DL and the Weather Forecasting Problem - Precipitation Nowcasting

Zhihan Gao, Xingjian Shi, Hao Wang, Dit-Yan Yeung,
Wang-chun Woo, Wai-Kin Wong

Feb. 29th, 2020

Chapter 18

Deep Learning and the Weather Forecasting Problem - Precipitation Nowcasting

18.1 Introduction

Precipitation nowcast refers to the forecasting of rainfall and other types of precipitation up to 6 hours ahead (as defined by the World Meteorological Organization). Since rainfall can be localized and highly-changeable, users of precipitation nowcast typically demand to know the exact time, location and intensity of rainfall. It is therefore necessary to make very high resolution, both spatially and temporally, precipitation nowcast products in a timely manner, typically in the order of minutes. The most important use of precipitation nowcast is to support the operations of rainstorm warning systems managed by meteorological services around the world. Rainstorm warning systems provide early alerts to the public, disaster risk reduction agencies, government departments in particular those related to public security and works, as well as managers of infrastructures and facilities. Upon the issuance of rainstorm warnings, these parties take actions according to their own standard operating procedures with a view to saving lives and protecting properties. It has tremendous impact on various areas from aviation service, public safety, to people's daily life. For example, commercial airlines rely on precipitation nowcasting to predict extreme weather events and ensures flight safety. On land, heavy rainfall severely can affect the road conditions and increases the risk of traffic accidents, which can be avoided with the help of precipitation nowcasting. For local businesses, the number of customers and their feedback about a restaurant are largely related to the weather [7], especially the rain rate. Thus, accurate and timely prediction

of rainfall helps restaurants predict and adjust their sales strategies. There-
fore, the past years have seen an ever-growing need for real-time, large-scale,
and fine-grained precipitation nowcasting [42, 43, 31, 31, 1]. Due to the in-
herent complexities of the atmosphere and relevant dynamical processes, the
problem imposes new challenge to the meteorological community [47].

Traditionally, precipitation nowcasting is approached by either *Optical
Flow* (OF) based methods [34, 39] or numerical methods [56, 47, 6]. OF
based methods first estimate the flow field, which represents the convective
motion of the precipitation, with the observed weather data (e.g., the *Con-
stant Altitude Plan Position Indicator* (CAPPI) radar echo maps [15]) and
then use the flow field for extrapolation [57]. The numerical methods build
mathematical models of the atmosphere on top of the physical principles
such as the dynamic and thermodynamic laws. Future rainfall intensities
are predicted by numerically solving partial differential equations within the
mathematical models. However, both approaches have deficiencies that limit
their success. The OF based methods attempt to identify the convective mo-
tion of the cloud, but they fail to represent cloud initiation or decay and lack
the ability of expressing strong nonlinear dynamics. In addition, the flow
field estimation step and the radar echo extrapolation step are separated,
making it challenging to determine the best model parameters. Numerical
methods can provide reliable forecast but require meticulous simulation of
the physical equations. The inference time of numerical models usually take
several hours and they are therefore not suitable for generating fine-grained
predictions required by precipitation nowcasting.

Recently, a new approach: *Deep Learning for Precipitation Nowcasting*
has emerged in the area and shown promising results. Shi et al. [42] first
formulated precipitation nowcasting as a spatiotemporal sequence forecast-
ing problem and proposed a DL based model, dubbed *Convolutional Long
Short-Term Memory* (ConvLSTM), to directly predict the future rainfall in-
tensities based on the past radar echo maps. The model is learned end-to-end
with a large amount of historical weather data and performs substantially
better than the OF based algorithm in the operational *Short-range Warning
of Intense Rainstorms in Localized System* (SWIRLS) developed by *Hong
Kong Observatory* (HKO) [34, 57]. After this seminal work, researchers
start to explore DL based methods for precipitation nowcasting and have
built models with state-of-the-art performance [18, 43, 37, 51, 31, 10, 1]. In
essence, precipitation nowcasting is well-suited for DL due to three reasons.
Firstly, the problem satisfies the big data requirement of DL. Numerous
amount of weather data are generated on a daily basis and can be used to
train the nowcasting model. For example, in *National Oceanic and Atmo-
spheric Administration* (NOAA), tens of terabytes of data are generated in
a single day [49]. Secondly, DL is suitable for modeling complex dynamical
systems [17]; a single-hidden-layer *Multi-Layer Perceptron* (MLP), which
is the most basic form of DL models, is a universal functional approxima-

tor [13]. Thirdly, the inference speed of DL models is faster than numerical methods [1]. Moreover, in the inference stage, we can dynamically update the DL model with the newly observed data [43], making the model more adaptive to the emerging weather patterns.

In this chapter, we introduce the current progress of DL based methods for precipitation nowcasting. In Section 18.2, we describe how to mathematically formulate precipitation nowcasting as a spatiotemporal sequence forecasting problem. In Section 18.3, we review the high-level strategies for constructing and learning DL models for precipitation nowcasting; because precipitation nowcasting requires predicting rainfall intensities for multiple timestamps ahead, we introduce various strategies to learn such a multi-step forecasting model. In Section 18.4.1 and Section 18.4.2, we introduce the DL models in two categories: *Feed-forward Neural Network* (FNN) based models, and *Recurrent Neural Network* (RNN) based models. In Section 18.5, we describe the first systematic benchmark of the DL models for precipitation nowcasting, the HKO-7 benchmark. We conclude this chapter and discuss the potential future works along this area in Section 18.6.

18.2   Formulation

Precipitation nowcasting can be formulated as a spatiotemporal sequence forecasting problem. Suppose that the meteorological system is defined over an $M \times N$ grid, in which there are M rows and N columns. Within each cell (i, j) of the grid, there are D measurements that vary over time. By taking snapshots of the system at timestamps $t_1, t_2, ..., t_T$, we get a spatiotemporal sequence that can be denoted as a sequence of tensors $\mathbf{X}_{t_1:t_T} = [\mathbf{X}_{t_1}, \mathbf{X}_{t_2}, ..., \mathbf{X}_{t_T}]$. Here, $\mathbf{X}_{t_i} \in \mathbb{R}^{D \times M \times N}$ is the observed meteorological data at timestamp t_i. In most DL models for precipitation nowcasting, the meteorological observations $\mathbf{X}_{t_i}$s are $2D$ CAPPI radar echo maps [42, 43], satellite images [31], or data from another *Quantitative Precipitation Estimation* (QPE) product [1, 62]. Thus, in most cases, the grid is regular and each pixel covers a region in the geographical map, e.g., a 1km × 1km local area. Some works [18, 37] mainly explore the impact of multi-modal meteorological data without considering the spatial correlations. In this case, $M = N = 1$ and the spatiotemporal sequence forecasting problem degenerates to sequence forecasting problem, in which $\mathbf{X}_{t_i}$ degenerates to a vector $\boldsymbol{x}_{t_i} \in \mathbb{R}^D$. In addition, in most scenarios, the time difference between two consecutive snapshots, i.e. $t_{i+1} - t_i$, is always the same. Thus, we are able to simplify the definition and denote $\mathbf{X}_{t_i}$ as $\mathbf{X}_i$.

The spatiotemporal sequence forecasting problem is to predict the most likely length-L sequence in the future given the previous J observations including the current one [42]. The mathematical definition is given in equation 18.1, in which $\tilde{\mathbf{X}}_{t+1:t+L}$ are the predictions, $\mathbf{X}_{t-J+1:t}$ are the ob-

servations, and $p(\mathbf{X}_{t+1:t+L} \mid \mathbf{X}_{t-J+1:t})$ is the model. In the terminology of precipitation nowcasting, the goal is to use the previously observed sequence to predict short term rainfalls of a local region (e.g., Hong Kong, Shanghai, New York, or Tokyo) in the future. In most nowcasting systems, radar echo map is the mainstay of the meteorological observations due to its high spatial and temporal resolutions. The radar maps are usually taken from the weather radar every 6-10 minutes and predictions are given for the following 0-6 hours. If we record one radar frame every six minutes, the task is to predict for 0-60 frames ahead:

$$\tilde{\mathbf{X}}_{t+1:t+L} = \underset{\mathbf{X}_{t+1:t+L}}{\arg\max}\, p(\mathbf{X}_{t+1:t+L} \mid \mathbf{X}_{t-J+1:t}). \tag{18.1}$$

The spatiotemporal sequence forecasting problem is different from the conventional multi-variate time series forecasting problem because the prediction target of our problem is a sequence that contains both spatial and temporal structures. Because the number of possible sequences grows exponentially with respect to both the spatial and temporal dimensionality, we have to, in practice, exploit the structure of the spatiotemporal space to reduce the dimensionality and hence make the problem tractable.

18.3   Learning Strategies

Precipitation nowcasting is intrinsically a multi-step forecasting problem. Learning a model for multi-step forecasting is challenging because the elements in the predicted sequence $\hat{\mathbf{X}}_{t+1:t+L}$ are not i.i.d.. Nevertheless, predicting the rainfall for multiple timestamps ahead is a crucial requirement of precipitation nowcasting and DL based methods adopt different ways to solve the issue. In this section, we introduce the learning strategies for multi-step forecasting. We first explain and compare two basic strategies called *Iterative Multi-step Estimation* (IME) and *Direct Multi-step Estimation* (DME) [11] and then introduce one extension called *Scheduled Sampling* (SS) that bridges the gap between IME and DME.

