TRU-NET: a deep learning approach to high resolution prediction of rainfall

Rilwan A. Adewoyin^{1,2} · Peter Dueben³ · Peter Watson⁴ · Yulan He¹ · Ritabrata Dutta⁵

Received: 29 August 2020 / Revised: 14 May 2021 / Accepted: 15 June 2021 /
Published online: 7 July 2021
© The Author(s) 2021

Abstract

Climate models (CM) are used to evaluate the impact of climate change on the risk of floods and heavy precipitation events. However, these numerical simulators produce outputs with low spatial resolution that exhibit difficulties representing precipitation events accurately. This is mainly due to computational limitations on the spatial resolution used when simulating multi-scale weather dynamics in the atmosphere. To improve the prediction of high resolution precipitation we apply a Deep Learning (DL) approach using input data from a reanalysis product, that is comparable to a climate model's output, but can be directly related to precipitation observations at a given time and location. Further, our input excludes local precipitation, but includes model fields (weather variables) that are more predictable and generalizable than local precipitation. To this end, we present TRU-NET (Temporal Recurrent U-Net), an encoder-decoder model featuring a novel 2D cross attention mechanism between contiguous convolutional-recurrent layers to effectively model multi-scale spatio-temporal weather processes. We also propose a non-stochastic variant of the conditional-continuous (CC) loss function to capture the zero-skewed patterns of rainfall. Experiments show that our models, trained with our CC loss, consistently attain lower RMSE and MAE scores than a DL model prevalent in precipitation downscaling and outperform a state-of-the-art dynamical weather model. Moreover, by evaluating the performance of our model under various data formulation strategies, for the training and test sets, we show that there is enough data for our deep learning approach to output robust, high-quality results across seasons and varying regions.

Keywords Climate modelling · Precipitation downscaling · Attention mechanism · Rainfall forecasting · Recurrent neural networks

Editor: Gustavo Batista.

✉ Rilwan A. Adewoyin
rilwan.adewoyin@warwick.ac.uk

Extended author information available on the last page of the article

1 Introduction and background

Across the globe, society is becoming increasingly prone to extreme precipitation events due to climate change. The United Nations stated that flooding was the most dominant weather-related disaster over the 20 years to 2015, affecting 2.3 billion people and accounting for \$1.89 trillion in reported economic losses (Wallemacq and Herden 2015). With the increase in the monetary and societal risk posed by flooding (Shukla et al. 2019), the CM predictions for extreme precipitation events are an important resource in guiding the decision of policy-makers.

State-of-the-art regional climate models (RCM) typically run at horizontal resolutions of $\sim 2\text{--}25$ km for simulations. These simulations provide an approach to get local detail, but they are computationally expensive and must be developed separately for each climate model (CM) (May 2004; IPCC 2007).

Quantile mapping is a well researched bias correction method (Wood et al. 2002; Panjwani et al. 2021) which aims to learn a mapping from quantiles of the empirical distribution of a CM's rainfall predictions to quantiles of the corresponding observed rainfall distribution. Multivariate quantile mapping techniques (Cannon 2018; Bárdossy and Pegram 2012), improve upon their univariate counterparts by leveraging the correlation between rainfall and other weather variables and learning a mapping to observed rainfall from the empirical multivariate distribution of a CM's predictions for several weather variables. However, since quantile mapping methods use precipitation as an input we can expect their performance to weaken when using input data that is not from the CM during training. This is because precipitation simulation is a major challenge in climate modelling, and each climate model's precipitation simulation contains a unique precipitation bias profile. These bias profiles include factors such as the degrees of: overestimation of number of days with rain; underestimation of extreme rainfall events; or prematurity of simulated seasonal weather patterns related to rain.

In contrast, our approach performs downscaling while only utilising input weather variables that are generally less sensitive to the choice of climate model used to generate them. As such, our approach is expected to be more robust to the climate model it is trained on.

The **main aim** of this paper is to create a model that can produce high resolution predictions for daily rainfall across the UK. During training, the input to our model is the low resolution model field data, excluding rainfall, from a reanalysis product. Precipitation observations for the related time and location are used as the target. This reanalysis input data is comparable to the output of a climate model and includes weather variables which, unlike precipitation, are well simulated by climate models. When trained, our model can then be used on the output of any CM simulation for the future, allowing us to produce computationally cheap long-term forecasts of high resolution precipitation into the future, helping to diagnose changes in precipitation events due to climate change.

Our task can be interpreted as a combination of the two DL tasks of Image Super-Resolution and Sequence Transduction. DeepSD (Vandal et al. 2018) extended a popular Image Super-Resolution model to downscale 2D precipitation images up to a factor of 4x. They also incorporated an optimization scheme utilising a probabilistic conditional-continuous (CC) loss function for improving the modelling of zero-skewed precipitation events. In the related sequence transduction task, precipitation nowcasting, Shi et al. (2015) introduced the Convolutional Long Short-Term Memory (ConvLSTM) cell to simultaneously model the behaviour of weather dynamics in space and time, using convolutions to incorporate the surrounding flow-fields and a recurrent structure to model the temporal dependencies

of weather. Extending upon this, the encoder-decoder ConvLSTM (Shi et al. 2015) is also able to represent weather dynamics defined on multiple scales in **space** due to the use of successive convolution based layers, each layer to model larger scale dynamics.

The two deep learning approaches we propose are the hierarchical Convolutional Gated Recurrent Unit model (HCGRU) and TRU-NET. Our TRU-NET model extends an encoder-decoder ConvLSTM by adding the ability to represent weather dynamics in multiple scales in **time** and includes the following three features:

1. A novel Fused Temporal Cross Attention (FTCA) mechanism to improve upon existing methods (Jauhar et al. 2018; Zhao et al. 2019; Liu et al. 2019b) to model multiple temporal scales of weather dynamics using a stacked recurrent structure.
2. A encoder-decoder structure adapting the U-NET (Ronneberger et al. 2015) structure, by contracting or expanding the temporal dimensions as opposed to the spatial dimensions.
3. A non probabilistic adaptation of the conditional continuous (Husak et al. 2007; Vandal et al. 2018) as a training loss objective which improves the prediction of extreme precipitation events, by decoupling the modelling of low intensity precipitation events and high intensity precipitation events.

We present HCGRU, TRU-NET and the conditional continuous loss in Sect. 3 and discuss model training details in Sect. 4. In Sect. 5, we perform experiments to compare TRU-NET to baselines and investigate TRU-NET's performance on out-of-sample forecasting tasks. Finally, we perform an ablation study to compare our proposed FTCA against alternative existing methods. The results of these experiments show that:

- Our novel model, TRU-NET, achieves a lower RMSE and MAE than both a state-of-the-art dynamical weather forecasting model's coarse precipitation prediction and U-NET.
- The quality of TRU-NET's predictions are stable when evaluated on out-of-sample weather predictions formed by time periods outside of the training set.
- Our proposed FTCA outperforms existing methods to decrease the temporal scale modelled by successive recurrent layers in a stacked recurrent structure.

2 Data

We use the ERA5 reanalysis dataset (Hersbach et al. 2020) as an analogue for CM output. Reanalysis data is based on a weather forecast model – that is running at a similar resolution to a CM – to which observations are constantly assimilated to yield the best estimate of the weather state. We take these historical weather state estimates and use it to predict the high resolution precipitation observations, made available by the E-OBS rainfall dataset (Cornes et al. 2019). The E-OBS rainfall dataset is created from the observational records of a network of weather stations. In order to get an estimate for precipitation at points on a structured geographic grid, E-obs uses a pre-trained spatial correlation function. Generally, E-Obs is a high quality dataset, having been found to have strong correlation with the highest quality UK precipitation dataset (Hofstra et al. 2009). However, inspection revealed that the E-obs precipitation values had significant bias towards smaller rainfall values, especially under extreme precipitation events (Hofstra et al. 2009).

