  $K + 1$ . The RNN models are composed of dilated LSTM layers with additional  
 936 residual connections. The M4-winning entry used slightly different architectures for the different  
 937 type of time series in the competition.

938

939

### 3.7 Intermittent Time Series

940

941

942

943

944

945

946

947

948

949

950

951

952

953

954

955

956

957

958

959

960

961

962

963

964

965

966

967

968

969

970

971

972

973

974

975

976

977

978

979

980

We noted in the introduction that deep forecasting models had a major impact on operational forecasting problems. In these large-scale problem, intermittent time series occur regularly [25]. Accordingly, research on NNs for intermittent time series forecasting has been an active area. Salinas et al. [156] propose a standard RNN architecture with a negative binomial likelihood to handle intermittent demand similar to [171] in classical methods. To the best of our knowledge, other likelihoods that have been proposed for intermittent time series in classical models, e.g., by [160], have not yet been carried over to NNs. However, some initial work is available via more standard likelihoods [93, 156].

In the seminal paper on intermittent demand forecasting [41], Croston separates the data in a sequence of observed non-zero demands and a sequence of time intervals between positive demand observations, and runs exponential smoothing separately on both series. A comparison of NNs to classical models for intermittent demand first appeared in Gutierrez et al. [73], where the authors compare the performance of a shallow and narrow MLP with Croston's method. They find NNs to outperform classical methods by a significant margin.

Kourentzes [105] proposes two MLP architectures for intermittent demand, taking demand sizes and intervals as inputs. As in Gutierrez et al. [73], the networks are shallow and narrow by modern standards, with only a single hidden layer and three hidden units. The difference between the two architectures is in the output. In one case interval times and non-zero occurrences are output separately, while in the other a ratio of the two is computed. The approach proposed by Kourentzes [105] outperforms other approaches primarily with respect to inventory metrics, but not forecasting accuracy metrics, challenging previous results in [73]. It is unclear whether the models are used as global or local. However, given the concern around overfitting and regularization, we assume that these models were primarily used as local models in the experiments.

Both approaches of [73, 105] only offer point forecasts. This shortcoming is addressed by [180, 181], where the authors propose renewal processes as natural models for intermittent demand forecasting. Specifically, they use RNNs to modulate both discrete time and continuous time renewal processes, using the simple analogy that RNNs can replace exponential smoothing in [41].

Finally, a recent trend in sequence modelling employs NNs in modelling discrete event sequences observed in continuous time [48, 133, 164, 165, 182, 194] and [166] for an overview. Notably, Xiao et al. [195] use two RNNs to parametrize a probabilistic "point process" model. These networks consume data from asynchronous event sequences and uniformly sampled time series observations respectively. Their model can be used in forecasting tasks where time series data can be enriched with discrete event observations in continuous time.

973

974

### 3.8 Generalized Adversarial Networks

975

976

977

978

979

980

Additionally to the approaches mentioned in Sections 2.4 and 2.4.3, the recent literature contains further examples for density estimation, most prominently via Generalized Adversarial Networks (GANs) [71]. While GANs have received much attention in the overall deep learning literature [52, 53, 98, 121, 199], this has not been reflected in forecasting. We speculate that this is because a discriminator network can be replaced by metrics such as CRPS which measure the quality of

981 generated samples. We therefore only provide a brief overview here and mention that, while they  
982 rely on the buildings blocks discussed in Section 2, they typically require architectures that are  
983 more complex than then ones discussed here and lead to involved optimization problems.

984 Despite the comparably less attention that GANs have received in forecasting, they have been  
985 recently applied to the time series domain [53, 199] to synthesize data [53, 177] or to employ  
986 an adversarial loss in forecasting tasks [192]. Many time series GAN architectures use recurrent  
987 networks to model temporal dynamics [53, 134, 199]. Modelling long-range dependencies and  
988 scaling recurrent networks to higher lengths is inherently difficult and limits the application of  
989 time series GANs to short sequence lengths [53, 199]. One way to achieve longer realistic synthetic  
990 time series is by employing convolutional [11, 62, 183] and self-attention architectures [184].

991 Convolutional architectures are able to learn relevant features from the raw time series data [11,  
992 62, 183], but are ultimately limited to local receptive fields and can only capture long-range  
993 dependencies via many stacks of convolutional layers. Self-attention can bridge this gap and  
994 allow for modelling long-range dependencies from convolutional feature maps, which has been a  
995 successful approach in the image [203] and time series forecasting domain [116]. Another technique  
996 to achieve long sample sizes is progressive growing, which successively increases the resolution by  
997 adding layers to a GAN generator and discriminator during training [97]. A recent proposal [21]  
998 synthesizes progressive growing with convolutions and self-attention into a novel architecture  
999 particularly geared towards time series.

### 1000 3.9 Summary and Practical Guidelines

1001 In Section 2 and this section, we introduced a large number of deep forecasting models. We  
1002 summarize the main approaches in Table 2. The list below provide keys to reading the table.

- 1003 • *Forecast* distinguishes between probabilistic (*Prob*) and *Point* forecasts.
- 1004 • *Horizon* indicates whether the model does one-step predictions (noted 1) in which case  
1005 multi-step forecasts are obtained recursively, or if it directly predicts a whole sequence  
1006 ( $\geq 1$ ).
- 1007 • *Loss* and *Metrics* specifies the loss used for training and metrics used for evaluation. Here,  
1008 we only provide an explanation of the acronyms and not the definition of each metric  
1009 which can be easily found in the corresponding papers: negative log-likelihood (NLL),  
1010 quantile loss (QL), continuous ranked probability score (CRPS), (normalized) (root) mean  
1011 squared error (NRMSE, RMSE, MSE), root relative squared error (RRSE), relative geometric  
1012 RMSE (RGRMSE), weighted absolute percentage error (WAPE), normalized deviation (ND),  
1013 mean absolute deviation (MAD), mean absolute error (MAE), mean relative error (MRE),  
1014 (weighted) mean absolute percentage error (wMAPE, MAPE), mean absolute scaled error  
1015 (MASE), overall weighted average (OWA), mean scaled interval score (MSIS), Kullback-  
1016 Leibler divergence (KL), Value-at-Risk (VaR), expected shortfall (ES), empirical correlation  
1017 coefficient (CORR), area under the receiver operating characteristic (AUROC), percentage  
1018 best (PB).

1019 While Table 2 serves to illustrate the wealth of deep forecasting methods now available, their sheer  
1020 number may be slightly overwhelming. Furthermore, empirical evidence on the effectiveness of the  
1021 different architectures has so far not revealed a clearly superior approach [4]. In this, forecasting  
1022 differs from other domains, e.g., natural language processing where Transformer-based models [184]  
1023 dominate overall. Also, deep forecasting methods seem to differ from other model families, such  
1024 as tree-based methods where LightGBM [99] or XGBoost [33] dominate (as in the recent M5  
1025 forecasting competition [92]). We speculate that this diffuse picture is in part due to the practical  
1026 reasons, the relative immaturity of the field and the corresponding software implementations and in  
1027  
1028  
1029

part due to fundamental reason as a natural consequence of the breadth and diversity of forecasting problems.

So, choosing the appropriate architecture for a problem at hand can be a daunting task. In the following, we therefore attempt to provide guidelines for a more informed deep forecasting model selection. These are largely based on our own experience in working with practical forecasting problem and they should primarily be taken as a non-exhaustive guidance on where to start model exploration.

*3.9.1 Baseline methods and standard mode of deployment.* At the start of any in-depth model exploration, considering a baseline model is commonly accepted best practice. To the best of our knowledge, the most mature deep forecasting models are DeepAR [156] and MQCNN [190] which exist in a number of open-source and commercial implementations.<sup>3</sup> As a practical guideline, we recommend to start model exploration using at least these methods as baselines. Other candidates we would consider are N-BEATS [138], WaveNet [183] and a Transformer-based model. The relative performance of these methods compared with other methods should give reasonable, directional evidence whether the problem at hand is amenable to deep forecasting methods. We note that AutoML approaches for forecasting are available<sup>4</sup> but while promising are in their infancy. At least in the M5 competition, they are still outperformed by the aforementioned more specialized deep forecasting models.

Our typical suggestion is to employ NNs as global models since, given enough data, global methods outperform classical local methods when dealing with groups of similar time series.<sup>5</sup> Interestingly, recent empirical evidence have shown that global models can achieve a state-of-the-art performance even in heterogeneous groups of time series. This is supported by the M4 [128] and M5 competitions where the top performing models had some form of globality. This suggests a more general applicability of global methods with a high impact on practical application where a general automated forecasting mechanism is required.

*3.9.2 Data characteristics.* The amount of data available is among the easiest dimensions in choosing a deep forecasting model. First, NNs require a minimum amount of data to be effective in comparison to other, more parsimoniously parametrized models. This is perhaps the most important factor in successful applications of NNs in forecasting. How much data does one need for a given application? Several important points should be discussed on this question. First, the amount of data is often misunderstood as the *number of time series* but in reality the amount of data typically relates to the *number of observations*. For instance, one may have only one time series but many thousands of observations, as in the case of a time series from a real-time sensor where measurements happen every second for a year, allowing to fit a complex NN [2]. Second, it is probably better to see the amount of data in terms of *information quantity*. For instance, in finance the amount of information of many millions of hourly transactions is limited given the very low signal-to-noise ratio in contrast to a retailer whose products follow clear seasonality and patterns, making it easier to apply deep learning methods. The more structured the data is (e.g., via strong seasonality or knowledge about the underlying process) the better deep forecasting models that incorporate these structures will fare. On the contrary, if the time series are more irregular or short, a more data-driven approach (e.g., via Transformer-based models) will often be preferable. The importance of covariate information for the forecasting problem at hand can further help

<sup>3</sup><https://aws.amazon.com/blogs/machine-learning/now-available-in-amazon-sagemaker-deepar-algorithm-for-more-accurate-time-series-forecasting/>

<sup>4</sup><http://ai.googleblog.com/2020/12/using-automl-for-time-series-forecasting.html>

<sup>5</sup>This is a more generally applicable fact beyond NN. Montero-Manso and Hyndman [135] show favorable theoretical and empirical properties for global over local models.

