DeepLSF: Fusing Knowledge and Data for Time Series Forecasting

This paper was downloaded from TechRxiv (<https://www.techrxiv.org>).

LICENSE

CC BY 4.0

SUBMISSION DATE / POSTED DATE

12-11-2021 / 17-11-2021

CITATION

Chattha, Muhammad Ali (2021): DeepLSF: Fusing Knowledge and Data for Time Series Forecasting. TechRxiv. Preprint. <https://doi.org/10.36227/techrxiv.17003944.v1>

DOI

[10.36227/techrxiv.17003944.v1](https://doi.org/10.36227/techrxiv.17003944.v1)

DeepLSF: Fusing Knowledge and Data for Time Series Forecasting

Muhammad Ali Chattha, Muhammad Imran Malik, Ludger van Elst, Andreas Dengel, and Sheraz Ahmed

Abstract—Deep Neural Networks (DNNs) rise to fame has been meteoric. However, there has been an increased realization that relying purely on data to find solutions can be sub-optimal because, in doing so, we ignore one very important modality i.e. knowledge. As a result, techniques based on fusing knowledge with DNNs are becoming common. However, most of the current fusion schemes revolve around knowledge distillation, where student network mimics predictions of a more complex teacher network. We argue that each modality, be it data-driven or knowledge-driven, has its own fortes and a fusion technique should combine individual strengths of each of the modalities rather than trying to copy. Motivated by this, we present Deep Latent Space Fusion (DeepLSF) for time series forecasting that uses features from both knowledge and data domains in a complementary manner. We test DeepLSF on three real-world time series datasets belonging to different domains. We show that predictions produced by DeepLSF are not only better than the predictions made by the constituent modalities but are also better than the state-of-the-art for each of the datasets by a significant margin. Leveraging the knowledge domain to complement the performance of neural networks will prove useful in many real-world application scenarios.

Index Terms—Time Series Forecasting, Deep Neural Networks, Knowledge Driven Systems, Knowledge Fusion

1 INTRODUCTION

Time series forecasting has always remained a vital problem due to its direct impact on many real-world applications. Be it within the ambit of supply chain management, resource planning or service and maintenance, etc., having an accurate estimate of prospects can yield higher profits, minimize losses, and can enable optimal resource utilization. As a result, any improvement in performance in forecasting methods is highly valuable.

Methods involving time series forecasting can be broadly categorized into two domains: data-driven Deep Neural Networks (DNNs) and knowledge-driven Knowledge Driven Systems (KDS). In a typical DNN setting, the DNN is trained on historical data where it learns to extract useful features and patterns from the data that are effective in reaching the solution. Recent advances in computational hardware and DNN architectures have enabled DNNs to achieve state-of-the-art (SOTA) performance not only in time series forecasting tasks [1] but also in other domains such as image classification [2], natural language processing [3], speech recognition [4], visual question answering (VQA) [5], image captioning [6], etc. The performance of DNNs across different domains speaks volumes about their dexterity in extracting useful features from the data. However, it is a common observation that knowledge about the problem is also important for the solution and although DNNs excel in processing data, they completely ignore domain knowledge which can be equally useful. Domain knowledge encompasses information obtained by experts who have a considerable understanding of the problem through relevant education and experience. As a result, where DNN is only limited to information present in his-

torical data, knowledge is not and in fact, may also contain information that is not present in the past data. Methods that rely on knowledge for prediction generally encapsulate the required knowledge in a KDS. A KDS usually consists of a knowledge acquisition module that extracts knowledge from domain experts and formulates it into a set of predefined rules which are then used to come up with a solution. For time series forecasting, these rules are usually in the form of logic rules [7], [8] or statistical methods [9], [10] where knowledge, be it defined by logical or arithmetic operations, is used to come up with predictions. Statistical methods have shown considerable prowess in terms of modeling time series [11], [12], [13]. Here knowledge is represented in the form of arithmetic operations that are formulated by experts and are used in a predefined manner. Statistical methods such as Auto-Regressive Integrated Moving Average (ARIMA) [9] and Exponential moving average [10] are still widely used by domain experts for forecasting tasks.

However, relying on only one modality, knowledge or data, can be sub-optimal since each modality has its separate strengths. DNNs are quite proficient in extracting useful features from data while KDS can model the underlying process very well. A natural step forward is to build a framework where individual strengths of each of these modalities are combined and their weaknesses are suppressed. However, most of the hybrid methods that revolve around combining knowledge with DNNs use knowledge distillation [14] or posterior regularization [15] as their underlying knowledge sharing mechanism. These methods work by minimizing Kullback-Leibler (KL) divergence between probability distributions of the output of the two modalities in an attempt to make the two distributions similar and hence, work in a teacher-student setting where the student network tries to be more like the teacher network in terms

• M. A. Chattha is affiliated with DFKI, NUST, TU KL, 67659 KL.
E-mail: muhammad_ali.chattha@dfki.de

of the output. We argue that in an optimal setting one network should complement the other instead of copying. Moreover, in time series forecasting or regression tasks, the outputs are numeric estimates of future values instead of a distribution. As a result, KL divergence based knowledge sharing schemes can not be directly employed as there is no underlying probability distribution. Currently, there exists a huge gap between data-driven and knowledge-driven modalities for time series forecasting tasks, and given its importance, it is imperative that a knowledge sharing framework for forecasting is developed.

On this front, we introduce a novel latent-space fusion technique, Deep Latent Space Fusion (DeepLSF), for time series forecasting that utilizes information contained in both data and knowledge modalities to come up with predictions. DeepLSF combines features drawn from the knowledge stream with those obtained by DNN at the correct level of abstraction, essentially offsetting the information missing from one modality with the one present in the other. We test DeepLSF on three real-world benchmark forecasting datasets belonging to different applications and having different temporal characteristics. DeepLSF not only produces better forecasts when compared to predictions made by either DNN or KDS individually, highlighting the effectiveness of its knowledge sharing capabilities, but also achieves SOTA results on all three datasets, cementing its prediction efficacy and robustness irrespective of the application domain. In particular, the contributions of this paper are:

- We introduce Latent-Space fusion network that automatically encodes knowledge into appropriate representation from predictions made by KDS and projects it into the intermediate layer(s) of the DNN by aligning the subspace of knowledge representations with that of intermediate DNN layer(s).
- We perform extensive experimentation and ablation studies to highlight the effectiveness of combining information from both knowledge and data modalities in multi horizon forecasting setting by using the proposed fusion network. DeepLSF achieves state-of-the-art results on every dataset tested. On average, it outperforms SOTA by **21.5%** on every evaluation metric used.

The rest of the paper is organized as follows. We start by doing a review of the work that has already been presented in the literature in Section 2. We briefly explain the forecasting problem in Section 3. In section 4, we provide a detailed description of DeepLSF architecture along with a thorough explanation of each of the components involved in it. This includes the KDS, the DNN architecture, and the latent space fusion framework. We then present experimental settings along with a brief description of the dataset and results obtained for each of the datasets in Section 5. This section also includes an extensive ablation study, where we try out different settings of DeepLSF and evaluate their results. Finally, we conclude the paper in Section 6.

2 RELATED WORK

In this section, we introduce relevant work and background studies pertaining to knowledge sharing and multi-modal

learning approaches. Integrating domain knowledge or any sort of extra information to boost DNN's performance has been an active area of research over recent years with the rationale that introducing useful human priors in the learning process can increase the performance of DNNs.

