

AN EFFICIENT NETWORK FOR PREDICTING TIME-VARYING DISTRIBUTIONS

Anonymous authors

Paper under double-blind review

ABSTRACT

While deep neural networks have achieved groundbreaking prediction results in many tasks, there is a class of data where existing architectures are not optimal – sequences of probability distributions. Performing forward prediction on sequences of distributions has many important applications. However, there are two main challenges in designing a network model for this task. First, neural networks are unable to encode distributions compactly as each node encodes just a real value. A recent work of Distribution Regression Network (DRN) solved this problem with a novel network that encodes an entire distribution in a single node, resulting in improved accuracies while using much fewer parameters than neural networks. However, despite its compact distribution representation, DRN does not address the second challenge, which is the need to model time dependencies in a sequence of distributions. In this paper, we propose our Recurrent Distribution Regression Network (RDRN) which adopts a recurrent architecture for DRN. The combination of compact distribution representation and shared weights architecture across time steps makes RDRN suitable for modeling the time dependencies in a distribution sequence. Compared to neural networks and DRN, RDRN achieves the best prediction performance while keeping the network compact.

1 INTRODUCTION

Deep neural networks have achieved state-of-the-art results in many tasks by designing the network architecture according to the data type. For instance, the convolutional neural network (CNN) uses local filters to capture the features in an image and max pooling to reduce the image representation size. By using a series of convolution and max pooling layers, CNN extracts the semantic meaning of the image. The recurrent architecture of recurrent neural networks (RNN) when unrolled, presents a shared weight structure which is designed to model time dependencies in a data sequence. However, among the major network architectures, the multilayer perceptron, convolutional neural network and recurrent neural network, there is no architecture suitable for representing sequences of probability distributions. Specifically, we address the task of forward prediction on distribution sequences.

There are two main challenges in designing a network for sequences of probability distributions. First, conventional neural networks are unable to represent distributions compactly. Since each node encodes only a real value, a distribution has to be decomposed to smaller parts that are represented by separate nodes. When the distribution has been decomposed into separate nodes, the notion of distribution is no longer captured explicitly. Similarly, for image data, the fully-connected multilayer perceptron (MLP), unlike convolutional neural networks, fails to capture the notion of an image. A recently proposed network, Distribution Regression Network (DRN) (Kou et al., 2018), has solved this problem. DRN uses a novel representation of encoding an entire distribution in a single node, allowing DRN to use more compact models while achieving superior performance for distribution regression. It has been shown that DRN can achieve better accuracies with 500 times fewer parameters compared to MLP. However, despite the strengths of DRN, it is a feedforward network and hence it does not address a second problem, which is the need to model time dependencies in a distribution sequence.

We address these two challenges and propose a recurrent extension of DRN, named the Recurrent Distribution Regression Network (RDRN). In the hidden states of RDRN, each node represents a distribution, thus containing much richer information while using fewer weights compared to the

real-valued hidden states in RNN. This compact representation consequently results in better generalization performance. Compared to DRN, the shared weights in RDRN captures time dependencies better and results in better prediction performance. By having both compact distribution representations and modeling of time dependencies, RDRN is able to achieve superior prediction performance compared to the other methods.

2 RELATED WORK

Performing forward prediction on time-varying distributions has many important applications. Many real-world systems are driven by stochastic processes. For such systems, the Fokker-Planck equation (Risken, 1996) has been used to model the time-varying distribution, with applications in astrophysics (Noble & Wheatland, 2011), biological physics (Gu erin et al., 2011), animal population studies (Butler & King, 2004) and weather forecasting (Palmer, 2000). In these applications, it is very useful to predict the future state of the distribution. For example, the Ornstein-Uhlenbeck process, which is a specific case of the Fokker-Planck equation, has been used to model and predict commodity prices (Schwartz & Smith, 2000). Extrapolating a time-varying distribution into the future has also been used for predictive domain adaptation, where a classifier is trained on data distribution which drifts over time (Lampert, 2015).

Various machine learning methods have been proposed for distribution data, ranging from distribution-to-real regression (P oczoz et al., 2013; Oliva et al., 2014) to distribution-to-distribution regression (Oliva et al., 2015; 2013). The Triple-Basis Estimator (3BE) has been proposed for the task of function-to-function regression. It uses basis representations of functions and learns a mapping from Random Kitchen Sink basis features (Oliva et al., 2015). The authors have applied 3BE for distribution regression, showing improved accuracy and speed compared to an instance-based learning method (Oliva et al., 2013). More recently, Kou et al. (2018) proposed the Distribution Regression Network which extends the neural network representation such that an entire distribution is encoded in a single node. With this compact representation, DRN showed better accuracies while using much fewer parameters than conventional neural networks and 3BE (Oliva et al., 2015).