Study	Structure	Forecast	Horizon	Loss	Metrics	Data Types	Comments
DeepAR [156]	RNN	Prob	1	NLL	Coverage, QL, ND, NRMSE	demand, traffic, electricity	Learns parametric distributions
Toubeau et al. [179]	RNN/CNN	Prob	1	NLL/QL	RMSE, price	electricity	Nonparametric copula to capture multivariate dependence
Salinas et al. [155]	RNN	Prob	1	NLL	QL, MSE	electricity, traffic, exchange rate, solar, taxi, wiki	Learns multivariate model via low-rank Gaussian copula processes
ARMDN [136]	RNN	Prob	1	NLL	wMAPE	demand	Like [156], but using mixture of Gaussian's and domain specific feature processing
QARNN [196]	MLP	Prob	1	QL	VaR, ES	finance	Conditional quantile function estimation
SQF-RNN [64]	RNN	Prob	1	CRPS	QL, MSIS, NRMSE, OWA	demand, traffic, count data, finance, M4	Models non-parametric distributions with splines
LSTNet [107]	CNN + RNN + MLP	Point	1	$\ell_1$	RRSE, CORR	traffic, solar, electricity, exchange rate	Extracts short and long temporal patterns with a CNN and RNN, respectively
Zhu and Laptev [205]	RNN + MLP	Prob	1	-	sMAPE, calibration	daily trips	Fits an encoder (RNN) that constructs an embedding state, which is fed to a prediction network (MLP)
Laptev et al. [111]	RNN	Prob	1	MSE	sMAPE	traffic, M3	LSTM as feature extractor
Qiu et al. [144]	MLP + SVR	Point	1	$\ell_2$ for MLP, SVR objective	RMSE, MAPE	energy, housing	Ensemble of DBNs where their output is fed to an SVR
A-LSTM [82]	RNN + MLP	Point	1	$\ell_2, \ell_2$ regularizer	RMSE	electricity consumption	Combination of LSTM with autoencoders

1079  
1080  
1081  
1082  
1083  
1084  
1085  
1086  
1087  
1088  
1089  
1090  
1091  
1092  
1093  
1094  
1095  
1096  
1097  
1098  
1099  
1100  
1101  
1102  
1103  
1104  
1105  
1106  
1107  
1108  
1109  
1110  
1111  
1112

								1113
								1114
								1115
								1116
								1117
								1118
								1119
								1120
								1121
								1122
								1123
								1124
								1125
								1126
								1127
								1128
								1129
								1130
								1131
								1132
								1133
								1134
								1135
								1136
								1137
								1138
								1139
								1140
								1141
								1142
								1143
								1144
								1145
								1146
Borovykh et al. [24]	CNN	Point	1	$\ell_1, \ell_2$ regularizer	RMSE, MASE, HITS	index forecasting, exchange rate	WaveNet [183] based model adjusted for time series forecasting	
SAnD [173]	MLP + Attention	Point	1	$\ell_2$ , cross-entropy, multi-label classification loss	AUROC, MASE, MSE	clinical	Transformer [116] based model adjusted for time series forecasting	
Zhang [202]	MLP	Point	1	MSE	MSE, MAD	sunspot, lynx, exchange rate	Hybrid local model that uses ARIMA to capture the linear component and a NN for the nonlinear residuals	
Khashei and Bijari [100]	MLP	Point	1	MSE	MAE, MSE	sunspot, lynx, exchange rate	Hybrid local model that uses ARIMA and a NN for trend correction	
Deep State Space [146]	RNN + State Space	Prob	$\geq 1$	NLL	P50, P90 quantile loss	traffic, electricity, tourism, M4	RNN parametrized linear Gaussian SSM	
NKF [42]	RNN + State Space + Normalizing Flow (NF)	Prob	$\geq 1$	NLL	QL	traffic, electricity, exchange rate, solar, wiki	RNN parametrized linear Gaussian SSM combined with normalizing flow, which acts as an emission model to handle non-Gaussian data	
ARSGLS [106]	Recurrent Switching State Space + NN	Prob	$\geq 1$	NLL	QL	traffic, electricity, exchange rate, solar, wiki	Recurrent Switching State Space combined with decoder-type NN, which acts as an emission model to handle non-Gaussian data	
MQ-RNN/CNN [190]	RNN/CNN + MLP	Prob	$\geq 1$	QL	QL, calibration, sharpness	demand	Learns pre-specified grid of quantiles	
Wen and Torkkola [189]	CNN + MLP	Prob	$\geq 1$	QL, inverse reconstruction loss, NLL	QL, quantile crossing, QL over sum of future intervals	demand	Combines model in [190] with Gaussian copula	
DeepTCN [35]	CNN + MLP	Prob	$\geq 1$	QL	QL	retail demand	Learns pre-specified grid of quantiles	



								1147
								1148
								1149
								1150
								1151
								1152
								1153
								1154
								1155
								1156
								1157
								1158
								1159
								1160
								1161
								1162
								1163
								1164
								1165
								1166
								1167
								1168
								1169
								1170
								1171
								1172
								1173
								1174
								1175
								1176
								1177
								1178
								1179
								1180
N-BEATS [138]	MLP	Point	$\geq 1$	sMAPE, MASE, MAPE	sMAPE, MASE, OWA	M4	Deep, residual MLP that learns interpretable trend and seasonality function	
Lv et al. [126]	Stacked autoencoder	Point	$\geq 1$	MSE, KL sparsity constraint	MAE, MRE, RMSE	traffic	Stacked autoencoders with logistic regression output layer	
DCRNN [118]	RNN	Point	$\geq 1$	NLL	MAE, MAPE, RMSE	traffic	Diffusion convolution for spatial and RNN for temporal dependencies	
Asadi and Regan [6]	CNN + RNN	Point	$\geq 1$	$\ell_2$	MAE, RMSE	traffic	Decomposition-based model for spatio-temporal forecasting	
Bandara et al. [14]	RNN + Classical Decomposition	Point	$\geq 1$	-	sMAPE	CIF2016, NN5	Clusters time series based on set of features and train one model per cluster	
LSTM-MSNet [13]	RNN + Classical Decomposition	Point	$\geq 1$	$\ell_1$	sMAPE, MASE	M4, energy	Decomposition based model with multiple seasonal patterns	
Cinar et al. [39]	RNN + Attention	Point	$\geq 1$	$\ell_2, \ell_2$ regularizer	MSE, sMAPE	energy, max temperature, CPU usage, air quality	Attention mechanism on top of RNN	
Deep Factors [188]	RNN + GP	Prob	$\geq 1$	NLL	QL, MAPE	electricity, traffic, taxi, uber	Global RNN and a local GP	
DeepGLO [162]	CNN	Point	$\geq 1$	$\ell_2$	WAPE, MAPE, sMAPE	electricity, traffic, wiki	Global matrix factorization regularized by a deep leveled network	
ES-RNN [169]	RNN	Point	$\geq 1$	QL	MASE, sMAPE, MSIS	M4	Locally estimated seasonality and trend and global RNN	
Kourentzes [105]	MLP	Point	1	$\ell_2$	ME, MAE, service level	intermittent demand	MLP-based intermittent demand model	

Attentional Twin RNN [195]	RNN	Prob	1	NLL	MAE	point process data	Event sequence prediction	1181
Gutierrez et al. [73]	MLP	Point	1	$\ell_2$	MAPE, RGRMSE, PB	intermittent demand	MLP-based intermittent demand model	1182
Deep Renewal Process [181]	RNN	Prob	$\geq 1$	NLL	P50, P90 quantile loss	intermittent demand	RNN-based intermittent demand model inspired by point processes	1183
WaveNet [4, 183]	CNN	Prob	$\geq 1$	NLL	mean opinion score	traffic, electricity, M4	Diluted causal convolutions	1184
Transformer [116]	MLP	Point	1	NLL	QL	electricity, traffic, wind, M4, solar	Transformer with causal convolutions and sparse attention	1185
AttnAR [198]	CNN	Point	1	$\ell_2$	RMSE	electricity, traffic, solar, exchange rate	Multivariate forecasting	1186
LSTM MAF [151]	RNN	Prob	$\geq 1$	NLL	RMSE	electricity, traffic, solar, exchange rate	Multivariate forecasting using normalizing flows	1187
TimeGrad [150]	RNN	Prob	$\geq 1$	NLL	RMSE	electricity, traffic, solar, exchange rate, taxi, wikipedia	Multivariate forecasting using diffusion models.	1188
TFT [119]	LSTM, MLP	Prob	$\geq 1$	QL	P50, P90 quantile loss	electricity, traffic, retail, volatility	Modified transformer architecture for improved interpretability	1189
MQ-Transformer [51]	CNN, MLP	Prob	$\geq 1$	QL	P50, P90, LT-SP	electricity, traffic, retail, volatility, retail demand (proprietary)	Architectural improvements on MQ-RNN/CNN for multi-step forecasting	1190
Informer [204]	CNN, MLP	Point	$\geq 1$	MSE	MSE, MAE	electricity, weather, sensor data	Sparse and computationally efficient transformer architecture	1191
								1192
								1193
								1194
								1195
								1196
								1197
								1198
								1199
								1200
								1201
								1202
								1203
								1204
								1205
								1206
								1207
								1208
								1209
								1210
								1211
								1212
								1213
								1214

Table 2. Summary of modern deep forecasting models.

1215 determine the correct method. Some NN architectures need extensions to include such information  
1216 while others readily accept them.