Ghazvininejad et al. (2018) [16] presented a knowledge-based conversation model powered by neural networks. In addition to utilizing data containing the conversation history, they also conditioned the output of their model on external non-conversational information, relevant to the current context of the conversation, by using a multi-task learning approach. Their model was able to produce realistic responses that were found to be more informative and appropriate by human evaluation. Similarly, Venugopalan et al. (2016) [17] proposed a neural network based language model that also employed linguistic knowledge mined, from a considerably large text corpus, in the form of semantics to produce video descriptions. They reported significant improvement in both the grammar as well as the overall description quality generated by the model. They experimented with different fusion techniques namely early, late, and deep fusion. The fusion stage was defined as the point at which they concatenated the hidden states from both the video to text network as well as the LSTM language model. Another domain where features from different modalities are used in a multi-modal setting is Visual Question Answering (VQA) systems. Ma et. al.(2018) [18] proposed a VQA system that utilizes neural networks to learn joint embeddings from visual and textual domains representing image and question pairs. In addition to learning joint embeddings, they also employed a co-attention mechanism that gives weights to every feature learned from visual and also textual domains highlighting the relevant features in both domains. They claimed that the framework achieves SOTA results on two benchmarks VQA datasets. All of these systems are highly dependent on expert performance, which in turn is dependent on the quality of additional corpus. Moreover, trying to learn similar embeddings from two different modalities although, is useful in extracting common information, however, we believe that in time series forecasting setting, focusing on features learned separately from each modality is equally important to fill in information missing from one modality.

Towell et. al.(1994) [19] proposed the Knowledge-based Artificial Neural Networks (KBANN) architecture where they utilized propositional rules for knowledge representation. These logic rules were structured hierarchically and neural network architecture is constructed in a way where it has a direct correspondence with elements in the ruleset. Each element in the ruleset is represented by a neuron in the neural network architecture and the connections between the neurons are determined by the relations between these elements defined in the ruleset. To ensure that the model has some excess capacity to learn new features, additional neurons were also introduced whose weights were learned during the training phase. A similar approach has been employed in Tran et al.(2018) [20]. Here too, a set of logic rules is employed in conjunction with the neural network. These approaches have the advantage of incorporating the knowledge from the knowledge base directly into the architecture of the neural network. However, this limits the

flexibility of the model as the architecture of the neural network is strictly defined by the ruleset. This renders most of the current SOTA DNN models incompatible with the approach and thus, limits the ability to use off-the-shelf architectures for the task at hand.

Buda et al. (2018) [21] leveraged the capabilities of statistical models to aid the neural network in producing accurate forecasting results. These forecasting results were ultimately used for anomaly detection. To further boost performance, they used a family of different statistical models whose predictions were combined into a homogeneous setup. The predictions from all the models were compared against the available ground-truth to select the most plausible prediction based on the lowest Root Mean Squared Error (RMSE) which they termed as "single-step merge". They also explored an alternate approach where only a single model was selected for all the predictions based on the lowest RMSE value. Since the model considers all of the predictions in isolation, it is unable to take advantage of both streams simultaneously as opposed to DeepLSF.

The use of statistical methods as a source of knowledge in modeling time series is not uncommon. Chattha et al. (2019a) [22] proposed a residual learning scheme, called Knowledge Integrated Neural Network(KINN), where they incorporated expert knowledge in the form of prediction in the network by adding it to the network's output. Despite its advantages, the approach cannot be directly scaled to multi-step predictions. Moreover, the performance offered by KINN did not scale to different datasets. This was primarily because KINN was unable to model time series that had a strong trend present in the past sequences. They improved upon KINN in DeepEx (2019) [23] where they used a separate predictor to model trends in the data. Similarly, Munir et al. (2019) [24] also utilized the power of statistical models to enhance the performance of their DNN which was ultimately used for the task of anomaly detection. The proposed model offered better performance in anomaly detection tasks.

Hu et al. (2016) [25] harvested expert knowledge in the form of first-order logic rules. To transfer the knowledge to the network parameters, they utilized an iterative knowledge distillation technique. They used their expert model as the teacher to train a student network comprised of a DNN based architecture. The DNN attempted to replicate the predictions made by the expert i.e. the teacher network. They updated both the teacher as well as the student at each iteration during the learning process. The main attempt was to find a teacher network that can match the ruleset in terms of predictions while not diverging significantly from the predictions of the student. Additionally, they also minimized the KL-divergence between the distributions of the prediction made by the teacher and the student to make the two distributions similar. With the proposed framework, they were able to achieve SOTA performance in classification tasks on the evaluated datasets. Following the similar steed, Xie et al. (2019) [26] used a noisy student-teacher learning paradigm. They initially train the teacher model on the ImageNet dataset which is an expert in their context. This teacher/expert network is then used to generate predictions on a larger dataset, i.e., JFT-300m which is comprised of 300 million images. These pseudo

labeled images along with the original ImageNet images are passed onto the student network for training. This process is repeated where the student is replaced as the teacher and a new more powerful student network is spawned. With this self-training scheme, they improved ImageNet top-1 accuracy by 1%. Since the task of the student network in this framework is to emulate the predictions made by the teacher, this leads to the strength of the student network being ignored by the system. In contrast, DeepLSF does not force the strengths of one modality on the other but rather works in a complementary fashion where individual strengths of each modality are retained.

3 PROBLEM FORMALIZATION

In this section, we mathematically formalize the forecasting problem and explain the terminologies that will be used throughout this paper. In a typical forecasting setting using DNNs, current value of the time series along with its lagged versions are used to predict values in the future. Formally, this is represented by a list $\mathbf{X}_p$ containing values $x_t, x_{t-1}, \dots, x_{t-p}$. Here p represents the total number of past value used for prediction and x_t represents the value of time series at time t . The aim of time series forecasting network is to learn a parametric model which maps $\mathbf{X}_p$ values from the past to $\hat{\mathbf{Y}}_h$. Here $\hat{\mathbf{Y}}_h$ represents a list containing predicted values $\hat{x}_{t+1}, \hat{x}_{t+2}, \dots, \hat{x}_{t+h}$ and h represents *horizon* i.e. the number of future values to be predicted. This parametric mapping can be mathematically expressed as:

$$[\hat{x}_{t+1}, \hat{x}_{t+2}, \dots, \hat{x}_{t+h}] = \Phi([x_t, x_{t-1}, \dots, x_{t-p}]; \mathcal{W}) \quad (1)$$

$$\hat{\mathbf{Y}}_h = \Phi(\mathbf{X}_p; \mathcal{W}) \quad (2)$$

where $\mathcal{W} = \{W_l, b_l\}_{l=1}^L$ encapsulates the parameters of the network which is comprised of L layers and $\Phi : R^{p+1} \mapsto R^h$ defines the mapping from the input space to the output space. The optimal parameters of the mapping function $\mathcal{W}^*$ are computed in a recursive manner using gradient descent. After each iteration, the loss is calculated that estimates how far predictions given by the model are from the ground truth, and the parameters of the model are adjusted to minimize the loss. This process is repeated until the loss plateaus. Typically, Mean Squared Error (MSE) is used as a loss function, and hence, the optimization problem for regression can be mathematically stated as:

$$\mathcal{W}^* = \arg \min_{\mathcal{W}} \frac{1}{h} \sum_{i=1}^h \|Y_{t+i} - \Phi([x_t, \dots, x_{t-p}]; \mathcal{W})\|_2^2 \quad (3)$$

$$\mathcal{W}^* = \arg \min_{\mathcal{W}} \frac{1}{h} \sum_{i=1}^h \|Y_{t+i} - \hat{Y}_{t+i}\|_2^2 \quad (4)$$

where Y_{t+i} and $\hat{Y}_{t+i}$ denotes the ground truth and the predicted value at time $t+i$ respectively.