The above methods are for general distribution regression. For predicting the future state of a time-varying distribution, it is important to model the time dependencies in the distribution sequence. Lampert (2015) proposed the Extrapolating the Distribution Dynamics (EDD) method which predicts the future state of a time-varying distribution given a sequence of samples from previous time steps. EDD uses the reproducing kernel Hilbert space (RKHS) embedding of distributions and learns a linear mapping to model the dynamics of how the distribution evolves between adjacent time steps. EDD is shown to work for a few variants of synthetic data, but the performance deteriorates for tasks where the dynamics is non-linear. Since the regression is performed with just one input time step, it is unclear how EDD can be extended for more complex trajectories that require multiple time steps of history. Another limitation is that the EDD can only learn a single trajectory of the distribution and not from multiple trajectories.

3 RECURRENT DISTRIBUTION REGRESSION NETWORK (RDRN)

3.1 TIME-SERIES DISTRIBUTION REGRESSION

We address the task of forward prediction from a time-varying distribution: Given a series of distributions with T equally-spaced time steps, $X^{(1)}, X^{(2)}, \dots, X^{(T)}$, we want to predict $X^{(T+k)}$, ie. the distribution at k time steps later. We assume the distributions to be univariate. The input at each time step may consist of more than one distribution, for instance, when tracking multiple independent distributions over time. In this case, the input distribution sequence is denoted as $(X_1^{(1)}, \dots, X_n^{(1)}), \dots, (X_1^{(T)}, \dots, X_n^{(T)})$, where there are n data distributions per time step.

Performing prediction on distribution sequences requires both compact distribution representations and modeling of time dependencies. While the recurrent neural network works well for time series data, it has no efficient representation for distributions. As for DRN, although it has a compact representation for distributions, the feedforward architecture does not capture the time dependencies in the distribution sequence. Hence, we propose our Recurrent Distribution Regression Network (RDRN) which is a recurrent extension of DRN.

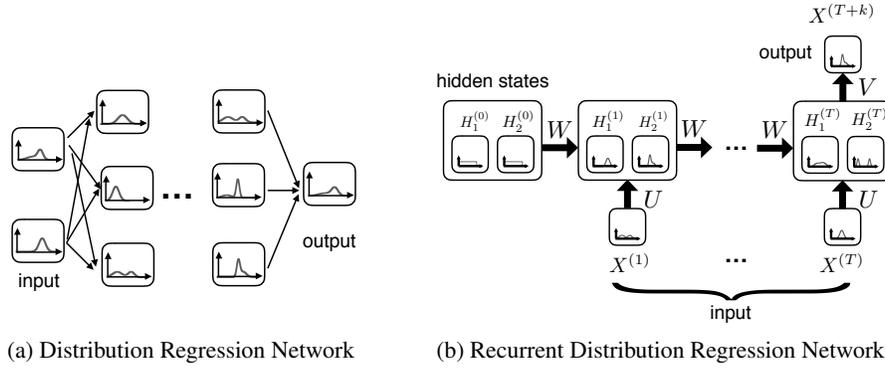


Figure 1: (a) A schematic of Distribution Regression Network (DRN) which performs distribution regression by encoding each node with an entire distribution. (b) An example network for Recurrent Distribution Regression Network (RDRN) which takes in T time steps of distributions to predict the distribution at $T + k$. The arrows represent fully-connected weights. The input-hidden weights U and the hidden-hidden weights W are shared across time steps. V is the weights between the final hidden state and the output distribution. The hidden states at $t = 0$ are initialized as uniform distributions.

3.2 BACKGROUND: DISTRIBUTION REGRESSION NETWORK (DRN)

Neural network models work well if the network architecture is designed according to the data type. Convolutional neural networks are suited for image data as they employ convolution to capture local features from neighboring image pixels. Such important data domain knowledge is not built in the fully-connected multilayer perceptron. For analysis of distributions, there are no conventional neural network architectures like what CNN does for images. To that end, Kou et al. (2018) proposed the Distribution Regression Network (DRN) for the task of distribution-to-distribution regression. To cater to distribution data, DRN has two main innovations: 1) each network node encodes an entire distribution and 2) the forward propagation is specially designed for propagating distributions, with a form inspired by statistical physics.