Iterative Multi-step Estimation   The IME strategy trains a single-step forecasting model and iteratively feeds the generated samples to the forecaster to get multi-step-ahead predictions. The IME model can either be deterministic or probabilistic. Here, we denote the model as $p(\mathbf{X}_{t+1} \mid \mathbf{X}_{1:t}; \boldsymbol{\theta})$ to cover both cases, in which the deterministic model has a delta distribution. In the terminology of precipitation nowcasting, $\mathbf{X}_{1:t}$ is the sequence of past weather data, $\mathbf{X}_{t+1}$ is the rainfall intensity that the model will predict at timestamp $t + 1$, and $\boldsymbol{\theta}$ is the parameter of the model. To train the model, we factorize the distribution $p(\mathbf{X}_{t+1:t+L} \mid \mathbf{X}_{t-J+1:t})$ as $\prod_{i=1}^{L} p(\mathbf{X}_{t+i} \mid \mathbf{X}_{t-J+1:t+i-1}; \boldsymbol{\theta})$. The optimal paramter $\boldsymbol{\theta}^{\star}$ can be estimated

by maximizing the likelihood:

$$\boldsymbol{\theta}^{\star} = \arg\max_{\phi} \mathbb{E}_{\hat{p}_{\text{data}}} \left[\sum_{i=1}^{L} \log p(\mathbf{X}_{t+i} \mid \mathbf{X}_{t-J+1:t+i-1}; \boldsymbol{\theta}) \right]. \tag{18.2}$$

There are two advantages of the IME approach: 1) The objective function in equation 18.2 is easy to train because it only requires optimizing for the one-step-ahead forecasting error and 2) we can predict for an arbitrary horizons in the future by recursively applying the basic forecaster. However, there is an intrinsic discrepancy between training and testing in IME. In the training phase, we use the ground-truths from $t+1$ to $t+i-1$ to predict the regional rainfall at timestamp $t+i$, which is also known as teacher-forcing [17]. While in the testing phase, we feed the model predictions instead of the ground-truths back to the forecaster. This makes the model prone to accumulative errors in the forecasting process [4]. Usually, the optimal forecaster for timestamp $t+i$, which is obtained by maximizing $\mathbb{E}_{\hat{p}_{\text{data}}}[\log p(\mathbf{X}_{t+i} \mid \mathbf{X}_{t-J+1:t}; \boldsymbol{\theta})]$, is not the same as recursively applying the optimal one-step-ahead forecaster when the model is nonlinear. This is because the forecasting error at earlier timestamps will propagate to later timestamps [35].

Direct Multi-step Estimation The main motivation behind DME is to avoid the error drifting problem in IME by directly minimizing the long-term prediction error. Instead of training a single model, DME trains a different model $p(\mathbf{X}_{t+i} \mid \mathbf{X}_{t-J+1:t}; \boldsymbol{\theta}_i)$ for each forecasting horizon i, in which $\boldsymbol{\theta}_i$ is the parameter. There can thus be L models in the DME approach. The set of optimal parameters $\{\boldsymbol{\theta}_1^{\star}, \ldots, \boldsymbol{\theta}_L^{\star}\}$ can be estimated from the following optimization problem:

$$\boldsymbol{\theta}_1^{\star}, \ldots, \boldsymbol{\theta}_L^{\star} = \arg\max_{\boldsymbol{\theta}_1, \ldots, \boldsymbol{\theta}_L} \mathbb{E}_{\hat{p}_{\text{data}}} \left[\sum_{i=1}^{L} \log p(\mathbf{X}_{t+i} \mid \mathbf{X}_{t-J+1:t}; \boldsymbol{\theta}_i) \right] \tag{18.3}$$

To disentangle the model size from the number of forecasting steps L, we can also construct $p(\mathbf{X}_{t+i} \mid \mathbf{X}_{t-J+1:t}; \boldsymbol{\theta}_i)$ by recursively applying the single-step forecaster $p(\mathbf{X}_{t+1} \mid \mathbf{X}_{1:t}; \boldsymbol{\theta})$. In this case, the model parameters $\{\boldsymbol{\theta}_1, ..., \boldsymbol{\theta}_L\}$ are shared. For example, when the single-step forecasting model is deterministic and predicts $\tilde{\mathbf{X}}_{t+1}$ as $m(\mathbf{X}_{1:t}; \boldsymbol{\theta})$, we can obtain the second step prediction by feeding in the predicted rainfall intensity, i.e., $\tilde{\mathbf{X}}_{t+2} = m(\mathbf{X}_1, \mathbf{X}_2, \ldots, \mathbf{X}_t, m(\mathbf{X}_{1:t}; \boldsymbol{\theta}); \boldsymbol{\theta})$. By repeating the process for L times, we obtain the predictions $\tilde{\mathbf{X}}_{t+1:t+L}$. The optimal parameter $\boldsymbol{\theta}^{\star}$ can be estimated by minimizing the distance between the prediction and the ground-truth, i.e., $\boldsymbol{\theta}^{\star} = \arg\min_{\boldsymbol{\theta}} \mathbb{E}_{\hat{p}_{\text{data}}} d(\tilde{\mathbf{X}}_{t+1:t+L}, \mathbf{X}_{t+1;t+L})$, in which $d(\cdot, \cdot)$ is a distance function. We need to emphasize here that the aforementioned objective function directly optimizes the multi-step-ahead forecasting error

and is different from equation 18.2, which only minimizes the one-step-ahead forecasting error.

Scheduled Sampling According to [11], DME leads to more accurate predictions when 1) the model is misspecified, 2) the sequences are non-stationary, or 3) the training set is too small. However, DME is more computationally expensive than IME. For DME, if the $\boldsymbol{\theta}_h$s are not shared, we need to store and train L models. If the $\boldsymbol{\theta}_h$s are shared, we need to recursively apply the basic forecasting model for $O(L)$ steps [11, 4, 30]. Both cases require larger memory storage or longer running time than solving the IME objective. Overall, IME is easier to train but less accurate for multi-step forecasting, while DME is more difficult to train but more accurate.

Schedules Sampling (SS) [4] tries to bridge the gap between IME and DME. The idea of SS is to first train the model with IME and then gradually replace the ground-truths in the objective function with samples generated by the model itself. When all ground-truth samples are replaced with model-generated samples, the training objective falls back into the DME objective. The generation process of SS is described in equation 18.4:

$$
\begin{aligned}
&\forall 1 \leq i \leq L, \\
&\tilde{\mathbf{X}}_{t+i} \sim p(\mathbf{X}_{t+i} \mid \hat{\mathbf{X}}_{t-J+1:t}, \bar{\mathbf{X}}_{t+1:t+i-1}; \boldsymbol{\theta}), \\
&\bar{\mathbf{X}}_{t+i} = (1 - \tau_{t+i})\hat{\mathbf{X}}_{t+i} + \tau_{t+i}\tilde{\mathbf{X}}_{t+i}, \\
&\tau_{t+i} \sim \text{Binomial}(1, \epsilon_k).
\end{aligned}
\tag{18.4}
$$

Here, $\tilde{\mathbf{X}}_{t+i}$ and $\hat{\mathbf{X}}_{t+i}$ are correspondingly the generated sample and the ground-truth at timestamp $t + i$. $p(\mathbf{X}_{t+i} \mid \hat{\mathbf{X}}_{t-J+1:t}, \bar{\mathbf{X}}_{t+1:t+i-1}; \boldsymbol{\theta})$ is the basic single-step forecasting model. Meanwhile, τ_{t+h} is generated from a binomial distribution and controls whether to use the ground-truth or the generated sample. ϵ_k is the probability of choosing the ground-truth at the kth iteration. In the training phase, SS minimizes the distance between $\tilde{\mathbf{X}}_{t+1:t+L}$ and $\hat{\mathbf{X}}_{t+1:t+L}$. In the testing phase, τ_{t+i}s are fixed to 0, meaning that the model-generated samples are always used.

SS lies in the mid-ground between IME and DME. If ϵ_k equals to 1, the ground-truths are always chosen, and the objective function will be the same as in the IME strategy. If ϵ_k is 0, the generated samples are always chosen, and the optimization objective will be the same as in the DME strategy. In practice [4, 53], ϵ_k is gradually decayed during the training phase to make the optimization objective shift smoothly from IME to DME, which is a type of *curriculum learning* [5].

When applied for precipitation nowcasting, existing DL models adopt either of these three learning strategies. We will introduce the detailed architectures of these two types of models in Section 18.4.1 and Section 18.4.2 and give an overview of the learning strategy that each model uses in Section 18.6.

18.4 Models

18.4.1 FNN-based Models

FNN refers to deep learning models that construct the mapping $\mathbf{Y} = f(\mathbf{X}; \theta)$ by stacking various basic blocks such as the *Fully-Connected* (FC) layer, the convolution layer, the deconvolution layer, and the activation layer. Common types of FNNs include *Multilayer Perceptron* (MLP) [41, 17], which stacks multiple FC layers and nonlinear activations and *Convolutional Neural Network* (CNN) [29], which stacks multiple convolution layers, pooling layers, deconvolution layers, FC layers, activation layers, normalization layers [23, 58] and other transformations. The parameters of FNN are usually estimated by minimizing the loss function plus some regularization terms, i.e., $\theta^\star = \arg\min_\theta \mathbb{E}_{\hat{p}_{\text{data}}}[l(\mathbf{Y}, f(\mathbf{X}; \theta))] + \Omega(\boldsymbol{\theta})$ where $l(\cdot, \cdot)$ is the loss function and $\Omega(\cdot)$ is the regularization function such as the L_2 or L_1 loss [17]. Usually, the optimization problem is solved via stochastic-gradient-based methods [17], in which the gradient is computed by backpropagation [36].

The convolution layer takes advantage of the translational invariance property of image data. The convolution layer computes the output by scanning over the input and applying the same set of linear filters. Although the input can have an arbitrary dimensionality [50], we mainly focus on 2D convolution, since for precipitation nowcasting, the convolution layer is mainly used for extracting the spatial correlation in meteorological images. For input $\mathbf{X} \in \mathbb{R}^{C_i \times H_i \times W_i}$, the output of the convolution layer $\mathbf{H} \in \mathbb{R}^{C_o \times H_o \times W_o}$, which is also known as feature map. Ayzel et al. [2] proposed DozhdyaNet for precipitation nowcasting. DozhdyaNet consists of only convolution layers and is a type of all convolutional network [44]. The input of the model is a sequence of radar images $\mathbf{X}_{t-J+1:t}$. To facilitate the 2D convolution layer to deal with the sequence of 3D tensors, the author concatenates all frames along the temporal dimension and treats them as different channels:

$$\mathbf{X}_{in} = \text{concat}(\mathbf{X}_{t-J+1}, \mathbf{X}_{t-J+2}, ...\mathbf{X}_t), \tag{18.5}$$

