

Streaming Adaptation of Deep Forecasting Models using Adaptive Recurrent Units

Prathamesh Deshpande
IIT Bombay
pratham@cse.iitb.ac.in

Sunita Sarawagi
IIT Bombay
sunita@iitb.ac.in

ABSTRACT

We present ARU, an Adaptive Recurrent Unit for streaming adaptation of deep globally trained time-series forecasting models. The ARU combines the advantages of learning complex data transformations across multiple time series from deep global models, with per-series localization offered by closed-form linear models. Unlike existing methods of adaptation that are either memory-intensive or non-responsive after training, ARUs require only fixed sized state and adapt to streaming data via an easy RNN-like update operation. The core principle driving ARU is simple — maintain sufficient statistics of conditional Gaussian distributions and use them to compute local parameters in closed form. Our contribution is in embedding such local linear models in globally trained deep models while allowing end-to-end training on the one hand, and easy RNN-like updates on the other. Across several datasets we show that ARU is more effective than recently proposed local adaptation methods that tax the global network to compute local parameters.

KEYWORDS

Local-global models; Online Adaptation; Time-series Forecasting

ACM Reference Format:

Prathamesh Deshpande and Sunita Sarawagi. 2019. Streaming Adaptation of Deep Forecasting Models using Adaptive Recurrent Units. In *The 25th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD '19)*, August 4–8, 2019, Anchorage, AK, USA. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3292500.3330996>

1 INTRODUCTION

Time series forecasting is a critical analysis tool in retail [9, 15], finance, and commerce, and has been extensively researched across communities for multiple decades [1, 5, 6, 11, 13, 21, 27] (See [8] for a recent survey). However, with the remarkable success of deep learning in other classically hard tasks like speech recognition, the time series forecasting problem also requires revisiting. We consider typical forecasting settings where we are given multiple series along different points in time. Each series at time t is characterized by a real-valued output y_t^i and a vector of input features \mathbf{x}_t^i . Our goal is

to predict the y for K future points of time for which we are given inputs $\mathbf{x}_{T+1}^i, \dots, \mathbf{x}_{T+K}^i$.

Classical methods of time series analysis like ARIMA, exponential smoothing, and other state space models [13], train separate parameters θ^i for each time series i and can be considered as local models. In contrast recent deep learning methods [11, 18, 29] train shared parameters θ for predicting each y as a function of a summary of each time series' previous observation. Typically, a Recurrent Neural Networks (RNNs) is used for representing such sequential history as fixed dimensional vectors. For sequences representing time series, in addition to general purpose gated RNNs like LSTMs, recent proposals also include un-gated RNNs like Statistical Recurrent Units (SRUs [19]) and Fourier Recurrent Units (FRU [31]). During training, data from multiple time series teaches the RNN to capture enough context from the history relevant for making predictions at each of the future K points in time. Such models have been found to surpass purely local models like ARIMA and conventional globally trained models like boosted regression trees [21].

In general however such global model may be sub-optimal when the various series are heterogeneous, or when not all time series are observed during training. The only local influence in these models is the RNN state that evolves with time separately for each time series i . We take a different view and treat the per-series state as *local parameters* that map a future input along with its context to a prediction. We use a clever trick to update the local parameters in closed form, and in a streaming manner much like normal RNN parameters. We call our unit a Adaptive Recurrent Unit (ARU) whose finite-sized state can be used to compute per-series local parameters in closed form.

Recently, other methods have also been proposed to augment a global time series model with local parameters [12, 23]. However, their method of local adaptation requires the global model to be retrained for each update to a time series. More generally, localization of a globally trained model can be treated as a domain adaptation problem [3, 4] for which many methods have been proposed including model fine-tuning, meta-learning [10, 17, 24], and memory-augmented networks [26, 28]. These methods either require gradient-based iterative updates, or are memory-intensive, and/or rely on techniques like self-attention that is quadratic in T for large time-series. ARU does not entail iterative training of local parameters, its adaptation is light-weight and streaming, and only requires a time-series specific state of constant size. This makes it particularly efficient for long time-series since its storage and compute requirement is constant in the sequence length.

The core principle of ARU is to use RNN-like updates to incrementally maintain per-series sufficient statistics required for fitting

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](https://permissions.acm.org).

KDD '19, August 4–8, 2019, Anchorage, AK, USA

© 2019 Association for Computing Machinery.

ACM ISBN 978-1-4503-6201-6/19/08...\$15.00

<https://doi.org/10.1145/3292500.3330996>

a local conditional Gaussian distribution in closed form. Our contribution is in exploiting this classical trick in the context of modern deep learning models, and designing network architectures that provide the right local-global tradeoff. Surprisingly, we show that this strategy of local adaptation is more robust across different types of time-series than recently proposed more ambitious methods that tax the deep network to learn to generate local parameters from RNN states.

We compare our ARU-based adaptation to two state-of-art adaptation methods on five time-series datasets under various settings. We show that ARU is effective in reducing error by 10 to 20% compared to the baseline, and is generally better than existing adaptation methods. This reduction in error comes at very little overheads of running time, whereas existing state-of-the-art adaptation methods are up to a factor of four slower.

2 OUR MODEL

We start with a review of state-of-the-art global models and classical local models for time series forecasting.

Review: Global Models. Typical deep learning based global models for multi-horizon time series forecasting [11, 29] deploy the encoder-decoder architecture. First an input layer maps the input features \mathbf{x}_t^i to a real vector. This could include embedding lookups for categorical attributes and any rescaling for continuous attributes. Next the transformed input along with the previous output $[\mathbf{x}_t^i, y_{t-1}^i]$ is fed to one or more encoder RNN layers. The RNN could be either a generic LSTM or time-series specific units like SRU or FRU that capture context useful for future predictions. The output of the encoder is its final state \mathbf{g}_T^i at the end of T steps. This can be treated as a summary of the known y values that is relevant as a *context* for future predictions. The decoder initialized with \mathbf{g}_T^i is responsible for making the predictions on the K future time horizons as a function of respective inputs $\mathbf{x}_t^i : t = T + 1 \dots T + K$. The decoder could be auto-regressive, where the previous predicted y is fed as input to be next step, or independently make each of the K predictions. We found the independent model to provide higher accuracy for long-term forecasts than the auto-regressive model fed with noisy previous predictions. This is corroborated in [29]. The decoder takes each transformed future input $\mathbf{x}_t^i : t = T + 1 \dots T + K$ concatenated with \mathbf{g}_T^i and generates an output vector \mathbf{h}_t^i using an optional RNN and one or more feed-forward layers. Lastly, a Gaussian distribution is imposed on the output by using a linear layer to map \mathbf{h}_t^i into a mean and variance. The final equations driving the prediction at future times are:

$$\mathbf{g}_T^i = \text{RNN}([\mathbf{y}_{t-1}^i, \mathbf{x}_t^i] : t = 1 \dots T | \theta_{\text{enc}}) \quad (1)$$

$$\mathbf{h}_t^i = \text{FF}([\mathbf{g}_T^i, \mathbf{x}_t^i] : t = T + 1 \dots T + K | \theta_{\text{dec}}) \quad (2)$$

$$\mu_t^i = \theta_\mu[\mathbf{h}_t^i, 1], \sigma_t^i = \log(1 + \exp(\theta_\sigma[\mathbf{h}_t^i, 1])) \quad (3)$$

$$\Pr(y_t^i | \mathbf{x}_t^i, (\mathbf{x}_1^i, y_1^i), \dots, (\mathbf{x}_T^i, y_T^i)) = \mathcal{N}(\mu_t^i, \sigma_t^i) \quad (4)$$

where $\theta_{\text{enc}}, \theta_{\text{dec}}, \theta_\mu, \theta_\sigma$ are all parts of the global parameters θ . During training we are given data comprising of input-output pairs over several time series. $D = \{(\mathbf{x}_1^i, y_1^i), \dots, (\mathbf{x}_T^i, y_T^i) : i = 1 \dots N\}$. We simulate multiple encoder-decoder windows from this history using a sliding window of stride ℓ and compute the data likelihood

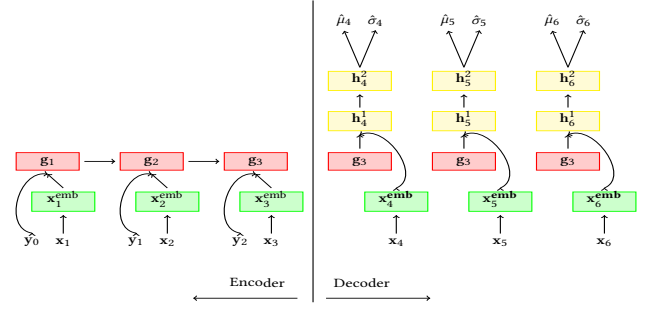


Figure 1: Diagram of the global model with encoder size of 3 and decoder size of 3. All the nodes/states with the same label are copies.

as follows:

$$\max_{\theta} \sum_{i=1}^N \sum_{T'=2:\ell}^{T-K} \sum_{t=T'+1}^{T'+K} \log \mathcal{N}(y_t^i; \mu_t^i, \sigma_t^i | \theta)$$

where \mathcal{N} denotes the Gaussian density function. The parameter θ includes all network parameters spanning the input-layer, the encoder RNN, the decoder RNN, and the last output layer that generates μ_t^i, σ_t^i . Figure 1 shows the global model.

The global model is driven by parameters trained across multiple time-series, and the only local influence in the final equation that outputs $\Pr(y_t^i | \mathbf{x}_t^i, \{(\mathbf{x}_s^i, y_s^i)\}_{s=1}^T)$ is the RNN state \mathbf{g}_T^i that serves as context. The RNN parameters are trained end-to-end to find the most relevant context.

Review: Local Models. A local model would train separate parameters θ^i for each i . However, unless the length of each time-series is very large, training parameters of a complex network like a multi-layer deep network would not work. Hence, local models have traditionally been simple models such as linear state space models that use linear parameters to transition from one state to the next, and make state-ful predictions. Let \mathbf{s}_t^i denote the local state at time t of series i . Parameters θ_{tr}^i control transition from one state to the other via affine transforms, and parameter θ_μ^i controls the linear transform to generate the mean output from the local state. The variance is a fixed learned parameter, and that makes the local model as:

$$\mathbf{s}_t^i = \theta_{\text{tr}}^i[\mathbf{s}_{t-1}^i, \mathbf{x}_t^i, 1] \quad (5)$$

$$\mu_t^i = \theta_\mu^i[\mathbf{s}_t^i, 1], \sigma_t^i = \theta_\sigma^i \quad (6)$$

where the parameters θ_{tr}^i and $\theta_\mu^i, \theta_\sigma^i$ are learned locally for each time-series using its known values up to T that is $(\mathbf{x}_1^i, y_1^i), \dots, (\mathbf{x}_T^i, y_T^i)$. Assuming T is large, local models can capture peculiarities of individual time-series in ways that cannot be approximated by local RNN state of global models.

2.1 The ARU RNN

Our model not only attempts to capture the best of the local and global models, but does so in a manner that allows the local parameters to be computed recurrently in a streaming manner. We call our

unit Adaptive Recurrent Unit (ARU) that views the per-series state as local predictive parameters that are fitted on the fly. Additionally, ARU is designed to capture local predictive parameters without expensive parameter retraining.

We use the shared global network to learn the complex transformation of the history and the input into a final output vector \mathbf{h}_t^i . However, instead of the global last linear layer (Equation 3 in the global model), we use a local model. We perform the local fit in a clever way to allow easy end-to-end training of the global network parameters without inner loop of local updates as in existing methods like [10]. The main idea behind the ARU is to use a RNN to provide the best possible least square fit locally for each time series based on observed $\mathbf{x}_t, \mathbf{h}_t$ values. The motivation behind this choice is that the optimal local parameters can be obtained in closed form using sufficient statistics that can be maintained in a streaming manner. In contrast local parameters that depend on iterative gradient-based updates, require multiple data passes and are difficult to embed within a larger global network that is trained via gradient descent. Although the ARU unit itself only models linear interaction between its inputs and outputs, it is embedded in a larger neural network that provides non-linear transformation of both the input to the ARU and its output. We describe the ARU-RNN in this section and then in Section 2.2 describe how we fit it in the global model.

2.1.1 The ARU Update Equations. Let \mathbf{h}_t^i be the output vector at time t from the global model. The ARU RNN has two modes: adapt and predict. In the adapt mode it is fed the input \mathbf{h}_t^i along with the true label y_t^i at that time and produces an updated state \mathbf{s}_t^i . In the predict mode denoted $\text{ARU}(\mathbf{s}_{t-1}^i, \mathbf{h}_t^i)$ it outputs the local prediction based on state \mathbf{s}_{t-1}^i on the input \mathbf{h}_t^i .