We give a brief description of DRN following the notations of Kou et al. (2018). Figure 1a illustrates the propagation in DRN. Similar to MLP, DRN consists of multiple fully-connected layers connecting the data input to the output in a feedforward manner, where each connection has a real-valued weight. The novelty of DRN is that each node in the network encodes an entire probability distribution. The distribution at each node is computed using the distributions of the incoming nodes, the weights and the bias parameters. Let $P_k^{(l)}$ represent the probability density function (pdf) of the k^{th} node in the l^{th} layer where $P_k^{(l)}(s_k^{(l)})$ is the density of the pdf when the node variable is $s_k^{(l)}$. The unnormalized distribution $\tilde{P}_k^{(l)}$ is computed by marginalizing over the product of the unnormalized conditional probability $\tilde{Q}(s_k^{(l)} | s_1^{(l-1)}, \dots, s_n^{(l-1)})$ and the incoming probabilities.

$$\tilde{P}_k^{(l)}(s_k^{(l)}) = \int_{s_1^{(l-1)}} \dots \int_{s_n^{(l-1)}} \tilde{Q}(s_k^{(l)} | s_1^{(l-1)}, \dots, s_n^{(l-1)}) P_1^{(l-1)}(s_1^{(l-1)}) \dots P_n^{(l-1)}(s_n^{(l-1)}) ds_1^{(l-1)} \dots ds_n^{(l-1)} \quad (1)$$

$$\tilde{Q}(s_k^{(l)} | s_1^{(l-1)}, \dots, s_n^{(l-1)}) = e^{-E(s_k^{(l)} | s_1^{(l-1)}, \dots, s_n^{(l-1)})} \quad (2)$$

E is the energy for a given set of node variables,

$$E(s_k^{(l)} | s_1^{(l-1)}, \dots, s_n^{(l-1)}) = \sum_i^n w_{ki}^{(l)} \left(\frac{s_k^{(l)} - s_i^{(l-1)}}{\Delta} \right)^2 + b_{q,k}^{(l)} \left(\frac{s_k^{(l)} - \lambda_{q,k}^{(l)}}{\Delta} \right)^2 + b_{a,k}^{(l)} \left| \frac{s_k^{(l)} - \lambda_{a,k}^{(l)}}{\Delta} \right|, \quad (3)$$

where $w_{ki}^{(l)}$ is the weight connecting the i^{th} node in layer $l - 1$ to the k^{th} node in layer l . $b_{q,k}^{(l)}$ and $b_{a,k}^{(l)}$ are the quadratic and absolute bias terms acting on positions $\lambda_{q,k}^{(l)}$ and $\lambda_{a,k}^{(l)}$ respectively. Δ is the support length of the distribution. After obtaining the unnormalized probability, the distribution from Eq. (1) is normalized. Forward propagation is performed layer-wise obtain the output prediction. With such a propagation method, DRN exhibits useful propagation properties such as peak spreading, peak shifting, peak splitting and the identity mapping (Kou et al., 2018). Due to space constraints, we refer the readers to Kou et al. (2018) for a more detailed description.

3.3 EXTENSION TO RECURRENT ARCHITECTURE

Since DRN is a feedforward network, it does not explicitly capture the time dependencies in distribution sequences. In this work, we introduce our Recurrent Distribution Regression Network (RDRN) which is a recurrent extension of DRN. The input data is a distribution sequence as described in Section 3.1. Figure 1b shows an example network for RDRN, where the network takes in T time steps of distributions to predict the distribution at $T + k$. The hidden state at each time step may consist of multiple distributions. The arrows represent fully-connected weights. The input-hidden weights U and the hidden-hidden weights W are shared across the time steps. V represents the weights between the final hidden state and the output distribution. The bias parameters for the hidden state nodes are also shared across the time steps. The hidden state distributions at $t = 0$ represents the ‘memory’ of all past time steps before the first input and can be initialized with any prior information. In our experiments, we initialize the $t = 0$ hidden states as uniform distributions as we assume no prior information is known.

We formalize the propagation for the general case where there can be multiple distributions for each time step in the data input layer and the hidden layer. Let n and m be the number of distributions per time step in the data layer and hidden layers respectively. Propagation starts from $t=1$ and performed through the time steps to obtain the hidden state distributions. $X_i^{(t)}(r_i^{(t)})$ represents the input data distribution at node i and time step t , when the node variable is $r_i^{(t)}$. $H_k^{(t)}(s_k^{(t)})$ represents the density of the pdf of the k^{th} hidden node at time step t when the node variable is $s_k^{(t)}$. $\tilde{H}_k^{(t)}(s_k^{(t)})$ represents the unnormalized form. The hidden state distributions at each time step is computed from the hidden state distributions from the previous time step and the input data distribution from the current time step.