where $\mathbf{X}_{in} \in \mathbb{R}^{C_i' \times H_i \times W_i}$ is the network input and the number of channels C_i' is the product of the number of observations within each local grid and the input sequence length, i.e., $C_i' = C_i \times J$. Notice that in this manner, the channel C_i' is determined by the input length, hence unlike the RNN-based models, which will be explained in detail in Section 18.4.2, the input length must be fixed for the FNN-base models. The radar image at the next timestamp $\mathbf{X}_{t+1}$ is predicted by feeding the preprocessed input sequence $\mathbf{X}_{in}$ into the FNN: $\mathbf{X}_{t+1} = f(\mathbf{X}_{in}; \boldsymbol{\theta})$. Also, to predict for multiple steps ahead, the author adopted the IME strategy by feeding the predicted radar image back to the network in the inference phase. Also, the author compared different transformation techniques for preprocessing the 2D radar images and used two radar images taken with the 5min interval as the input to
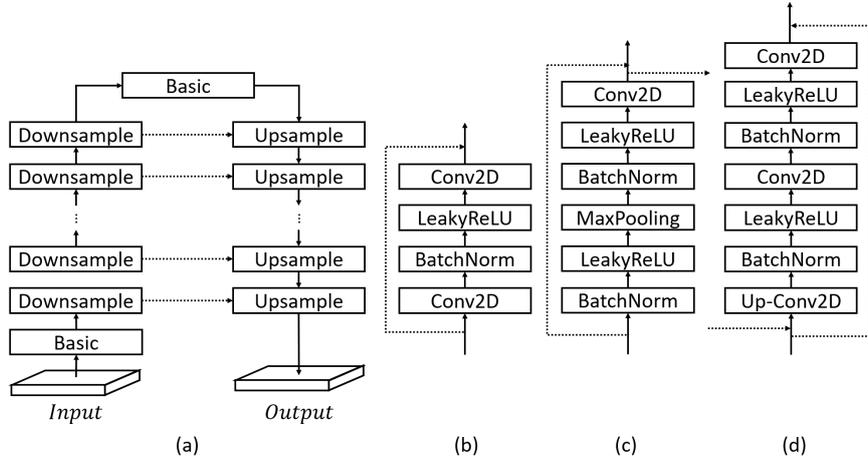
Figure 18.1: (a) The overall structure of the U-NET in [1]. Solid lines indicate input connections between layers. Dashed lines indicate skip connections. (b) The operations within the basic layer. (c) The operations within our down-sample layers. (d) The operations within the up-sample layers.

predict the next radar image.

Agrawal et al. [1] also concatenates the input images in the temporal dimension and used the U-Net [40] architecture for prediction. U-Net combines down-sampling, up-sampling and skip connections to learn better hidden representations. Fig. 18.1 illustrates how these building blocks are organized. The iterative down-sampling part extracts more global and more abstract representations and the up-sampling part gradually refines the representation and adds the finer details to the generated output. The skip connection helps preserve high-resolution details and facilitates gradient backpropagation. The author used QPE data from the *Multi-Radar Multi-Sensor* (MRMS) system [62] for training and testing the model. The input is a sequence of radar images taken with 2min interval for one hour and the output is the sequence of radar images for the next several hours. Experiments show that the U-Net based DL model outperforms OF based model and the HRRR model by NOAA [6].

Klein et al. [28] designed the dynamic convolution layer to replace the conventional convolution layer. Instead of using a data-independent filter, the dynamic convolution layer generates both the feature maps and the filters from the input and convolves the filters with the feature maps to get the output. The feature maps and the filters are obtained from the input with two sub-networks. Because the filters are depend on the input, they will vary from one sample to another in the testing phase. The author concate-
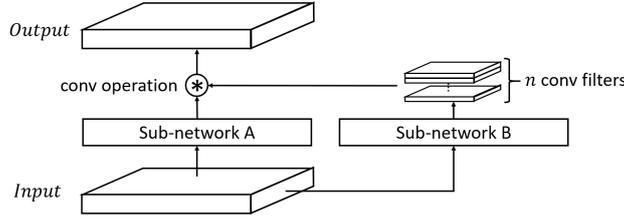
Figure 18.2: The dynamic convolutional layer. The input is fed into two sub-networks. The features are the result of sub-network A while the convolution filters are obtained from sub-network B. The final output of the dynamic convolution layer is computed by convolving the filters from sub-network B across the features from sub-network A.

nates four radar images as the input to predict the next radar image. Also, the author proposed a patch-by-patch synthesis technique which predicts a 10×10 patch in the output from a sequence of 70×70 patches in the input. Fig. 18.2 illustrates the workflow of the dynamic convolution layer. Notice that this layer is different from the dynamic filter [25] layer that is introduce in Section 18.4.2. In dynamic convolution layer, the filter is shared for all locations in the input, while they are adaptively selected in the dynamic filter layer.

Besides the radar images, satellite images are also commonly used as input in FNN-based models. In [31], satellite images and the observations from Global Forecast System (GFS) [9] are combined and used as the input. These two types of data are in different modalities and are misaligned with regard to spatial and temporal resolution. Thus, the author remapped them into the same spatial and temporal grid by interpolation. Lebedev et al. [31] also applied the U-Net architecture.

Similar to [31], Hernández et al. [18] and Qiu et al. [37] also deals with meteorological data from multiple modalities, including temperature, humidity, wind speed, barometric pressure, Dew point, etc. However, they do not consider the spatial dimension of these data. FNNs with 1D convolution layers and FC layers are built for the $1 \times D$ input data. The weather nowcasting problem is formulated as learning a deterministic mapping $\boldsymbol{Y}_{t+1} = f(\boldsymbol{X}_t)$ that maps the current meteorological observation $\boldsymbol{X}_t$ to the precipitation at next step $\boldsymbol{Y}_{t+1}$. Since the formulation has not fully utilized the spatiotemporal structure of the data, we will not go into the details here.

18.4.2 RNN-based Models

As mentioned in Section 18.2, precipitation nowcasting can be formulated as a spatiotemporal sequence forecasting problem with the sequence of past radar maps as input and the sequence of future radar maps as output. In

the advancement of DL, RNN-based architectures, such as *Gated Recurrent Unit* (GRU) [12] and LSTM [19], are proven to be effective for modeling sequential data [48, 26, 38, 45, 59]. Different from FNN-based models, which are designed for modeling inputs with static shapes, RNN-based models are designed for modeling dynamic systems. In this section, we introduce the RNN-based models for precipitation nowcasting. We first introduce the encoder-forecaster structure which is the common approach for constructing RNN-based models for spatiotemporal sequence forecasting. Then we introduce the *Convolutional LSTM* (ConvLSTM) network [42], which combines the advantage of CNN and RNN and is the first DL-based model for precipitation nowcasting. After that, we introduce other RNN-based models like the ConvLSTM with star-shaped bridge [8, 10], *Predictive RNN* (PredRNN) [52], *Memory In Memory* (MIM) Network [53], and the *Trajectory GRU* (TrajGRU) [43], which improves upon ConvLSTM from different directions.

18.4.3   Encoder-Forecaster Structure

The *Encoder-Forecaster* (EF) structure [45, 42] is a widely-used neural network architecture for sequence forecasting. It first encodes the observations into a state with an encoder. The state can be a single vector, multiple vectors, or other mathematical objects. Based on the state, it generates the predictions with a forecaster. Following the same notation as in Section 18.2, we can formulate the EF structure as follows:

$$\mathbf{H} = f(\mathbf{X}_{t-J+1:t}; \boldsymbol{\theta}_1), \quad \hat{\mathbf{X}}_{t+1:t+L} \quad = g(\mathbf{H}; \boldsymbol{\theta}_2). \tag{18.6}$$

Here, $f(\cdot; \boldsymbol{\theta}_1)$ is the encoder parameterized by $\boldsymbol{\theta}_1$, $g(\cdot; \boldsymbol{\theta}_2)$ is the forecaster parameterized by $\boldsymbol{\theta}_2$, $\hat{\mathbf{X}}_{t+1:t+L}$ are the predictions.

18.4.4   Convolutional LSTM

In the DL community, *Fully-Connected LSTM* (FC-LSTM) is a type of RNN with gates and memory cells for dealing with the vanishing gradient problem [17]. The formula of FC-LSTM is given as follows:

$$\begin{aligned}
\boldsymbol{i}_t &= \sigma(\boldsymbol{W}_{xi}\boldsymbol{x}_t + \boldsymbol{W}_{hi}\boldsymbol{h}_{t-1} + \boldsymbol{w}_{ci} \odot \boldsymbol{c}_{t-1} + \boldsymbol{b}_i), \\
\boldsymbol{f}_t &= \sigma(\boldsymbol{W}_{xf}\boldsymbol{x}_t + \boldsymbol{W}_{hf}\boldsymbol{h}_{t-1} + \boldsymbol{w}_{cf} \odot \boldsymbol{c}_{t-1} + \boldsymbol{b}_f), \\
\boldsymbol{c}_t &= \boldsymbol{f}_t \odot \boldsymbol{c}_{t-1} + \boldsymbol{i}_t \odot \tau_h(\boldsymbol{W}_{xc}\boldsymbol{x}_t + \boldsymbol{W}_{hc}\boldsymbol{h}_{t-1} + \boldsymbol{b}_c), \\
\boldsymbol{o}_t &= \sigma(\boldsymbol{W}_{xo}\boldsymbol{x}_t + \boldsymbol{W}_{ho}\boldsymbol{h}_{t-1} + \boldsymbol{w}_{co} \odot \boldsymbol{c}_t + \boldsymbol{b}_o), \\
\boldsymbol{h}_t &= \boldsymbol{o}_t \odot \tau_o(\boldsymbol{c}_t),
\end{aligned} \tag{18.7}$$

in which $\boldsymbol{i}_t, \boldsymbol{f}_t, \boldsymbol{o}_t$ are correspondingly the input gate, forget gate, and output gate. The $\boldsymbol{c}_{t-1}, \boldsymbol{c}_t$ are the memory cells. $\tau_h(\cdot)$ and $\tau_o(\cdot)$ are the activations, e.g., the 'tanh' function. $\boldsymbol{x}_t$s and $\boldsymbol{h}_t$s are input vectors and the hidden states.
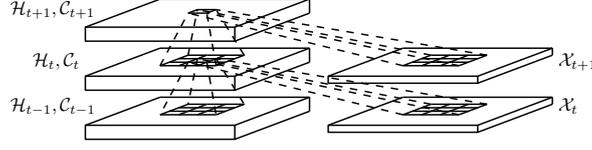
Figure 18.3: Inner structure of ConvLSTM. Source: [42].

However, FC-LSTM is not suitable for precipitation nowcasting in which the input and output are spatiotemporal sequences. The major drawback of FC-LSTM in handling spatiotemporal data is its usage of full-connections in input-state and state-state transitions that loses the spatial structure.