4 DEEPLSF: THE PROPOSED METHOD

In this section, we introduce the proposed DeepLSF network and discuss in detail individual components of the architecture. The overall architecture of the network is shown in Fig. 1. DeepLSF consists of (i) a KDS, (ii) a DNN, and (iii) a latent space fusion network. KDS utilizes the underlying knowledge to come up with predictions. The DNN is a traditional convolutional neural network and the fusion network consists of an Encoder-Decoder based network that encodes predictions made from the KDS into knowledge representations. The fusion network also consists of a projection network that projects knowledge representations extracted using feature extractor network into the intermittent layers of DNN. Since the intermediate layers of DNN after fusion consists of features drawn from both KDS and DNN, the final predictions made by DeepLSF use information from both the knowledge and the data modality to come up with predictions. In the following subsections, we further elaborate on each of the constituent networks.

4.1 Knowledge Driven System

The term ‘knowledge’ is tricky to explain and equally difficult to store and manifest in a computer algorithm. However, knowledge can be considered as a set of rules that are to be followed in order to reach a reasonable solution and in the case of time series forecasting, rules and methods drawn from the statistical theory have shown impressive results [11], [12], [13]. Hence, we also employ concepts used in one of the statistical methods in our KDS. Specifically, we use 4Theta¹ method.

4Theta draws its root from Theta [27] method where time series is decomposed into theta lines. Theta line θ at point t can be obtained by the Eq. 5

$$\begin{aligned} Y_T^\theta &= \theta Y_T'' = \theta(Y_t - 2Y_{t-1} + Y_{t+2}) \\ &= \theta Y_T + (1 - \theta)(b + at) \end{aligned} \quad (5)$$

where Y'' is the second difference of the data and b and a are the intercept and slope of the regression at time Y^0 . The parameter θ controls the curvature of theta lines with $\theta = 0$ modeling the long term linear trend in the time series and higher order θ modeling the deviations from the linear trend by magnifying the local curvatures. For forecasting, these lines are extrapolated and combined. For simplicity, we mathematically represent a theta model comprising of only two theta lines as follows

$$Y_t = \omega_0 Y_t^0 + \omega_\theta Y_t^\theta \quad (6)$$

where ω_0 and ω_θ are the weights of Theta lines modeling the linear and curvature of the original data. The parameter θ can be optimized to attain forecasts with the lowest errors. 4Theta improves upon the classical Theta method by automatically adjusting for linear and non-linear trends by using multiple Theta models and selecting the best one. This offers a more generalized forecasting framework that is capable of handling more complex time series [28].

Although We have used the 4Theta method as our knowledge based model, DeepLSF is agnostic to underlying

knowledge used in the KDS. DeepLSF only takes predictions made by the KDS into account regardless of its internal architecture or implementation details. Hence, any other technique or statistical method that is capable of coming up with predictions can be employed in DeepLSF. This will be further explained in subsection 4.3 where we explain the knowledge fusion technique in detail.

4.2 Deep Neural Network

We employ Convolutional Neural Networks (CNN) as our data-driven method. Although the aim of this research is to develop a framework that allows for knowledge sharing between two different modalities, we still spent considerable compute effort to find optimal network parameters for DNN through extensive grid-search over reasonable hyperparameter search space. One interesting finding of architecture search was that although Long Short-Term Memory (LSTM) networks [29] are considered the go-to architecture for time series modeling, owing to their significant proficiency in modeling long-term dependencies in the data, their performance, however, was slightly inferior or comparable to the performance given by CNN for the datasets employed in our experiments. These discoveries also align with the findings of [30]. The final architecture for all the datasets consisted of three convolutional layers and a fully connected layer which were then followed by the prediction layer responsible for giving the final predictions. We used Rectified Linear Activation Unit (ReLU) as the activation function for convolutional and fully-connected layers. For the traffic dataset, the architecture has 16, 16, 32, and 64 units while for energy and NASDAQ datasets, the number of units is 4, 32, 64, and 128. Details of datasets will be given in detail in section 5.1.

4.3 Latent Space Fusion

When processing information, humans tend to find relatable templates and features that match the information they have already seen or experienced [31]. DNNs, inspired by the human neural system, also attempt to learn general features from the data which can be then further voted to reach a final prediction. However, unlike DNNs whose flexibility is limited by the training set, humans can use information from other domains to aid in their decision making process [32]. Sharing knowledge from different modalities and sources of information is perhaps one of the strongest feats of humans and we aim to enable DNNs to have a similar ability. Motivated by this, we propose a novel fusion mechanism that allows for knowledge sharing across different modalities in DNNs. This is done by utilizing two specialized networks (i) feature extractor and (ii) projection network as shown in Fig. 1. The job of the feature extractor network is to obtain relevant features from the KDS predictions that can be further used to supplement features obtained by DNN. For this, we employ an Encoder-Decoder based approach. In a typical Encoder-Decoder setting, the encoder is trained to learn compressed latent features and the decoder is trained to reconstruct the original input from these compressed features. This way the latent features learned by the encoder are most representative of the input data, where redundant information present in the data is

1. <https://tinyurl.com/vtwfzm7p>

Fig. 1. DeepLSF architecture. Historical data is given to both DNN and KDS. Latent space fusion network takes prediction made by the KDS and encodes it into suitable knowledge representation which is then projected into intermediate layer of the DNN using projection network. DNN takes features drawn from both historical data and knowledge representation to come up with predictions.

mitigated. We employ the Encoder-Decoder network in a similar fashion, where the predictions obtained from KDS, sec 4.1, are first fed into the Encoder-Decoder network which learns compressed knowledge representations. This training is done in isolation, separate from end-to-end DeepLSF optimization. The representation learned by the encoder consists of salient information contained in the knowledge driven method since the decoder is trained to minimize reconstruction loss. Once the reconstruction loss has plateaued, the decoder is not utilized further and is discarded. The weights of the encoder are frozen and are not further optimized. This can be represented mathematically as:

$$(\mathcal{W}_E^{Exp*}, \mathcal{W}_D^{Exp*}) = \arg \min_{\mathcal{W}_E^{Exp}, \mathcal{W}_D^{Exp}} \frac{1}{|\mathcal{X}|} \sum_{\mathbf{x} \in \mathcal{X}} \|\mathbf{x} - D(E(\mathbf{x}; \mathcal{W}_E^{Exp}); \mathcal{W}_D^{Exp})\|_2^2 + \lambda (\|\mathcal{W}_E^{Exp}\|_2^2 + \|\mathcal{W}_D^{Exp}\|_2^2) \quad (7)$$

where $\mathcal{W}_E^{Exp}$ and $\mathcal{W}_D^{Exp}$ denotes the weights of the encoder and the decoder respectively trained using KDS predictions. Similarly, D and E denotes the activation function of the decoder and encoder. The optimal weights obtained after training the models are denoted by $*$. λ controls the weight of the regularization term as compared to the reconstruction loss.