$$\tilde{H}_k^{(t)}(s_k^{(t)}) = \int_{r_1^{(t)}, \dots, r_n^{(t)}, s_1^{(t-1)}, \dots, s_m^{(t-1)}} \tilde{Q}(s_k^{(t)} | r_1^{(t)}, \dots, s_1^{(t-1)}, \dots) \quad (4)$$

$$X_1^{(t)}(r_1^{(t)}) \dots X_n^{(t)}(r_n^{(t)}) H_1^{(t-1)}(s_1^{(t-1)}) \dots H_m^{(t-1)}(s_m^{(t-1)}) dr_1^{(t)} \dots dr_n^{(t)} ds_1^{(t-1)} \dots ds_m^{(t-1)}$$

$$\tilde{Q}(s_k^{(t)} | r_1^{(t)}, \dots, s_1^{(t-1)}, \dots) = e^{-E(s_k^{(t)} | r_1^{(t)}, \dots, s_1^{(t-1)}, \dots)} \quad (5)$$

The energy function is given by

$$E(s_k^{(t)} | r_1^{(t)}, \dots, s_1^{(t-1)}, \dots) = \sum_i^n u_{ki} \left(\frac{s_k^{(t)} - r_i^{(t)}}{\Delta} \right)^2 + \sum_j^m w_{kj} \left(\frac{s_k^{(t)} - s_j^{(t-1)}}{\Delta} \right)^2 \quad (6)$$

$$+ b_{q,k} \left(\frac{s_k^{(t)} - \lambda_{q,k}}{\Delta} \right)^2 + b_{a,k} \left| \frac{s_k^{(t)} - \lambda_{a,k}}{\Delta} \right|,$$

where for each time step, u_{ki} is the weight connecting the i^{th} input distribution to the k^{th} hidden node. Similarly, for the hidden-hidden connections, w_{kj} is the weight connecting the j^{th} hidden node in the previous time step to the k^{th} hidden node in the current time step. As in DRN, the hidden node distributions are normalized before propagating to the next time step. At the final time step, the output distribution is computed from the hidden state distributions, through the weight vector V and bias parameters at the output node.

3.4 NETWORK COST AND OPTIMIZATION

Following Kou et al. (2018), the cost function for the forward prediction task is measured by the Jensen-Shannon (JS) divergence (Lin, 1991) between the label and output distributions. Optimization is performed by backpropagation through time. We adopt the same parameter initialization method as Kou et al. (2018), where the network weights and bias are randomly initialized following a uniform distribution and the bias positions are uniformly sampled from the support length of the data distribution. The integrals in Eq. (4) are performed numerically, by partitioning the distribution into q bins, resulting in a discrete probability mass function.

4 EXPERIMENTS

We conducted experiments on four datasets which involve prediction of time-varying distributions. To evaluate the effectiveness of the recurrent structure in RDRN, we compare with DRN where the input distributions for all time steps are concatenated at the input layer. We also compare with conventional neural network architectures and other distribution regression methods. The benchmark methods are DRN, RNN, MLP and 3BE. For the final dataset, we also compare with EDD as the data involves only a single trajectory of distribution. Among these methods, RNN and EDD are designed to take in the inputs sequentially over time while for the rest the inputs from all T time steps are concatenated.

4.1 BENCHMARK METHODS

Distribution Regression Network (DRN) Since DRN is a feedforward network, the distributions for all input time steps are concatenated and fed in together. The architecture consists of fully-connected layers, where each node encodes an entire distribution. DRN is optimized using JS divergence.

Recurrent Neural Network (RNN) At each time step, the distribution is discretized into bins and represented by separate input nodes. The RNN architecture consists of a layer of hidden states, where the number of nodes is chosen by cross validation. The input-hidden and hidden-hidden weights are shared across time steps. The final hidden state is transformed by the hidden-output weights and processed by a softmax layer to obtain the output distribution. The cost function is the mean squared error between the predicted and output distribution bins.

Multilayer Perceptron (MLP) The input layer consists of the input distributions for all time steps and each distribution is discretized into bins that are represented by separate nodes. Hence, for T input time steps and discretization size of q , there will be $T \times q$ input nodes in MLP. MLP consists of fully-connected layers and a final softmax layer, and is optimized with the mean squared error.

Triple-Basis Estimator (3BE) For 3BE, each distribution is represented by its sinusoidal basis coefficients. The number of basis coefficients and number of Random Kitchen Sink basis functions are chosen by cross validation.

Extrapolating the Distribution Dynamics (EDD) Since EDD learns from a single trajectory of distribution, it is unsuitable for most of the datasets. We performed EDD for the final dataset which has only a single distribution trajectory. For the RKHS embedding of the distributions, we use the radial basis function kernel, following Lampert (2015).