To overcome the problem of FC-LSTM, ConvLSTM [42] extends FC-LSTM by having convolutional structures in both the input-state and state-state transitions. The key equations of ConvLSTM is given in equation 18.8. A distinguishing feature of the design is that all the inputs $\mathbf{X}_t$s, cell states $\mathbf{C}_t$s, hidden states $\mathbf{H}_t$s, and gates $\mathbf{I}_t, \mathbf{F}_t, \mathbf{O}_t$ of the ConvLSTM are 3D tensors whose last two dimensions are spatial dimensions (rows and columns). To get a better picture of the inputs and states, we may imagine them as vectors standing on a spatial grid. The ConvLSTM determines the future state of a certain cell in the grid by the inputs and past states of its local neighbors. Fig. 18.3 illustrates the connection structure of ConvLSTM.

$$
\begin{aligned}
\mathbf{I}_t &= \sigma(\mathbf{W}_{xi} * \mathbf{X}_t + \mathbf{W}_{hi} * \mathbf{H}_{t-1} + \mathbf{W}_{ci} \odot \mathbf{C}_{t-1} + \boldsymbol{b}_i) \\
\mathbf{F}_t &= \sigma(\mathbf{W}_{xf} * \mathbf{X}_t + \mathbf{W}_{hf} * \mathbf{H}_{t-1} + \mathbf{W}_{cf} \odot \mathbf{C}_{t-1} + \boldsymbol{b}_f) \\
\mathbf{C}_t &= \mathbf{F}_t \odot \mathbf{C}_{t-1} + \mathbf{I}_t \odot \tau_h(\mathbf{W}_{xc} * \mathbf{X}_t + \mathbf{W}_{hc} * \mathbf{H}_{t-1} + \boldsymbol{b}_c) \qquad (18.8) \\
\mathbf{O}_t &= \sigma(\mathbf{W}_{xo} * \mathbf{X}_t + \mathbf{W}_{ho} * \mathbf{H}_{t-1} + \mathbf{W}_{co} \odot \mathbf{C}_t + \boldsymbol{b}_o) \\
\mathbf{H}_t &= \mathbf{O}_t \odot \tau_o(\mathbf{C}_t)
\end{aligned}
$$

In order to ensure that each state has the same number of rows and columns, the author uses zero-padding in the convolution operator and views it as initializing the state of the outside world to be all zero. Also, the traditional FC-LSTM can be viewed as a special case of ConvLSTM with all features standing on a single cell.

The author adopts ConvLSTM as the building block for the EF architecture. The initial states and cell outputs of the forecasting network are copied from the last state of the encoding network. Both encoder and forecaster are formed by stacking several ConvLSTM layers. All states in the forecasting network are concatenated and fed into a 1×1 convolutional layer to generate the final prediction.

The nowcasting model based on ConvLSTM is compared with the OF-based ROVER [57] algorithm operated in HKO and the FC-LSTM based model on a 97-day radar echo data in Hong Kong. Experiments show that ConvLSTM outperforms both two baselines. Also, the results showed that
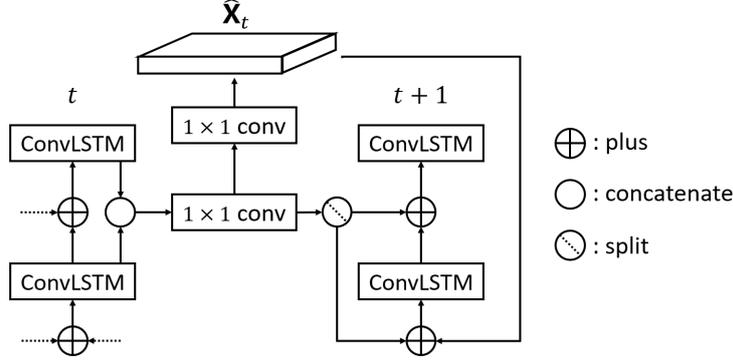
Figure 18.4: Connection structure of the star-shaped bridge.

setting the kernel size of the state-state convolution to be larger than 1 is essential for the final performance.

18.4.5   ConvLSTM with Star-shaped Bridge

To make the feature flow in multi-layer ConvLSTM more robust, Cao et al. [8] proposed the connection structure called star-shaped bridge. In this structure, the states of all ConvLSTM layers at timestamp t are concatenated and passed to a convolution layer with kernel size 1×1 to obtain a global state. The global state has residual connections to all ConvLSTM states at timestamp $t + 1$. The detailed structure is shown in Figure 18.4.

Apart from the star-shaped bridge, the author also inserts *Group Normalization* (GN) [58] between ConvLSTM layers. Ablation study shows that the best performance is obtained by combining ConvLSTM, star-shaped bridge, and GN. Also, experiments on 4-year radar echo data from Shanghai, China showed that the learning-based model outperforms the conventional COTREC method [10].

18.4.6   Predictive RNN

With memory cells being updated every time step inside the ConvLSTM block, the encoder-forecaster architecture is able to model the underlying temporal dynamics. However, the memory cells across different layers lack mutual communication, which are hence not powerful enough for capturing and memorizing spatial correlations. If we directly stack multiple ConvL-STM layers, the information goes only upwards and makes the features more and more abstract. However, since the network needs to predict a spatiotemporal sequence with fine details, information from the lower-level features, including the raw inputs, should be maintained. To solve the issue, Wang et al. [52] proposed *Spatiotemporal LSTM* (ST-LSTM) which keeps an extra
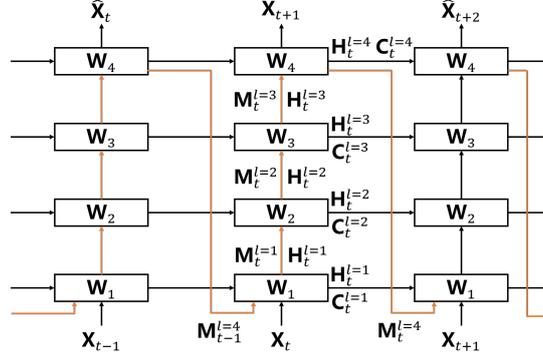
Figure 18.5: Connection structure of PredRNN. The orange arrows in PredRNN denote the flow of the spatiotemporal memory $\mathbf{M}_t^l$.

memory cell $\mathbf{M}_t^l$ to enhance the memory capacity. The external memory is updated in a zigzag direction illustrated in Figure 18.5 and the author named the whole multi-layer architecture as *Predictive RNN* (PredRNN). The updating rule of ST-LSTM is formulated as follows:

$$
\begin{aligned}
\mathbf{I}_t &= \sigma(\mathbf{W}_{xi} * \mathbf{X}_t + \mathbf{W}_{hi} * \mathbf{H}_{t-1}^l + \boldsymbol{b}_i), \\
\mathbf{F}_t &= \sigma(\mathbf{W}_{xf} * \mathbf{X}_t + \mathbf{W}_{hf} * \mathbf{H}_{t-1}^l + \boldsymbol{b}_f), \\
\mathbf{C}_t^l &= \mathbf{F}_t \odot \mathbf{C}_{t-1}^l + \mathbf{I}_t \odot \tanh(\mathbf{W}_{xc} * \mathbf{X}_t + \mathbf{W}_{hc} * \mathbf{H}_{t-1}^l + \boldsymbol{b}_c), \\
\mathbf{I}_t' &= \sigma(\mathbf{W}_{xi}' * \mathbf{X}_t + \mathbf{W}_{mi} * \mathbf{M}_t^{l-1} + \boldsymbol{b}_i'), \\
\mathbf{F}_t' &= \sigma(\mathbf{W}_{xf}' * \mathbf{X}_t + \mathbf{W}_{mf} * \mathbf{M}_t^{l-1} + \boldsymbol{b}_f'), \\
\mathbf{M}_t^l &= \mathbf{F}_t \odot \mathbf{M}_t^{l-1} + \mathbf{I}_t \odot \tanh(\mathbf{W}_{xm} * \mathbf{X}_t + \mathbf{W}_{hm} * \mathbf{M}_t^{l-1} + \boldsymbol{b}_m), \\
\mathbf{O}_t &= \sigma(\mathbf{W}_{xo} * \mathbf{X}_t + \mathbf{W}_{ho} * \mathbf{H}_{t-1}^l + \mathbf{W}_{co} \odot \mathbf{C}_t + \mathbf{W}_{mo} \odot \mathbf{M}_t^l + \boldsymbol{b}_o), \\
\mathbf{H}_t &= \mathbf{O}_t \odot \tanh(\mathbf{W}_{1 \times 1} * [\mathbf{C}_t^l, \mathbf{M}_t^l]).
\end{aligned}
\tag{18.9}
$$

Here, $\mathbf{I}_t, \mathbf{F}_t, \mathbf{H}_t, \mathbf{O}_t$ have the same meaning as in equation 18.8. $\mathbf{C}_t^l$ means the cell state at layer l at timestamp t, and $\mathbf{M}_t^l$ means the external memory cell that will be updated in a zigzag order. For the bottom ST-LSTM with $l = 1$, the memory cell from the previous layer is defined as $\mathbf{M}_t^{l-1} = \mathbf{M}_{t-1}^L$, which results in a zigzag update flow. Experiments show that PredRNN outperforms the ConvLSTM structure in precipitation nowcasting. The experiment is conducted on a dataset with 10,000 consecutive radar observations recorded every 6 minutes in Guangzhou, China. 10 frames are used as the input to predict the future 10 frames.
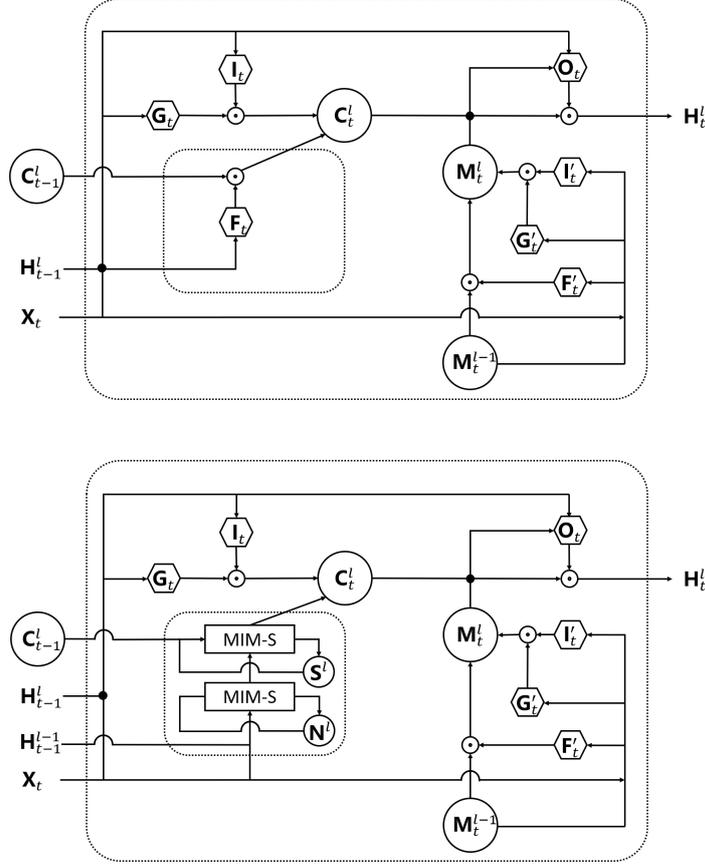
Figure 18.6: ST-LSTM block (up) and Memory In Memory block (bottom). For brevity, $\mathbf{G}_t = \tanh(\mathbf{W}_{xc} * \mathbf{X}_t + \mathbf{W}_{hc} * \mathbf{H}^l_{t-1} + \boldsymbol{b}_c)$, $\mathbf{G}'_t = \tanh(\mathbf{W}_{xm} * \mathbf{X}_t + \mathbf{W}_{hm} * \mathbf{M}^{l-1}_t + \boldsymbol{b}_m)$. MIM is designed to introduce two recurrent modules (yellow squares) to replace the forget gate (dashed box) in ST-LSTM. MIM-N is the non-stationary module and MIM-S is the stationary module.