Fig. 1 Illustration of Input to TRU-NET and HCGRU: The input datum $X_{tj}, j \in [1-4]$ is obtained from ERA5 reanalysis dataset, while the target datum Y_t is from the E-OBS dataset. An input datum is formed as a stack of 6 model fields including longitudinal and latitudinal components of wind velocity at 850 hPa, specific humidity, geopotential at 500 hPa, air temperature and total column water vapour in the entire vertical column. These model fields are defined at 6-hourly resolution on a 20×21 grid over the United Kingdom. A target datum Y_t is defined at a daily resolution on a 100×140 grid over the UK. The grids in the figure are not drawn to relative scale

More concretely, our input data is formed as a timeseries of length t days ($t \in [1, T]$) containing 6 key model fields (air temperature, specific humidity, longitudinal and latitudinal components of wind velocity at 850 hPa, geopotential height at 500 hPa and total column water vapour in the entire vertical column), each defined on a (20×21) grid representing the UK at approximately 65 km spatial resolution chosen to match that used in the UK Climate Projections datasets (Murphy et al. 2018). By stacking together the six model fields, as in Fig. 1, we have a $(20, 21, 6)$ matrix, $X_{tj}, j \in [1-4]$, representing the UK weather state at 6-h intervals. Our model will therefore take as input a sequence of daily model field observations $X_t \in \mathbb{R}^{4 \times 20 \times 21 \times 6}$, from which it will output a prediction for the true daily total precipitation (mm), $Y_t \in \mathbb{R}^{100 \times 140}$, defined on a $(100, 140)$ grid over the UK with approximately 8.5 km spatial resolution and depicted in Fig. 1.

2.1 Experimental settings

For training, we use the 16×16 stencils surrounding sixteen locations to form our training and validation sets, namely, Cardiff, London, Glasgow (G'gow), Birmingham (B'ham), Lancaster (Lanc.), Manchester (Manc.), Liverpool (L'pool), Bradford (B'ford), Edinburgh (Edin), Leeds, Dublin, Truro, Newry, Norwich, Plymouth (P'mth) and Bangor.

These locations were chosen as they are important population centres that sample a wide breadth of locations across the UK. Further, collectively these locations possess varied meteorological profiles, depicted in Fig. 2. For example, percentage of days with rainfall > 10 mm (R10) ranges from 2.4 to 11.9% and average rainfall conditional on an R10 event is ranging from 13.8 to 16.5 mm.

Fig. 2 Precipitation profiles: Here we present three statistics for each of the regions used in the training set for the experiment in Sect. 5.1 where we evaluate our approaches against Baselines

During testing, we either test on the whole UK, region by region, or test on a the region around a single location such as a city.

3 Models

In the following we initially introduce our deep learning approach, TRU-NET as well as the Conditional Continuous loss used to train it. After this, we introduce the Hierarchical Convolutional Gated Recurrent Unit model.

3.1 Temporal recurrent U-NET

Our TRU-NET model, visualised in Fig. 3, maps the 6-hourly low resolution model fields, $\mathcal{X}_t \in \mathbb{R}^{4 \times 20 \times 21 \times 6}$, to a representation capturing variability on 6-hourly, daily and weekly time scales, and then uses these representations to output a prediction, $\hat{Y}_t \in \mathbb{R}^{100 \times 140}$, for daily total rainfall.

As a first step within TRU-NET, we map the input data of the coarse grid onto the fine grid using bi-linear upsampling.¹

The encoder contains a stack of 3 bi-directional ConvGRU layers. Within the encoder, these 3 layers map the input into coarser spatial/temporal scales, from six-hourly/8.5 km, to daily/34 km, and to weekly/136 km. To achieve this reduction in the temporal scales modelled by contiguous encoder layers, we propose a novel Fused Temporal Cross Attention mechanisms (FTCA) as shown in Fig. 4. These scales are aligned to the timescales associated with extreme rainfall events in the UK (Burton 2011).

The decoder maps the latent representation captured at the weekly scale back to the daily scale before feeding it to an output layer for daily rain prediction.

Due to memory constraints we do not input the full $(28 \cdot 4 \times 100 \times 140 \times 6)$ dimensional model fields at once. In space, we extract stencils of 16×16 grid-points for the input to predict precipitation over the stencil of 4×4 grid-points in the centre of input stencil. TRU-Net processes 28 days worth of information at a time, generating an output of total daily precipitation for all of the 28 days for each application of TRU-NET:

$$\hat{Y}_t, \dots, \hat{Y}_{t-J} = f(\mathcal{X}_t, \mathcal{X}_{t-1}, \dots, \mathcal{X}_{t-J}) \quad (1)$$

with $J = 28$. This will naturally generate a lack of information on the past for the first timesteps ($J = 1, 2, 3 \dots$) and a lack of information on the future for the last timesteps ($J = \dots 26, 27, 28$). However, this could be avoided by a stream of input data that only makes predictions for the time-steps in the centre of the time-series in future studies.

In the following, we describe each of the main components of TRU-NET in more detail.

3.1.1 Encoder

The encoder of our TRU-NET model, as shown in Fig. 3, has $L = 3$ ConvGRU layers, where the l -th layer decreases the sequence length by a factor m_l : $1 \rightarrow 4 \rightarrow 7$. This results in the number of units in each ConvGRU based layer decreasing in the manner: $112 \rightarrow 28 \rightarrow 4$, corresponding to six-hourly, daily and weekly temporal resolutions.

The conventional ConvGRU is a recurrent neural network designed to model spatial-temporal information. In a conventional ConvGRU Layer, each unit i shares its trainable weight matrices $\{W_k, U_k, b_k : k \in [z, r, \tilde{A}]\}$ with other units in the layer, and collectively they are described as having *tied weights*. Each unit i takes two inputs, namely the previous state A_{i-1} and the input in the current time step $\hat{B}_i^{(h_b, w_b, c_b)}$, and outputs a state $A_i^{(h_a, w_a, c_a)}$, as detailed below. Here, z_i is the update gate, r_i is the reset gate, $\tilde{A}$ is the cell state, $\bullet$ and $*$ denote the Hadamard product and convolution, respectively.

$$\begin{aligned} z_i &= \sigma(\hat{B}_i * W_z + A_{i-1} * U_z + b_z) \\ \tilde{A}_i &= \tanh(\hat{B}_i * W_{\tilde{A}} + r_i \bullet A_{i-1} * U_{\tilde{A}} + b_{\tilde{A}}) \\ r_i &= \sigma(\hat{B}_i * W_r + A_{i-1} * U_r + b_r) \\ A_i &= z_i \bullet A_{i-1} + (1 - z_i) \bullet \tilde{A}_i \end{aligned} \quad (2)$$

¹ Please note that an increase in resolution is called *upsampling* in machine learning but *down-scaling* in meteorology literature.