4.2 SHIFTING GAUSSIAN

For the first experiment, we chose a dataset where the output distribution has to be predicted from multiple time steps of past distributions. Specifically, we track a Gaussian distribution whose mean varies in the range $[0.2, 0.8]$ sinusoidally over time while the variance is kept constant at 0.01. Given a few consecutive input distributions taken from time steps spaced $\Delta t = 0.2$ apart, we predict the next time step distribution. Figure 2a illustrates how the mean changes over time. It is apparent that we require more than one time step of past distributions to predict the future distribution. For instance, at two different time points, the distribution means can be the same, but one has increasing mean while the other has a decreasing mean.

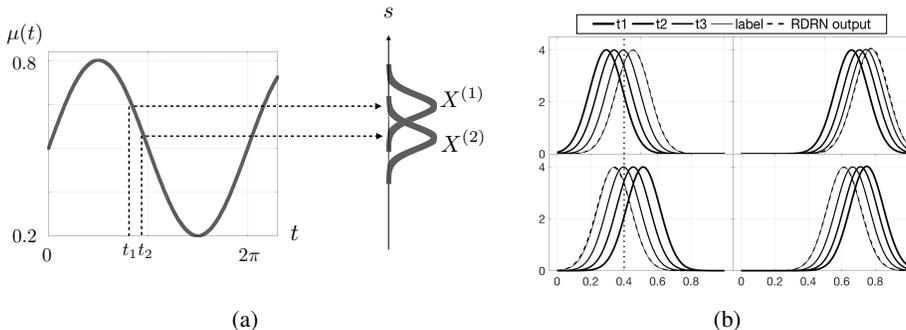


Figure 2: (a) Shifting Gaussian dataset: The mean of the Gaussian distribution varies sinusoidally with time, hence forward prediction requires more than one time step of past distributions. (b) RDRN’s prediction for four test data for the shifting Gaussian dataset shows a good fit with the labeled output. The top and bottom left data have the same mean at $t = 3$, but are moving in opposite directions, showing that more than one input time steps are required for this task.

	Shifting Gaussian			Climate Model		
	$L2(10^{-2})$	T	N_p	$L2(10^{-2})$	T	N_p
RDRN	4.55(0.42)	3	59	11.98(0.13)	5	59
DRN	4.90(0.46)	3	224	12.27(0.34)	3	44
RNN	17.50(0.89)	3	2210	13.29(0.59)	5	12650
MLP	10.32(0.41)	3	1303	13.52(0.25)	3	22700
3BE	22.30(1.89)	3	6e+5	14.18(1.29)	5	2.2e+5

Table 1: Performance of RDRN and other methods for the shifting Gaussian and climate model datasets. $L2$ denotes the $L2$ loss, T is the optimal number of input time steps and N_p is the number of model parameters used. Lower loss values reflect better regression accuracies. The number in the parentheses is the standard error of the performance measures, over repeated runs.

To create the dataset, for each data we randomly sample the first time step from $[0, 2\pi]$. The distributions are truncated with support of $[0, 1]$ and discretized with $q=100$ bins. We found that for all methods, a history length of 3 time steps is optimal. Following Oliva et al. (2014) the regression performance is measured by the $L2$ loss. Table 1 shows the regression results, where lower $L2$ loss is favorable. 20 training data was sufficient for RDRN and DRN to give good predictions. RDRN’s regression accuracy is the best, followed by DRN. Figure 2b shows four test data, where the input distributions at $t=1, 2, 3$ are shown, along with the label output for $t=4$ and RDRN’s prediction. We observe good fit for the predictions. Additionally, the top and bottom left data shows that two data can have the same mean at $t=3$, but are moving in opposite directions. Hence, to predict the next distribution at $t=4$, multiple time steps in history are required as input and the model has to determine the direction of movement from the history of distributions. Since RDRN is designed to model time dependencies in the distribution sequence, it is able to infer the direction of the mean movement well. In contrast, the neural network counterparts of RNN and MLP showed considerable overfitting which is likely due to the fact that excessive number of nodes are used to represent the distributions, resulting in many model parameters.

4.3 CLIMATE MODELS

In the field of climate modeling, variability of climate measurements due to noise is an important factor to consider (Hasselmann, 1976). The next experiment is based on the work of Lin & Koshyk (1987), where they model the heat flux at the sea surface as a time-varying one-dimensional distribution. Specifically, the evolution of the heat flux over time obeys the stochastic Ornstein-Uhlenbeck (OU) process (Uhlenbeck & Ornstein, 1930), and the diffusion and drift coefficients are determined from real data measurements obtained by Oort & Rasmusson (1971).