The ARU state \mathbf{s}_t at each time t keeps four types of sufficient statistics $\mathbf{s}_t = [\mathbf{sxx}_t^i, \mathbf{sxy}_t^i, \mathbf{sn}_t^i, \mathbf{ss}_t^i]$ to provide a local least square fit between the \mathbf{h} and y s. A vector α of J aging factors maintains this statistics for varying amounts of aging of the old data. Initially, all states are zero and at each time-step it is updated as follows:

$\text{ARU_Update}(\mathbf{s}_{t-1}^i, \mathbf{h}_t^i, y_t^i)$:

$$\mathbf{sxx}_t^i = \alpha \mathbf{sxx}_{t-1}^i + [\mathbf{h}_t^i \ 1]^T [\mathbf{h}_t^i \ 1]$$

$$\mathbf{sxy}_t^i = \alpha \mathbf{sxy}_{t-1}^i + [\mathbf{h}_t^i \ 1]^T (y_t^i)$$

$$\mathbf{sn}_t^i = \alpha \mathbf{sn}_{t-1}^i + 1$$

$$\mathbf{ss}_t^i = \alpha \mathbf{ss}_{t-1}^i + (y_t^i - \text{ARU}(\mathbf{s}_{t-1}^i, \mathbf{h}_t^i))^2$$

$$\text{Return } \mathbf{s}_t^i = [\mathbf{sxx}_t^i, \mathbf{sxy}_t^i, \mathbf{sn}_t^i, \mathbf{ss}_t^i]$$

From these update equations, it is easy to interpret what each of the ARU state components represent. The state \mathbf{sxx}^i represents age-weighted sum of pairwise feature product of the input vector, \mathbf{sxy}^i represents age-weighted sum of product of input features and output y , \mathbf{sn}_t^i represents the age adjusted count so far, and \mathbf{ss}_t^i represents the accumulated fitting noise that will serve as the variance from the local predictions. In this equation $\text{ARU}(\mathbf{s}_{t-1}^i, \mathbf{h}_t^i)$ denotes the mean local prediction as explained below. The ARU state is updated in a streaming manner every time a true y_t^i is known.

We next provide the equations of ARU in the predict mode. $\text{ARU}(\mathbf{s}_{t-1}^i, \mathbf{h}_t^i)$ represents the predicted values from the ARU at




Figure 2: ARU cell combined with the decoder of the global model. Showing only the mean computation here.

state \mathbf{s}_{t-1}^i on the input \mathbf{h}_t^i . The predicted output using only ARU states can be computed by exploiting the closed form of the least squares solution. First, the local parameter is calculated in closed form using the sufficient statistics. The λI term regularizes the local parameters. The local parameter is used to compute the mean and variance of the local prediction.

$\text{ARU_Predict}(\mathbf{s}_t^i, \mathbf{h}_t^i)$:

$$\theta_{t,\mu}^i = (\mathbf{sxx}_t^i + \lambda I)^{-1} \mathbf{sxy}_t^i, \quad \theta_{t,\sigma}^i = \mathbf{ss}_t^i / \mathbf{sn}_t^i \quad (7)$$

$$\mathbf{m}_t^i = \theta_{t,\mu}^i [\mathbf{h}_t^i \ 1], \quad \mathbf{a}_t^i = \theta_{t,\sigma}^i \quad (8)$$

where \mathbf{m}_t^i denotes the mean local prediction and \mathbf{a}_t^i denotes the local variance. The ARU fits local parameters, which can be used to compute the mean prediction (and variance) much like in the local model in Equation 5. However, unlike in local models, we do not output these directly as predictions. Instead, we exploit the availability of multiple time series to further combine these local predictions with global parameters as follows.

2.2 ARU in Global Model

We concatenate the local mean and variance predictions with the final vector \mathbf{h}_t^i and use a two-layer feed forward network to transform them as follows.

$$\mu_t^i = \theta_\mu(\text{FF2}[\mathbf{h}_t^i, \mathbf{m}_t^i, 1], \sigma_t^i = \log(1 + \exp(\theta_\sigma(\text{FF2}[\mathbf{h}_t^i, \mathbf{a}_t^i, 1]))) \quad (9)$$

Figure 2 shows how ARU is embedded in the global model for one decoder step. We only show mean computation in the figure. The variance computation can also be handled similarly. The output \mathbf{m}_4 obtained from ARU is combined with hidden layer and passed further to a two layer feed-forward network. The FF2 network further evaluates the importance of the global output \mathbf{h}_t^i and local ARU output \mathbf{m}_t^i and makes the final predictions.

3 RELATED WORK

Before we present our empirical evaluation we discuss related work that could have served as alternatives to our method of adaptation.

Generate Local Parameters from RNN State. Recently [12] and [23] propose to use the power of deep learning to directly compute local parameters θ^i from the local state of the RNN. As a state-of-art representative of this class of methods, we discuss the DeepState method of [23]. In DeepState the RNN state \mathbf{g}_t^i computed from global parameters (Eq 1) is passed through feed-forward networks to directly output the local parameters θ_{tr}^i and $\theta_{\mu}^i, \theta_{\sigma}^i$, that is,

$$[\theta_{tr}^i, \theta_{\mu}^i, \theta_{\sigma}^i] = FF(\mathbf{g}_t^i; \theta_{meta}) \quad (10)$$

where θ_{meta} are additional global parameters. The computed local parameters are used to generate the mean and variance of the output y_s using a linear state-space model (e.g. Equation 5). The global parameters are trained via joint likelihood on all training examples. Unlike our approach, these methods cannot adapt to new time series since local y -s are only transferred via the joint likelihood based training. In our experiment section we show that these models do not train well unless the length of each time-series is long.

Dedicated Local parameters jointly trained with Global. An approach often used for localization of global models is to dedicate per-locale parameters that play specific role in the global model and are trained jointly end-to-end. For example [25] uses per-document parameter along with global parameters to develop better embeddings. In our case, we would train time-series specific parameters θ^i along with the rest of the global parameters. This method is only applicable when all time-series are known during training time. A downside of this approach is that it cannot be easily evolved when new values are observed for a time series. Also, each time series has to be long enough for the local parameters to train well.