1217 From a practical perspective, NNs have been reported to outperform demand forecasting baselines  
1218 starting from 50000 observations in [156] and from a few hundred observations in load-forecasting  
1219 [146, 188]. Understanding better these limitation, both theoretically and empirically, is an area of  
1220 current research and is not yet well understood. See [131] for some current theoretical work on  
1221 sample complexity of global-local approaches for instance and [23] for empirical work.

1222  
1223 *3.9.3 Problem characteristics.* The characteristics of the forecasting problem to be solved are  
1224 natural important decision points. We list a few dimensions to consider here.

1225 One important aspect of a model is its forecast nature, i.e., if it produces probabilistic or point  
1226 forecasts. The choice of this is highly dependent on the underlying application. To illustrate this we  
1227 can examine two different forecasting use cases: product demand and CPU utilization. In the former  
1228 use case one wishes to forecast the future demand of a product in order to take a more informed  
1229 decision about the stock that is required to have in a warehouse or to optimize the labour planning  
1230 based on the traffic that is expected. In the latter, the forecast of CPU utilization could be used to  
1231 identify in a timely manner if a process will fail in order to proactively resolve associated issues, or  
1232 to detect possible anomalous behaviours that could trigger some root cause analysis and system  
1233 improvements. Although in both applications a forecast is required, the end goal is different, which  
1234 changes the requirements of the chosen forecasting model. For example, for product demand the  
1235 whole distribution of the future demand might be important: one cannot rely on a single forecast  
1236 value since the variance in the forecast plays an important role to avoid out of stock issues or  
1237 under/over planning the expected required labour. Therefore, in this application it is important  
1238 to use a model that focuses on predicting accurately the whole distribution. On the other hand,  
1239 for CPU utilization one might be interested in the 99-th percentile, since everything below that  
1240 threshold might not be of particular interest or does not produce any actionable alarm. In this case,  
1241 a model that focuses on a particular quantile of importance is of higher interest than a model that  
1242 predicts the whole distribution with possibly worse accuracy on the selected quantile.

1243 It is observed [42, 106, 146, 155, 156] empirically that autoregressive models are superior in  
1244 performance (in terms of forecast accuracy) compared to state space models, especially when the  
1245 data is less noisy and the forecast horizon is not too long. This is not surprising given that the  
1246 autoregressive models directly use past observations as input features and treat own predictions as  
1247 lag inputs in the multi-step forecast setting. A general rule of thumb is that if one knows details  
1248 such as the forecast horizon, the quantile to query or the exact goals of the forecasting problem  
1249 in advance and these are unlikely to change, then a discriminative model is often a good default  
1250 choice. Conversely, state space models proved to be robust when there are missing and/or noisy  
1251 observations [42]. Moreover, if the application-specific constraints can be incorporated in the latent  
1252 state, then state space models usually perform better even in the low-data regimes [146].

1253 The length of the forecast horizon relative to the history or, more generally speaking, the  
1254 importance of the historic values for future values must further be taken into account. For example,  
1255 very long forecast horizons may require to control (e.g., via differential equations) the exponential  
1256 growth in the target. A canonical example for this is forecasting of a pandemic. This example  
1257 further clarifies the importance of being able to produce counterfactuals for what-if analysis (e.g.,  
1258 the incorporation of intervention). Not all deep forecasting models allow for this.

1259  
1260 *3.9.4 Other Aspects.* A number of other aspects can further help to narrow the model exploration  
1261 space. For example, computational constraints (how much time/money for training is available,  
1262 are there constraints on the latency during inference) can favor “simpler” NNs, see e.g., [23] for a  
1263

1264 discussion on multi-objective forecasting model selection. Another aspect to consider could be CNN  
1265 over RNN-based architectures. The skill set of the research team available is an important factor.  
1266 For example, probabilistic models often are more sensitive towards parametrization and identifying  
1267 reasonable parameter ranges requires in-depth knowledge. On the other extreme, troubleshooting  
1268 Transformer-based models requires deep learning experience that not every research team may  
1269 possess. The time budget available for model development and the willingness to extend existing  
1270 models are further factors.

1271

1272

#### 4 CONCLUSIONS AND AVENUES FOR FUTURE WORK

1273

1274 This article has attempted to provide an introduction to and an overview of NNs for forecasting  
1275 or deep forecasting. We began by providing a panorama of some of the core concepts in the  
1276 modern literature on NNs chosen by their degree of relevance for forecasting. We then reviewed  
1277 the literature on recent advances in deep forecasting models.

1278 Deep forecasting methods have received considerable attention in the literature because they excel  
1279 at addressing forecasting problems with many related time series and at extracting weak signals  
1280 and complex patterns from large amounts of data. From a practical perspective, the availability of  
1281 efficient programming frameworks helps to alleviate many of the pain points that practitioners  
1282 experience with other forecasting methods such as manual feature engineering or the need to  
1283 derive gradients. However, NNs are not a silver bullet. For many important classes of forecasting  
1284 problems such as long-range macro-economic forecasts or other problems requiring external domain  
1285 knowledge not learnable from the data, deep forecasting methods are not the most appropriate  
1286 choice and will likely never be. Still, it is our firm belief that NNs belong to the toolbox of every  
1287 forecaster, in industry and academia.

1288 Building onto the existing promising work in NNs for forecasting, many challenges remain  
1289 to be solved. We expect that the current trends of hybridizing existing time series techniques  
1290 with NNs [64, 146, 169, 181] and bringing innovations from other related areas or general purpose  
1291 techniques to forecasting [71, 183, 184] will continue organically. Typical general challenges for  
1292 NNs, such as data effectiveness, are important in forecasting and likely need a special treatment (see  
1293 [59] for an approach in time series classification with transfer learning). Other topics of general ML  
1294 interest such as interpretability, explainability and causality (e.g., [19, 122, 158]) are of particular  
1295 practical importance in the forecasting setting. It is our hope that original methods such as new  
1296 NN architectures will be pioneered in the time series prediction sector (e.g., [138]) and that those  
1297 will then feed back into the general NN literature to help solve problems in other disciplines.

1298 Beyond such organic improvements, we speculate that another area in which NNs have had  
1299 tremendous impact [167, 168] may become important for forecasting, namely deep reinforcement  
1300 learning. In contrast to current practice, where forecasting merely serves as input to downstream  
1301 decision problems (often mixed-integer nonlinear stochastic optimization problems), for example  
1302 to address problems such as restocking decisions, reinforcement learning allows to directly learn  
1303 optimal decisions in business context [89]. It will be interesting to see whether reinforcement based  
1304 approaches can improve decision making – and how good forecasting models could help improve  
1305 reinforcement approaches.

1306 As methodology advances, so will the applicability. Many potential applications of forecasting  
1307 methods are under-explored. To pick areas that are close to the authors' interests, in database  
1308 management, cloud computing, and system operations a host of applications would greatly benefit  
1309 from the use of principled forecasting methods (see e.g., [9, 22, 60]). Forecasting can also be used  
1310 to improve core ML tasks such as hyperparameter optimization (e.g., [47]) and we expect more  
1311 applications to open up in this area.