18.4.7   Memory In Memory Network

Temporal dynamics of spatiotemporal processes are usually non-stationary. Most RNN-based models approximate the non-stationary dynamics in a stationary manner. To learn a better representation of the underlying high-order non-stationary structure, *Memory In Memory* (MIM) [53] extends the ST-LSTM by replacing the forget gate with another two embedded long short-term memories. It leverages the differential information between neighboring hidden states in the recurrent paths, and can gradually stationarize the spatiotemporal process by stacking multiple MIM blocks.

As shown in Figure 18.6, two cascaded temporal memory recurrent modules are designed to replace the temporal forget gate $\mathbf{F}_t$ in ST-LSTM. The
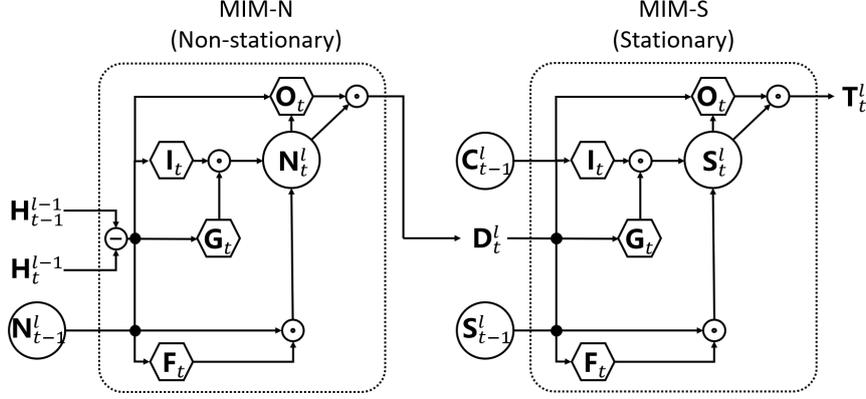
Figure 18.7: The non-stationary module (MIM-N) and the stationary module (MIM-S), which are interlinked in a cascaded structure in the MIM block. Non-stationarity is modeled by differencing.

first module additionally taking $\mathbf{H}_{t-1}^{l-1}$ as input is used to capture the non-stationary variations based on the difference $\mathbf{H}_t^{l-1} - \mathbf{H}_{t-1}^{l-1}$ between two consecutive hidden representations. Thus, it is named as the non-stationary module (shown as MIM-N in Figure 18.7). It generates differential features $\mathbf{D}_t^l$ based on the difference-stationary assumption. The other recurrent module takes the output $\mathbf{D}_t^l$ of the MIM-N module and the outer temporal memory $\mathbf{C}_{t-1}^l$ as inputs to capture the approximately stationary variations in spatiotemporal sequences. Thus, it is named as the stationary module (shown as MIM-S in Figure 18.7). By replacing the forget gate with the final output $\mathbf{T}_t^l$ of the cascaded non-stationary and stationary modules, the non-stationary dynamics can be captured more effectively. The complete formula of MIM is given as follows:

$$
\begin{aligned}
\mathbf{I}_t &= \sigma(\mathbf{W}_{xi} * \mathbf{X}_t + \mathbf{W}_{hi} * \mathbf{H}_{t-1}^l + \boldsymbol{b}_i), \\
\mathbf{F}_t &= \sigma(\mathbf{W}_{xf} * \mathbf{X}_t + \mathbf{W}_{hf} * \mathbf{H}_{t-1}^l + \boldsymbol{b}_f), \\
\mathbf{D}_t^l &= \mathrm{MIM\text{-}N}(\mathbf{H}_t^{l-1}, \mathbf{H}_{t-1}^{l-1}, \mathbf{N}_{t-1}^l), \\
\mathbf{T}_t^l &= \mathrm{MIM\text{-}S}(\mathbf{D}_t^l, \mathbf{C}_{t-1}^l, \mathbf{S}_{t-1}^l), \\
\mathbf{C}_t^l &= \mathbf{F}_t \odot \mathbf{T}_t^l + \mathbf{I}_t \odot \tanh(\mathbf{W}_{xc} * \mathbf{X}_t + \mathbf{W}_{hc} * \mathbf{H}_{t-1}^l + \boldsymbol{b}_c), \\
\mathbf{I}_t' &= \sigma(\mathbf{W}_{xi}' * \mathbf{X}_t + \mathbf{W}_{mi} * \mathbf{M}_t^{l-1} + \boldsymbol{b}_i'), \\
\mathbf{F}_t' &= \sigma(\mathbf{W}_{xf}' * \mathbf{X}_t + \mathbf{W}_{mf} * \mathbf{M}_t^{l-1} + \boldsymbol{b}_f'), \\
\mathbf{M}_t^l &= \mathbf{F}_t \odot \mathbf{M}_t^{l-1} + \mathbf{I}_t \odot \tanh(\mathbf{W}_{xm} * \mathbf{X}_t + \mathbf{W}_{hm} * \mathbf{M}_t^{l-1} + \boldsymbol{b}_m), \\
\mathbf{O}_t &= \sigma(\mathbf{W}_{xo} * \mathbf{H}_t^{l-1} + \mathbf{W}_{ho} * \mathbf{H}_{t-1}^l + \mathbf{W}_{co} \odot \mathbf{C}_t + \mathbf{W}_{mo} \odot \mathbf{M}_t^l + \boldsymbol{b}_o), \\
\mathbf{H}_t &= \mathbf{O}_t \odot \tanh(\mathbf{W}_{1\times 1} * [\mathbf{C}_t^l, \mathbf{M}_t^l]),
\end{aligned}
\tag{18.10}
$$
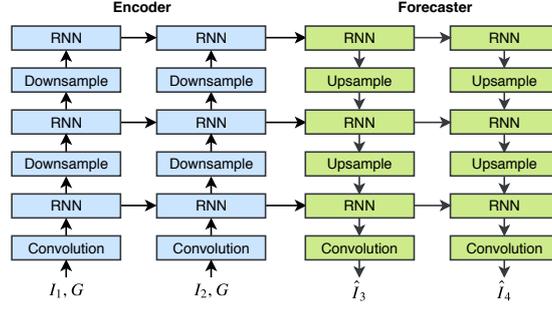
Figure 18.8: Encoder-forecaster architecture adopted in [43]. Source: [43].

where **S** and **N** denote the horizontally-transited memory cells in the non-stationary module (MIM-N) and stationary module (MIM-S) respectively; $\mathbf{D}_t^l$s are the differential features learned by MIM-N; $\mathbf{T}_t^l$ is the memory passing the virtual "forget gate". MIM-N is a ConvLSTM with $\mathbf{H}_t^{l-1} - \mathbf{H}_{t-1}^{l-1}$ as the hidden state input. MIM-S is a ConvLSTM with $\mathbf{D}_t^l$ as the hidden state input. The detailed formula are omitted here and readers can refer to [53] for more details.

18.4.8 Trajectory GRU

Shi et al. [43] proposed a U-Net-like modification to the EF architecture. As Figure 18.8 illustrates, the order of the forecaster network is reversed comparing to that in [42]. There are downsampling and upsampling layers between the RNNs, which are implemented by strided convolution and deconvolution. In this structure, the encoder adopts a local-to-global feature extraction process while the decoder adopts a coarse-to-fine generation process. There are "skip-connections" between the encoder and the forecaster to preserve the details from the raw inputs.

Along with the new EF structure, Shi et al. [43] also pointed out a side-effect of the convolution operation in ConvLSTM. In essence, the location-invariant convolution filters are inefficient to capture location-variant spatiotemporal relationship. To overcome this problem, Shi et al. [43] proposed the *Trajectory GRU* (TrajGRU) model which uses a sub-network to output the state-state connection structures before state transitions. TrajGRU extends upon *Convolutional GRU* (ConvGRU), which is a variant of ConvLSTM, and allows the state to be aggregated along some learned trajectories.

The formula of ConvGRU is given as follows:

$$
\begin{aligned}
\mathbf{Z}_t &= \sigma(\mathbf{W}_{xz} * \mathbf{X}_t + \mathbf{W}_{hz} * \mathbf{H}_{t-1} + \boldsymbol{b}_z), \\
\mathbf{R}_t &= \sigma(\mathbf{W}_{xr} * \mathbf{X}_t + \mathbf{W}_{hr} * \mathbf{H}_{t-1} + \boldsymbol{b}_r), \\
\mathbf{H}'_t &= f(\mathbf{W}_{xh} * \mathbf{X}_t + \mathbf{R}_t \odot (\mathbf{W}_{hh} * \mathbf{H}_{t-1} + \boldsymbol{b}_h)), \\
\mathbf{H}_t &= (1 - \mathbf{Z}_t) \odot \mathbf{H}'_t + \mathbf{Z}_t \odot \mathbf{H}_{t-1}.
\end{aligned}
\tag{18.11}
$$

Here, $\mathbf{H}_t, \mathbf{R}_t, \mathbf{Z}_t, \mathbf{H}'_t \in \mathbb{R}^{C_h \times H \times W}$ are the memory state, reset gate, update gate, and new information, respectively.