Fig. 3 TRU-NET architecture: This figure depicts the Conditional Continuous variant, which outputs predictions for rain level, $\hat{Y}_t$, and rain probability, $\hat{r}_t$, for 28 consecutive days. Each input tensor $\mathcal{X}_t$ is the concatenation of four time sequential $\mathcal{X}_{t_j}$ matrices, presented in Fig. 1. The Sequence Length of 3D tensors between layers contracts/expands through the encoder/decoder. This relates to an increasing/decreasing of the temporal scales modelled. The horizontal direction from left to right indicates time

Fig. 4 ConvGRU with Fused Temporal Cross Attention (FTCA): **a** illustrates a ConvGRU with FTCA layer and **b** illustrates an individual unit within the layer. The grey box in **(b)** shows our adaptation of the generic ConvGRU through the addition of an FTCA operation (grey box) that outputs $\hat{B}_i$

When mapping an input from one time scale to another, e.g. generating the daily time scale tensor for day t from a sequence of 4 corresponding six-hourly time scale tensors, a simple approach is to average the 4 six-hourly tensors. However, such a simple aggregation strategy ignores the influence of the daily time scale tensor from the previous day $t - 1$. We instead propose *Fused Temporal Cross Attention* (FTCA), as a better aggregation strategy based on the cross attention mechanism.

In the final two ConvGRU layers of the encoder, FTCA is fused into the ConvGRU in order to aggregate the inputs from the previous layer to generate a representation for the current layer. The ConvGRU with FTCA is illustrated in Fig. 4 and explained in the following subsection.

3.1.2 Convolutional gated recurrent unit with fused temporal cross attention (ConvGRU w/FTCA)

In the conventional ConvGRU, the i th unit of the l th layer, denoted as $D_{l,i}$, takes two inputs, the previous state A_{l-1} and the input in the current time step $\hat{B}_i$. In our setup here, however, we stack ConvGRU layers with different temporal scales. As such, the input in the current time step to $D_{l,i}$ is no longer a single tensor, but instead, an ordered sequence of tensors, $\mathcal{B}_i \equiv B_{1:T_b}^{(h_b, w_b, c_b)}$, as shown in Fig. 4a, where the input $\mathcal{B}_{l,i}$ consists of T_b time-aligned outputs from the $(l-1)$ -th ConvGRU layer, i.e., $\mathcal{B}_{l,i} \equiv \{A_{l-1,j:j=1,\dots,T_b}\}$. For example, if the l th layer has the daily time resolution, then the $(l-1)$ th layer would have the six-hourly time resolution, $T_b = 4$ and $\mathcal{B}_{l,i} \equiv \{A_{l-1,1}, A_{l-1,2}, A_{l-1,3}, A_{l-1,4}\}$.

Given $\mathcal{B}_i \equiv B_{1:T_b}^{(h_b, w_b, c_b)}$, we propose a Fused Temporal Cross Attention (FTCA) mechanism to calculate a weighted average $\hat{B}_i$. Here, we use $\mathcal{A}_{l-1}$ to derive a query tensor and $\mathcal{B}_i$ to derive both a key tensor and a value tensor. The query tensor is compared with the key tensor to generate weights which are used to aggregate various elements in the value tensor to obtain the final aggregated representation of $\hat{B}_i$.

Afterwards, the ConvGRU operations in Eq. 2 are resumed. The FTCA related operations for unit i have been decomposed into the following three steps:

- *Downscaling representations:* On $\mathcal{A}_i$ and $\mathcal{B}_i$, we first perform a 3D average pooling² (3DAP) with a pool size of $M \times M \times 1$ and transform them to matrices $\mathcal{A}_i^{PF}$ and $\mathcal{B}_i^{PF}$ of dimensions $(1, d_a)$ and (T_b, d_b) respectively, via matrix-reshaping, where $d_a = h_a \times w_a \times c_a \times M^{-2}$ and $d_b = h_b \times w_b \times c_b \times M^{-2}$.

$$\begin{aligned}\mathcal{A}_i^{PF} &= \text{Reshape}(\text{3DAP}(\mathcal{A}_i))^{(h_a, w_a, c_a) \rightarrow (1, d_a)} \\ \mathcal{B}_i^{PF} &= \text{Reshape}(\text{3DAP}(\mathcal{B}_i))^{(T_b, h_b, w_b, c_b) \rightarrow (T_b, d_b)}\end{aligned}\quad (3)$$

- *Similarity calculation using relative attention score (RAS):* We transform $\mathcal{A}_i^{PF}$ and $\mathcal{B}_i^{PF}$ to $Q = \mathcal{A}_i^{PF} \circ W_Q$ and $K^{(T_b, d_o)} = \mathcal{B}_i^{PF} \circ W_K$ through matrix multiplication, $\circ$, with two trainable weight matrices, $W_Q^{(d_a, d_o)}$ and $W_K^{(d_b, d_o)}$. We then compute a matrix of weights $S^{(1, T_b)}$, corresponding to the T_b vectors in K , as follows:

$$S = \text{softmax}\left(\frac{Q \circ (K + a^K)^T}{\sqrt{d_o}}\right)\quad (4)$$

Note here we use the relative attention score (RAS) function (Shaw et al. 2018) to compute the similarity in Eq. 4. Generally to calculate the similarity scores between Q and each vector K_b , the inner product function is used (Vaswani et al. 2017). RAS

² The use of 3D average pooling is motivated by the high spatial correlation within a given feature map due to the spatially correlated nature of weather and to reduce the computational expense of the matrix multiplication.

extends this inner product scoring function by considering the relative position of each vector K_b to one another. In our case, this position relates to the temporal position of K_b relative to other members of K . To facilitate this, we also learn vectors a_b^K which encode the relative position of each K_b .

- *Informative representation*: Finally the new informative representation $\hat{\mathcal{B}}$ is learnt using two trainable convolution weight matrices with c_f filters, $W_{V_1}^{(c_f, 4, A, c_b)}$ and $W_{V_2}^{(c_f, 3, 3, c_b)}$ and a set of trainable vectors $a_i^V \in a^V$, encoding the relative position of each vector $V_i \in V$ as following:

$$V = \mathcal{B} * W_{V_1} \quad \hat{\mathcal{B}} = (\mathcal{S}o(V + a^V)) * W_{V_2} \quad (5)$$

We also use Multi-Head Attention (MHA) which allows the attention mechanism to encode multiple patterns of information by using H heads, $\{W_Q, W_K, W_{V_1}\}$, and performing H parallel cross-attention calculations. The different values of $\{W_Q, W_K, W_{V_1}\}$ across the heads capture different pattern/relationship in data, whereas simply using one head will lead to less diverse or informative patterns captured.

$$\begin{aligned} Q^h &= \mathcal{A}_{PF} \circ W_Q^h & Q &= \text{Concat}(Q_{1:H}^h) \\ K^h &= \mathcal{B}_{PF} \circ W_K^h & K &= \text{Concat}(K_{1:H}^h) \\ V^h &= \mathcal{B} * W_{V_1}^h & V_i &= \text{Concat}(V_{1:H}^h) \end{aligned} \quad (6)$$

3.1.3 Decoder

The decoder is composed of one Dual State ConvGRU (dsConvGRU) layer and an output layer which outputs predictions $\hat{Y}_{t:t+28}$ for the rain level for 28 consecutive days. If the conditional-continuous framework is in use, a second output layer outputs the corresponding predictions for the probability of rainfall $\hat{r}_{t:t+28}$ as illustrated in Fig. 3.

dsConvGRU: As illustrated in Fig. 3, the inputs to the dsConvGRU layer comes from the 2nd and the 3rd Encoder layers, while the output of the dsConvGRU layer is a sequence of 28 tensors which form a latent representation for the 28 days of the target observed daily precipitation $Y_{t:t+28}$.

As the dsConvGRU layer contains 28 units, we must expand the 3rd Encoder layer's output from sequence length 4 to sequence length 28. To do this, we repeat every element in the sequence of length 4, 7 times, as in Tai et al. (2015). As such, each unit in the dsConvGRU layer receives an input from the temporally aligned unit in the 3rd Encoder layer.