1312

## REFERENCES

- [1] Martin Abadi, Paul Barham, Jianmin Chen, Zhifeng Chen, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Geoffrey Irving, Michael Isard, Manjunath Kudlur, Josh Levenberg, Rajat Monga, Sherry Moore, Derek G. Murray, Benoit Steiner, Paul Tucker, Vijay Vasudevan, Pete Warden, Martin Wicke, Yuan Yu, and Xiaoqiang Zheng. 2016. TensorFlow: A system for large-scale machine learning. In *12th USENIX Symposium on Operating Systems Design and Implementation (OSDI 16)*. 265–283.
- [2] Subutai Ahmad, Alexander Lavin, Scott Purdy, and Zuha Agha. 2017. Unsupervised real-time anomaly detection for streaming data. *Neurocomputing* 262 (2017), 134–147.
- [3] Amr Ahmed, Moahmed Aly, Joseph Gonzalez, Shrvan Narayanamurthy, and Alexander J Smola. 2012. Scalable inference in latent variable models. In *Proceedings of the fifth ACM International Conference on Web Search and Data Mining*. ACM, 123–132.
- [4] Alexander Alexandrov, Konstantinos Benidis, Michael Bohlke-Schneider, Valentin Flunkert, Jan Gasthaus, Tim Januschowski, Danielle C Maddix, Syama Sundar Rangapuram, David Salinas, Jasper Schulz, et al. 2020. GluonTS: Probabilistic and Neural Time Series Modeling in Python. *Journal of Machine Learning Research* 21, 116 (2020), 1–6.
- [5] Abdul Fatir Ansari, Konstantinos Benidis, Richard Kurlle, Ali Caner Turkmen, Harold Soh, Alexander J Smola, Bernie Wang, and Tim Januschowski. 2021. Deep Explicit Duration Switching Models for Time Series. *Advances in Neural Information Processing Systems* 34 (2021).
- [6] Reza Asadi and Amelia C Regan. 2020. A Spatial-Temporal Decomposition Based Deep Neural Network for Time Series Forecasting. *Applied Soft Computing* 87 (2020), 105963.
- [7] George Athanasopoulos, Roman A Ahmed, and Rob J Hyndman. 2009. Hierarchical forecasts for Australian domestic tourism. *International Journal of Forecasting* 25, 1 (2009), 146–166.
- [8] George Athanasopoulos, Rob J Hyndman, Nikolaos Kourentzes, and Fotios Petropoulos. 2017. Forecasting with temporal hierarchies. *European Journal of Operational Research* 262, 1 (2017), 60–74.
- [9] Fadhel Ayed, Lorenzo Stella, Tim Januschowski, and Jan Gasthaus. 2020. Anomaly Detection at Scale: The Case for Deep Distributional Time Series Models. arXiv:2007.15541 [cs.LG]
- [10] Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2014. Neural Machine Translation by Jointly Learning to Align and Translate. arXiv preprint arXiv:1409.0473 (2014).
- [11] Shaojie Bai, J Zico Kolter, and Vladlen Koltun. 2018. An empirical evaluation of generic convolutional and recurrent networks for sequence modeling. arXiv preprint arXiv:1803.01271 (2018).
- [12] Luca Vincenzo Ballestra, Andrea Guizzardi, and Fabio Palladini. 2019. Forecasting and trading on the VIX futures market: A neural network approach based on open to close returns and coincident indicators. *International Journal of Forecasting* 35, 4 (2019), 1250 – 1262.
- [13] Kasun Bandara, Christoph Bergmeir, and Hansika Hewamalage. 2020. LSTM-MSNet: Leveraging forecasts on sets of related time series with multiple seasonal patterns. *IEEE Transactions on Neural Networks and Learning Systems* (2020).
- [14] Kasun Bandara, Christoph Bergmeir, and Slawek Smyl. 2017. Forecasting across time series databases using long short-term memory networks on groups of similar series. arXiv preprint arXiv:1710.03222 8 (2017), 805–815.
- [15] Kasun Bandara, Peibei Shi, Christoph Bergmeir, Hansika Hewamalage, Quoc Tran, and Brian Seaman. 2019. Sales demand forecast in e-commerce using a long short-term memory neural network methodology. In *International Conference on Neural Information Processing*. Springer, 462–474.
- [16] David Barber. 2012. *Bayesian reasoning and machine learning*. Cambridge University Press.
- [17] Souhaib Ben Taieb, James W Taylor, and Rob J Hyndman. 2017. Coherent probabilistic forecasts for hierarchical time series. In *International Conference on Machine Learning*. 3348–3357.
- [18] Marin Biloš, Johanna Sommer, Syama Sundar Rangapuram, Tim Januschowski, and Stephan Günnemann. 2021. Neural Flows: Efficient Alternative to Neural ODEs. *Advances in Neural Information Processing Systems* 34 (2021).
- [19] Alexander Binder, Sebastian Bach, Gregoire Montavon, Klaus-Robert Müller, and Wojciech Samek. 2016. Layer-wise relevance propagation for deep neural network architectures. In *Information Science and Applications (ICISA)*. Springer, 913–922.
- [20] Toby Bischoff and Austin Gross. 2019. Wavenet & Dropout: An efficient setup for competitive forecasts at scale. In *Proceedings of the International Symposium on Forecasting*.
- [21] Michael Bohlke-Schneider, Paul Jeha, Pedro Mercado, Shubham Kapoor, Jan Gasthaus, and Tim Januschowski. 2022. PSA-GAN: Progressive Self Attention GANs for Synthetic Time Series. *International Conference on Learning Representations (ICLR)* (2022).
- [22] Michael Bohlke-Schneider, Shubham Kapoor, and Tim Januschowski. 2020. Resilient Neural Forecasting Systems. In *Proceedings of the Fourth International Workshop on Data Management for End-to-End Machine Learning* (Portland, OR, USA) (DEEM’20). Association for Computing Machinery, New York, NY, USA, Article 4, 5 pages.

- 1362 [23] Oliver Borchert, David Salinas, Valentin Flunkert, Tim Januschowski, and Stephan Günnemann. 2022. Multi-Objective  
 1363 Model Selection for Time Series Forecasting. *arXiv preprint arXiv:2202.08485* (2022).
- 1364 [24] Anastasia Borovykh, Sander Bohte, and Cornelis W Oosterlee. 2017. Conditional time series forecasting with  
 1365 convolutional neural networks. *arXiv preprint arXiv:1703.04691* (2017).
- 1366 [25] Joos-Hendrik Böse, Valentin Flunkert, Jan Gasthaus, Tim Januschowski, Dustin Lange, David Salinas, Sebastian  
 1367 Schelter, Matthias Seeger, and Yuyang Wang. 2017. Probabilistic demand forecasting at scale. *Proceedings of the VLDB*  
*Endowment* 10, 12 (2017), 1694–1705.
- 1368 [26] Bernhard E. Boser, Isabelle M. Guyon, and Vladimir N. Vapnik. 1992. A Training Algorithm for Optimal Margin  
 1369 Classifiers. In *Proceedings of the 5th Annual Workshop on Computational Learning Theory (COLT '92)*. ACM, New  
 1370 York, NY, USA, 144–152.
- 1371 [27] Sofiane Brahim-Belhouari and Amine Bermak. 2004. Gaussian process for nonstationary time series prediction.  
 1372 *Computational Statistics & Data Analysis* 47, 4 (2004), 705–712.
- 1373 [28] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan,  
 1374 Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *Advances in Neural*  
 1375 *Information Processing Systems* 33 (2020), 1877–1901.
- 1376 [29] Laurent Callot, Mehmet Caner, A Özlem Önder, and Esra Ulaşan. 2019. A Nodewise Regression Approach to Estimating  
 1377 Large Portfolios. *Journal of Business & Economic Statistics* (2019), 1–12.
- 1378 [30] Laurent AF Callot, Anders B Kock, and Marcelo C Medeiros. 2017. Modeling and forecasting large realized covariance  
 1379 matrices and portfolio choice. *Journal of Applied Econometrics* 32, 1 (2017), 140–158.
- 1380 [31] Nicolas Chapados. 2014. Effective Bayesian modeling of groups of related count time series. In *International Conference*  
 1381 *on Machine Learning*. PMLR, 1395–1403.
- 1382 [32] Ricky TQ Chen, Yulia Rubanova, Jesse Bettencourt, and David K Duvenaud. 2018. Neural ordinary differential  
 1383 equations. *Advances in Neural Information Processing Systems* 31 (2018).
- 1384 [33] Tianqi Chen and Carlos Guestrin. 2016. XGBoost: A Scalable Tree Boosting System. In *Proceedings of the 22nd ACM*  
 1385 *SIGKDD International Conference on Knowledge Discovery and Data Mining (San Francisco, California, USA) (KDD*  
 1386 *'16)*. ACM, 785–794.
- 1387 [34] Tianqi Chen, Mu Li, Yutian Li, Min Lin, Naiyan Wang, Minjie Wang, Tianjun Xiao, Bing Xu, Chiyuan Zhang, and  
 1388 Zheng Zhang. 2015. Mxnet: A flexible and efficient machine learning library for heterogeneous distributed systems.  
 1389 *NeurIPS Workshop on Machine Learning Systems* (2015).
- 1390 [35] Yitian Chen, Yanfei Kang, Yixiong Chen, and Zizhuo Wang. 2020. Probabilistic forecasting with temporal convolutional  
 1391 neural network. *Neurocomputing* 399 (2020), 491–501.
- 1392 [36] KyungHyun Cho, Bart van Merriënboer, Dzmitry Bahdanau, and Yoshua Bengio. 2014. On the Properties of Neural  
 1393 Machine Translation: Encoder-Decoder Approaches. *CoRR abs/1409.1259* (2014).
- 1394 [37] Jan Chorowski, Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2014. End-to-end continuous speech  
 1395 recognition using attention-based recurrent NN: First results. *arXiv preprint arXiv:1412.1602* (2014).
- 1396 [38] Jan K Chorowski, Dzmitry Bahdanau, Dmitriy Serdyuk, Kyunghyun Cho, and Yoshua Bengio. 2015. Attention-based  
 1397 models for speech recognition. In *Advances in Neural Information Processing Systems*. 577–585.
- 1398 [39] Yagmur Gizem Cinar, Hamid Mirisaee, Parantapa Goswami, Eric Gaussier, Ali Ait-Bachir, and Vadim Strijov. 2017.  
 1399 Position-based content attention for time series forecasting with sequence-to-sequence RNNs. In *International*  
 1400 *Conference on Neural Information Processing*. 533–544.
- 1401 [40] Michael J Crawley. 2012. Mixed-effects models. *The R Book, Second Edition* (2012), 681–714.
- 1402 [41] J. D. Croston. 1972. Forecasting and Stock Control for Intermittent Demands. *Journal of the Operational Research*  
 1403 *Society* 23, 3 (01 Sep 1972), 289–303.
- 1404 [42] Emmanuel de Bézenac, Syama Sundar Rangapuram, Konstantinos Benidis, Michael Bohlke-Schneider, Richard  
 1405 Kurl, Lorenzo Stella, Hilaf Hasson, Patrick Gallinari, and Tim Januschowski. 2020. Normalizing Kalman Filters for  
 1406 Multivariate Time Series Analysis. *Advances in Neural Information Processing Systems* 33 (2020).
- 1407 [43] Ailin Deng and Bryan Hooi. 2021. Graph neural network-based anomaly detection in multivariate time series. In  
 1408 *Proceedings of the 35th AAAI Conference on Artificial Intelligence, Vancouver, BC, Canada*. 2–9.
- 1409 [44] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of Deep Bidirectional  
 1410 Transformers for Language Understanding. In *Proceedings of the 2019 Conference of the North American Chapter*  
 1411 *of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*.  
 1412 Association for Computational Linguistics, Minneapolis, Minnesota, 4171–4186.
- 1413 [45] I. Dimoukas, P. Mazidi, and L. Herre. 2019. Neural networks for GEFCom2017 probabilistic load forecasting.  
 1414 *International Journal of Forecasting* 35, 4 (2019), 1409 – 1423.
- 1415 [46] Laurent Dinh, Jascha Sohl-Dickstein, and Samy Bengio. 2017. Density estimation using Real NVP. In *5th International*  
 1416 *Conference on Learning Representations, ICLR 2017*.