After the encoder has been trained, the compressed feature representations learned by it are fed into the projection network which fuses these features with intermediate layers of the DNN. The projection network serves two

purposes. First, it projects the compressed feature vectors into equivalent subspace as that of latent features of the intermediate layer of the DNN, where these features from the knowledge base are being added. Second, it matches the dimensionality of the compressed feature vector with that of the hidden layer of the DNN. The projection network is trained in conjunction with the training of DNN in an end-to-end manner. This way the projection network learns the optimal projection that gives the least prediction errors while training. We believe that integrating features from the knowledge domain at the intermediate layers of the network will prove beneficial since it constrains the optimization problem without relying on the end-to-end learning scheme to extract useful features from KDS predictions. This also avoids the trivial case where the network only learns to leverage information from one stream due to poor learning of the latent space mapping. The final output of the prediction network is based on features drawn from the data using DNN and as well as on features drawn from the KDS. Mathematically the prediction given by the prediction network can be represented as:

$$\mathbf{y}' = \Phi_{k+1:N} \left(\left[\Phi_{1:k}(\mathbf{x}; \mathcal{W}_{1:k}^*), P(E(\mathbf{y}^{exp}; \mathcal{W}_E^{Exp*}); \mathcal{W}_P^*) \right]; \mathcal{W}_{k+1:N}^* \right) \quad (8)$$

where $\mathbf{y}'$ is the forecast given by the DeepLSF prediction network. Φ is the final predictor with N layers and k is the layer at which compressed knowledge representations from KDS predictions are embedded. $\mathcal{W}_E^{Exp*}$ and E denotes

the optimal weights and activation function of the encoder already trained using expert predictions in Eq. 7. y^{exp} are the expert predictions and W_P^* and P denotes final weights and activation of the projection network. $W_{k+1:N}^*$ denotes the optimal weights of the DNN. The optimal weights are computed by minimizing the mean squared error between predicted and desired forecasts, which can be mathematically represented by the following equation:

$$\begin{aligned} (W^*, W_P^*) = \arg \min_{W, W_P} \frac{1}{|\mathcal{X}|} \sum_{(x, y)} \|y - \Phi_{k+1:N} \\ ([\Phi_{1:k}(x; W_{1:k}), P(E(y^{exp}; W_E^{Exp*}); W_P)]; W_{k+1:N})\|_2^2 \\ + \lambda (\|W\|_2^2 + \|W_P\|_2^2) \quad (9) \end{aligned}$$

where x is the input while y is the desired forecast, Φ is the final predictor with N layers.

As can be seen from Eq 8, the final output given by the network makes use of the information contained in knowledge and as well as the data streams by combining features obtained from both of these modalities. The parameter k is a hyperparameter that dictates the layer of DNN at which knowledge representations are incorporated. We also study the impact of using different k in Section 5.6. For the rest of the experiments, we always fuse the features at the first layer, i.e. $k = 1$ since it gave the best results on datasets evaluated in this paper.

5 EXPERIMENTS

In this section, we evaluate DeepLSF with extensive experimentation. For evaluation, we have used three widely used real-world forecasting benchmarking datasets. Firstly, we will briefly discuss the characteristics of the datasets which will be followed by training details for each of the datasets. After which we will demonstrate results achieved by DeepLSF in making short and long-term predictions and will compare them with the results of current SOTA and other baseline methods present in the literature. Next, we will validate the effectiveness of the proposed knowledge fusion mechanism by comparing the results of the architecture with and without knowledge fusion. Finally, we will present an ablation study in which we will evaluate and analyze different configurations of DeepLSF and comment about the configuration of DeepLSF proposed in this paper

5.1 Datasets and Data Preprocessing

In order to determine the efficacy of DeepLSF, we employ three real-world publicly available datasets for evaluation. Each dataset belongs to a different application domain. The datasets utilized are as follows.

- **Traffic [33]:** This dataset consists of vehicular traffic data obtained from Caltrans Performance Measurement System (PeMS). PeMS has been deployed by California Department of Transportation consisting of over 39,000 sensor stations installed on California state highways that collect traffic information every 30 seconds. For our experiments, we use PeMSD7(M)

[33] dataset corresponding to District 7 of California. The dataset aggregates traffic into 5 minute intervals and contains weekdays readings of 228 sensors from May to June of 2012.

- **Energy [34]:** This dataset contains measurement of 26 different attributes relating to appliances energy consumption in a single household. The dataset contains samples recorded every 10 minute for 4.5 months and is obtained from UCI appliances energy database.
- **Stocks [35]:** This dataset contains stock prices of 81 corporations and the index value of NASDAQ 100 which is recorded every minute for 105 days.

In Table 1 we list some statistics of each of the datasets used along with train, test and validation split.

TABLE 1
Statistics of Datasets

Dataset	Traffic	Energy	Stocks
Total samples	17544	19735	40560
No. of time series	228	26	82
Sampling rate	5 min	10 min	1 min
Train size	84%	80%	90%
Valid size	8%	10%	5%
Test size	8%	10%	5%

Please note that the train, test and validation splits used in this study are the same with other studies that we compare against in section 5.4 in order to keep the comparison fair.

5.2 Training Details and Preprocessing

To get training samples from each of the time series, a rolling window approach is used where a window of size *inputWindowSize* and size *HorizonSize* is used to get an input-label pair from the time series respectively. The next training sample is obtained by shifting both of the windows by an index of 1 and this process is continued until the whole of the time series is covered. As a result, each time series gives $(TSLength - inputWindowSize - HorizonSize + 1)$, where *TSLength* is the length of the time series. Samples for validation and test set are obtained using the split mentioned in Table 1. An important factor that directly impacts the performance of DNN is preprocessing. Each input sample given to DNN can have a different scale and distribution which, in turn, may increase the difficulty of modeling the problem for DNNs since, in addition to learning the input to output mapping, the DNNs also have to learn the scale of each of the input. As a result normalization of input variables is essential for DNN based approaches to confine the scale of the input variables within a certain range. In our experiments, we employ Min-Max scaling as our normalization technique. Here each input-label pair given by the rolling window approach is normalized using Min-Max scaling, which scales the range of input values to be between 0 and 1. Normalization is done for all of the datasets except for the traffic dataset, where results without normalization were far better. This can be attributed to the fact that values for this particular dataset were always within a confined range which made it possible for the network to model the time series efficiently.

Once all of the time series in a given dataset are converted into input label pairs and normalized, where applicable, the training samples from each of the time series are concatenated and DeepLSF is trained in a univariate setting, where one DeepLSF model is trained for all of the time series in the dataset. Feature extractor network is also trained in a similar fashion where a single encoder model is trained for all the time series of a particular dataset. Here too, the inputs are normalized using Min-Max scaling. During the training of DeepLSF, the initial learning rate is set to 10^{-3} which is decreased by 10^{-1} every time the loss plateaus. The set of parameters that give the lowest validation loss are saved and used for the evaluation of the test set.

5.3 Evaluation Metrics and Baseline Methods

To evaluate the performance of DeepLSF we employ two widely used evaluation metrics namely, Root Mean Squared Error (RMSE) and Mean Absolute Percentage Error (MAE). RMSE and MAE are calculated for each of the time series in the dataset which is then averaged to form one representative value for the entire dataset. This can be mathematically expressed as

$$RMSE = \sqrt{\frac{1}{n} \sum_{j=1}^n \frac{1}{s} \sum_{i=1}^s (Y_{t+i,j} - \hat{Y}_{t+i,j})^2}$$

$$MAE = \frac{1}{n} \sum_{j=1}^n \frac{1}{s} \sum_{i=1}^s \|Y_{t+i,j} - \hat{Y}_{t+i,j}\|$$

where n is the number of time series in the dataset, Y_{t+i} is the ground truth and $\hat{Y}_{t+i}$ is the prediction of the model at the future moment $t + i$. s is the number of samples in the test set.