Extending Equations 2, the dsConvGRU augments the conventional ConvGRU by replacing the input $\hat{\mathcal{B}}_i$ with two separate inputs $\hat{\mathcal{B}}_{i(1)}$ and $\hat{\mathcal{B}}_{i(2)}$, each possessing the same dimensions as $\hat{\mathcal{B}}_i$. Further, the i -th unit of the dsConvGRU layer takes three inputs, A_{i-1} , $\hat{\mathcal{B}}_{i(1)}$ and $\hat{\mathcal{B}}_{i(2)}$, and outputs a state A_i .

Finally, referring to Equations 2, we calculate two sets of the values, $\{z_{i(j)}, r_{i(j)}, \tilde{A}_{i(j)}, A_{i(j)}\}_{j \in [1,2]}$, corresponding to the use of $\hat{\mathcal{B}}_{(1)}$ or $\hat{\mathcal{B}}_{(2)}$ in place of $\hat{\mathcal{B}}$. Finally, A_i is calculated as the average of $A_{i(1)}$ and $A_{i(2)}$.

Output Layer: As we need to output two sequences of values, rainfall probabilities $\hat{r}_{t:t+28}$ and rainfall values $\hat{Y}_{t:t+28}$, for the conditional-continuous framework which will be discussed in Sect. 3.1.4, our model contains a separate output layer stacked over the dual-state ConvGRU layer for each output. Each output layer contains two 2D convolution layers, with 32 and 1 filters respectively and a kernel shape of (3,3).

3.1.4 Conditional continuous (CC) augmentation

To reflect the zero-skewed nature of rainfall data, due to many days without rainfall, a probabilistic conditional continuous (CC) distribution (Husak et al. 2007; Stern and Coe 1984) is often used to model precipitation. These distributions can be interpreted as the composition of a discrete component and a continuous distribution to jointly model the occurrence and intensity of rainfall:

$$\delta(Y_t) \approx \begin{cases} 1, & Y_t = 0 \\ 0, & Y_t \neq 0 \end{cases} \quad (7)$$

$$p(Y_t; \gamma) = (1 - r_t)\delta(Y_t) + r_t \cdot g(Y_t) \cdot (1 - \delta(Y_t)) \quad (8)$$

where δ is the Dirac function such that $\int_{-\infty}^{\infty} \delta(x)dx = 1$, r_t is the probability of rain at t -th day and $g(\cdot)$ is a Gaussian distribution with unit variance and predicted rainfall $\hat{Y}_t$ as mean. Therefore $(1 - r_t)\delta(y_t)$ models the no rain events, while $r_t \cdot g(Y_t) \cdot (1 - \delta(Y_t))$ handles the rain events.

Vandal et al. (2018) used a probabilistic CC loss for precipitation upscaling by using a Bayesian Neural Network (BNN) and performing multiple stochastic forward passes to generate a sample of predictions for each input. From this sample, estimates for the first and second moments are used to calculate the predictive loss and perform one back propagation step. In contrast our model avoids the computational expense of multiple forward passes by outputting a prediction, $\hat{r}_t$, for the probability of rain occurring as well as a prediction, $\hat{Y}_t$, for the level of rainfall conditional on day t being a rainy day.

To facilitate the requirement of two outputs, $\hat{Y}_t$ and $\hat{r}_t$, we augment the decoder to contain a second identical output layer. In this case, the TRU-NET model has a branch like structure, with $\hat{r}_t$ and $\hat{Y}_t$ the respective outputs of each of these branches.

During training, we sample one set of $[\hat{Y}_t, \hat{r}_t]$ per prediction and use the following loss function. This can be observed as a combination of the binary cross entropy on predictions for whether or not it rained (the first term) and a squared error term on the predicted unconditional rainfall intensity (the second term).

$$\begin{aligned} L(Y_t, [\hat{Y}_t, \hat{r}_t]) = & \frac{1}{T} \left[\sum_{t=1}^T \left(\mathbb{1}_{y_t > 0} \cdot \log(\hat{r}_t) + (\mathbb{1}_{y_t = 0}) \cdot \log(1 - \hat{r}_t) \right) \right. \\ & \left. - \sum_{t=1}^T \|Y_t - \hat{Y}_t\|^2 \right] \end{aligned} \quad (9)$$

3.1.5 Monte Carlo model averaging (MCMA)

When training with dropout, each of the n weights in the neural network has a probability p of being masked. As such, there are 2^n possible models, defined by the combination of weights that can be masked. When sampling predictions from the model, it is infeasible to sample from each of the 2^n variations. Instead, we can form a sample of predictions from a random selection of the 2^n possible models, and calculate the average of the sample. More formally, MCMA is the process of using dropout during training and testing. During training, dropout is performed with a fixed probability p of masking weights. During testing we draw n samples, from our model for each prediction. To do this we use n different

Fig. 5 Illustration of the conditional-continuous variant of the HCGRU model used as a baseline in experiments

dropout masks on the model's weights. Each dropout mask uses the same masking probability, p , on the model's weight as was used during training. We then calculate the mean of these samples to arrive at a *model averaged* prediction. Experiments in Srivastava et al. (2014), §7.5, show this method is effective to sample from neural networks trained with dropout.

During inference, we use the MCMA framework to produce $i \in I$ samples $[\hat{r}_t^i, \hat{Y}_t^i]$ for each observed rainfall Y_t . For each observation, we calculate a final prediction $\hat{Y}_t$ for Y_t :

$$\hat{Y}_t = \frac{1}{I} \sum_{i=1}^I \mathbb{1}_{\hat{r}_t^i > 0.5} \cdot \hat{Y}_t^i \quad (10)$$

3.2 Hierarchical convolutional GRU (HCGRU)

The general structure of a stack of convolutional layers has been used successfully in precipitation nowcasting (Shi et al. 2015, 2017) wherein it outperformed an Optical Flow algorithm (Woo and Wong 2017) produced by the Hong Kong Observatory. In the aforementioned work, the task was sequence to point forecasting. For this work, we adapt the structure to the task of sequence to sequence transduction and propose the structure illustrated in Fig. 5. Our implementation contains 4 ConvGRU layers and an output layer, matching the number of layers in our TRU-NET model. Prior to the first layer, we reduce the input sequence from length 112 to 28, by concatenating blocks of 4 sequential elements. Each of the 4 ConvGRU layers contain 28 recurrent units, with each recurrent unit in each layer

Table 1 Model size: Here we present the sizes of models used in our experiments. We choose similar sizes for the neural network architectures to ensure that any difference in performance is primarily due to modeling capability of the architecture as opposed to model size

	TRU-NET	HCGRU	U-NET
Parameter Count	6,890,290	6,609,602	7,709,888

containing convolutional operations with 80 filters and a kernel shape of (4, 4). Skip connections exists over the final 3 ConvGRU layers and a final skip connection exists from the output of the first ConvGRU layer to the input of the output layer. The output layer follows the same formulation as in TRU-NET, with two 2D Convolution layers. Similar to TRU-NET, we create two variants of HCGRU. One variant uses the Conditional Continuous loss and therefore has two output layers for probability of rain and conditional rain intensity. The other variant only outputs a prediction for rain intensity and as such only has one output layer.

4 Experimental setup

This section describes baseline models used for comparison, hardware setup.