- 1411 [47] Tobias Domhan, Jost Tobias Springenberg, and Frank Hutter. 2015. Speeding Up Automatic Hyperparameter  
1412 Optimization of Deep Neural Networks by Extrapolation of Learning Curves. In *Proceedings of the 24th International*  
1413 *Conference on Artificial Intelligence* (Buenos Aires, Argentina) (*IJCAI'15*). AAAI Press, 3460–3468.
- 1414 [48] Nan Du, Hanjun Dai, Rakshit Trivedi, Utkarsh Upadhyay, Manuel Gomez-Rodriguez, and Le Song. 2016. Recurrent  
1415 marked temporal point processes: Embedding event history to vector. In *Proceedings of the 22nd ACM SIGKDD*  
*International Conference on Knowledge Discovery and Data Mining*. ACM, 1555–1564.
- 1416 [49] James Durbin and Siem Jan Koopman. 2012. *Time series analysis by state space methods*. Vol. 38. Oxford University  
1417 Press.
- 1418 [50] Elena Ehrlich, Laurent Callot, and François-Xavier Aubet. 2021. Spliced Binned-Pareto Distribution for Robust  
1419 Modeling of Heavy-tailed Time Series. *arXiv preprint arXiv:2106.10952* (2021).
- 1420 [51] Carson Eisenach, Yagna Patel, and Dhruv Madeka. 2020. MQTransformer: Multi-Horizon Forecasts with Context  
1421 Dependent and Feedback-Aware Attention. *arXiv preprint arXiv:2009.14799* (2020).
- 1422 [52] Jesse Engel, Kumar Krishna Agrawal, Shuo Chen, Ishaan Gulrajani, Chris Donahue, and Adam Roberts. 2018.  
1423 GANSynth: Adversarial Neural Audio Synthesis. In *International Conference on Learning Representations*.
- 1424 [53] Cristóbal Esteban, Stephanie L Hyland, and Gunnar Rätsch. 2017. Real-valued (medical) time series generation with  
1425 recurrent conditional gans. *arXiv preprint arXiv:1706.02633* (2017).
- 1426 [54] Fotios Petropoulos et al. 2020. Forecasting: theory and practice. *International Journal of Forecasting* (2020).
- 1427 [55] Christos Faloutsos, Valentin Flunkert, Jan Gasthaus, Tim Januschowski, and Yuyang Wang. 2019. Forecasting Big  
1428 Time Series: Theory and Practice. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge*  
*Discovery & Data Mining, KDD 2019, Anchorage, AK, USA, August 4-8, 2019*.
- 1429 [56] Christos Faloutsos, Valentin Flunkert, Jan Gasthaus, Tim Januschowski, and Yuyang Wang. 2020. Forecasting Big  
1430 Time Series: Theory and Practice. In *Companion Proceedings of the Web Conference 2020* (Taipei, Taiwan) (*WWW '20*).  
1431 Association for Computing Machinery, 320–321.
- 1432 [57] Christos Faloutsos, Jan Gasthaus, Tim Januschowski, and Yuyang Wang. 2018. Forecasting big time series: old and  
1433 new. *Proceedings of the VLDB Endowment* 11, 12 (2018), 2102–2105.
- 1434 [58] Christos Faloutsos, Jan Gasthaus, Tim Januschowski, and Yuyang Wang. 2019. Classical and Contemporary Ap-  
1435 proaches to Big Time Series Forecasting. In *Proceedings of the 2019 International Conference on Management of Data*  
*(SIGMOD '19)*. ACM, New York, NY, USA.
- 1436 [59] Hassan Ismail Fawaz, Germain Forestier, Jonathan Weber, Lhassane Idoumghar, and Pierre-Alain Muller. 2018.  
1437 Transfer learning for time series classification. In *IEEE International Conference on Big Data, Big Data 2018, Seattle,*  
*WA, USA, December 10-13, 2018*. 1367–1376.
- 1438 [60] Valentin Flunkert, Quentin Rebjock, Joel Castellon, Laurent Callot, and Tim Januschowski. 2020. A simple and  
1439 effective predictive resource scaling heuristic for large-scale cloud applications. *arXiv preprint arXiv:2008.01215*  
1440 (2020).
- 1441 [61] Marco Fraccaro, Simon Kamronn, Ulrich Paquet, and Ole Winther. 2017. A disentangled recognition and nonlinear  
1442 dynamics model for unsupervised learning. *Advances in Neural Information Processing Systems* 30 (2017).
- 1443 [62] Jean-Yves Franceschi, Aymeric Dieuleveut, and Martin Jaggi. 2019. Unsupervised scalable representation learning for  
1444 multivariate time series. *Advances in Neural Information Processing Systems* 32 (2019).
- 1445 [63] Victor Garcia Satorras, Syama Sundar Rangapuram, and Tim Januschowski. 2022. Multivariate Time Series Forecasting  
1446 with Latent Graph Inference. *arXiv preprint* (2022).
- 1447 [64] Jan Gasthaus, Konstantinos Benidis, Yuyang Wang, Syama Sundar Rangapuram, David Salinas, Valentin Flunkert,  
1448 and Tim Januschowski. 2019. Probabilistic Forecasting with Spline Quantile Function RNNs. In *The 22nd International*  
*Conference on Artificial Intelligence and Statistics*. 1901–1910.
- 1449 [65] Andrew Gelman, John B Carlin, Hal S Stern, David B Dunson, Aki Vehtari, and Donald B Rubin. 2013. *Bayesian data*  
*analysis*. CRC press.
- 1450 [66] John Geweke. 1977. The dynamic factor analysis of economic time series. *Latent variables in socio-economic models*  
1451 (1977).
- 1452 [67] Agathe Girard, Carl Edward Rasmussen, Joaquin Quinonero Candela, and Roderick Murray-Smith. 2003. Gaussian  
1453 process priors with uncertain inputs application to multiple-step ahead time series forecasting. In *Advances in Neural*  
*Information Processing Systems*. 545–552.
- 1454 [68] Xavier Glorot, Antoine Bordes, and Yoshua Bengio. 2011. Deep sparse rectifier neural networks. In *Proceedings of the*  
*14th International Conference on Artificial Intelligence and Statistics*. 315–323.
- 1455 [69] Tilmann Gneiting, Fadoua Balabdaoui, and Adrian E Raftery. 2007. Probabilistic forecasts, calibration and sharpness.  
1456 *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 69, 2 (2007), 243–268.
- 1457 [70] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. 2016. *Deep Learning*. MIT Press. <http://www.deeplearningbook.org>.

- 1460 [71] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and  
 1461 Yoshua Bengio. 2014. Generative adversarial nets. *Advances in Neural Information Processing Systems* 27 (2014).
- 1462 [72] Adèle Gouttes, Kashif Rasul, Mateusz Koren, Johannes Stephan, and Tofiqh Naghibi. 2021. Probabilistic Time Series  
 1463 Forecasting with Implicit Quantile Networks. *arXiv preprint arXiv:2107.03743* (2021).
- 1464 [73] Rafael S. Gutierrez, Adriano O. Solis, and Somnath Mukhopadhyay. 2008. Lumpy demand forecasting using neural  
 1465 networks. *International Journal of Production Economics* 111, 2 (Feb. 2008), 409–420.
- 1466 [74] Hilaf Hasson, Bernie Wang, Tim Januschowski, and Jan Gasthaus. 2021. Probabilistic Forecasting: A Level-Set  
 1467 Approach. In *Advances in Neural Information Processing Systems*. Curran Associates, Inc.
- 1468 [75] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In  
 1469 *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 770–778.
- 1470 [76] Hansika Hewamalage, Christoph Bergmeir, and Kasun Bandara. 2021. Recurrent neural networks for time series  
 1471 forecasting: Current status and future directions. *International Journal of Forecasting* 37, 1 (2021), 388–427.
- 1472 [77] Geoffrey E. Hinton. 2002. Training Products of Experts by Minimizing Contrastive Divergence. *Neural Computation*  
 1473 14, 8 (08 2002), 1771–1800.
- 1474 [78] Geoffrey E Hinton, Simon Osindero, and Yee-Whye Teh. 2006. A fast learning algorithm for deep belief nets. *Neural  
 1475 Computation* 18, 7 (2006), 1527–1554.
- 1476 [79] Tin Kam Ho. 1995. Random decision forests. In *Proceedings of 3rd International Conference on Document Analysis and  
 1477 Recognition*, Vol. 1. IEEE, 278–282.
- 1478 [80] Sepp Hochreiter. 1998. The vanishing gradient problem during learning recurrent neural nets and problem solutions.  
 1479 *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems* 6, 02 (1998), 107–116.
- 1480 [81] Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural Computation* 9, 8 (1997), 1735–1780.
- 1481 [82] Daniel Hsu. 2017. Time series forecasting based on augmented long short-term memory. *arXiv preprint  
 1482 arXiv:1707.00666* (2017).
- 1483 [83] M. J. C. Hu and Halbert E. Root. 1964. An adaptive data processing system for weather forecasting. *Journal of  
 1484 Applied Meteorology* (1964).
- 1485 [84] Rob J Hyndman and George Athanasopoulos. 2018. *Forecasting: principles and practice*. OTexts.
- 1486 [85] Rob J Hyndman and Anne B Koehler. 2006. Another look at measures of forecast accuracy. *International Journal of  
 1487 Forecasting* (2006), 679–688.
- 1488 [86] Rob J Hyndman, Anne B Koehler, J Keith Ord, and Ralph D Snyder. 2008. *Forecasting with Exponential Smoothing: the  
 1489 State Space Approach*. Springer.
- 1490 [87] Rob J Hyndman, Earo Wang, and Nikolay Laptev. 2015. Large-scale unusual time series detection. In *IEEE International  
 1491 Conference on Data Mining Workshop (ICDMW)*. 1616–1619.
- 1492 [88] Tim Janke, Mohamed Ghanmi, and Florian Steinke. 2021. Implicit Generative Copulas. *Advances in Neural Information  
 1493 Processing Systems* 34 (2021).
- 1494 [89] Tim Januschowski, Jan Gasthaus, Yuyang Wang, Syama Sundar Rangapuram, and Laurent Callot. 2018. Deep Learning  
 1495 for Forecasting: Current Trends and Challenges. *Foresight: The International Journal of Applied Forecasting* 51 (2018),  
 1496 42–47.
- 1497 [90] Tim Januschowski, Jan Gasthaus, Yuyang Wang, David Salinas, Valentin Flunkert, Michael Bohlke-Schneider, and  
 1498 Laurent Callot. 2019. Criteria for classifying forecasting methods. *International Journal of Forecasting* (2019).
- 1499 [91] Tim Januschowski and Stephan Kolassa. 2019. A Classification of Business Forecasting Problems. *Foresight: The  
 1500 International Journal of Applied Forecasting* 52 (2019), 36–43.
- 1501 [92] Tim Januschowski, Yuyang Wang, Hilaf Hasson, Timo Erkkila, Kari Torkkila, and Jan Gasthaus. 2021. Forecasting  
 1502 with Trees. *International Journal of Forecasting* (2021).
- 1503 [93] Yunho Jeon and Sihyeon Seong. 2021. Robust recurrent network model for intermittent time-series forecasting.  
 1504 *International Journal of Forecasting* (2021).
- 1505 [94] Michael I. Jordan. 1986. *Serial Order: A parallel, distributed processing approach*. Technical Report. Institute for  
 1506 Cognitive Science, University of California, San Diego.
- 1507 [95] Michael I. Jordan. 1989. Serial Order: A Parallel, Distributed Processing Approach. In *Advances in Connectionist  
 1508 Theory: Speech*. Erlbaum.
- [96] Kelvin Kan, François-Xavier Aubet, Tim Januschowski, Youngsuk Park, Konstantinos Benidis, Lars Ruthotto, and  
 Jan Gasthaus. 2022. Multivariate Quantile Function Forecaster. In *The 25th International Conference on Artificial  
 Intelligence and Statistics*.
- [97] Tero Karras, Timo Aila, Samuli Laine, and Jaakko Lehtinen. 2017. Progressive growing of gans for improved quality,  
 stability, and variation. *arXiv preprint arXiv:1710.10196* (2017).
- [98] Tero Karras, Samuli Laine, and Timo Aila. 2019. A Style-Based Generator Architecture for Generative Adversarial  
 Networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.