Fine-tune Global Parameters for each series. Another approach is to treat localization as a problem of domain adaptation, for which a huge literature exists [3, 4]. A well-known solution is to fine-tune the global parameters on labeled data of each time-series using gradient descent on the loss over the limited labeled data for each series [24]. This method would require storage of separate local parameters for each series. Also, parameter fine-tuning could lead to unpredictable local performance. This has led to an explosion of meta-learning methods that exploit multiple local datasets during training so as to learn the adaptation process [10, 17]. Another method of adaptation is using memory [22, 28] that combines parameter fine-tuning with memory-based recall. Of these a recent state-of-art approach is the SNAIL model that uses a simple method of learned deep self-attention on local data for adaptation. As a state-of-art representative of learned adaptation methods, we will compare with this method in our experiments. Our approach is closed-form self-attention that exploits the special regression form of our prediction function. Self-attention based models for time series forecasting are also explored in [20].

For theoretical insights on the learning efficiency of global vs local models see [14].

Special RNNs for Time Series. We designed ARU so that it can be embedded in a global network much like any other RNN. Recently, [19] and [31] also propose special RNNs for time-series. Oliva et al. [19] proposed an un-gated alternative to LSTMs, called the Statistical Recurrent Unit (SRU). SRUs maintain moving averages of

Dataset	N	T	K	Enc.Len	#Features
Rossmann	1115	1600	16	16	39
Walmart	3331	143	8	8	16
Electricity	370	44000	24	168	5
Traffic	963	2100	24	168	3
Parts	2246	52	8	8	1

Table 1: Statistics of the datasets used in our experiments. N is the number of time series, T is the number of values in a series, K is the forecast horizon we used in our experiments, Enc.Len is the length of the encoder in each prediction instance.

summary statistics in their hidden states and generalize exponential moving average like statistics. However, they treat the SRU state like 'context' to be used for predictions in place of LSTM states. Fourier Recurrent Units FRUs [31] is another such un-gated time-series unit. In contrast, we convert the ARU state to local parameters by exploiting the closed-form fit from the the least square sufficient statistics.

4 EXPERIMENTS

We now empirically compare our ARU-based streaming adaptation method to state-of-the-art deep learning based global models as baselines, and to two recent approaches of adapting them to local time-series. In addition to five real-life public time-series datasets, we also present controlled experiments on synthetic datasets to gain insights on different methods of localizing deep global models. Our code and experiments can be found at ¹.

This section is organized as follows: Section 4.1 describes the datasets, Section 4.2 describes the five methods we compare with, Section 4.3 describes detailed setup of the experiments, Section 4.4 presents experiments on the synthetic dataset to qualitatively understand the difference between DeepState and ARU, and Section 4.5 presents anecdotes, accuracy, and running times on the real datasets.

4.1 Datasets

We perform our experiments on five publicly available time-series datasets. The first two datasets we collected from Kaggle contains rich x features at each point and are more aligned with typical retail forecasting setups where the deep models have gained most traction recently. The remaining three contain only derived time features and are included for comparison with the existing literature. A summary of the datasets appears in Table 1.

Rossmann The Rossmann dataset² contains daily sales history of 1115 Rossmann stores. The data is collected between 1st Jan. 2013 to 31st July 2015, i.e., roughly 900 values per series. Each time series is associated with various co-variate features. These include predetermined store-specific features such as store type and distance to the nearest competing store and time-varying features such as whether the store is running any promotion, and external

¹<https://github.com/pratham16/ARU.git>

²<https://www.kaggle.com/c/rossmann-store-sales/data>

Set	# RNN units	Hidden layer sizes
Small	8	8,6,6
Medium	16	16,15,10
Large	50	32,20,15

Table 2: Hyperparameter sets

features such as weather conditions, is current day a school holiday or a state holiday, etc.

Walmart The Walmart³ dataset maintains weekly sales in each department of 45 walmart stores from Feb 2010 to Oct 2010. Here, (store, department) tuple uniquely identifies each time-series, resulting in 3331 time-series. Each time-series contains roughly 143 values. This dataset also provides promotional features and external features such as temperature, unemployment, consumer price index, major holidays in a given week etc.

Electricity The electricity⁴ dataset contains hourly energy consumption of 370 houses (in kW) from 1st Jan. 2011 to 31st Dec. 2015, a total of 44,000 values per time-series.

Traffic The Traffic⁵ dataset contains hourly occupancy rates of 963 car lanes in San Francisco bay area. The occupancy rates are in the range [0, 1]. The data is collected between 1st Jan. 2008 to 30th March 2009, i.e., 10560 values per series.

Parts This dataset [7] contains monthly sales information of 2246 car parts over 52 months. This is a small dataset compared to others in our collection.

Features. For each dataset, we extract time-related features in addition to any features available in the data (e.g. in Rossman and Walmart). For Walmart, we use month of the year and week of the year features. For Rossman, we use day of the month and month of the year features. Since Parts is a monthly dataset, we use month of the year feature for Parts. Since Electricity and Traffic are hourly datasets, we use hour of the day and day of the week features. In addition to these, we use month of the year feature for electricity dataset. We embed categorical features into real vectors and rescale the real valued attributes between 0 and 1 across the time-series.

4.2 Methods

We compare ARU against two state-of-the-art global deep learning models as baselines and two most recent local adaptation methods. We drop comparison with classical local-only models such as ARIMA because the DeepState and DeepAR methods that we compare with have been shown to conclusively surpass them.

Baseline As a baseline we use the globally trained model where the decoder makes independent predictions for each of the future K points in time. Note that the encoder used to compute the per-series state \mathbf{g}_T^i continues to be auto-regressive. Our experiments and [29] have found the independent decoder to yield better performance than an auto-regressive decoder. This baseline is part of our code base and uses exactly the same setup as ARU.

³<https://www.kaggle.com/c/walmart-recruiting-store-sales-forecasting/data>

⁴<https://archive.ics.uci.edu/ml/datasets/ElectricityLoadDiagrams20112014>

⁵<https://archive.ics.uci.edu/ml/datasets/PEMS-SF>

DeepAR As another global model, we use DeepAR, a production quality deep forecasting model that is built-in Amazon’s SageMaker machine learning service⁶. As described in [11], DeepAR is also an encoder-decoder model but their decoder is auto-regressive. Since the source code is not public, further details of specific features, scaling, etc that they deploy is unknown. To the best of our ability we tried to match the features and experiment setup with our global baseline but we cannot be sure without access to the source code.

SNAIL The SNAIL method proposed in [17] is a local adaptation model that uses interleaving of dilated causal convolution and self-attention layers to capture long range dependencies and for localization respectively. We use two convolution layers with dilation rates 2 and 4 respectively, and two self-attention layers.