These metrics are chosen primarily because of their widespread usage in related literature, which makes it easier to compare the performance of different techniques. For traffic dataset, we mainly compared with Spatio-Temporal Graph Convolutional Networks (STGCN) [33], Diffusion Convolutional Recurrent Neural Network (DCRNN) [36], Graph-WaveNet [37] and Spatial-Temporal Transformer Networks (STTN) [1]. Almost all of these techniques revolve around fusing graph networks with traditional neural networks. This allows the merging of spatial information from all the time series in the dataset with temporal information obtained from the individual time series being predicted. For datasets relating to vehicular traffic, spatial information proves useful since roads in a particular area are connected and hence information of traffic on the adjacent or connecting road eventually leads to better prediction. Currently, STTN [1] is SOTA for PeMSD7(M) dataset. Along with these recent models, we also compare our results with traditional statistical and DNN methods including Linear Support Vector Regression (LSVR), Auto Regressive Moving Average (ARIMA), Feed-Forward Neural Network (FNN), and Fully-Connected LSTM (FC-LSTM).

For Energy and NASDAQ datasets, we compared with Auto-Encoder Convolutional and Recurrent Neural Network (AECRNN) [38], Long- and Short-term time series network (LSTNet) [39] and Multi-level Construal Neural Network (MLCNN) [40]. MLCNN [40] revolves around

extracting multi-level abstract representations of the raw data by using convolutional networks. These multiple representations are extracted for both near and distant future predictions which are then fused by using the Encoder-Decoder approach and scaled by using traditional Autoregression models. Currently, MLCNN [40] is SOTA on both of these datasets. Apart from these methods we also compared with classical LSTM [41] and Multi-Task CNN (MTCNN) [42] models along with Vector AutoRegressive (VAR) method.

5.4 Results

To keep the comparison with other techniques impartial, we use the same learning settings as those used in current best performing work in the literature. We use the same train, validation, and test split along with the same input window and horizon size. Table 2 summarizes the results of DeepLSF and other baseline techniques mentioned in section 5.3 for traffic PeMS7(M) dataset. All the models are evaluated on predictions made for 15, 30, and 45 minutes ahead forecasts. As the sampling rate of the data is 5 mins, this translates to the horizon of 3, 6, and 9 respectively i.e. $\{t+3, t+6, t+9\}$.

DeepLSF outperforms all of the recent methods by significant margins. In particular, DeepLSF improves over current best performing model STTN [1] by **11.7%, 6.3%, 32.7%** on MAE and by **12.6%, 11.4%, 33.6%** on RMSE for 15, 30 and 45 min ahead prediction respectively. Although STTN [1] takes advantage of long term temporal information along with spatial dependencies in the data, the results show that utilizing knowledge allows DeepLSF to model time series more effectively.

Table 3 shows results on energy dataset. Here, the models are evaluated for predictions made for the horizon of 3, 6, and 12 i.e. $\{t+3, t+6, t+12\}$ corresponding to values for 30, 60, and 120 minutes in the future. Similar to the results in the traffic dataset, DeepLSF outperforms current state-of-the-art and other baseline methods in every setting. This demonstrates the effectiveness of DeepLSF in predicting both near and distant future values. Specifically, DeepLSF improves the SOTA by **3.2%, 15.1%, 22.4%** on MAE and **5.4%, 5.1%, 6.6%** on RMSE.

Table 4 shows the results of all of the evaluated models on the NASDAQ dataset. Here, the models are evaluated for predictions made for the horizon of 3, 6, and 12 which correspond to future values for 3, 6, and 12 minutes in this case since the sampling rate of the dataset is 1 minute. Here as well, DeepLSF outperforms every baseline model including the SOTA. DeepLSF outperforms SOTA method MLCNN [40] by **30.8%, 35.4%, 39.8%** on MAE and **36.7%, 38.2%, 40.6%** on RMSE. DeepLSF yields the biggest improvement on

DeepLSF establishes new SOTA on all of the three datasets in every setting evaluated. On average, DeepLSF improves the SOTA by **16.9%** on MAE and by **19.2%** on RMSE for traffic dataset, by **13.6%** and **5.7%** for energy dataset and by **35.3%** and **38.5%** for stocks dataset as shown in Fig. 2. This demonstrates the efficacy of DeepLSF in a multi-horizon prediction setting where the network is evaluated for predicting both near and distant values in the future. Moreover, it also highlights the robustness of DeepLSF in modeling different time series belonging to different

TABLE 2
Results on Traffic dataset for 15, 30 and 45 min ahead forecasts

Models	MAE			RMSE		
	15 mins	30 mins	45 mins	15 mins	30 mins	45 mins
LSVR	2.50	3.63	4.54	4.55	6.67	8.28
ARIMA	5.55	5.86	6.27	9.00	9.13	9.38
FNN	2.74	4.02	5.04	4.75	6.98	8.58
FC-LSTM	3.57	3.94	4.16	6.20	7.03	7.51
DCRNN [39] (2018)	2.37	3.31	4.01	4.21	5.96	7.13
STGCN [33] (2018)	2.25	3.03	3.57	4.04	5.70	6.77
DeepGLO [43] (2019)	-	-	3.53	-	-	6.49
Graph-WaveNet [37] (2019)	2.14	2.80	3.19	4.01	5.48	6.25
STTN [1] (2021)	2.14	2.70	3.03	4.04	5.37	6.05
DeepLSF	1.89	2.53	2.04	3.53	4.76	4.02

TABLE 3
Results on Energy Dataset for 30, 60 and 120 mins ahead forecasts

Models	MAE			RMSE		
	30 mins	60 mins	120 mins	30 mins	60 mins	120 mins
VAR	2.90	3.32	3.87	15.51	16.25	16.95
RNN-LSTM	2.73	3.05	3.67	15.82	16.76	17.29
MTCNN [42] (1998)	3.42	3.90	4.31	15.84	16.55	17.48
AECRNN [38] (2018)	2.27	3.01	3.40	15.71	16.26	17.17
LSTNet [39] (2018)	1.82	2.39	3.11	15.51	15.80	16.89
MLCNN [40] (2020)	1.88	2.38	3.04	15.13	15.99	16.78
DeepLSF	1.82	2.02	2.36	14.31	15.17	15.67

TABLE 4
Results on Stocks Dataset for 3, 6 and 12 mins ahead forecasts

Models	MAE			RMSE		
	3 mins	6 mins	12 mins	3 mins	6 mins	12 mins
VAR	1.83	2.08	2.01	2.73	3.05	3.05
LSTM	2.21	2.34	2.45	4.53	4.95	5.35
AECRNN [38] (2018)	4.37	4.50	4.37	9.79	9.89	9.73
MTCNN [42] (1998)	2.43	2.38	2.40	4.20	3.93	4.34
LSTNet [39] (2018)	0.093	0.135	0.195	0.366	0.522	0.754
MLCNN [40] (2020)	0.091	0.130	0.186	0.365	0.516	0.739
DeepLSF	0.063	0.082	0.112	0.231	0.316	0.439

Fig. 2. Percentage improvement offered by DeepLSF over previous SOTA for every dataset evaluated. Orange and blue bar represents percentage gain over SOTA for MAE and RMSE respectively

domains and having different temporal characteristics. The biggest percentage improvement over the previous SOTA is on the NASDAQ dataset, where DeepLSF outperforms the previous best model by more than 30% on each of the metrics and for each of the forecasting horizon. Generally, as

the forecasting horizon becomes larger the accuracy of any forecasting model decreases since the model has to predict more values in the future which involves more uncertainty. However, the percentage improvement offered by DeepLSF increases with the increase in horizon size. Fig. 3 shows the average improvement given by DeepLSF for predictions made for near, medium, and far horizons.

Getting more accurate predictions for values that are more distant in the future is one of the significant advantages of DeepLSF which will prove useful in almost every domain. Specifically in domains like stocks and health etc, where more accurate estimates of the distant future can lead to more efficient planning which can potentially yield higher profits and save human lives.