4.1 Baseline models

We compare TRU-NET with the following baselines:

Integrated Forecast System (IFS): The IFS is a computationally expensive numerical weather prediction system which is solving the physical equations of atmospheric motion. IFS is used for operational weather predictions at the European Centre for Medium-Range Weather Forecasts (ECMWF). It is also used to generate the ERA5 reanalysis data which is used as input data for TRU-NET. While the input fields are a product of the data assimilation process of ERA5, there are also data for precipitation predictions available which are diagnosed from short-term forecast simulations with IFS which use ERA5 as initial conditions. There are two forecast simulations started each day at 6 a.m. and 6 p.m. We extract the precipitation fields for the first 12 hours of each simulation to reproduce daily precipitation—this is presently the optimal way to derive meaningful precipitation predictions from a dynamical model that is consistent with the large-scale fields in the ERA5 reanalysis data. The ERA5 and precipitation data is available on a grid with 31 km resolution. However, our target is to use model fields from climate models as input which are typically run at coarser resolution. We therefore map the ERA5 data onto the grid that is used in the HadGEM3 climate model (Murphy et al. 2018).

U-NET: U-NET (Ronneberger et al. 2015) is a popular convolution based network architecture initially proposed for biomedical imaging segmentation, but recently used in image super-resolution and precipitation downscaling (Agrawal et al. 2019; Trebing and Mehrkanoon 2020). Similar to TRU-NET and HCGRU, U-NET features an encoder and a decoder. However, unlike HCGRU and TRU-NET, U-NET does not feature any structure to capture temporal relationships. As U-NET is a simpler architecture, we have utilised a larger structure with 4 sub-layers in the Encoder and 4 sub-layers in the decoder. This ensures that the

Table 2 Hyper-parameter tuning values: Here we present the values used for hyper-parameter grid search tuning. Sub-table (a) shows TRU-NET CC's values and sub-table (b) shows HCGRU's values. The combination of these hyper-parameter options resulted in a set of 64 models for each of TRU-NET CC and HCGRU CC. d.o. is short-hand for dropout

β_2	Learning rate	Input d.o.	Recurrent d.o.	Clip norm	Attention d.o.
(a) TRU-NET CC					
0.900	1e-3	0.150	0.150	4.500	0.150
0.990	1e-4	0.350	0.350	5.500	0.350
β_1	β_2	Learning rate	Input dropout	Recurrent dropout	
(b) HCGRU CC					
0.750	0.900	1e-3	0.100	0.100	
0.900	0.990	1e-4	0.225	0.225	
			0.350	0.350	

number of parameters present in U-NET is similar to that of TRU-NET and HCGRU, as shown in Table 1.

4.2 Hyperparameter settings

In this section we describe the hyper-parameter tuning process we used for the TRU-NET and HCGRU variants. We focus on the hyper-parameter tuning process for TRU-NET CC. An identical process was followed for the other models.

Initial training runs were used to guide our decision to tune the following hyper-parameters: input dropout, recurrent dropout, attention dropout, learning rate, β_1 and β_2 and clip norm. Input dropout and recurrent dropout are hyper-parameters used by the convolutional-recurrent based structures in the HCGRU and TRU-NET variants. Attention dropout is the hyper-parameter used by TRU-NET's attention based structures. The Rectified Adam optimizer's (Liu et al. 2019a) hyper-parameters we chose to optimize were learning rate, β_1 and β_2 . We also included clip norm for the gradients. Table 2 show the values used for the hyper-parameter tune grid search. The dataset used for tuning includes the following cities: Cardiff, London, Glasgow, Birmingham, Lancaster, Manchester, Liverpool, Bradford, Edinburgh, Leeds. The training set covered data from 1999 till 2008 and the test set contained data from 2009 till 2014.

Our initial training runs indicated that TRU-NET CC's performance was significantly sensitive to clip norm. A clip norm below 6.5 and linear warmup were required to ensure stable training.

Table 3 displays the hyper-parameter settings corresponding to the top 5 performing TRU-NET CC and HCGRU CC models, ranked by R10 RMSE. For the TRU-NET CC models, the hyper-parameters with the most significant effect on R10 RMSE were input dropout, recurrent dropout and attention dropout. For these hyper-parameters Fig. 6 uses boxplots to illustrate the distribution of R10 RMSE scores conditional on hyper-parameter values. We see that values of 0.35, 0.15 and 0.15 for the attention dropout, input dropout and recurrent dropout optimize TRU-NET CC performance. Given the results presented in Table 3, we settle on using the following hyper-parameter values for TRU-NET CC during experiments: learning rate = 1e-4, β_2 = 0.90, clip norm = 4.5, input dropout = 0.15, recurrent dropout = 0.15 and attention dropout = 0.35. A similar procedure was followed

Table 3 Grid search hypertuning results: Here we present the five hyper-parameter settings that resulted in the best R10 RMSE scores for TRU-NET CC and HCGRU CC. Sub-Table (a) shows the results for TRU-NET CC models trained for 40 epochs and Sub-Table (b) shows the results for HCGRU CC models trained for 30 epochs. The training set consisted of data between 1999 to 2014 and the testing set consists of data between 2014 to 2019. d.o. is short-hand for dropout

β_2	Learning rate	Input d.o.	Recurrent d.o.	Clip norm	Attention d.o.	R10 RMSE
(a) TRUNET CC						
0.900	1e-4	0.350	0.150	4.500	0.350	6.410
0.990	1e-4	0.350	0.150	5.500	0.350	6.571
0.900	1e-4	0.150	0.150	5.500	0.350	6.744
0.990	1e-4	0.150	0.150	4.500	0.350	6.859
0.990	1e-3	0.150	0.350	5.500	0.350	6.894
β_1	β_2	Learning rate	Input d.o.	Recurrent d.o.	Attention d.o.	R10 RMSE
(b) HCGRU CC						
0.900	0.900	1e-3	0.225	0.350		7.593
0.900	0.990	1e-3	0.225	0.350		7.607
0.750	0.990	1e-3	0.350	0.225		7.619
0.750	0.990	1e-3	0.100	0.225		7.630
0.900	0.990	1e-3	0.100	0.100		7.642

(a) R10 RMSE to attention dropout

(b) R10 RMSE to input dropout

(c) R10 RMSE to recurrent dropout

Fig. 6 Hyperparameter tuning results: Here we present the results of the 64 TRU-NET CC models that were trained. The sub-figures show the distribution of R10 RMSE values relative to the hyper-parameter values for the 64 TRU-NET models, judged by R10 RMSE. The units for R10 RMSE are mm/day

Table 4 Computational Expense: Comparison of cpu core-hours required by the IFS model and gpu-hours required by TRUNET, HCGRU and U-NET during prediction. These estimates are based on the runs of IFS on Cray XC40 supercomputer at ECMWF (with 250 Intel E5-2695v4 Broadwell processors with 36 cores each) and of TRUNET, HCGRU and U-NET on a single core machine with 4 GPUs (GTX 1080 Ti Cards)

	Prediction Time Length	IFS (CPU)	TRUNET (GPU)	HCGRU (GPU)	U-NET (GPU)
Prediction	28 days	~12.5	6.32e−3	5.89e−3	4.91e−3
	1 year	~163.0	8.5e−2	7.68e−2	6.40e−2
	5 years	~814.0	4.1e−1	3.84e−1	3.20e−1

for our HCGRU model to arrive at an optimal parameter range of learning rate = 0.001, $\beta_1 = 0.90$, $\beta_2 = 0.99$, input dropout = 0.225 and recurrent dropout = 0.35.