- [99] Guolin Ke, Qi Meng, Thomas Finley, Taifeng Wang, Wei Chen, Weidong Ma, Qiwei Ye, and Tie-Yan Liu. 2017. LightGBM: A highly efficient gradient boosting decision tree. *Advances in Neural Information Processing Systems* 30 (2017).
- [100] Mehdi Khashei and Mehdi Bijari. 2011. A novel hybridization of artificial neural networks and ARIMA models for time series forecasting. *Applied Soft Computing* 11, 2 (2011), 2664–2675.
- [101] Diederik P. Kingma and Prafulla Dhariwal. 2018. Glow: Generative Flow with Invertible 1x1 Convolutions. In *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018*. 10236–10245.
- [102] Thomas Kipf, Ethan Fetaya, Kuan-Chieh Wang, Max Welling, and Richard Zemel. 2018. Neural relational inference for interacting systems. In *International Conference on Machine Learning*. PMLR, 2688–2697.
- [103] Roger Koenker. 2005. *Quantile Regression*. Cambridge University Press.
- [104] Stephan Kolassa. 2020. Why the “best” point forecast depends on the error or accuracy measure. *International Journal of Forecasting* 36, 1 (2020), 208–211.
- [105] Nikolaos Kourentzes. 2013. Intermittent demand forecasts with neural networks. *International Journal of Production Economics* 143, 1 (2013), 198–206.
- [106] Richard Kurlle, Syama Sundar Rangapuram, Emmanuel de Bézenac, Stephan Günnemann, and Jan Gasthaus. 2020. Deep Rao-Blackwellised Particle Filters for Time Series Forecasting. *Advances in Neural Information Processing Systems* 33 (2020).
- [107] Guokun Lai, Wei-Cheng Chang, Yiming Yang, and Hanxiao Liu. 2018. Modeling long- and short-term temporal patterns with deep neural networks. In *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval*. ACM, 95–104.
- [108] F. Laio and S. Tamea. 2007. Verification tools for probabilistic forecasts of continuous hydrological variables. *Hydrology and Earth System Sciences* 11, 4 (2007), 1267–1277.
- [109] Alex M Lamb, Anirudh Goyal, Ying Zhang, Saizheng Zhang, Aaron C Courville, and Yoshua Bengio. 2016. Professor forcing: A new algorithm for training recurrent networks. *Advances in Neural Information Processing Systems* 29 (2016).
- [110] Martin Långkvist, Lars Karlsson, and Amy Loutfi. 2014. A review of unsupervised feature learning and deep learning for time-series modeling. *Pattern Recognition Letters* 42 (2014), 11–24.
- [111] Nikolay Laptev, Jason Yosinski, Li Li Erran, and Slawek Smyl. 2017. Time-series Extreme Event Forecasting with Neural Networks at Uber. In *ICML Time Series Workshop*.
- [112] Yann LeCun. 1989. Generalization and network design strategies. In *Connectionism in perspective*.
- [113] Yann LeCun and Yoshua Bengio. 1995. Convolutional networks for images, speech, and time series. *The handbook of brain theory and neural networks* 3361, 10 (1995).
- [114] Yann LeCun, Sumit Chopra, Raia Hadsell, Fu Jie Huang, and et al. 2006. A tutorial on energy-based learning. In *Predicting Structured Data*. MIT Press.
- [115] Yann LeCun, L.D. Jackel, Leon Bottou, A. Brunot, Corinna Cortes, J. S. Denker, Harris Drucker, I. Guyon, U.A. Muller, Eduard Sackinger, Patrice Simard, and V. Vapnik. 1995. Comparison of learning algorithms for handwritten digit recognition. In *International Conference on Artificial Neural Networks*, Vol. 60. Perth, Australia, 53–60.
- [116] Shiyang Li, Xiaoyong Jin, Yao Xuan, Xiyu Zhou, Wenhu Chen, Yu-Xiang Wang, and Xifeng Yan. 2019. Enhancing the locality and breaking the memory bottleneck of transformer on time series forecasting. *Advances in Neural Information Processing Systems* 32 (2019).
- [117] Xuerong Li, Wei Shang, and Shouyang Wang. 2019. Text-based crude oil price forecasting: A deep learning approach. *International Journal of Forecasting* 35, 4 (2019), 1548 – 1560.
- [118] Yaguang Li, Rose Yu, Cyrus Shahabi, and Yan Liu. 2018. Diffusion Convolutional Recurrent Neural Network: Data-Driven Traffic Forecasting. In *International Conference on Learning Representations*.
- [119] Bryan Lim, Sercan Ö Arık, Nicolas Loeff, and Tomas Pfister. 2021. Temporal fusion transformers for interpretable multi-horizon time series forecasting. *International Journal of Forecasting* 37, 4 (2021), 1748–1764.
- [120] Bryan Lim and Stefan Zohren. 2021. Time-series forecasting with deep learning: a survey. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 379, 2194 (Feb 2021).
- [121] Kevin Lin, Dianqi Li, Xiaodong He, Zhengyou Zhang, and Ming-Ting Sun. 2017. Adversarial ranking for language generation. *Advances in Neural Information Processing Systems* 30 (2017).
- [122] Zachary C. Lipton. 2018. The Mythos of Model Interpretability. *Queue* 16, 3 (2018), 30:31–30:57.
- [123] Yucheng Low, Deepak Agarwal, and Alexander J Smola. 2011. Multiple domain user personalization. In *Proceedings of the 17th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. ACM, 123–131.
- [124] Rui Luo, Weinan Zhang, Xiaojun Xu, and Jun Wang. 2018. A neural stochastic volatility model. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 32.