5.5 Effectiveness of the Knowledge Fusion Scheme

In this subsection, we evaluate and validate the effectiveness of using the proposed fusion scheme in modeling time series, which in turn contributes towards better forecasting performance.

In order to demonstrate this, we first train the DNN used in DeepLSF without the knowledge fusion scheme, which we refer to as vanilla DNN. Vanilla DNN has the exact

Fig. 3. Average percentage improvement given by DeepLSF over different forecasting horizons, when compared to SOTA. Here, h denotes the horizon size.

same specifications in terms of the number of layers, the number of units in each layer, activation functions, etc, as that of the DNN used in the DeepLSF architecture. Training and validation curves of both of the models, vanilla DNN and DeepLSF, are observed to determine the capability of each network in modeling the time series. Fig. 4 gives an overview of loss curves for both models for the traffic dataset. Fig 4 shows loss curves of each of the network for 50 epochs on traffic dataset.

Fig. 4. Train and Validation loss curves for traffic dataset

DeepLSF achieves lower loss on the training set when compared to the loss achieved by vanilla DNN. This can be seen by the solid blue line in Fig 4. Similarly, DeepLSF also achieves lower validation loss which shows that overall DeepLSF is able to model the time series more effectively as it has a better in-sample fit without overfitting.

However, achieving lower loss in training only paints half of the picture and it is also imperative that the model performs equally well on unseen data in order to make sure that the model learned by the architecture can work in a real-world setting. Moreover, it is also important to demonstrate that the improvement offered by DeepLSF can actually be attributed to knowledge fusion scheme. To show this, we compare predictions given by DeepLSF with those given by the constituent modalities, knowledge-driven and data-

driven, when evaluated in isolation without the knowledge fusion scheme. For this, we computed RMSE and MAE on predictions made by the vanilla DNN and, as well as, the KDS on the test set for each of the datasets. The performance metrics of these networks are then compared with the ones given by DeepLSF to validate the effectiveness of the fusion mechanism in improving the prediction efficacy of the model. Table 5 shows the results given by vanilla DNN, KDS, and DeepLSF. An additional column is also added that shows the improvement in performance achieved by DeepLSF in percentage over vanilla DNN and KDS.

DeepLSF consistently gives better forecasting performance on all of the three datasets, when compared to the individual predictions given by simple DNN and KDS. On average DeepLSF performed around 16% better in MAE and 11% better in RMSE. The results validate that the knowledge fusion mechanism employed in DeepLSF is capable of fusing information from both streams in a way where information from one stream offsets the information missing from the other i.e. both modalities are used in a complementary manner rather than one modality trying to copy strengths of the other. This is evident, as it gives better overall forecasting performance when compared to the individual performance given by the constituent networks. Combining the strengths of the data and knowledge-driven modalities was one of the goals that we initially set towards achieving a more natural fusion mechanism.

5.6 Ablation Studies

In this section, we evaluate different architectural configurations of DeepLSF in order to support the architecture proposed in DeepLSF.

DeepLSF uses a latent space fusion network to combine information from the knowledge stream at the first layer of the architecture. We believe that fusing features in the initial layers of DNN allows the network to focus on information missing from the expert network while extracting features from the input layer. Moreover, it also gives the model enough capacity to learn and compute more complex non-linear relationships from features fused from the knowledge stream in later layers of the network. However, since the DNN used in DeepLSF has three convolutional layers, it is also interesting to investigate the impact on forecasting performance if the projection network is used to fuse information to other intermediary layers of the network. This translates to the value of k in Eq 8 to be greater than 1. We analyze the impact of using different intermediary layers for fusion, individually, along with using the fusion network to project knowledge features on all three layers of DNN at once. Table 6 shows results of these experiments evaluated on traffic dataset.

The results verify the arguments we gave above as fusing information at the first layer of the DNN gives the best overall forecasting performance. Fusing at all layers of the DNN did not improve the results, which intuitively makes sense since the network already has the information that is being fused at deeper layers. However, all the results obtained, irrespective of the layer at which knowledge is fused, were still significantly better than the one produced by vanilla DNN with no knowledge fusion.

TABLE 5

Comparison of MAE and RMSE values of Vanilla DNN and KDS with those of DeepLSF. Percentage improvement achieved by DeepLSF over constituent networks, DNN and KDS, is also given

Model	MAE	Percentage improvement in MAE	RMSE	Percentage improvement in RMSE
Traffic (15/30/45 mins)				
Vanilla DNN	1.95/2.61/3.04	3.1%/3.1%/32.9%	3.55/4.81/5.64	0.6%/1.0%/28.7%
KDS	2.13/2.87/3.44	11.3%/11.8%/40.7%	3.99/5.54/6.89	11.5%/14.1%/41.7%
DeepLSF	1.89/2.53/2.04	-	3.53/4.76/4.02	-
Energy (30/60/120 mins)				
Vanilla DNN	3.52/2.97/3.03	48.3%/32.0%/22.1%	18.31/18.52/18.60	21.8%/18.1%/15.8%
KDS	2.30/2.29/2.62	20.9%/11.8%/9.9%	17.19/19.59/17.59	16.8%/22.6%/10.9%
DeepLSF	1.82/2.02/2.36	-	14.31/15.17/15.67	-
NASDAQ (3/6/12 mins)				
Vanilla DNN	0.0633/0.090/0.158	1.3%/8.9%/29.1%	0.232/0.332/0.781	0.4%/4.8%/39.7%
KDS	0.078/0.095/0.123	19.2%/13.7%/8.9%	0.282/0.361/0.471	18.1%/12.5%/6.8%
DeepLSF	0.0625/0.082/0.112	-	0.231/0.316/0.439	-

TABLE 6

Comparison of RMSE and MAE values with knowledge fusion at different intermediary layers and with no fusion for the traffic dataset

Fusion Layer(s)	MAE	RMSE
No fusion (Vanilla DNN)	3.04	5.64
First (DeepLSF)	2.04	4.02
Second	2.12	4.09
Third	2.14	4.11
All	2.14	4.10

In addition to this, it is also interesting to evaluate the importance of Encoder-Decoder pre-training. Here, we have used an encoder-decoder based feature extractor in the latent space fusion module that learns to encode KDS predictions into appropriate knowledge representations. The encoder in the Encoder-Decoder network is pre-trained, separately from the end-to-end DeepLSF training as explained in Sec 4.3. However, when using two separate modalities the literature often leans towards latent space alignment techniques such as in visual question answering methods, where network from both modalities are trained together in an end-to-end optimization setting. This results in the latent space of both of the networks being aligned with each other. We also evaluate both of these settings. In the first setting, the encoder in the Encoder-Decoder network is pre-trained i.e. the setting used in DeepLSF while in the other, the encoder is trained in an end-to-end fashion along with the training of complete DeepLSF architecture. Table 7 shows results of both of these scenarios.

TABLE 7

Comparison of MAE and RMSE values of DeepLSF with and without encoder pre-training in latent space fusion network for the traffic dataset

Optimization Strategy	MAE	RMSE
DeepLSF (with encoder pre-training)	2.04	4.02
Without encoder pre-training	2.07	4.09

As evident from table 7, extracting features from the knowledge stream beforehand as proposed in DeepLSF produces better results on both of the evaluation metrics.

This may be due to the fact that extracting representative features beforehand makes the job of an end-to-end learning scheme easier since it does not have to learn to extract useful features from the KDS in addition to finding useful features from the raw data. This finding is also in line with the results presented by Palacio et. al. [44], where pre-training an Encoder-Decoder system at the input layer produced more robust results which were less prone to adversarial attacks in image classification tasks.