4.3 Computational cost

While conventional weather and climate prediction models do not require training, the testing and evaluation of these models still generate a significant computational cost which is often higher when compared to the computational cost for operational predictions and to develop and train a weather and climate prediction model at a grid-spacing as used in this paper requires the use of large supercomputers. On the other hand, for TRUNET the majority of all costs are due to the training phase while the evaluation phase contributes negligible cost. However, the training phase can still be performed on a single GPU and does not require supercomputing facilities.

If the absolute cost for a 28-day global weather predictions with the IFS model on the Cray XC40 supercomputer at ECMWF with 250 Intel E5-2695v4 Broadwell processors with 36 cores each is proportionally reduced to only cover the domain of the UK, this simulation will require 12.5 cpu core-hours. In comparison, TRUNET produces a 28-day precipitation prediction at a comparable grid-spacing for the UK in approximately 6.32e−3 gpu-hours on a machine with 4 GPUs (GTX 1080 Ti Cards). The computational cost is further illustrated in Table 4, where we report the computational cost for predicting precipitation for the UK at a 9 km or comparable grid-spacing correspondingly to a predictions of 28 days, 1 and 5 years.

5 Experiments

5.1 Comparison with baselines

5.1.1 Seasonal breakdown for all of the UK

We use the following metrics to evaluate the performance of each model: Root Mean Squared Error (RMSE), RMSE for days of observed rainfall over 10/mm (R10 RMSE) and Mean Absolute Error (MAE). We present these metrics for each season, where the seasons have been defined as Spring (March, April, May), Summer (June, July, August), Autumn (September, October, November) and Winter (December, January, February). The training

Table 5 Comparing TRU-NET with Baselines: Here, we present the results for models tested on the whole UK between the dates of 2014 and August 2019. The TRU-NET CC (TRU-NET.CC) model achieves the best RMSE and MAE scores across all seasons. The units for all metrics are mm/day

Model name	RMSE	R10 RMSE	MAE
(a) All Seasons			
IFS	3.627	9.001	1.976
U-NET	3.790	9.677	2.879
U-NET.CC	3.580	10.262	1.846
HCGRU	3.268	8.792	1.762
HCGRU.CC	3.266	8.671	1.739
TRU-NET	3.106	8.766	1.784
TRU-NET.CC	3.081	8.759	1.644
(b) Winter			
IFS	3.950	9.114	2.233
U-NET	4.194	9.929	3.021
U-NET.CC	4.144	10.651	2.252
HCGRU	3.740	8.879	2.135
HCGRU.CC	3.731	8.894	2.039
TRU-NET	3.613	9.138	2.126
TRU-NET.CC	3.570	9.096	1.978
(c) Spring			
IFS	3.135	8.455	1.692
U-NET	3.355	8.776	2.719
U-NET.CC	2.945	9.420	1.483
HCGRU	2.707	7.922	1.439
HCGRU.CC	2.710	7.832	1.419
TRU-NET	2.549	7.817	1.487
TRU-NET.CC	2.504	7.777	1.328
(d) Summer			
IFS	3.663	9.021	2.018
U-NET	3.646	9.593	2.850
U-NET.CC	3.406	9.940	1.761
HCGRU	3.210	9.056	1.718
HCGRU.CC	3.193	8.701	1.695
TRU-NET	3.001	8.764	1.749
TRU-NET.CC	2.991	8.800	1.616
(e) Autumn			
IFS	3.765	9.222	1.987
U-NET	3.941	10.011	2.939
U-NET.CC	3.755	10.532	1.9311
HCGRU	3.381	9.073	1.783
HCGRU.CC	3.398	8.923	1.773
TRU-NET	3.210	8.892	1.798
TRU-NET.CC	3.215	8.926	1.680

The bold values represent the evaluation score of the best performing model in a given season

set spans the years 1979 till 2008, the validation set spans the years 2009 till 2013 and the test set spans the time period 2014 till August 2019.

In Table 5(a) we observe that the TRU-NET CC model generally outperforms alternative models in terms of RMSE and MAE. Further, our CC variants of TRU-NET and HCGRU achieve a better R10 RMSE than their non conditional continuous counterparts.

We notice that while TRU-NET CC achieves a lower R10 RMSE than HCGRU CC, HCGRU CC achieves a slightly lower R10 RMSE than TRU-NET CC. However, the improvement (5.7%) in the TRU-NET CC RMSE score relative to HCGRU CC RMSE score is far larger than the improvement (1.1%) of the in the HCGRU CC R10 RMSE score relative to the TRU-NET CC RMSE score. As such we believe that a more exhaustive fine-tuning processing that tested more values for TRU-NET's dropout, in the FTCA structure, would lead to further improvements in the R10 RMSE for TRU-NET CC. Finally, we notice, that U-NET/U-NET CC achieves a significantly worse RMSE/R10 RMSE score than the other DL approaches. The poor relative performance of U-NET highlights the importance of using a a neural network structure that utilises the temporal information present in weather data.

5.1.2 City-wise breakdown

In the previous sub-section, we presented seasonal performance metrics for each model tested on the whole country. Here we focus on the predictive errors on 5 specific cities across the range of precipitation profiles displayed in Fig. 2. These cities chosen can be divided into two groups; those with lower rainfall characteristics (London, Birmingham and Manchester) and those with high rainfall characteristics (Cardiff and Glasgow). These locations have been chosen in order to discern whether the quality of predictions over a region is related to the region's precipitation profile. The following tables present the predictive scores of the TRU-NET CC model trained on data from 16 locations over the time span covering 1979 till 2013. The results are presented in Table 6 where we provide the performance of the IFS model as the second number in each cell.

We observe that both TRU-NET CC and IFS generally achieves lower RMSE scores during the Spring and Summer months with less rainfall. By observing the Mean Error (MAE) we notice our model generally under-predicts rainfall for cities with high average rainfall (Glasgow and Cardiff) and over-predicts rainfall for cities with low average rainfall (London, Birmingham, Manchester).

5.1.3 Distribution of predictions

Figure 7 illustrates TRU-NET CC's and IFS predicted rainfall values, for the whole UK, plotted against the true observed rainfall over the period 2014–2019.

When comparing TRU-NET's predictions to IFS predictions, we notice a significant number of cases wherein both TRU-NET and IFS predict rainfall higher than 0 mm, for days where observed rainfall is 0 mm. However, as can be seen by the vertical blue cloud of points to the left of each sub-figure, TRU-NET's log-transformed predictions for non-rainy days spread up to 2.75, while IFS performs worse and spread up to 3.4.

For observed rainfall events between 10 and 19 mm/day we notice that both TRU-NET and IFS slightly under-predict the observed rainfall by a similar amount. However, TRU-NET's predictions have less variance than the IFS predictions, which routinely produce predictions significantly below or above the $y = x$ line. This is highlighted by the large

Table 6 Seasonally disaggregated Performance Metrics for TRU-NET CC: This TRU-NET CC model was trained on data between 1979 and 2013 and tested on the $7.1 \times 10^3 \text{ km}^2$ region around 5 cities between 2014 and August 2019. The first/second number in each cell is the associated performance of TRU-NET/IFS. The left hand column contains the following abbreviations for Seasons: Winter = WNT, Spring = SPR, Summer = SUM, Autumn = AUT. The units for all metrics are mm/day. The scores in bold represent the best predictive performance for each metric for each seasons. We observe that both TRU-NET CC generally achieves lower RMSE scores during the Spring and Summer months with less rainfall, and under/over predicts rainfall for cities with high/low average rainfall