- 1558 [125] Helmut Lütkepohl. 2005. Vector Autoregressive Moving Average Processes. In *New Introduction to Multiple Time*  
 1559 *Series Analysis*. Springer, 419–446.
- 1560 [126] Yisheng Lv, Yanjie Duan, Wenwen Kang, Zhengxi Li, and Fei-Yue Wang. 2014. Traffic flow prediction with big data:  
 1561 A deep learning approach. *IEEE Transactions on Intelligent Transportation Systems* 16, 2 (2014), 865–873.
- 1562 [127] Danielle C Maddix, Yuyang Wang, and Alex Smola. 2018. Deep Factors with Gaussian Processes for Forecasting.  
*arXiv preprint arXiv:1812.00098* (2018).
- 1563 [128] Spyros Makridakis, Evangelos Spiliotis, and Vassilios Assimakopoulos. 2018. The M4 Competition: Results, findings,  
 1564 conclusion and way forward. *International Journal of Forecasting* 34, 4 (2018), 802–808.
- 1565 [129] Spyros Makridakis, Evangelos Spiliotis, and Vassilios Assimakopoulos. 2018. Statistical and Machine Learning  
 1566 forecasting methods: Concerns and ways forward. *PLOS ONE* 13 (03 2018).
- 1567 [130] Spyros Makridakis, Evangelos Spiliotis, Vassilios Assimakopoulos, Zhi Chen, Anil Gaba, Ilia Tsetlin, and Robert L  
 1568 Winkler. 2021. The M5 Uncertainty competition: Results, findings and conclusions. *International Journal of Forecasting*  
 1569 (2021).
- 1570 [131] Zelda Mariet and Vitaly Kuznetsov. 2019. Foundations of sequence-to-sequence modeling for time series. In *The 22nd*  
 1571 *International Conference on Artificial Intelligence and Statistics*. PMLR, 408–417.
- 1572 [132] James E. Matheson and Robert L. Winkler. 1976. Scoring Rules for Continuous Probability Distributions. *Management*  
 1573 *Science* 22, 10 (1976), 1087–1096.
- 1574 [133] Hongyuan Mei and Jason M. Eisner. 2017. The neural hawkes process: A neurally self-modulating multivariate point  
 1575 process. In *Advances in Neural Information Processing Systems*. 6754–6764.
- 1576 [134] Olof Mogren. 2016. C-RNN-GAN: Continuous recurrent neural networks with adversarial training. *arXiv preprint*  
 1577 *arXiv:1611.09904* (2016).
- 1578 [135] Pablo Montero-Manso and Rob J Hyndman. 2020. Principles and algorithms for forecasting groups of time series:  
 1579 Locality and globality. *arXiv preprint arXiv:2008.00444* (2020).
- 1580 [136] Srayanta Mukherjee, Devashish Shankar, Atin Ghosh, Nilam Tathawadekar, Pramod Kompalli, Sunita Sarawagi, and  
 1581 Krishnendu Chaudhury. 2018. ARMDN: Associative and recurrent mixture density networks for etail demand  
 1582 forecasting. *arXiv preprint arXiv:1803.03800* (2018).
- 1583 [137] Junier Oliva, Avinava Dubey, Manzil Zaheer, Barnabas Poczos, Ruslan Salakhutdinov, Eric Xing, and Jeff Schneider.  
 1584 2018. Transformation autoregressive networks. In *International Conference on Machine Learning*. PMLR, 3898–3907.
- 1585 [138] Boris N Oreshkin, Dmitri Carпов, Nicolas Chapados, and Yoshua Bengio. 2019. N-BEATS: Neural basis expansion  
 1586 analysis for interpretable time series forecasting. *arXiv preprint arXiv:1905.10437* (2019).
- 1587 [139] Boris N Oreshkin, Dmitri Carпов, Nicolas Chapados, and Yoshua Bengio. 2020. Meta-learning framework with  
 1588 applications to zero-shot time-series forecasting. *arXiv preprint arXiv:2002.02887* (2020).
- 1589 [140] George Papamakarios, Theo Pavlakou, and Iain Murray. 2017. Masked autoregressive flow for density estimation.  
 1590 *arXiv preprint arXiv:1705.07057* (2017).
- 1591 [141] Youngsuk Park, Danielle Maddix, François-Xavier Aubet, Kelvin Kan, Jan Gasthaus, and Yuyang Wang. 2022. Learning  
 1592 Quantile Functions without Quantile Crossing for Distribution-free Time Series Forecasting. In *The 25th International*  
 1593 *Conference on Artificial Intelligence and Statistics*.
- 1594 [142] Razvan Pascanu, Tomas Mikolov, and Yoshua Bengio. 2013. On the difficulty of training recurrent neural networks.  
 1595 In *Proceedings of the 30th International Conference on Machine Learning*. Vol. 28. 1310–1318.
- 1596 [143] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming  
 1597 Lin, Natalia Gimelshein, Luca Antiga, et al. 2019. Pytorch: An imperative style, high-performance deep learning  
 1598 library. *Advances in Neural Information Processing Systems* 32 (2019).
- 1599 [144] Xueheng Qiu, Le Zhang, Ye Ren, Ponnuthurai N Suganthan, and Gehan Amaratunga. 2014. Ensemble deep learning  
 1600 for regression and time series forecasting. In *Symposium on Computational Intelligence in Ensemble Learning (CIEL)*.  
 1601 IEEE, 1–6.
- 1602 [145] Stephan Rabanser, Tim Januschowski, Valentin Flunkert, David Salinas, and Jan Gasthaus. 2020. The Effectiveness of  
 1603 Discretization in Forecasting: An Empirical Study on Neural Time Series Models. *arXiv preprint arXiv:2005.10111*  
 1604 (2020).
- 1605 [146] Syama Sundar Rangapuram, Matthias W Seeger, Jan Gasthaus, Lorenzo Stella, Yuyang Wang, and Tim Januschowski.  
 1606 2018. Deep state space models for time series forecasting. In *Advances in Neural Information Processing Systems*.  
 7785–7794.
- [147] Syama Sundar Rangapuram, Shubham Shubham Kapoor, Rajbir Nirwan, Pedro Mercado, Tim Januschowski, Yuyang  
 Wang, and Michael Bohlke-Schneider. 2022. Coherent Probabilistic Forecasting for Temporal Hierarchies.
- [148] Syama Sundar Rangapuram, Lucien D Werner, Konstantinos Benidis, Pedro Mercado, Jan Gasthaus, and Tim  
 Januschowski. 2021. End-to-end learning of coherent probabilistic forecasts for hierarchical time series. In *In-*  
*ternational Conference on Machine Learning*. PMLR, 8832–8843.
- [149] Carl Edward Rasmussen and Christopher KI Williams. 2006. *Gaussian process for machine learning*. MIT press.

- 1607 [150] Kashif Rasul, Calvin Seward, Ingmar Schuster, and Roland Vollgraf. 2021. Autoregressive denoising diffusion  
1608 models for multivariate probabilistic time series forecasting. In *International Conference on Machine Learning*. PMLR,  
1609 8857–8868.
- 1610 [151] Kashif Rasul, Abdul-Saboor Sheikh, Ingmar Schuster, Urs Bergmann, and Roland Vollgraf. 2020. Multi-variate  
1611 probabilistic time series forecasting via conditioned normalizing flows. *arXiv preprint arXiv:2002.06103* (2020).
- 1612 [152] Frank Rosenblatt. 1957. *The perceptron, a perceiving and recognizing automaton Project Para*. Cornell Aeronautical  
1613 Laboratory.
- 1614 [153] David E Rumelhart, Geoffrey E Hinton, and Ronald J Williams. 1985. *Learning internal representations by error*  
1615 *propagation*. Technical Report. University of California San Diego, La Jolla Institute for Cognitive Science.
- 1616 [154] David E. Rumelhart, Geoffrey E. Hinton, and Ronald J. Williams. 1986. Learning Representations by Back-propagating  
1617 Errors. *Nature* 323, 6088 (1986), 533–536.
- 1618 [155] David Salinas, Michael Bohlke-Schneider, Laurent Callot, and Jan Gasthaus. 2019. High-Dimensional Multivariate  
1619 Forecasting with Low-Rank Gaussian Copula Processes. In *Advances in Neural Information Processing Systems*.
- 1620 [156] David Salinas, Valentin Flunkert, Jan Gasthaus, and Tim Januschowski. 2020. DeepAR: Probabilistic forecasting with  
1621 autoregressive recurrent networks. *International Journal of Forecasting* 36, 3 (2020), 1181–1191.
- 1622 [157] Harshit Saxena, Omar Aponte, and Katie T. McConky. 2019. A hybrid machine learning model for forecasting a  
1623 billing period’s peak electric load days. *International Journal of Forecasting* 35, 4 (2019), 1288 – 1303.
- 1624 [158] Bernhard Schölkopf. 2019. Causality for machine learning. *arXiv preprint arXiv:1911.10500* (2019).
- 1625 [159] Matthias Seeger. 2004. Gaussian processes for machine learning. *International Journal of Neural Systems* 14, 02 (2004),  
1626 69–106.
- 1627 [160] Matthias W Seeger, David Salinas, and Valentin Flunkert. 2016. Bayesian intermittent demand forecasting for large  
1628 inventories. In *Advances in Neural Information Processing Systems*. 4646–4654.
- 1629 [161] Artemios-Anargyros Semenoglou, Evangelos Spiliotis, Spyros Makridakis, and Vassilios Assimakopoulos. 2021.  
1630 Investigating the accuracy of cross-learning time series forecasting methods. *International Journal of Forecasting* 37,  
1631 3 (2021), 1072–1084.
- 1632 [162] Rajat Sen, Hsiang-Fu Yu, and Inderjit S Dhillon. 2019. Think globally, act locally: A deep neural network approach to  
1633 high-dimensional time series forecasting. *Advances in Neural Information Processing Systems* 32 (2019).
- 1634 [163] Chao Shang, Jie Chen, and Jinbo Bi. 2021. Discrete Graph Structure Learning for Forecasting Multiple Time Series. In  
1635 *International Conference on Learning Representations*.
- 1636 [164] Anuj Sharma, Robert Johnson, Florian Engert, and Scott Linderman. 2018. Point process latent variable models of  
1637 larval zebrafish behavior. In *Advances in Neural Information Processing Systems*. 10919–10930.
- 1638 [165] Oleksandr Shchur, Ali Caner Turkmen, Tim Januschowski, Jan Gasthaus, and Stephan Günnemann. 2021. Detecting  
1639 Anomalous Event Sequences with Temporal Point Processes. *Advances in Neural Information Processing Systems* 34  
1640 (2021).
- 1641 [166] Oleksandr Shchur, Ali Caner Türkmen, Tim Januschowski, and Stephan Günnemann. 2021. Neural temporal point  
1642 processes: A review. *arXiv preprint arXiv:2104.03528* (2021).
- 1643 [167] David Silver, Aja Huang, Chris J Maddison, Arthur Guez, Laurent Sifre, George Van Den Driessche, Julian Schrittwieser,  
1644 Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, et al. 2016. Mastering the game of Go with deep neural  
1645 networks and tree search. *nature* 529, 7587 (2016), 484–489.
- 1646 [168] David Silver, Thomas Hubert, Julian Schrittwieser, Ioannis Antonoglou, Matthew Lai, Arthur Guez, Marc Lanctot,  
1647 Laurent Sifre, Dhharshan Kumaran, Thore Graepel, Timothy Lillicrap, Karen Simonyan, and Demis Hassabis. 2018.  
1648 A general reinforcement learning algorithm that masters chess, shogi, and Go through self-play. *Science* 362, 6419  
1649 (2018), 1140–1144.
- 1650 [169] Slawek Smyl. 2020. A hybrid method of exponential smoothing and recurrent neural networks for time series  
1651 forecasting. *International Journal of Forecasting* 36, 1 (2020), 75–85.
- 1652 [170] Slawek Smyl and N. Grace Hua. 2019. Machine learning methods for GEFCom2017 probabilistic load forecasting.  
1653 *International Journal of Forecasting* 35, 4 (2019), 1424–1431.
- 1654 [171] Ralph D. Snyder, J. Keith Ord, and Adrian Beaumont. 2012. Forecasting the intermittent demand for slow-moving  
1655 inventories: A modelling approach. *International Journal of Forecasting* 28, 2 (2012), 485–496.
- [172] Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. 2015. Deep unsupervised learning  
using nonequilibrium thermodynamics. In *International Conference on Machine Learning*. PMLR, 2256–2265.
- [173] Huan Song, Deepta Rajan, Jayaraman J Thiagarajan, and Andreas Spanias. 2018. Attend and diagnose: Clinical time  
series analysis using attention models. In *Thirty-Second AAAI Conference on Artificial Intelligence*.
- [174] Yang Song and Diederik P Kingma. 2021. How to train your energy-based models. *arXiv preprint arXiv:2101.03288*  
(2021).
- [175] Kamile Stankeviciute, Ahmed M Alaa, and Mihaela van der Schaar. 2021. Conformal Time-series Forecasting. *Advances*  
*in Neural Information Processing Systems* 34 (2021).