As stated in [43], when used for capturing spatiotemporal correlations, the deficiency of ConvGRU and ConvLSTM is that the connection structure and weights are fixed for all the locations. The convolution operation basically applies a *location-invariant* filter to the input. If the inputs are all zero and the reset gates are all one, the author pointed out that the calculation process of $\mathbf{H}'_t$ at a specific location (i, j), i.e, $\mathbf{H}'_{t,:,i,j}$, can be rewritten as follows:

$$
\mathbf{H}'_{t,:,i,j} = f\left(\sum_{l=1}^{|\mathcal{N}^h_{i,j}|} \boldsymbol{W}^l_{hh} \mathbf{H}_{t-1,:,p(l,i,j),q(l,i,j)}\right),
\tag{18.12}
$$

in which $\mathcal{N}^h_{i,j}$ is the ordered neighborhood set at location (i, j) defined by the hyper-parameters of the state-state convolution such as kernel size, dilation and padding [61]. $(p(l, i, j), q(l, i, j))$ is the lth neighborhood location corresponding to position (i, j).

Based on this observation, TrajGRU uses the current input and previous state to generate the local neighborhood set for each location at each timestamp. The detailed formula is given in equation 18.13. Here, L is the number of allowed links. $\mathbf{U}_t, \mathbf{V}_t \in \mathbb{R}^{L \times H \times W}$ are the flow fields that store the local connection structure generated by $\gamma(\mathbf{X}_t, \mathbf{H}_{t-1})$. The $\mathbf{W}^l_{hz}, \mathbf{W}^l_{hr}, \mathbf{W}^l_{hh}$ are the weights for projecting the channels and were chosen as 1×1 convolutions in the paper. The warp($\mathbf{H}_{t-1}, \mathbf{U}_{t,l}, \mathbf{V}_{t,l}$) function selects the positions pointed out by $\mathbf{U}_{t,l}, \mathbf{V}_{t,l}$ from $\mathbf{H}_{t-1}$ via the bilinear sampling kernel [24, 21, 43].

$$
\begin{aligned}
\mathbf{U}_t, \mathbf{V}_t &= \gamma(\mathbf{X}_t, \mathbf{H}_{t-1}), \\
\mathbf{Z}_t &= \sigma\left(\mathbf{W}_{xz} * \mathbf{X}_t + \sum_{l=1}^{L} \mathbf{W}^l_{hz} * \text{warp}(\mathbf{H}_{t-1}, \mathbf{U}_{t,l}, \mathbf{V}_{t,l})\right), \\
\mathbf{R}_t &= \sigma\left(\mathbf{W}_{xr} * \mathbf{X}_t + \sum_{l=1}^{L} \mathbf{W}^l_{hr} * \text{warp}(\mathbf{H}_{t-1}, \mathbf{U}_{t,l}, \mathbf{V}_{t,l})\right), \\
\mathbf{H}'_t &= f\left(\mathbf{W}_{xh} * \mathbf{X}_t + \mathbf{R}_t \odot \left(\sum_{l=1}^{L} \mathbf{W}^l_{hh} * \text{warp}(\mathbf{H}_{t-1}, \mathbf{U}_{t,l}, \mathbf{V}_{t,l})\right)\right), \\
\mathbf{H}_t &= (1 - \mathbf{Z}_t) \odot \mathbf{H}'_t + \mathbf{Z}_t \odot \mathbf{H}_{t-1}.
\end{aligned}
\tag{18.13}
$$

(a) For convolutional RNN, the recurrent connections are fixed over time.

(b) For trajectory RNN, the recurrent connections are dynamically determined.
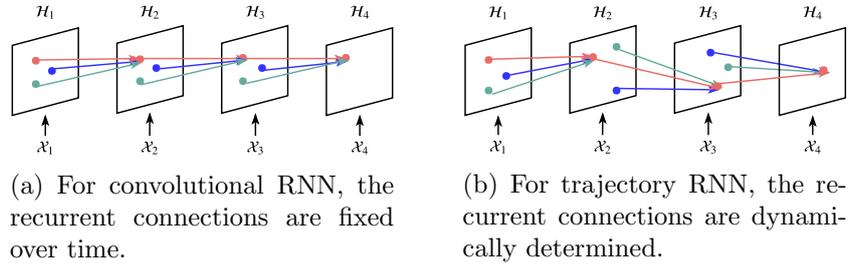
Figure 18.9: Comparison of the connection structures of convolutional RNN and trajectory RNN. Links with the same color share the same transition weights. (Best viewed in color). Source of figure: [43].

The advantage of such a structure is that it could learn the connection topology by learning the parameters of the subnetwork γ. γ has only a small number of parameters and thus adds nearly no cost to the overall computation. Compared to a ConvGRU with $K \times K$ state-state convolution, TrajGRU is able to learn a more efficient connection structure with $L < K^2$. For ConvGRU and TrajGRU, the number of model parameters is dominated by the size of the state-state weights, which is $O(L \times C_h^2)$ for TrajGRU and $O(K^2 \times C_h^2)$ for ConvGRU. If L is chosen to be smaller than K^2, the number of parameters of TrajGRU can also be smaller than the ConvGRU and the TrajGRU model is able to use the parameters more efficiently. Illustration of the recurrent connection structures of ConvGRU and TrajGRU is given in Figure 18.9.

Experiments in the paper showed that TrajGRU outperforms ConvGRU, 2D CNN, 3D CNN, and the ROVER algorithm in precipitation nowcasting.

18.5   Benchmark

Despite the rapid development of DL models in solving this problem, the way to evaluate the models has some deficiencies. Firstly, the deep learning models are only evaluated on relatively small dataset containing limited data frames. Secondly, different models report evaluation results on different criterias. As the needs of real-world precipitation nowcasting system diverge from indicating raining or not to rainstorms alert, single criterion is not sufficient for demonstrating the algorithm's overall performance. Thirdly, in the real-world scenario, the meteorological data arrive in a stream and the nowcasting algorithm should be able to actively adapt to the new-coming sequences. Considering this online setting is not less important than considering offline setting with fixed-length input. In fact, as the area *deep learning for precipitation nowcasting* is still in its early stages, it is not clear how models should be evaluated to meet the needs of real-world applications.

Shi et al. [43] proposed the large-scale HKO-7 benchmark for precipi-

tation nowcasting to address this problem. HKO-7 benchmark is built on the HKO-7 dataset containing radar echo data from 2009 to 2015 near Hong Kong. Since the radar echo maps arrive in a stream in the real-world scenario, the nowcasting algorithms can adopt online learning to adapt to the newly emerging patterns dynamically. To take this setting into account, there are two testing protocols in this benchmark: the offline setting in which the algorithm can only use a fixed window of the previous radar echo maps and the *online setting* in which the algorithm is free to use all the historical data and any online learning algorithm. Another issue for the precipitation nowcasting task is that the proportions of rainfall events at different rain-rate thresholds are highly imbalanced. Heavier rainfall occurs less often but has a higher real-world impact. *Balanced Mean Squared Error* (B-MSE) and *Balanced Mean Absolute Error* (B-MAE) measures are thus introduced for training and evaluation, which assign more weights to heavier rainfalls in the calculation of MSE and MAE. Empirical study showed that the balanced variants of the loss functions are more consistent with the overall nowcasting performance at multiple rain-rate thresholds than the original loss functions. Moreover, training with the balanced loss functions is essential for deep learning models to achieve good performance at higher rain-rate thresholds.

Using the new dataset, testing protocols, and training loss, there are seven models being extensively evaluated, including a simple baseline model which always predicts the last frame, two OF based models (ROVER and its nonlinear variant), and four representative deep learning models (TrajGRU, ConvGRU, 2D CNN, and 3D CNN). This large-scale benchmark for precipitation nowcasting is the first comprehensive benchmark of deep learning models for the precipitation nowcasting problem.

18.5.1   HKO-7 Dataset

The HKO-7 dataset used in the benchmark contains radar echo data from 2009 to 2015 collected by HKO. The radar CAPPI reflectivity images, which have resolution of 480×480 pixels, are taken from an altitude of 2km and cover a 512km $\times$ 512km area centered in Hong Kong. The data are recorded every 6 minutes and hence there are 240 frames per day. The raw logarithmic radar reflectivity factors are linearly transformed to pixel values via pixel $= \lfloor 255 \times \frac{\text{dBZ}+10}{70} + 0.5 \rfloor$ and are clipped to be between 0 and 255. The radar reflectivity values are converted to rainfall intensity values (mm/h) using the Z-R relationship: $\text{dBZ} = 10 \log a + 10 b \log R$ where R is the rain-rate level, $a = 58.53$ and $b = 1.56$. As rainfall events occur sparsely, the rainy days are selected based on the rain barrel information to form the final dataset, which has 812 days for training, 50 days for validation and 131 days for testing.

The raw radar echo images generated by Doppler weather radar are noisy

due to factors like ground clutter, sea clutter, anomalous propagation and electromagnetic interference [33]. To alleviate the impact of noise in training and evaluation, Noisy pixels are filtered out by generating the noise masks with a two-stage process.

18.5.2 Evaluation Methodology

As the radar echo maps arrive in a stream, nowcasting algorithms can apply online learning to adapt to the newly emerging spatiotemporal patterns. The evaluation protocol of HKO-7 benchmark consists of two settings: (1) the offline setting in which the algorithm always receives 5 frames as input and predicts 20 frames ahead, and (2) the online setting in which the algorithm receives segments of length 5 sequentially and predicts 20 frames ahead for each new segment received. The testing environment guarantees that the same set of sequences is tested in both the offline and online settings for fair comparison.

For both settings, models are evaluated according to the skill scores for multiple thresholds that correspond to different rainfall levels to give an all-round evaluation of the algorithms' nowcasting performance. Table 18.1 shows the distribution of different rainfall levels in HKO-7 dataset. The thresholds 0.5, 2, 5, 10, 30 are selected to calculate the CSI and Heidke Skill Score (HSS) [20]. For calculating the skill score at a specific threshold τ, which is 0.5, 2, 5, 10 or 30, the pixel values in prediction and ground-truth are first converted to 0/1 by thresholding with τ. Then calculate the TP (prediction=1, truth=1), FN (prediction=0, truth=1), FP (prediction=1, truth=0), and TN (prediction=0, truth=0). The CSI score is calculated as $\frac{TP}{TP+FN+FP}$ and the HSS score is calculated as $\frac{TP \times TN - FN \times FP}{(TP+FN)(FN+TN)+(TP+FP)(FP+TN)}$. During the computation, the masked noisy points are ignored.