	RMSE	R10 RMSE	ME
(a) Birmingham			
WNT	2.155/ 2.061	4.019 /4.583	0.685/ 0.075
SPR	1.812 /2.817	3.258 /8.193	0.135 /0.349
SUM	1.646 /2.884	3.368 /6.805	-0.095 /0.334
AUT	1.781 /2.661	2.998 /6.676	0.318/ 0.125
All	1.933 /2.621	5.215 /6.649	0.263/ 0.222
(b) Cardiff			
WNT	2.292 /3.486	4.737 /6.564	0.155/ 0.009
SPR	1.829 /3.313	3.516 /6.644	0.301 /0.727
SUM	2.026 /3.735	4.940 /7.769	-0.377/ 0.345
AUT	2.241 /3.701	4.517 /8.408	-0.185/ 0.143
All	2.215 /3.551	4.809 /7.396	-0.008 /0.315
(c) Glasgow			
WNT	3.752 /5.079	7.454 /9.479	-1.100/ -1.025
SPR	2.316 /3.023	5.733 /7.172	-0.312/ -0.008
SUM	2.455 /3.758	5.244 /7.413	-0.241 /0.266
AUT	3.022 /4.128	6.523 /7.812	-0.460 /-0.561
All	3.132 /4.059	1.783 /2.387	-0.531/ -0.337
(d) London			
WNT	2.159 /2.611	3.775 /7.768	0.244/ -0.031
SPR	1.935 /2.411	3.866 /6.829	0.309 /0.447
SUM	2.367 /3.156	6.824 /9.787	-0.040 /0.487
AUT	1.940 /2.771	4.461 /9.105	0.083/ -0.081
All	2.221 /2.735	7.850 /8.500	0.158 /0.210
(e) Manchester			
WNT	2.372 /3.212	4.252 /7.223	0.378/ 0.271
SPR	2.048 /3.342	4.189 /7.959	0.032 /0.585
SUM	1.998 /3.795	5.186 /9.244	-0.179 /0.618
AUT	2.253 /3.428	5.008 /7.783	-0.221/ -0.093
All	2.374 /3.428	6.517 /7.994	0.013 /0.353

vertical spread of IFS predictions, in Fig. 7b, between observed rainfall of 10 mm/day and 3.

For observed rainfall events above 20 mm/day, we notice that TRU-NET under-predicts rainfall events more than IFS. We believe that the rarity of rainfall > 20 events in the training set has negatively impacted TRU-NET's ability to learn these relationships, while IFS learns the underlying physical equations.

5.1.4 Cross correlation across predictions for cities

Here, we check the spatial structure of the predictions via cross-correlation plots for TRU-NET CC Normal predictions on the central point within pairs of cities. We use Leeds as our base location and compute pairwise cross-correlations with the following six locations; Bradford (13 km), Manchester (57 km), Liverpool (104 km), Edinburgh (261 km), London (273 km) and Cardiff (280 km), where the each bracketed number is the distance of

Fig. 7 Distribution of predictions: These figures illustrate the distribution of predicted rainfall against observed rainfall for the TRU-NET CC and IFS models from Sect. 5.1.1. The dashed red line shows the mean and standard deviation of predictions in 3 mm/day intervals of observed rainfall. The purple line indicates the boundary for rainfall events with at least 10 mm/day. For illustrative purposes, we sub-sample every 25th pair of prediction and observed value. The log transform used is $\log(y + 1)$

this location from Leeds. The cross correlations with comparison cities are ordered with increasing distance from Leeds. Linear de-trending was used.

Figures 8 and 9 illustrate the cross correlations between TRU-NET CC's predictions for the central points of pairs of cities. Figure 8 shows the cross-correlation function up to 28 days lag. As expected, we notice a strong correlation up to approximately 5 days. For all sets of figures, the relationships exhibited by TRU-NET CC's predictions (blue line) are approximately mirrored by the observed values (orange line) confirming that our model is producing sensible predictions. In Fig. 9, as expected, we observe that the Lag 0 cross-correlation between the predicted daily rainfall for cities decreases as the cities become increasingly distant from each other.

5.2 Investigation of TRU-NET's limitations

The high temporal correlation in weather data reduces the *effective sample size* and provides the risk that any neural network trained on N consecutive years will only learn a limited set of weather patterns. The reliability of the DL model's extrapolation to out of

Fig. 8 Cross Correlation across Predictions: These figures illustrate the Cross Correlation function (XCF) between rainfall predictions for Leeds and alternative cities. Here, we present XCF up to lag 28. The orange line provides the same statistics, except with the true observed rainfall values. The red line provides the 5% significance threshold, above which we can assume their is significant correlation

sample predictions (new weather patterns) is more doubtful because DL models do not aim to learn the underlying physical equations, unlike numerical weather algorithms.

The three experiments introduced below evaluate the robustness of TRU-NET's out of sample predictive ability.

Fig. 9 Lag 0 Cross-Correlation between Leeds and other cities: This figure shows the Cross-Correlation function (XCF) at Lag 0 for the daily rain predictions between Leeds and the 6 cities on the x-axis. The cities are ordered by increasing distance from Leeds. We use the rain prediction for the central point within the 16×16 stencil representing a city. We provide comparison to the XCF of the true observed rainfall values

(a) RMSE

(b) R10 RMSE

(c) MAE

Fig. 10 Varied time span: Here, we vary the size of TRU-NET CC's training set and observe the corresponding predictive performances on a test set spanning 2014 to August 2019. Sub-figures (a), (b) and (c) show the predictive performances evaluated by RMSE, R10 RMSE and MAE respectively. We observe that RMSE and MAE scores improve as the training set size increases. The units of all metrics here are mm/day

5.2.1 Varied time span experiment

Here, we fix the test set to span the years 2014 to August 2019 and vary the number of years, starting from 1979, used to train our TRU-NET CC model. We measure the training set size by years and by unique test datums. As our model operates on extracted temporal patches from the coarse grid, the amount of unique datums in a training set is proportional to the product of the number of years we choose to train on and the number

Fig. 11 Forecasting range: Here, we inspect the annually aggregated predictive performance for a TRU-NET CC model trained on data spanning 1979 to 1997. We notice no clear trend in the R10 RMSE predictive performance as the test year becomes further forward in time from the training set. However, the RMSE and MAE show a steady decline between 1998 and 2005, after which point the predictive scores stay steady. The units of all metrics here are mm/day

of locations included in the training set. In Fig. 10, we observe a downward trend in RMSE and MAE as the number of years and unique test datums increases. The fact that the RMSE is reaching the lowest value for the largest dataset indicates that an increase of our dataset by using more locations could achieve further improvements in our model's predictive ability.

5.2.2 Forecasting range evaluation experiment

Here, we evaluate the change in quality of predictions at increasingly larger temporal distances from the time covered by the training set. We train a TRU-NET CC model using data between 1979 and 1997 and then calculate annual RMSE, R10 RMSE and MAE metrics for each calendar year of predictions between 1998 and 2018. In the results, illustrated in Fig. 11, the R10 RMSE shows no clear upward or downward trend throughout the whole test period, while the RMSE and MAE slightly decline until 2005, after which the score remains steady. This indicates that our model's predictive ability is robust to at least 21 years of weather pattern changes due to climate change and natural decadal variability. We attribute the higher near-term RMSE and MAE errors to the increased number of R10 events between the years 1998 till 2002. Specifically, inspection of the E-obs precipitation dataset revealed that 8.3% of the days between 1998 till 2002 experience rainfall events over 10 mm/day whereas in the period 2003 till 2007 only 6.8% of days with rainfall events

Table 7 Varied time period predictive scores: Here we present the predictive scores for the experiment detailed in Sect. 5.2.3. We present results for RMSE (a), R10 RMSE (b) and MAE (c). We notice that M1 generally attains better predictive scores than M2, M3 and M4

	DS1	DS2	DS3	DS4
(a) RMSE				
M1	nan	3.365	3.393	3.324
M2	3.395	nan	3.385	3.314
M3	3.410	3.376	nan	3.315
M4	3.420	3.392	3.400	nan
(b) R10 RMSE				
M1	nan	9.319	9.399	9.139
M2	9.305	nan	9.389	9.085
M3	9.410	9.418	nan	9.210
M4	9.384	9.390	9.458	nan
(c) MAE				
M1	1.738	1.794	1.788	nan
M2	nan	1.800	1.796	1.812
M3	1.766	nan	1.805	1.835
M4	1.757	1.799	nan	1.821

over 10 mm/day. As such there are more R10 events to contribute to the RMSE in the former period.