6 CONCLUSION

This paper presents DeepLSF, a novel knowledge fusion technique for time series forecasting that provides a mechanism to utilize the strengths of knowledge and data domains by combining features from both of these modalities. DeepLSF achieves this by making use of a Latent Space Fusion network that extracts relevant knowledge representations from the KDS and projects them into equivalent subspace, as that of the intermediate layer of DNN. The prowess of the proposed DeepLSF architecture is tested by evaluating its forecasting performance on three real-world datasets belonging to different application domains. DeepLSF is not only able to model the time series more efficiently by achieving lower training and validation losses when compared to the baseline DNN but also outperforms SOTA across the board, i.e. on every dataset. This shows the efficacy of DeepLSF in time series forecasting and also advocates the robustness of the architecture when dealing with different types of time series data belonging to different domains. The ability of DeepLSF to automatically extract knowledge representations makes it agnostic to the underlying knowledge base used in KDS. Moreover, the fusion network is capable of working with any off-the-shelf DNN architecture. This flexibility of DeepLSF will prove beneficial in areas where only a certain type of knowledge base is available or the use of different DNN architecture is required. Leveraging knowledge to complement the features drawn by DNNs will prove useful in many critical domains especially where expert knowledge is extremely desirable such as the health and the financial sector.

ACKNOWLEDGMENT

This work is supported by National University of Science and Technology (NUST) Pakistan, "Continental Automotive GmbH" and BMBF project KI Wissen (Grant 19A20020G).

REFERENCES

- [1] M. Xu, W. Dai, C. Liu, X. Gao, W. Lin, G.-J. Qi, and H. Xiong, "Spatial-temporal transformer networks for traffic flow forecasting," *arXiv preprint arXiv:2001.02908*, 2020.
- [2] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778.
- [3] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "Bert: Pre-training of deep bidirectional transformers for language understanding," *arXiv preprint arXiv:1810.04805*, 2018.
- [4] G. Saon, G. Kurata, T. Sercu, K. Audhkhasi, S. Thomas, D. Dimitriadis, X. Cui, B. Ramabhadran, M. Picheny, L.-L. Lim *et al.*, "English conversational telephone speech recognition by humans and machines," *arXiv preprint arXiv:1703.02136*, 2017.
- [5] K. Yi, J. Wu, C. Gan, A. Torralba, P. Kohli, and J. Tenenbaum, "Neural-symbolic vqa: Disentangling reasoning from vision and language understanding," in *Advances in Neural Information Processing Systems* 31, S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, Eds. Curran Associates, Inc., 2018, pp. 1031–1042. [Online]. Available: <http://papers.nips.cc/paper/7381-neural-symbolic-vqa-disentangling-reasoning-from-vision-and-language-understanding.pdf>
- [6] M. Z. Hossain, F. Sohel, M. F. Shiratuddin, and H. Laga, "A comprehensive survey of deep learning for image captioning," *ACM Computing Surveys (CSUR)*, vol. 51, no. 6, pp. 1–36, 2019.
- [7] I. Moghram and S. Rahman, "Analysis and evaluation of five short-term load forecasting techniques," *IEEE Transactions on power systems*, vol. 4, no. 4, pp. 1484–1491, 1989.
- [8] R. Zhang, B. Ashuri, and Y. Deng, "A novel method for forecasting time series based on fuzzy logic and visibility graph," *Advances in Data Analysis and Classification*, vol. 11, no. 4, pp. 759–783, 2017.
- [9] G. E. Box, G. M. Jenkins, and G. Reinsel, "Time series analysis: forecasting and control holden-day san francisco," *BoxTime Series Analysis: Forecasting and Control Holden Day* 1970, 1970.
- [10] J. S. Hunter, "The exponentially weighted moving average," *Journal of quality technology*, vol. 18, no. 4, pp. 203–210, 1986.
- [11] S. Makridakis and M. Hibon, "The m3-competition: results, conclusions and implications," *International journal of forecasting*, vol. 16, no. 4, pp. 451–476, 2000.
- [12] S. Makridakis, E. Spiliotis, and V. Assimakopoulos, "The m4 competition: Results, findings, conclusion and way forward," *International Journal of Forecasting*, 2018.
- [13] S. B. Taieb, G. Bontempi, A. F. Atiya, and A. Sorjamaa, "A review and comparison of strategies for multi-step ahead time series forecasting based on the nn5 forecasting competition," *Expert systems with applications*, vol. 39, no. 8, pp. 7067–7083, 2012.
- [14] G. Hinton, O. Vinyals, and J. Dean, "Distilling the knowledge in a neural network," *arXiv preprint arXiv:1503.02531*, 2015.
- [15] K. Ganchev, J. Graça, J. Gillenwater, and B. Taskar, "Posterior regularization for structured latent variable models," *The Journal of Machine Learning Research*, vol. 11, pp. 2001–2049, 2010.
- [16] M. Ghazvininejad, C. Brockett, M.-W. Chang, B. Dolan, J. Gao, W.-t. Yih, and M. Galley, "A knowledge-grounded neural conversation model," in *AAAI*, 2018.
- [17] S. Venugopalan, L. A. Hendricks, R. Mooney, and K. Saenko, "Improving lstm-based video description with linguistic knowledge mined from text," *arXiv preprint arXiv:1604.01729*, 2016.
- [18] C. Ma, C. Shen, A. Dick, Q. Wu, P. Wang, A. van den Hengel, and I. Reid, "Visual question answering with memory-augmented networks," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 6975–6984.
- [19] G. G. Towell and J. W. Shavlik, "Knowledge-based artificial neural networks," *Artificial intelligence*, vol. 70, no. 1-2, pp. 119–165, 1994.
- [20] S. N. Tran and A. S. d. Garcez, "Deep logic networks: Inserting and extracting knowledge from deep belief networks," *IEEE transactions on neural networks and learning systems*, vol. 29, no. 2, pp. 246–258, 2018.
- [21] T. S. Buda, B. Caglayan, and H. Assem, "Deepad: A generic framework based on deep learning for time series anomaly detection," in *PAKDD*. Springer, 2018, pp. 577–588.
- [22] M. A. Chattha, S. A. Siddiqui, M. I. Malik, L. van Elst, A. Dengel, and S. Ahmed, "Kinn," *arXiv preprint arXiv:1902.05653*, 2019.
- [23] M. A. Chattha, S. A. Siddiqui, M. Munir, M. I. Malik, L. van Elst, A. Dengel, and S. Ahmed, "Deepex: Bridging the gap between knowledge and data driven techniques for time series forecasting," in *International Conference on Artificial Neural Networks*. Springer, 2019, pp. 639–651.
- [24] M. Munir, S. A. Siddiqui, M. A. Chattha, A. Dengel, and S. Ahmed, "FuseAD: Unsupervised Anomaly Detection in Streaming Sensors Data by Fusing Statistical and Deep Learning Models," *Sensors*, vol. 19, no. 11, 2019.
- [25] Z. Hu, X. Ma, Z. Liu, E. Hovy, and E. Xing, "Harnessing deep neural networks with logic rules," *arXiv preprint arXiv:1603.06318*, 2016.
- [26] Q. Xie, E. Hovy, M.-T. Luong, and Q. V. Le, "Self-training with noisy student improves imagenet classification," *arXiv preprint arXiv:1911.04252*, 2019.
- [27] V. Assimakopoulos and K. Nikolopoulos, "The theta model: a decomposition approach to forecasting," *International journal of forecasting*, vol. 16, no. 4, pp. 521–530, 2000.
- [28] E. Spiliotis, V. Assimakopoulos, and S. Makridakis, "Generalizing the theta method for automatic forecasting," *European Journal of Operational Research*, 2020.
- [29] S. Hochreiter and J. Schmidhuber, "Lstm can solve hard long time lag problems," in *Advances in neural information processing systems*, 1997, pp. 473–479.
- [30] F. A. Gers, D. Eck, and J. Schmidhuber, "Applying lstm to time series predictable through time-window approaches," in *Neural Nets WIRN Vietri-01*. Springer, 2002, pp. 193–200.
- [31] P. H. Lindsay and D. A. Norman, *Human information processing: An introduction to psychology*. Academic press, 2013.
- [32] B. Landau, L. B. Smith, and S. S. Jones, "The importance of shape in early lexical learning," *Cognitive development*, vol. 3, no. 3, pp. 299–321, 1988.
- [33] B. Yu, H. Yin, and Z. Zhu, "Spatio-temporal graph convolutional networks: A deep learning framework for traffic forecasting," *arXiv preprint arXiv:1709.04875*, 2017.
- [34] L. M. Candanedo, V. Feldheim, and D. Deramaix, "Data driven prediction models of energy use of appliances in a low-energy house," *Energy and buildings*, vol. 140, pp. 81–97, 2017.
- [35] Y. Qin, D. Song, H. Chen, W. Cheng, G. Jiang, and G. Cottrell, "A dual-stage attention-based recurrent neural network for time series prediction," *arXiv preprint arXiv:1704.02971*, 2017.
- [36] Y. Li, R. Yu, C. Shahabi, and Y. Liu, "Diffusion convolutional recurrent neural network," in *Data-driven traffic forecasting, ICLR 2018 Conference*, 2017, pp. 1–16.
- [37] Z. Pan, Y. Liang, W. Wang, Y. Yu, Y. Zheng, and J. Zhang, "Urban traffic prediction from spatio-temporal data using deep meta learning," in *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 2019, pp. 1720–1730.
- [38] R.-G. Cirstea, D.-V. Micu, G.-M. Muresan, C. Guo, and B. Yang, "Correlated time series forecasting using deep neural networks: A summary of results," *arXiv preprint arXiv:1808.09794*, 2018.
- [39] G. Lai, W.-C. Chang, Y. Yang, and H. Liu, "Modeling long-and short-term temporal patterns with deep neural networks," in *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval*, 2018, pp. 95–104.
- [40] J. Cheng, K. Huang, and Z. Zheng, "Towards better forecasting by fusing near and distant future visions." in *AAAI*, 2020, pp. 3593–3600.
- [41] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural computation*, vol. 9, no. 8, pp. 1735–1780, 1997.
- [42] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner, "Gradient-based learning applied to document recognition," *Proceedings of the IEEE*, vol. 86, no. 11, pp. 2278–2324, 1998.
- [43] R. Sen, H.-F. Yu, and I. Dhillon, "Think globally, act locally: A deep neural network approach to high-dimensional time series forecasting," *arXiv preprint arXiv:1905.03806*, 2019.
- [44] S. Palacio, J. Folz, J. Hees, F. Raue, D. Borth, and A. Dengel, "What do deep networks like to see?" in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 3108–3117.