As shown in Table 18.1, the frequencies of different rainfall levels are highly imbalanced. Using weighted loss function helps solve this problem. Specifically, a weight $w(x)$ is assigned to each pixel according to its rainfall intensity x: $w(x) = \begin{cases} 1, & x < 2 \\ 2, & 2 \le x < 5 \\ 5, & 5 \le x < 10 \\ 10, & 10 \le x < 30 \\ 30, & x \ge 30 \end{cases}$. Also, the masked pixels have weight 0. The resulting B-MSE and B-MAE scores are computed as B-MSE $= \frac{1}{N} \sum_{n=1}^{N} \sum_{i=1}^{480} \sum_{j=1}^{480} w_{n,i,j} (x_{n,i,j} - \hat{x}_{n,i,j})^2$ and B-MAE $= \frac{1}{N} \sum_{n=1}^{N} \sum_{i=1}^{480} \sum_{j=1}^{480} w_{n,i,j} |x_{n,i,j} - \hat{x}_{n,i,j}|$, where N is the total number of frames and $w_{n,i,j}$ is the weight corresponding to the (i,j)th pixel in the nth frame. For the conventional MSE and MAE measures, all the weights are simply set to 1 except the masked points.

Table 18.1: Rain rate statistics in the HKO-7 benchmark. Source: [43].

Rain Rate (mm/h)			Proportion (%)	Rainfall Level
$0 \leq$	x	< 0.5	90.25	No / Hardly noticeable
$0.5 \leq$	x	< 2	4.38	Light
$2 \leq$	x	< 5	2.46	Light to moderate
$5 \leq$	x	< 10	1.35	Moderate
$10 \leq$	x	< 30	1.14	Moderate to heavy
$30 \leq$	x		0.42	Rainstorm warning

18.5.3 Evaluated Algorithms

There are seven nowcasting algorithms for evaluation in HKO-7 benchmark, including the simplest model which always predicts the last frame, two optical flow based methods (ROVER and its nonlinear variant), and four deep learning methods (TrajGRU, ConvGRU, 2D CNN, and 3D CNN). Specifically in the online setting, models are fine-tuned using AdaGrad [16] with learning rate equal to 10^{-4}. The training objective during offline training and online fine-tuning is the sum of B-MSE and B-MAE. During the offline training process, all models are optimized by the Adam optimizer with learning rate equal to 10^{-4} and momentum equal to 0.5, with early-stopping on the sum of B-MSE and B-MAE. The ConvGRU model is also trained with the original MSE and MAE loss, which is named "ConvGRU-nobal" in the paper [43], to evaluate the improvement by training with the B-MSE and B-MAE loss.

18.5.4 Evaluation Results

The experiment results show that training with balanced loss functions is essential for good nowcasting performance of heavier rainfall. The ConvGRU model that is trained without balanced loss, which best represents the model in [42], has a worse nowcasting score than the optical flow based methods at the 10mm/h and 30mm/h thresholds. Also, all the deep learning models that are trained with the balanced loss outperform the optical flow based models. Among the deep learning models, TrajGRU performs the best and 3D CNN outperforms 2D CNN, which shows that an appropriate network structure is crucial to achieving good performance. The improvement of TrajGRU over the other models is statistically significant because the differences in B-MSE and B-MAE are larger than three times their standard deviation. Moreover, the performance with online fine-tuning is consistently better than that without online fine-tuning, which verifies the effectiveness of online learning at least for this task.

The results of the Kendall's τ coefficients [27] between the MSE, MAE,

B-MSE, B-MAE and the CSI, HSS at different thresholds show that B-MSE
and B-MAE have stronger correlations with the CSI and HSS in most cases.

18.6   Discussion

In this chapter, we reviewed the DL based methods for precipitation now-
casting. The architecture, building block, training objective function, met-
rics, and data source of the reviewed methods are summarized in Table 18.2.
Precipitation nowcasting is formulated as a spatiotemporal sequence fore-
casting problem from the machine learning perspective. Thanks to the in-
creased computational power and the amounts of data, the area is making
rapid progress. Machine learning, specifically deep learning, facilitates the
large amount of weather data and provides promising research directions
for better modeling and understanding of precipitation nowcasting problem.
Despite the success of DL based methods achieved on precipitation now-
casting, this problem is still challenging. Below we list several major future
research directions that are not solved or have not been explored:

- **Utilization of multi-source meteorological data**
  While existing DL based models mainly focus on single-source data,
  typically radar echo maps or satellite images, multi-source meteorolog-
  ical data have become avaiblable thanks to rapidly developing sensing
  techniques as well as increasing data storage. Although DL based mod-
  els extract spatiotemporal features effectively, precipitation nowcast-
  ing using only single-source data is essentially ill-posed. The models
  are not offered complete knowledge on the dynamics of the meteoro-
  logical system, hence fail to accurately model it and infer its future
  evolution. Multi-source data, in contrast, provide multi-modal and
  multi-scale meteorological information, giving the model a more holis-
  tic view of the system. Therefore exploring DL models that are able
  to jointly process complementary multi-source data can certainly help
  learn better representations of the observing systems.

- **Handling uncertainty**
  Precipitation nowcasting involves complex physics dynamics. Accord-
  ing to chaos theory, chaotic behaviors in a meteorological system make
  it unpredictable due to high degree of uncertainty. Learning to capture
  the internal uncertainty is one of the major challenge in modeling and
  understanding the latent dynamics. However, most DL models address
  precipitation nowcasting in deterministic manner, which averages all
  possible futures into a single output, without describing its whole dis-
  tribution. For some application scenarios such as rainstorm alert, not
  only the average and the most likely futures are concerned, but also
  some possible extreme cases should be considered. There are some

recent works [60, 3, 14, 32] that developed stochastic spatiotemporal models to predict different possible futures through variational inference. Though stochastic spatiotemporal models are not yet evaluated on precipitation nowcasting tasks, they are inspiring can be potential solutions for handling uncertainty in precipitation nowcasting.

- **Integration with numerical methods**
  Compared with theory-driven quantitative precipitation forecast (QPF) methods with clear physical meanings, deep learning models are data-driven and typically suffer from poor interpretability. Although theory-driven precipitation nowcasting models are derived from physical theories, they are essentially phenomenological models built by summarizing the empirical relationship of observations instead of deriving from first principles, which means theory-driven models are not entirely different from but in essence analogous to data-driven models. Theory-driven models consist of interpretable components to describe the observed data, while keeping consistent with physical laws including conservation of mass, momentum, energy, etc. They are determined by human experts and are hence hard to adjust according to different data from different distributions. On the contrary, data-driven models are equipped with high flexibility to adapt to different datasets since they directly learn parameters from data under few constraints. These two approaches are complementary in respect of interpretability and flexibility. Integrating theory-driven and data-driven approaches provides new opportunities in future precipitation nowcasting research, including but not limited to model calibration, recognizing unidentified observations, etc.

18.7 Acknowledgement

Table 18.2: Summary of reviewed methods. The first half are FNN-based models and the second half are RNN-based models.

Method	Building Block	Architecture	Objectives	Metrics	Data Sources
[28]	Dynamic CNN	Stacked CNN and dynamic CNN	MSE	MSE	Radar
[2]	CNN	DozhdyaNet	MAE, MSE, Logcosh	MAE, CSI	Radar
[1]	CNN	U-Net	Cross-entropy for three binary classifications	Precision, Recall	
[31]	CNN	U-Net	Cross-entropy for three binary classifications, Dice loss [46]	Accuracy, Precision, Recall, F1 score	Satellite
[18]	MLP	AE, MLP	MSE	MSE	47 observational features
[37]	1D-CNN	Stacked 1D-CNN and FC layer	MSE with Frobenius norm	MSE, CSI, Correlation	Multi-site observational features
[42]	ConvLSTM	Spatiotemporal encoder-forecaster	Cross-entropy	MSE, CSI, FAR, POD, Correlation	Radar
[43]	TrajGRU	Encoder-forecaster with downsampling and upsampling	MSE, MAE, B-MSE, B-MAE	HKO-7 benchmark	Radar
[51]	ConvRNN	encoder-forecaster with downsampling and upsampling	SSIM, MS-SSIM [55, 54]	MSE, MAE, SSIM, MS-SSIM, PCC [22]	Radar
[8]	ConvLSTM	Star-shaped Bridge	Multi-Sigmoid Loss	MSE, CSI	Radar
[52]	ST-LSTM	PredRNN	MSE, MAE	MSE	Radar
[53]	MIM block	MIM network	MSE	MSE, CSI	Radar

Bibliography

[1] Shreya Agrawal, Luke Barrington, Carla Bromberg, John Burge, Cenk Gazen, and Jason Hickey. Machine learning for precipitation nowcasting from radar images. *arXiv preprint arXiv:1912.12132*, 2019.

[2] G Ayzel, M Heistermann, A Sorokin, O Nikitin, and O Lukyanova. All convolutional neural networks for radar-based precipitation nowcasting. *Procedia Computer Science*, 150:186–192, 2019.

[3] Mohammad Babaeizadeh, Chelsea Finn, Dumitru Erhan, RH Campbell, and Sergey Levine. Stochastic variational video prediction. In *6th International Conference on Learning Representations, ICLR 2018*, 2018.

[4] Samy Bengio, Oriol Vinyals, Navdeep Jaitly, and Noam Shazeer. Scheduled sampling for sequence prediction with recurrent neural networks. In *NIPS*, 2015.

[5] Yoshua Bengio, Jérôme Louradour, Ronan Collobert, and Jason Weston. Curriculum learning. In *ICML*, 2009.

[6] Stanley G Benjamin, Stephen S Weygandt, John M Brown, Ming Hu, Curtis R Alexander, Tatiana G Smirnova, Joseph B Olson, Eric P James, David C Dowell, Georg A Grell, et al. A north american hourly assimilation and model forecast cycle: The rapid refresh. *Monthly Weather Review*, 144(4):1669–1694, 2016.

[7] Milos Bujisic, Vanja Bogicevic, HG Parsa, Verka Jovanovic, and Anupama Sukhu. It's raining complaints! how weather factors drive consumer comments and word-of-mouth. *Journal of Hospitality & Tourism Research*, 43(5):656–681, 2019.

[8] Yuan Cao, Qiuying Li, Lei Chen, Junping Zhang, and Leiming Ma. Video prediction for precipitation nowcasting. *arXiv preprint arXiv:1907.08069*, 2019.

[9] Environmental Modeling Center. The gfs atmospheric model. *National Centers for Environmental Prediction Office Note*, 442:14, 2003.

[10] Lei Chen, Yuan Cao, Leiming Ma, and Junping Zhang. A deep learning based methodology for precipitation nowcasting with radar. *Earth and Space Science*, page e2019EA000812, 2019.