	CarEvolution				Stock			
	<i>NLL</i>	<i>T</i>	N_p		<i>NLL</i> (1 day)	<i>NLL</i> (10 days)	<i>T</i>	N_p
RDRN	3.9660(3e-4)	2	12313	RDRN	-473.71(0.18)	-458.36(0.20)	3	37
DRN	3.9663(2e-5)	2	28676	DRN	-474.31(0.01)	-457.38 (0.001)	1	9
MLP	3.9702(6e-4)	2	1.2e+7	RNN	-471.93(0.04)	-457.98(0.19)	3	4210
3BE	3.9781(0.003)	2	1.2e+7	MLP	-471.47(0.15)	-454.73(3.75)	3	10300
EDD	4.0405	1	64	3BE	-466.69(0.84)	-385.43(8.03)	1	14000

(a)

(b)

Table 2: Performance of RDRN and other methods for the (a) CarEvolution and (b) stock dataset. *NLL*: negative log-likelihood, *T*: optimal number of input time steps, N_p : number of model parameters used. Lower loss values reflect better regression accuracies.

The OU process is described by a time-varying Gaussian distribution. With the long-term mean set at zero, the pdf has a mean of $\mu(t) = y \exp(-\theta t)$ and variance of $\sigma^2(t) = \frac{D(1-e^{-2\theta t})}{\theta}$. t represents time, y is the initial point mass position, and D and θ are the diffusion and drift coefficients respectively. For the energy balance climate model, $D = 0.0013$, $\theta = 2.86$, and each unit of the nondimensional time corresponds to 55 days (Lin & Koshyk, 1987). At $t = 0$, the distribution is a delta-function at position y . To create a distribution sequence, we first sample $y \in [0.02, 0.09]$. For each sampled y , we generate 6 Gaussian distributions at $t_0 - 4\delta$, $t_0 - 3\delta$, ..., t_0 and $t_0 + 0.02$, with $\delta = 0.001$ and t_0 sampled uniformly from $[0.01, 0.05]$. The Gaussian distributions are truncated with support of $[-0.01, 0.1]$.

The regression task is as follows: Given the distributions at $t_0 - 4\delta$, $t_0 - 3\delta$, ..., t_0 , predict the distribution at $t_0 + 0.02$. With different sampled values for y and t_0 , we created 100 training and 1000 test data. The regression performance is measured by the $L2$ loss. The regression results on the test set are shown in Table 1. RDRN’s regression accuracy is the best, followed by DRN. This is followed by the neural network architectures MLP and RNN. It is noteworthy that RDRN and RNN, which explicitly capture time dependencies in the architecture, perform better than their feedforward counterparts. In addition, the recurrent models perform best with more time steps ($T=5$) compared to the feedforward models ($T=3$), which may suggest that the recurrent architecture captures the time dependencies in the data sequence better than a feedforward one. In terms of model compactness, RDRN and DRN use at 2-3 orders of magnitude fewer model parameters compared to the other methods, owing to the compact distribution representation.

4.4 CAREVOLUTION DATA

RDRN can be used to track the distribution drift of image datasets. For the next experiment, we use the CarEvolution dataset (Rematas et al., 2013) which was used by Lampert (2015) for the domain adaptation problem. The dataset consists of 1086 images of cars manufactured from the years 1972 to 2013. We split the data into intervals of 5 years (ie. 1970-1975, 1975-1980, ..., 2010-2015) where each interval has an average of 120 images. This gives 9 time intervals and for each interval, we create the data distribution from the DeCAF(fc6) features (Donahue et al., 2014) of the car images using kernel density estimation. The DeCAF features have 4096 dimensions. Performing accurate density estimation in very high dimensions is challenging due to the curse of dimensionality (Gu et al., 2013). Here we make the approximation that the DeCAF features are independent, such that the joint probability is a product of the individual dimension probabilities.

The regression task is to predict the next time step distribution of features given the previous T time step distributions. We found $T=2$ to work best for all methods. The first 7 intervals were used for the train set while the last 3 intervals were used for the test set, giving 5 training and 1 test data. The regression performance is measured by the negative log-likelihood of the test samples following Oliva et al. (2013), where lower negative log-likelihood is favorable. The regression results are shown in Table 2a. RDRN has the best prediction performance, followed by DRN. RNN had difficulty in optimization possibly due to the high number of input dimensions, so the results are not presented. EDD has the fewest number of parameters as it assumes the dynamics of the distribution follows a linear mapping between the RKHS features of consecutive time steps (ie. $T=1$). However, as the results show, the EDD model may be too restrictive for this dataset.

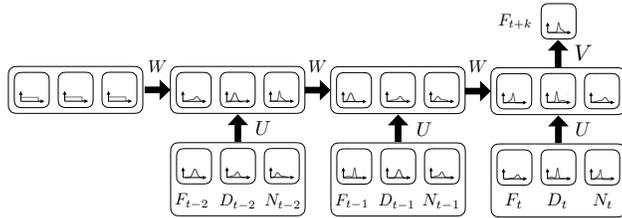


Figure 3: RDRN network for the stock dataset: past 3 days of distribution of returns of constituent companies in FTSE, DOW and Nikkei were used as inputs, to predict the next day’s distribution of returns for constituent companies in FTSE. One layer of hidden states is used, with 3 nodes per hidden state.