Muhammad Ali Chattha received the master's degree from Lahore University of Management Sciences, Pakistan in electrical engineering. He is currently pursuing the Ph.D. degree in computer science from Technische Universität Kaiserslautern with the German Research Center for Artificial Intelligence (DFKI GmbH) under the supervision of Prof. Dr. Prof. H. C. Andreas Dengel. His main area of research is focused on fusing traditional knowledge based systems with Deep Neural Networks. His research interests

include Time Series Analysis, Knowledge Fusion, Hybrid Networks and Forecasting.

Andreas Dengel received the Diploma degree in computer science from the University of Kaiserslautern and the Ph.D. degree from the University of Stuttgart. He is a Scientific Director with the German Research Center for Artificial Intelligence (DFKI GmbH), Kaiserslautern. In 1993, he became a Professor at the Computer Science Department, University of Kaiserslautern, where he holds the chair Knowledge-Based Systems. Since 2009, he has been a Professor (Kyakuin) with the Department of Computer Science and Information Systems, Osaka Prefecture University. He was also with IBM, Siemens, and Xerox Parc. Moreover, he has co-edited international computer science journals and has written or edited 12 books. He has authored more than 300 peer-reviewed scientific publications and has supervised more than 170 Ph.D. and master's theses. He is a member of several international advisory boards, has chaired major international conferences, and founded several successful start-up companies. He is an IAPR Fellow and has received prominent international awards. His main scientific emphasis is in the areas of pattern recognition, document understanding, information retrieval, multimedia mining, semantic technologies, and social media.

Muhammad Imran Malik received the bachelor's and the master's degrees in computer science from the University of Kaiserslautern, Germany, and the Ph.D. degree with the German Research Center for Artificial Intelligence, Germany, under the supervision of Dr. H. C. A. Dengel and Dr. H. M. Liwicki. In bachelor's thesis, he worked in the domains of real time object detection and image enhancement. In master's thesis, he primarily worked for the development of various systems for signature identification

and verification. His Ph.D. topic was automated forensic handwriting analysis on which he focused on both the perspectives of forensic handwriting examiners and pattern recognition researchers. He is currently an Assistant Professor with the National University of Sciences and Technology, Pakistan. He has more than 35 publications on the said and related topics including two journal papers.

Sheraz Ahmed received the master's degree in computer science from the Technische Universität Kaiserslautern, Germany, and the Ph.D. degree from the German Research Center for Artificial Intelligence, Germany, under the supervision of Prof. Dr. Prof. H. C. Andreas Dengel and Prof. Dr. habil. Marcus Liwicki. His Ph.D. topic is Generic Methods for Information Segmentation in Document Images. He is a Senior Researcher with the German Research Center for Artificial Intelligence (DFKI GmbH), Kaiserslautern, where he is leading the area of time series and document analysis.

From 2012 to 2013, he visited Osaka Prefecture University, Osaka, Japan, as a Research Fellow, supported by the Japanese Society for the Promotion of Science, and from 2014 to 2014, he visited the University of Western Australia, Perth, Australia, as a Research Fellow, supported by the DAAD, Germany and Go8, Australia. His research interests include document understanding, generic segmentation framework for documents, gesture recognition, pattern recognition, data mining, anomaly detection, and natural language processing. He has more than 30 publications on the said and related topics, including three journal papers and two book chapters. He is a frequent Reviewer of various journals and conferences, including Pattern Recognition Letters, Neural Computing and Applications, IJDAR, ICDAR, ICIFHR, and DAS.

Ludger van Elst is a researcher, consultant, and project manager for more than 20 years at the German Research Center for Artificial Intelligence (DFKI GmbH) in the Smart Data and Knowledge Services group of Prof. Andreas Dengel. He studied computer science, economics, and mathematics at the University of Kaiserslautern from where he also received his Diploma degree. His scientific background is in the fields of semantic technologies, especially formal knowledge representation, ontologies, knowledge acquisition, and knowledge management. He has authored, co-authored, and edited more than 60 scientific publications on these topics. His current interests and projects focus on semantic sensor data representation and analysis in Industrie 4.0 applications as well as in deep learning for applications in the automotive industry. His guiding vision is to develop technologies that better integrate data-driven and model-based approaches to create assistance systems that continuously learn and improve.

His current interests and projects focus on semantic sensor data representation and analysis in Industrie 4.0 applications as well as in deep learning for applications in the automotive industry. His guiding vision is to develop technologies that better integrate data-driven and model-based approaches to create assistance systems that continuously learn and improve.