[11] Guillaume Chevillon. Direct multi-step estimation and forecasting. *Journal of Economic Surveys*, 21(4):746–785, 2007.

[12] K. Cho, B. van Merrienboer, C. Gulcehre, F. Bougares, H. Schwenk, and Y. Bengio. Learning phrase representations using RNN encoder-decoder for statistical machine translation. In *EMNLP*, pages 1724–1734, 2014.

[13] Balázs Csanád Csáji et al. Approximation with artificial neural networks. *Faculty of Sciences, Etvs Lorand University, Hungary*, 24(48):7, 2001.

[14] Emily Denton and Rob Fergus. Stochastic video generation with a learned prior. In *International Conference on Machine Learning*, pages 1174–1183, 2018.

[15] RH Douglas. The stormy weather group (canada). In *Radar in Meteorology*, pages 61–68. Springer, 1990.

[16] John Duchi, Elad Hazan, and Yoram Singer. Adaptive subgradient methods for online learning and stochastic optimization. *Journal of Machine Learning Research*, 12(Jul):2121–2159, 2011.

[17] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. Deep learning. Book in preparation for MIT press, 2016.

[18] Emilcy Hernández, Victor Sanchez-Anguix, Vicente Julian, Javier Palanca, and Néstor Duque. Rainfall prediction: A deep learning approach. In *International Conference on Hybrid Artificial Intelligence Systems*, pages 151–162. Springer, 2016.

[19] S. Hochreiter and J. Schmidhuber. Long short-term memory. *Neural Computation*, 9(8):1735–1780, 1997.

[20] Robin J Hogan, Christopher AT Ferro, Ian T Jolliffe, and David B Stephenson. Equitability revisited: Why the "equitable threat score" is not equitable. *Weather and Forecasting*, 25(2):710–726, 2010.

[21] Eddy Ilg, Nikolaus Mayer, Tonmoy Saikia, Margret Keuper, Alexey Dosovitskiy, and Thomas Brox. Flownet 2.0: Evolution of optical flow estimation with deep networks. In *CVPR*, 2017.

[22] Deep Inamdar, George Leblanc, Raymond J Soffer, and Margaret Kalacska. The correlation coefficient as a simple tool for the localization of errors in spectroscopic imaging data. *Remote Sensing*, 10(2):231, 2018.

[23] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *ICML*, 2015.

[24] Max Jaderberg, Karen Simonyan, Andrew Zisserman, et al. Spatial transformer networks. In *NIPS*, 2015.

[25] Xu Jia, Bert De Brabandere, Tinne Tuytelaars, and Luc V Gool. Dynamic filter networks. In *Advances in Neural Information Processing Systems*, pages 667–675, 2016.

[26] A. Karpathy and L. Fei-Fei. Deep visual-semantic alignments for generating image descriptions. In *CVPR*, 2015.

[27] Maurice G Kendall. A new measure of rank correlation. *Biometrika*, 30(1/2):81–93, 1938.

[28] Benjamin Klein, Lior Wolf, and Yehuda Afek. A dynamic convolutional layer for short range weather prediction. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4840–4848, 2015.

[29] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. ImageNet classification with deep convolutional neural networks. In *NIPS*, 2012.

[30] Alex M Lamb, Anirudh Goyal ALIAS PARTH GOYAL, Ying Zhang, Saizheng Zhang, Aaron C Courville, and Yoshua Bengio. Professor forcing: A new algorithm for training recurrent networks. In *NIPS*. 2016.

[31] Vadim Lebedev, Vladimir Ivashkin, Irina Rudenko, Alexander Ganshin, Alexander Molchanov, Sergey Ovcharenko, Ruslan Grokhovetskiy, Ivan Bushmarinov, and Dmitry Solomentsev. Precipitation nowcasting with satellite imagery. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 2680–2688, 2019.

[32] Alex X Lee, Richard Zhang, Frederik Ebert, Pieter Abbeel, Chelsea Finn, and Sergey Levine. Stochastic adversarial video prediction. *arXiv preprint arXiv:1804.01523*, 2018.

[33] Hansoo Lee and Sungshin Kim. Ensemble classification for anomalous propagation echo detection with clustering-based subset-selection method. *Atmosphere*, 8(1):11, 2017.

[34] PW Li, WK Wong, KY Chan, and Edwin ST Lai. SWIRLS—an evolving nowcasting system. *Special Administrative Region Government*, 2000.

[35] Jin-Lung Lin and Clive WJ Granger. Forecasting from non-linear models in practice. *Journal of Forecasting*, 13(1):1–9, 1994.

[36] Jorge Nocedal and Stephen Wright. *Numerical optimization.* Springer Science & Business Media, 2006.

[37] Minghui Qiu, Peilin Zhao, Ke Zhang, Jun Huang, Xing Shi, Xiaoguang Wang, and Wei Chu. A short-term rainfall prediction model using multi-task convolutional neural networks. In *2017 IEEE International Conference on Data Mining (ICDM)*, pages 395–404. IEEE, 2017.

[38] MarcAurelio Ranzato, Arthur Szlam, Joan Bruna, Michael Mathieu, Ronan Collobert, and Sumit Chopra. Video (language) modeling: a baseline for generative models of natural videos. *arXiv preprint arXiv:1412.6604*, 2014.

[39] Maarten Reyniers. *Quantitative precipitation forecasts based on radar observations: Principles, algorithms and operational systems.* Institut Royal Météorologique de Belgique Brussel, Belgium, 2008.

[40] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer, 2015.

[41] Frank Rosenblatt. *Perceptions and the theory of brain mechanisms.* Spartan books, 1962.

[42] Xingjian Shi, Zhourong Chen, Hao Wang, Dit-Yan Yeung, Wai-kin Wong, and Wang-chun Woo. Convolutional LSTM network: A machine learning approach for precipitation nowcasting. In *NIPS*, 2015.

[43] Xingjian Shi, Zhihan Gao, Leonard Lausen, Hao Wang, Dit-Yan Yeung, Wai-kin Wong, and Wang-chun Woo. Deep learning for precipitation nowcasting: A benchmark and a new model. In *Advances in neural information processing systems*, pages 5617–5627, 2017.

[44] Jost Tobias Springenberg, Alexey Dosovitskiy, Thomas Brox, and Martin Riedmiller. Striving for simplicity: The all convolutional net. *arXiv preprint arXiv:1412.6806*, 2014.

[45] Nitish Srivastava, Elman Mansimov, and Ruslan Salakhutdinov. Unsupervised Learning of Video Representations using LSTMs. In *ICML*, 2015.

[46] Carole H Sudre, Wenqi Li, Tom Vercauteren, Sebastien Ourselin, and M Jorge Cardoso. Generalised dice overlap as a deep learning loss function for highly unbalanced segmentations. In *Deep learning in medical*

image analysis and multimodal learning for clinical decision support, pages 240–248. Springer, 2017.

[47] Juanzhen Sun, Ming Xue, James W Wilson, Isztar Zawadzki, Sue P Ballard, Jeanette Onvlee-Hooimeyer, Paul Joe, Dale M Barker, Ping-Wah Li, Brian Golding, et al. Use of nwp for nowcasting convective precipitation: Recent progress and challenges. *Bulletin of the American Meteorological Society*, 95(3):409–426, 2014.

[48] I. Sutskever, O. Vinyals, and Q. V. Le. Sequence to sequence learning with neural networks. In *NIPS*, pages 3104–3112, 2014.

[49] Katelyn Szura. The big data project: Enhancing public access to noaa's open data. In *AGU Fall Meeting Abstracts*, 2018.

[50] Du Tran, Lubomir Bourdev, Rob Fergus, Lorenzo Torresani, and Manohar Paluri. Learning spatiotemporal features with 3d convolutional networks. In *ICCV*, 2015.

[51] Quang-Khai Tran and Sa-kwang Song. Computer vision in precipitation nowcasting: Applying image quality assessment metrics for training deep neural networks. *Atmosphere*, 10(5):244, 2019.

[52] Yunbo Wang, Mingsheng Long, Jianmin Wang, Zhifeng Gao, and S Yu Philip. Predrnn: Recurrent neural networks for predictive learning using spatiotemporal lstms. In *Advances in Neural Information Processing Systems*, pages 879–888, 2017.

[53] Yunbo Wang, Jianjin Zhang, Hongyu Zhu, Mingsheng Long, Jianmin Wang, and Philip S Yu. Memory in memory: A predictive neural network for learning higher-order non-stationarity from spatiotemporal dynamics. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 9154–9162, 2019.

[54] Zhou Wang and Alan C Bovik. Mean squared error: Love it or leave it? a new look at signal fidelity measures. *IEEE signal processing magazine*, 26(1):98–117, 2009.

[55] Zhou Wang, Eero P Simoncelli, and Alan C Bovik. Multiscale structural similarity for image quality assessment. In *The Thrity-Seventh Asilomar Conference on Signals, Systems & Computers, 2003*, volume 2, pages 1398–1402. Ieee, 2003.

[56] Morris L Weisman, Christopher Davis, Wei Wang, Kevin W Manning, and Joseph B Klemp. Experiences with 0–36-h explicit convective forecasts with the wrf-arw model. *Weather and Forecasting*, 23(3):407–437, 2008.

[57] Wang-chun Woo and Wai-kin Wong. Operational application of optical flow techniques to radar-based rainfall nowcasting. *Atmosphere*, 8(3):48, 2017.

[58] Yuxin Wu and Kaiming He. Group normalization. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 3–19, 2018.

[59] K. Xu, J. Ba, R. Kiros, A. Courville, R. Salakhutdinov, R. Zemel, and Y. Bengio. Show, attend and tell: Neural image caption generation with visual attention. In *ICML*, 2015.

[60] Tianfan Xue, Jiajun Wu, Katherine Bouman, and Bill Freeman. Visual dynamics: Probabilistic future frame synthesis via cross convolutional networks. In *Advances in neural information processing systems*, pages 91–99, 2016.

[61] Fisher Yu and Vladlen Koltun. Multi-scale context aggregation by dilated convolutions. In *ICLR*, 2016.

[62] Jian Zhang, Kenneth Howard, Carrie Langston, Brian Kaney, Youcun Qi, Lin Tang, Heather Grams, Yadong Wang, Stephen Cocks, Steven Martinaitis, et al. Multi-radar multi-sensor (mrms) quantitative precipitation estimation: Initial operating capabilities. *Bulletin of the American Meteorological Society*, 97(4):621–638, 2016.