DeepState This model uses an RNN to predict the parameters of a local state-space model [23] and is currently the best known local adaptation model for time-series forecasting. The source code is not publicly available, so we only compare with their published numbers.

4.3 Experiment setup

We use a single RNN layer in the encoder. The decoder has three ReLU layers on a concatenation of the encoder state \mathbf{g}_T^i and input features \mathbf{x}_t^i . A skip connection from \mathbf{x}_{T+k}^i to the second ReLU layer is added. The batch size for all experiments is set to 64. Adam optimizer is used with learning rate 0.0001. ARU regularization λ and aging vector α are chosen based on validation loss. In all our experiments, $\alpha \in \{1.0, 0.99, 0.95, 0.9\}^J$. In all experiments, best performing value of α turns out to be [1.0], except for Traffic dataset in Table 5, where best performing $\alpha = [1.0, 0.95, 0.9]$. For number of RNN units and number of hidden units, we use three sets of hyper-parameter settings for baseline and ARU – Small, Medium, and Large as given in Table 2. We use the Small configuration for Parts since it contains time-series of length 52. For Rossman and Walmart, we use the Medium configuration. For Electricity and Traffic datasets, we use the Large configuration in Table 5 and the Medium configuration in the rest of the experiments. We rescale y values of each time-series with its average y value as follows:

$$\bar{y}_t^i = 1 + \frac{1}{T} \sum_{t'=1}^T y_{t'}^i$$

This rescaling is also used by [11]. On the Walmart dataset since each time-series is small, we perform above rescaling on a batch instead of the entire time-series i.e. $\bar{y}_t^i = 1 + \frac{1}{E} \sum_{t'=1}^E y_{t'}^i$, where E is encoder length.

We also experimented with SRU [19], which is an un-gated alternative to LSTMs. Since we obtained similar results with SRUs, we omit those numbers from our experiments.

In addition to RMSE, we report Normalized Deviation (ND), which is used in [11, 30] as an evaluation metric. It is defined as:

$$\text{ND}(y_t^i, \hat{y}_t^i) = \left(\sum_{i=1}^N \sum_{t=T+1}^{T+K} |y_t^i - \hat{y}_t^i| \right) / \left(\sum_{i=1}^N \sum_{t=T+1}^{T+K} |y_t^i| \right)$$

⁶<https://docs.aws.amazon.com/sagemaker/latest/dg/deepar.html>

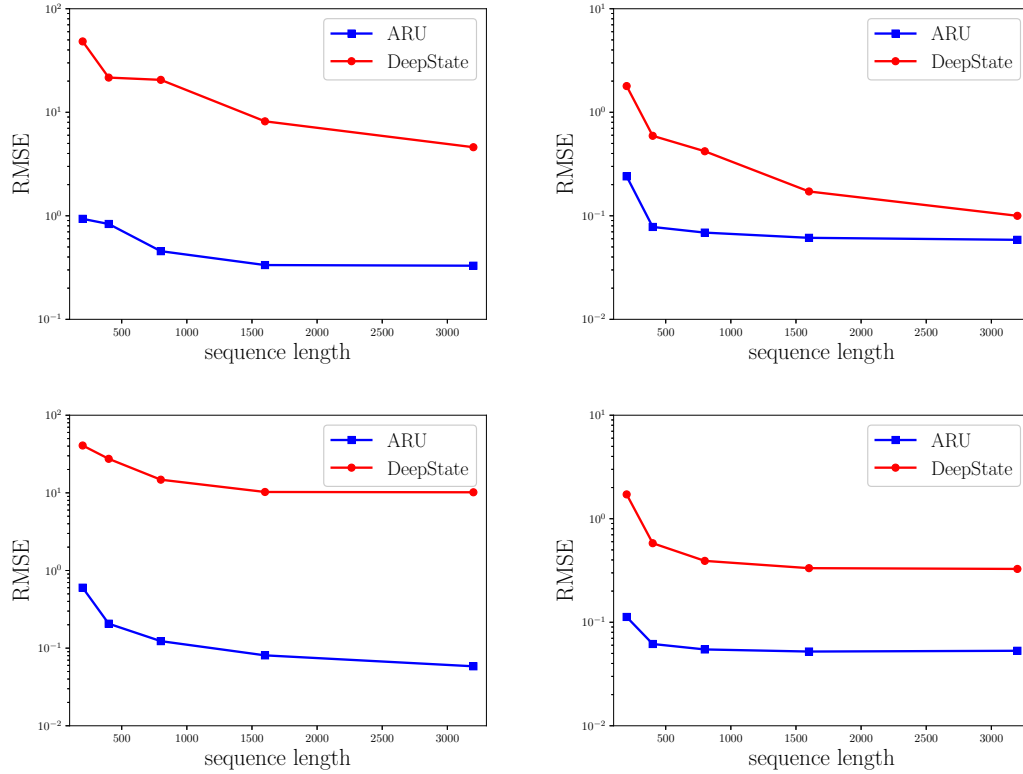


Figure 3: Performance of ARU and DeepState on synthetic data. The columns correspond to different values of γ ($\gamma = 20$ first column, $\gamma = 1$ second column). The first row are with time-series ids and second row are without time-series ids.

4.4 Qualitative Comparison on Synthetic Data

In order to get an insight into the working of ARU vis-a-vis the most recent DeepState method of local adaptation, we first present a qualitative comparison using a very simple synthetic dataset.

We generate 10 synthetic time-series of hourly data. The y values for these time-series are obtained as a linear combination of x features and local parameters θ_μ^i , which are uniformly sampled in the range $[-\gamma, \gamma]$. We use only hour of day and day of week as x features and generate time-series of different lengths by sampling $y \sim \theta_\mu^i[x_t \ 1] + \epsilon$ where $\epsilon \sim \mathcal{N}(0, 1)$ is a Gaussian noise with unit variance. We compute RMSE for predictions on last 24 hour forecasts in each time-series. This non-auto-regressive model might look too simplistic but it helps us bring out the essence of different approaches. Also, after conditioning on the encoder-state, we found non-auto-regressive decoder to be superior to the auto-regressive decoder even on real-life datasets. In this simple case, the state-space model would need to learn to generate θ_μ^i from its RNN state g_t^i in order to match the generative process exactly. In contrast, the ARU would compute θ_μ^i from the data whose sufficient statistics is summarized in the ARU state. The global network just needs to learn to provide the right input to the ARU cell and also correctly integrate its output with the global output. We used the Medium setting (Table 2) for the global network without any special customization for the synthetic experiments. Thus, the global network

needs to learn the correct meta-learning strategy for both methods of adaptation. When id i is part of \mathbf{x}_t , in theory, a powerful enough network can memorize each set of the 10 local parameter in its feed-forward layers, making them both identical. We therefore compare these models under two settings: one where each time-series i feeds its unique id i to the network, and second where it does not.