- 1656 [176] Souhaib Ben Taieb, James W Taylor, and Rob J Hyndman. 2021. Hierarchical probabilistic forecasting of electricity  
1657 demand with smart meter data. *J. Amer. Statist. Assoc.* 116, 533 (2021), 27–43.
- 1658 [177] Shuntaro Takahashi, Yu Chen, and Kumiko Tanaka-Ishii. 2019. Modeling financial time-series with generative  
1659 adversarial networks. *Physica A: Statistical Mechanics and its Applications* 527 (2019), 121261.
- 1660 [178] Filotas Theodosiou and Nikolaos Kourentzes. 2021. Forecasting with Deep Temporal Hierarchies. Available at SSRN:  
1661 <https://ssrn.com/abstract=3918315> or <http://dx.doi.org/10.2139/ssrn.3918315> (2021).
- 1662 [179] Jean-François Toubeau, Jérémie Bottieau, François Vallée, and Zacharie De Grève. 2018. Deep learning-based  
1663 multivariate probabilistic forecasting for short-term scheduling in power markets. *IEEE Transactions on Power Systems*  
1664 34, 2 (2018), 1203–1215.
- 1665 [180] Ali Caner Turkmen, Tim Januschowski, Yuyang Wang, and Ali Taylan Cemgil. 2021. Forecasting intermittent and  
1666 sparse time series: A unified probabilistic framework via deep renewal processes. *PlosOne* (2021).
- 1667 [181] Ali Caner Turkmen, Yuyang Wang, and Tim Januschowski. 2019. Intermittent demand forecasting with deep renewal  
1668 processes. *arXiv preprint arXiv:1911.10416* (2019).
- 1669 [182] Ali Caner Türkmen, Yuyang Wang, and Alexander J. Smola. 2019. FastPoint: Scalable Deep Point Processes. In *Joint  
1670 European Conference on Machine Learning and Knowledge Discovery in Databases*.
- 1671 [183] Aäron Van Den Oord, Sander Dieleman, Heiga Zen, Karen Simonyan, Oriol Vinyals, Alex Graves, Nal Kalchbrenner,  
1672 Andrew W Senior, and Koray Kavukcuoglu. 2016. WaveNet: A generative model for raw audio. *SSW* 125 (2016).
- 1673 [184] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia  
1674 Polosukhin. 2017. Attention is all you need. In *Advances in Neural Information Processing Systems*. 5998–6008.
- 1675 [185] François-Xavier Vialard, Roland Kwitt, Suan Wei, and Marc Niethammer. 2020. A Shooting Formulation of Deep  
1676 Learning. In *Advances in Neural Information Processing Systems*.
- 1677 [186] Chris S Wallace and David L Dowe. 2000. MML clustering of multi-state, Poisson, von Mises circular and Gaussian  
1678 distributions. *Statistics and Computing* 10, 1 (2000), 73–83.
- 1679 [187] Rui Wang, Danielle Maddix, Christos Faloutsos, Yuyang Wang, and Rose Yu. 2021. Bridging physics-based and  
1680 data-driven modeling for learning dynamical systems. In *Learning for Dynamics and Control*. PMLR, 385–398.
- 1681 [188] Yuyang Wang, Alex Smola, Danielle Maddix, Jan Gasthaus, Dean Foster, and Tim Januschowski. 2019. Deep factors  
1682 for forecasting. In *International Conference on Machine Learning*. 6607–6617.
- 1683 [189] Ruofeng Wen and Kari Torkkola. 2019. Deep Generative Quantile-Copula Models for Probabilistic Forecasting. *arXiv  
1684 preprint arXiv:1907.10697* (2019).
- 1685 [190] Ruofeng Wen, Kari Torkkola, Balakrishnan Narayanaswamy, and Dhruv Madeka. 2017. A multi-horizon quantile  
1686 recurrent forecaster. *arXiv preprint arXiv:1711.11053* (2017).
- 1687 [191] Shanika L Wickramasuriya, George Athanasopoulos, Rob J Hyndman, et al. 2015. Forecasting hierarchical and  
1688 grouped time series through trace minimization. *Department of Econometrics and Business Statistics, Monash University*  
1689 (2015).
- 1690 [192] Sifan Wu, Xi Xiao, Qianggang Ding, Peilin Zhao, Ying Wei, and Junzhou Huang. 2020. Adversarial sparse transformer  
1691 for time series forecasting. *Advances in Neural Information Processing Systems* 33 (2020), 17105–17115.
- 1692 [193] Zonghan Wu, Shirui Pan, Guodong Long, Jing Jiang, Xiaojun Chang, and Chengqi Zhang. 2020. Connecting the dots:  
1693 Multivariate time series forecasting with graph neural networks. In *Proceedings of the 26th ACM SIGKDD International  
1694 Conference on Knowledge Discovery & Data Mining*. 753–763.
- 1695 [194] Shuai Xiao, Mehrdad Farajtabar, Xiaojing Ye, Junchi Yan, Le Song, and Hongyuan Zha. 2017. Wasserstein learning of  
1696 deep generative point process models. In *Advances in Neural Information Processing Systems*. 3247–3257.
- 1697 [195] Shuai Xiao, Junchi Yan, Mehrdad Farajtabar, Le Song, Xiaokang Yang, and Hongyuan Zha. 2017. Joint modeling of  
1698 event sequence and time series with attentional twin recurrent neural networks. *arXiv preprint arXiv:1703.08524*  
1699 (2017).
- 1700 [196] Qifa Xu, Xi Liu, Cuixia Jiang, and Keming Yu. 2016. Quantile autoregression neural network model with applications  
1701 to evaluating value at risk. *Applied Soft Computing* 49 (2016), 1–12.
- 1702 [197] Xing Yan, Weizhong Zhang, Lin Ma, Wei Liu, and Qi Wu. 2018. Parsimonious quantile regression of financial asset  
1703 tail dynamics via sequential learning. *Advances in Neural Information Processing Systems* 31 (2018).
- 1704 [198] Jaemin Yoo and U Kang. 2021. Attention-Based Autoregression for Accurate and Efficient Multivariate Time Series  
Forecasting. In *Proceedings of the 2021 SIAM International Conference on Data Mining (SDM)*. SIAM, 531–539.
- [199] Jinsung Yoon, Daniel Jarrett, and Mihaela Van der Schaar. 2019. Time-series generative adversarial networks. *Advances  
in Neural Information Processing Systems* 32 (2019).
- [200] Hsiang-Fu Yu, Rao N., and I.S. Dhillon. 2016. Temporal regularized matrix factorization for high-dimensional time  
series prediction. *Advances in Neural Information Processing Systems* (2016), 847–855.
- [201] Guoqiang Zhang, B Eddy Patuwo, and Michael Y Hu. 1998. Forecasting with artificial neural networks: The state of  
the art. *International journal of forecasting* 14, 1 (1998), 35–62.

- 1705 [202] G Peter Zhang. 2003. Time series forecasting using a hybrid ARIMA and neural network model. *Neurocomputing* 50  
1706 (2003), 159–175.
- 1707 [203] Han Zhang, Ian Goodfellow, Dimitris Metaxas, and Augustus Odena. 2019. Self-attention generative adversarial  
1708 networks. In *International Conference on Machine Learning*. PMLR, 7354–7363.
- 1709 [204] Haoyi Zhou, Shanghang Zhang, Jieqi Peng, Shuai Zhang, Jianxin Li, Hui Xiong, and Wancai Zhang. 2020. Informer:  
1710 Beyond Efficient Transformer for Long Sequence Time-Series Forecasting. *arXiv preprint arXiv:2012.07436* (2020).
- 1711 [205] Lingxue Zhu and Nikolay Laptev. 2017. Deep and confident prediction for time series at Uber. In *IEEE International  
1712 Conference on Data Mining Workshops (ICDMW)*. 103–110.
- 1713 [206] Daniel Zügner, François-Xavier Aubet, Victor Garcia Satorras, Tim Januschowski, Stephan Günnemann, and Jan  
1714 Gasthaus. 2021. A study of joint graph inference and forecasting. *arXiv preprint arXiv:2109.04979* (2021).
- 1715
- 1716
- 1717
- 1718
- 1719
- 1720
- 1721
- 1722
- 1723
- 1724
- 1725
- 1726
- 1727
- 1728
- 1729
- 1730
- 1731
- 1732
- 1733
- 1734
- 1735
- 1736
- 1737
- 1738
- 1739
- 1740
- 1741
- 1742
- 1743
- 1744
- 1745
- 1746
- 1747
- 1748
- 1749
- 1750
- 1751
- 1752
- 1753