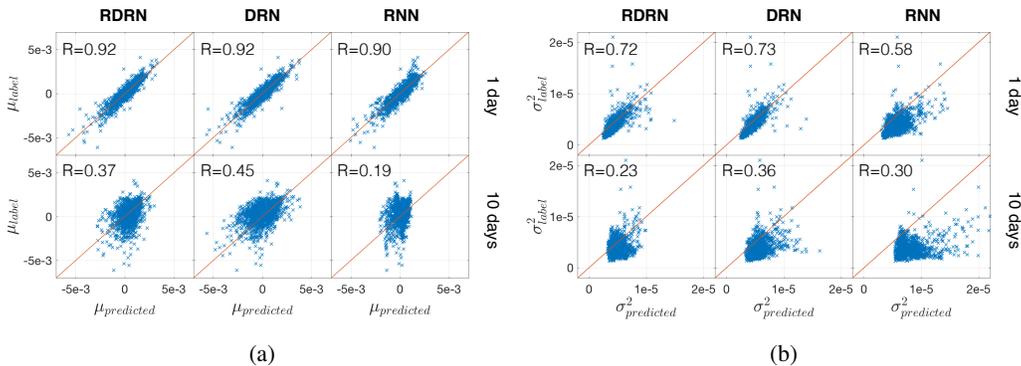


Figure 4: Comparison of the (a) mean and (b) variance of the label and predicted distributions for 1 and 10 days ahead stock prediction. The diagonal line represents a perfect fit. R represents the correlation coefficient.

4.5 STOCK PREDICTION

We show RDRN is useful for forward prediction of price movements in stock markets. We adopt a similar experimental setup as Kou et al. (2018). There have been studies that show that movement of indices of stock markets in the world correlate with each other, providing a basis for predicting future stock returns (Hamao et al., 1990; Chong et al., 2008). Specifically, the previous day stock returns of the Nikkei and Dow Jones Industrial Average (Dow) are found to be good predictors of the FTSE return (Vega & Smolarski, 2012). Furthermore, predicting the entire distribution of stock returns has been found to be more useful for portfolio selection compared to just a single index value (Cenesizoglu & Timmermann, 2008).

Following the setup in Kou et al. (2018), our regression task is as follows: given the past T days’ distribution of returns of constituent companies in FTSE, Dow and Nikkei, predict the distribution of returns for constituent companies in FTSE k days later. We used 9 years of daily returns from 2007 to 2015 and performed exponential window averaging on the price series following common practice (Murphy, 1999). The regression performance is measured by the negative log-likelihood of the test samples. The RDRN architecture used is shown in Figure 3, where the data input consists past 3 days of distribution returns and one layer of hidden states with 3 nodes per time step is used.

We tested on forward prediction of 1 and 10 days ahead. Table 2b shows the regression results. As before, RDRN and DRN’s performance surpasses the other methods by a considerable margin. For 1 day ahead prediction, RDRN’s performance is slightly below DRN, but for 10 days ahead, RDRN’s performance is better. This may suggest that the 1 day ahead task is simpler and does not involve long time dependencies. On the other hand, predicting 10 days ahead is a more challenging task which may benefit from having a longer history of stock movements.

We further visualize the results by comparing the mean and variance of the predicted and the label distributions, as shown in Figure 4. Each point represents one test data and we show the correlation coefficients between the predicted and labeled moments. As expected, the regression for all methods deteriorates for the 10 days ahead prediction. RDRN and DRN have the best regression performance as the points lie closest to the diagonal line. For the 10 days ahead task, the predicted distributions for RDRN are much better predicted than the other methods, showing RDRN’s strength in predicting with longer time steps ahead. On the other hand, RNN shows some sign of regression to the mean, as the means of the output distributions are limited to a small range about zero.

5 DISCUSSION

Neural network models work well by designing the architecture according to the data type. However, among the conventional neural network architectures, there is none that is designed for time-varying probability distributions. There are two key challenges in learning from distribution sequences. First, we require a suitable representation for probability distributions. Conventional neural networks, however, do not have suitable representations for distributions. As each node encodes only a real value, the distribution has to be split into smaller parts which are then represented by independent nodes. Hence, the neural network is agnostic to the distribution nature of the input data. A recently proposed Distribution Regression Network (DRN) addresses this issue. DRN has a novel network representation where each node encodes a distribution, showing improved accuracies compared to neural networks. However, a second challenge remains, which is to model the time dependencies in the distribution sequence. Both the recurrent neural network (RNN) and the Distribution Regression Network address only either one of the challenges. In this work, we propose our Recurrent Distribution Regression Network (RDRN) which extends DRN with a recurrent architecture. By having an explicit distribution representation in each node and shared weights across time steps, RDRN performs forward prediction on distribution sequences most effectively, achieving better prediction accuracies than RNN, DRN and other regression methods.