We compare the two methods against increasing length of time-series in Figure 3 for two different values of γ (20 and 1) both with and without time-id. When γ is large (first column) the local parameters differ a lot from one time series to the next and the DeepState method of computing them from the RNN state is not at all effective. With increasing sequence length, errors of both methods reduce but ARU converges much faster to error values less than 0.1. When local time-series are more similar to each other ($\gamma = 1$), DeepState starts matching ARU only after seeing many samples per time-series ($\gamma = 1$, with time-series id). Also, parameter generation methods like DeepState also crucially exploit the time-series specific identifiers to generate local parameters (compare red plots across the two rows). This implies that they memorize local parameters of the finite number of training time-series, and cannot perform adaptation of a series not seen during training. Also, they cannot exploit streaming availability of labeled data after training.

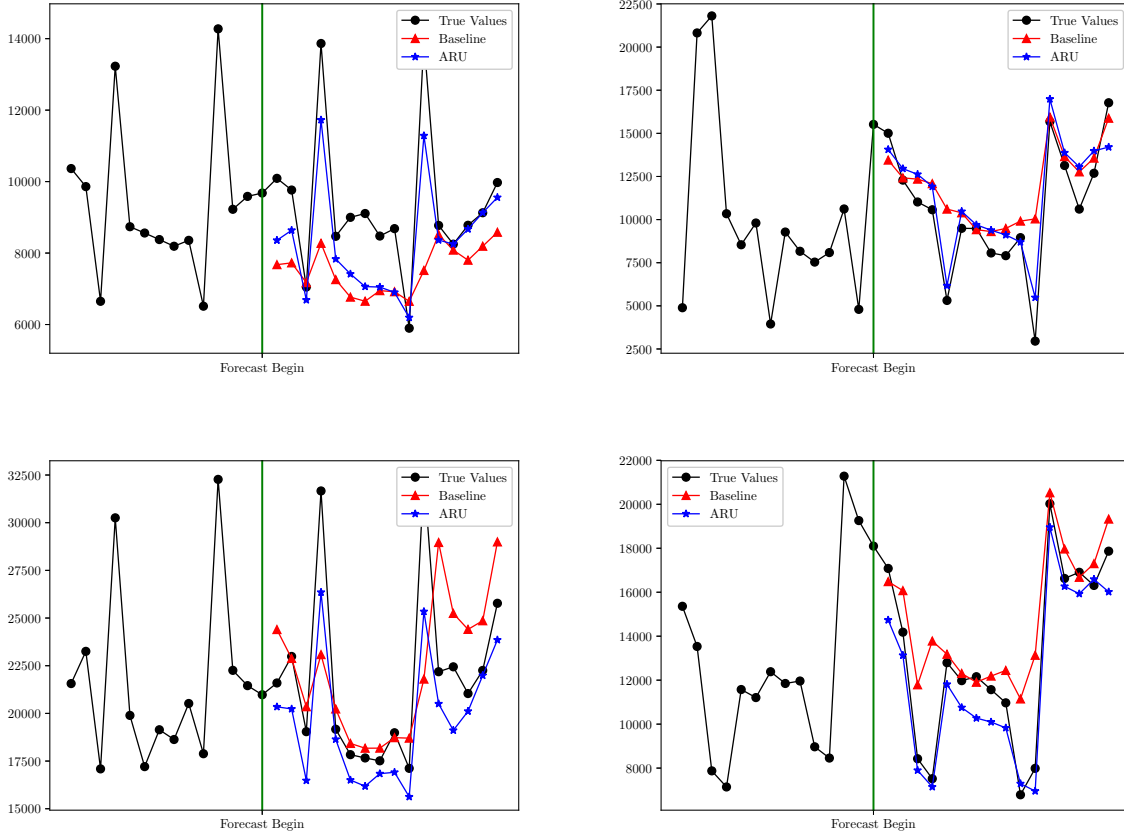


Figure 4: Examples of four time series from the Rossman dataset showing how well our ARU is able to track the local patterns of each series compared to the global baseline.

4.5 Comparisons on Real Data

We start by showing some anecdotes, then compare different methods on prediction error in two different settings, and finally compare them on running time.

4.5.1 Anecdotes. In Figure 4, we show how ARU enables a global model to adapt to the characteristics of each time series from the Rossman dataset. The vertical line denotes the prediction horizon. The ARU model is able to match the peaks of the data much better than the global model. The location and pattern of peaks in the actual data varies a lot across the four time-series and a single global model is unlikely to be able to capture them. The ARU by intervening at the last layer, with local sufficient statistics, is able to learn the relation between the inputs and peaks more accurately. For example, look at the two peaks after the green line in the first plot. For the time-series in the second row, first column we see how the global model’s peak prediction is offset by one step whereas ARU aligns perfectly with the true y s. In the last time-series (second row, second column), we see how nicely the gradual decline is matched by ARU, whereas the global baseline is needlessly jittery.

4.5.2 Quantitative comparison on prediction error. Next we move to a quantitative comparison. We first compare in a fixed horizon setting where we predict for a fixed horizon $T + 1, \dots, T + K$ that is known during training (T and K values summarized in Table 1). In Table 3 we compare different methods for a fixed horizon on the five datasets on four methods. The DeepState method is compared separately since we could not get access to the DeepState code base. We observe that on feature-rich datasets like Rossman and Walmart ARU is significantly better than all existing methods. Another property of ARU is that even when it does not improve beyond the baseline it is not much worse. This is because we meta-learn how to combine ARU’s local predictions with the global predictions.

We next compare different methods in a streaming setting where we provide the true E (encoder length) forecasts after the training time period T , and ask for prediction for the next K time periods. We do not retrain model parameters to simulate a true online deployment setting. For the electricity and traffic datasets, we use the rolling-window forecast setting used in [23] and [30]. Hourly values of 7 days are input and we make prediction for the next day. Then the forecast window is shifted by one day and we repeat this for the next 7 days. Since the frequency of the Walmart dataset

Dataset	Method	Normalized Deviation (ND)	RMSE
Rossman	Baseline	0.0987	1020.7
	DeepAR	0.2273	2197.4
	SNAIL	0.0963	993.4
	ARU	0.0851	908.6
Walmart	Baseline	0.1187	3585.8
	DeepAR	0.2215	5894.9
	SNAIL	0.1082	3465.0
	ARU	0.1058	3413.3
Electricity	Baseline	0.1264	377.3
	DeepAR	0.1838	530.9
	SNAIL	0.1307	363.2
	ARU	0.1260	344.6
Traffic	Baseline	0.1991	0.0374
	DeepAR	0.1894	0.0367
	SNAIL	0.2187	0.0380
	ARU	0.1964	0.0372
Parts	Baseline	1.5763	1.300
	DeepAR	1.4440	1.268
	SNAIL	1.4910	1.276
	ARU	1.5409	1.285

Table 3: Performance of different methods on the datasets in the fixed mode.

is weekly, we make rolling-window forecasts for 16 weeks, with window size 8. For the Rossman data, forecast window size is 16 days, and we make forecasts for 32 days.

Dataset	Method	Normalized Deviation (ND)	RMSE
Rossman	Baseline	0.094	983.2
	DeepAR	0.245	2389.0
	SNAIL	0.093	972.6
	ARU	0.089	934.9
Walmart	Baseline	0.137	4548.6
	DeepAR	0.233	6704.0
	SNAIL	0.144	4690.5
	ARU	0.114	3938.6
Electricity	Baseline	0.136	416.7
	DeepAR	0.172	544.0
	SNAIL	0.135	400.6
	ARU	0.127	396.4
Traffic	Baseline	0.170	0.0224
	DeepAR	0.145	0.0216
	SNAIL	0.165	0.0227
	ARU	0.161	0.0220

Table 4: Performance of different methods in the streaming mode.

In Table 4, we summarize the results on four datasets. The car parts dataset was dropped since each time-series is too small for streaming. We observe that even in this mode, ARU provides significant gains over all existing methods. A close second is the SNAIL

Dataset	Setting Enc Len-K	DeepAR	DeepState	ARU
Parts	12-12	1.245	1.470	1.318
Electricity	336-168	0.199	0.087	0.116
Traffic	336-168	0.148	0.168	0.159
Tourism-monthly	48-24	0.107	0.138	0.096

Table 5: ND loss for datasets on which DeepState reported results. The setting denotes the respective encoder and decoder lengths (prediction horizon).

method of adaptation but we will see later that SNAIL incurs significant runtime overheads.

4.5.3 Comparison with DeepState. Since the DeepState code is not publicly available, we compare ARU with DeepState on exactly those dataset and settings for which we were able to approximately match the publicly available DeepAR numbers in the DeepState paper. Here we also include a sixth dataset Tourism-monthly since it was used in the DeepState paper. This is a small dataset of 366 time-series containing monthly tourism demand for approximately 15 years [2]. In this dataset, the start timestamp as well as length of each time-series are different. Since this is a monthly dataset, we use month of the year feature. We observe that ARU is better than DeepState in 3 out of 4 datasets. Both DeepAR and DeepState models use a time-series id, whereas we do not. Perhaps these datasets do not show too much local variation, making it possible for large deep models to memorize local patterns in its parameters as we showed in our synthetic experiments.

4.5.4 ARU integration. Another way in which we differ from DeepState is that instead of directly yielding the local predictions, we integrate them with global predictions using more learned global parameters. This allows us to safeguard against the high variance in local predictions and to meta-learn the best fall-back strategy. In Table 6 we compare ARU with ARU-direct whose local predictions are output directly. We observe that ARU-Direct is much worse than ARU and also worse than the Baseline.

4.5.5 Running time comparison. In Table 7, we compare running times of different models of inference on the test data. The running time of SNAIL is larger, up to a factor of four higher than the baseline, due to its costly self-attention mechanism. In contrast, the running time of ARU is at most 1.5 times more than the baseline. Also, since ARU stores a fixed sized state per time-series, its storage overhead is also very small.

These experiments demonstrate that ARU is an effective lightweight method of streaming adaptation that can be easily plugged into any existing global model.

5 CONCLUSION

In this paper we presented ARU, an adaptive recurrent unit that provides streaming local adaptation of globally trained deep time-series models. The core principle in ARU is simple — exploit the kernel trick to obtain a closed-form optimal solution to linear least square regression loss. While this trick is classically known, our contribution is in recognizing the imminently practical use of this

Dataset	Method	Normalized Deviation (ND)	RMSE
Rossman	Baseline	0.0987	1020.7
	ARU	0.0851	908.6
	ARU-Direct	0.1358	1835.68
Walmart	Baseline	0.1187	3585.8
	ARU	0.1058	3413.3
	ARU-Direct	0.1271	4041.88
Electricity	Baseline	0.1264	377.3
	ARU	0.1260	344.6
	ARU-Direct	0.1389	398.11
Traffic	Baseline	0.1991	0.0374
	ARU	0.1964	0.0372
	ARU-Direct	0.2195	0.0379
Parts	Baseline	1.5763	1.300
	ARU	1.5409	1.285
	ARU-Direct	1.6555	1.5332

Table 6: Comparison of performance of ARU-Direct method with Baseline and ARU

Dataset	Baseline	ARU	SNAIL
Electricity	1.0458	1.1788	3.0754
Traffic	2.0740	2.3383	5.7673
Walmart	0.7034	0.9434	1.2693
Rossmann	0.4379	0.6717	2.2837

Table 7: Inference time (in seconds)

trick in modern deep network settings and designing a method of effectively integrating global-local patterns. Unlike existing self-attention or memory-augmented models with linearly increasing memory requirements, ARU allocates only constant sized state per time-series. Unlike methods that learn to generate local parameters, ARU can adapt to streaming data. Experiments on five real-life and synthetic datasets establish that ARU is an effective light-weight method of streaming adaptation of global deep forecasting models. In future, we plan to apply the ARU method of adaptation to linear layers in other parts of a deep network.

Acknowledgements: We thank Flipkart for sponsoring the project, Srayanta Mukherjee and P Kompalli for getting us interested in the problem, and M Prashanth for performing initial experiments.