REFERENCES

- Marguerite A Butler and Aaron A King. Phylogenetic comparative analysis: a modeling approach for adaptive evolution. *The American Naturalist*, 164(6):683–695, 2004.
- Tolga Cenesizoglu and Allan G Timmermann. Is the distribution of stock returns predictable? *Available at SSRN 1107185*, 2008.
- Terence Tai-Leung Chong, Ying-Chiu Wong, and Isabel Kit-Ming Yan. International linkages of the Japanese stock market. *Japan and the World Economy*, 20(4):601–621, 2008.
- Jeff Donahue, Yangqing Jia, Oriol Vinyals, Judy Hoffman, Ning Zhang, Eric Tzeng, and Trevor Darrell. Decaf: A deep convolutional activation feature for generic visual recognition. In *International conference on machine learning*, pp. 647–655, 2014.
- Chong Gu, Yongho Jeon, and Yi Lin. Nonparametric density estimation in high-dimensions. *Statistica Sinica*, pp. 1131–1153, 2013.
- T Guérin, J Prost, and J-F Joanny. Bidirectional motion of motor assemblies and the weak-noise escape problem. *Physical Review E*, 84(4):041901, 2011.
- Yasushi Hamao, Ronald W Masulis, and Victor Ng. Correlations in price changes and volatility across international stock markets. *Review of Financial studies*, 3(2):281–307, 1990.
- Klaus Hasselmann. Stochastic climate models part i. theory. *tellus*, 28(6):473–485, 1976.
- C. Kou, H. K. Lee, and T. K. Ng. A Compact Network Learning Model for Distribution Regression. *arXiv preprint arXiv:1804.04775v3*, April 2018.
- Christoph H Lampert. Predicting the future behavior of a time-varying probability distribution. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 942–950, 2015.

- Charles A Lin and John N Koshyk. A nonlinear stochastic low-order energy balance climate model. *Climate dynamics*, 2(2):101–115, 1987.
- Jianhua Lin. Divergence measures based on the Shannon entropy. *IEEE Transactions on Information theory*, 37(1):145–151, 1991.
- John J Murphy. Technical analysis of the futures markets: A comprehensive guide to trading methods and applications, New York Institute of Finance, 1999.
- PL Noble and MS Wheatland. Modeling the sunspot number distribution with a fokker-planck equation. *The Astrophysical Journal*, 732(1):5, 2011.
- Junier Oliva, William Neiswanger, Barnabás Póczos, Eric Xing, Hy Trac, Shirley Ho, and Jeff Schneider. Fast function to function regression. In *Artificial Intelligence and Statistics*, pp. 717–725, 2015.
- Junier B Oliva, Barnabás Póczos, and Jeff G Schneider. Distribution to distribution regression. In *ICML (3)*, pp. 1049–1057, 2013.
- Junier B Oliva, Willie Neiswanger, Barnabás Póczos, Jeff G Schneider, and Eric P Xing. Fast distribution to real regression. In *AISTATS*, pp. 706–714, 2014.
- Abraham H Oort and Eugene M Rasmusson. *Atmospheric circulation statistics*, volume 5. US Government Printing Office, 1971.
- Tim N Palmer. Predicting uncertainty in forecasts of weather and climate. *Reports on progress in Physics*, 63(2):71, 2000.
- Barnabás Póczos, Aarti Singh, Alessandro Rinaldo, and Larry A Wasserman. Distribution-free distribution regression. In *AISTATS*, pp. 507–515, 2013.
- Konstantinos Rematas, Basura Fernando, Tatiana Tommasi, and Tinne Tuytelaars. Does evolution cause a domain shift? In *Proceedings VisDA 2013*, pp. 1–3, 2013.
- Hannes Risken. Fokker-planck equation. In *The Fokker-Planck Equation*, pp. 63–95. Springer, 1996.
- Eduardo Schwartz and James E Smith. Short-term variations and long-term dynamics in commodity prices. *Management Science*, 46(7):893–911, 2000.
- George E Uhlenbeck and Leonard S Ornstein. On the theory of the Brownian motion. *Physical review*, 36(5):823, 1930.
- Jose G Vega and Jan M Smolarski. Forecasting FTSE Index using global stock markets. *International Journal of Economics and Finance*, 4(4):3, 2012.