REFERENCES

- [1] Miguel Araújo, Pedro Ribeiro, and Christos Faloutsos. 2018. Tensorcast: Forecasting Time-evolving Networks with Contextual Information. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence (IJCAI'18)*.
- [2] George Athanasopoulos, Rob J Hyndman, Haiyan Song, and Doris C Wu. 2011. The tourism forecasting competition. *International Journal of Forecasting* 27, 3 (2011), 822–844.
- [3] Shai Ben-David, John Blitzer, Koby Crammer, and Fernando Pereira. 2007. Analysis of Representations for Domain Adaptation. In *Advances in Neural Information Processing Systems 20*. MIT Press, Cambridge, MA.
- [4] John Blitzer, Ryan McDonald, and Fernando Pereira. 2006. Domain adaptation with structural correspondence learning. In *Proceedings of the 2006 conference on empirical methods in natural language processing*. Association for Computational Linguistics, 120–128.
- [5] G.E.P Box and D.R. Cox. 1964. An analysis of transformations. *Journal of Royal Statistical Society. Series B (Methodological)* 26, 2 (1964), 211–252.
- [6] G.E.P Box and Gwilym M. Jenkins. 1968. Some recent advances in forecasting and control. *Journal of Royal Statistical Society. Series C (Applied Statistics)* 17, 2 (1968), 91–109.
- [7] Nicolas Chapados. 2014. Effective Bayesian modeling of groups of related count time series. *arXiv preprint arXiv:1405.3738* (2014).
- [8] Christos Faloutsos, Jan Gasthaus, Tim Januschowski, and Yuyang Wang. 2018. Forecasting Big Time Series: Old and New. *Proc. VLDB Endow.* 11, 12 (Aug. 2018), 2102–2105.
- [9] Kris Johnson Ferreira, Bing Hong Alex Lee, and David Simchi-Levi. 2015. Analytics for and online retailer: Demand forecasting and price optimization. *Manufacturing and Service Operations Management* 18, 1 (2015), 69–88.
- [10] Chelsea Finn, Pieter Abbeel, and Sergey Levine. 2017. Model-Agnostic Meta-Learning for Fast Adaptation of Deep Networks. In *Proceedings of the 34th International Conference on Machine Learning*. 1126–1135.
- [11] Valentin Flunkert, David Salinas, and Jan Gasthaus. 2017. DeepAR: Probabilistic Forecasting with Autoregressive Recurrent Networks. *CoRR abs/1704.04110* (2017).
- [12] Hardik Goel, Igor Melnyk, and Arindam Banerjee. 2017. R2N2: Residual Recurrent Neural Networks for Multivariate Time Series Forecasting. *CoRR abs/1709.03159* (2017).
- [13] R. Hyndman, A.B. Koehler, J.K. Ord, and R.D. Snyder. 2008. *Forecasting with exponential smoothing: The state space approach*. Springer.
- [14] Vitaly Kuznetsov and Zelda Mariet. 2019. Foundations of Sequence-to-Sequence Modeling for Time Series. *AISTATS* (2019).
- [15] Paul D. Larson, David Simchi-Levi, Philip Kaminsky, and Edith Simchi-Levi. 2001. Designing and managing the supply chain. *Journal of Business Logistics* 22, 1 (2001), 259–261.
- [16] Lei Li, B. Aditya Prakash, and Christos Faloutsos. 2010. Parsimonious Linear Fingerprinting for Time Series. *Proc. VLDB Endow.* 3, 1-2 (Sept. 2010).
- [17] Nikhil Mishra, Mostafa Rohaninejad, Xi Chen, and Pieter Abbeel. 2018. A Simple Neural Attentive Meta-Learner. In *International Conference on Learning Representations*.
- [18] Srayanta Mukherjee, Devashish Shankar, Atin Ghosh, Nilam Tathawadekar, Pramod Kompalli, Sunita Sarawagi, and Krishnendu Chaudhury. 2018. ARMDN: Associative and Recurrent Mixture Density Networks for eRetail Demand Forecasting. *CoRR abs/1803.03800* (2018).
- [19] Junier B. Oliva, Barnabás Póczos, and Jeff G. Schneider. 2017. The Statistical Recurrent Unit. In *ICML*. 2671–2680.
- [20] Yao Qin, Dongjin Song, Haifeng Chen, Wei Cheng, Guofei Jiang, and Garrison W. Cottrell. 2017. A Dual-Stage Attention-Based Recurrent Neural Network for Time Series Prediction. In *IJCAI*. 2627–2633.
- [21] J.R. Quinlan. 1992. Learning with continuous classes. *Proceedings of the 5th Australian Joint Conference on Artificial Intelligence* (1992), 343–348.
- [22] Jack W Rae, Chris Dyer, Peter Dayan, and Timothy P Lillicrap. 2018. Fast Parametric Learning with Activation Memorization. *arXiv preprint arXiv:1803.10049* (2018).
- [23] Syama Sundar Rangapuram, Matthias W Seeger, Jan Gasthaus, Lorenzo Stella, Yuyang Wang, and Tim Januschowski. 2018. Deep State Space Models for Time Series Forecasting. In *Advances in Neural Information Processing Systems 31*, S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett (Eds.), 7796–7805.
- [24] Sachin Ravi and Hugo Larochelle. 2017. Optimization as a model for few shot learning. In *ICLR*.
- [25] Marek Rei. 2015. Online Representation Learning in Recurrent Neural Language Models. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*. <http://aclweb.org/anthology/D/D15/D15-1026.pdf>
- [26] Adam Santoro, Sergey Bartunov, Matthew Botvinick, Daan Wierstra, and Timothy P. Lillicrap. 2016. Meta-Learning with Memory-Augmented Neural Networks. In *ICML*. 1842–1850.
- [27] Matthias Seeger, David Salinas, and Valentin Flunkert. 2016. Bayesian Intermittent Demand Forecasting for Large Inventories. In *Proceedings of the 30th International Conference on Neural Information Processing Systems (NIPS'16)*.
- [28] Shiv Shankar and Sunita Sarawagi. 2018. Labeled Memory Networks for Online Model Adaptation. In *AAAI*.
- [29] Ruofeng Wen, Kari Torkkola, and Balakrishnan Narayanaswamy. 2017. A Multi-Horizon Quantile Recurrent Forecaster. *arXiv preprint arXiv:1711.11053* (2017).
- [30] Hsiang-Fu Yu, Nikhil Rao, and Inderjit S Dhillon. 2016. Temporal regularized matrix factorization for high-dimensional time series prediction. In *Advances in neural information processing systems*. 847–855.
- [31] Jiong Zhang, Yibo Lin, Zhao Song, and Inderjit S. Dhillon. 2018. Learning Long Term Dependencies via Fourier Recurrent Units. In *Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholm, Sweden, July 10–15, 2018*. 5810–5818.