5.2.3 Varied time period experiment

To judge the extent to which TRU-NET's predictive ability is dependent on the time period it is trained on, we divide the 40 year dataset into 4 sub-datasets of 10 years each. The first sub-dataset (DS1) corresponds to years 1979–1988, the second sub-dataset (DS2) to the years 1989–1998, the third sub-dataset (DS3) to the years 1999–2008 and the fourth sub-dataset (DS4) to the years 2009–2018. We set-up a K-fold cross validation based experiment by training a separate TRU-NET CC model on each of DS1, DS2, DS3 and DS4, creating four models M1, M2, M3 and M4.

Table 7 shows the results from testing each model on the out-of-sample datasets. For each evaluation metric we perform a Tukey HSD test. A Tukey HSD test is used to accept or reject a claim that the means of two groups of values are not significantly different from each other. In our case we use the models (M1–4) as treatment groups and each models predictive scores form a groups of observations. The Tukey HSD test, then compares two models for a significant difference between the mean of each models reported predictive scores.

The Tukey HSD results for each evaluation metrics and all pairs of Models is presented in Table 8(a,b,c). The 1st two columns indicate the models under comparison, the 3rd the mean difference in their predictive scores. The rightmost column confirms whether or not there is a significant difference (sig. diff) between the performance of the corresponding pair of models. We can observe that the predictive performance between each pair of models is not significantly different. This implies our TRU-NET CC model is fairly invariant to the period of data it is trained on.

Table 8 Varied Time Period-Tukey HSD test: We train four TRU-NET CC models (M 1–4) on 4 training sets, labelled (DS 1–4), which cover mutually exclusive 10 year time spans. Each model is then tested on all time spans, except that of its training set. We evaluate the predictions using three evaluation metrics (RMSE, R10 RMSE, MAE) and then for each evaluation metric perform a Tukey HSD test on the results. The final column of each table confirms that there is no statistically significant difference (sig. diff) between the mean performance of the corresponding two models. This implies that the performance of our TRU-NET CC model is invariant to the time period it is trained on, provided all the time periods have the same time length. The units of all metrics here are mm/day

1st model	2nd model	Mean diff.	<i>p</i> value	Sig. diff.
(a) RMSE				
M1	M2	0.004	0.900	False
M1	M3	0.007	0.900	False
M1	M4	0.044	0.524	False
M2	M3	0.003	0.900	False
M2	M4	0.040	0.591	False
M3	M4	0.037	0.635	False
(b) R10 RMSE				
M1	M2	−0.026	0.900	False
M1	M3	0.060	0.900	False
M1	M4	0.125	0.595	False
M2	M3	0.087	0.796	False
M2	M4	0.151	0.460	False
M3	M4	0.065	0.900	False
(c) MAE				
M1	M2	0.029	0.612	False
M1	M3	0.029	0.625	False
M1	M4	0.019	0.839	False
M2	M3	−0.001	0.900	False
M2	M4	−0.010	0.900	False
M3	M4	−0.010	0.900	False

Table 9 Ablation study: Here, we evaluate the predictive performance of alternative methods to achieve the multi-scale hierarchical structure in the TRU-NET CC Encoder. We evaluate these TRU-NET CC variants using a training set consisting of the years 1979 till 2013. The test set is composed of data between the dates 2013 till August 2019 for the whole UK. We observe that our proposed Temporal Cross Attention (T Cross-Attn) with 8 heads outperforms other methods. The units of all metrics here are mm/day

TRU-NET CC	RMSE	R10 RMSE	MAE
T Cross-Attn 8 heads	3.072	8.729	1.634
T Cross-Attn 4 heads	3.100	8.985	1.656
T Cross-Attn 1 heads	3.100	8.893	1.656
Concatenation	3.147	9.146	1.661
Last Element	3.098	8.861	1.641
T Self-Attn	3.124	9.085	1.646

5.3 Ablation: fused temporal cross attention

In this section, we investigate the efficacy of our FTCA relative to other methods for achieving the multi-scale hierarchical structure in the TRU-NET Encoder. More concretely, we replace the temporal fused cross attention with concatenation, last element method (Jauhar et al. 2018; Zhao et al. 2019) and temporal self attention (Liu et al. 2019b). We

examine how the effect of changing the number of heads in FTCA. Table 9 shows that our model achieves lower RMSE than other methods of achieving the multi-scale hierarchical structure. Furthermore, we notice strong performance relative to the self attention variant which has the same model size. This highlights the importance of using information from higher spatio-temporal scales to guide the aggregation of information from lower spatio-temporal scales in our TRU-NET model.

6 Conclusion and future work

In this work we present TRU-NET, designed to predict high resolution precipitation from a climate model's coarse-resolution simulation of other weather variables. TRU-NET features a novel Fused Temporal Cross Attention mechanism to improve the modelling of processes defined on multiple spatio-temporal scales. We also provide a non probabilistic adaptation of the conditional-continuous loss function to obtain predictions for zero-skewed rainfall events.

For the prediction of local precipitation over the UK, using inputs defined on a coarse resolution grid, our model achieves a 10% lower RMSE than a hierarchical ConvGRU model and a 15% lower RMSE than a dynamical weather forecasting model (IFS) initialised 0–24 h before each precipitation prediction. After further analysis, we observe that TRU-NET attains lower RMSE scores than IFS when predicting rainfall events up to and including 20 mm/day, which comprises the majority of rainfall events. However, after this point TRU-NET under-predicts rainfall events to a higher degree than IFS.

We address concerns regarding the suitability of DL approaches to precipitation prediction (Rasp et al. 2020, 2018), given the limited amount of training data. We show that the current amount of data available is sufficient for a DL approach to produce quality predictions at 10 km resolution with an input of coarse-resolution variables from a numerical simulation.

The current work used deterministic models and readily available reanalysis data as an analogue for climate model output. Future works, could utilise probabilistic neural network methods, such as Monte Carlo Dropout (Gal and Ghahramani 2015) or Horseshoe Prior (Ghosh et al. 2018), as well as data from climate simulations to simulate risks of severe weather under varying climate scenarios. Further, methods combining Extreme Value Theory and machine learning (Ding et al. 2019) could be used to improve TRU-NET's ability to predict rainfall events over 20 mm/day.

Acknowledgements The project was funded by Alan Turing Institute under Climate Action Pilot Projects Call. We also acknowledge support from the HPC Facilities at the University of Warwick (Orac cluster). The code and data used to train and evaluate our models can be downloaded from <https://github.com/Akann196/TRUNET>. We acknowledge the ERA5 (Hersbach et al. 2020, 2017) and the E-OBS (Cornes et al. 2019) dataset from the Copernicus Climate Change Service, and the EU-FP6 project UERRA (<http://www.uerra.eu>) and the data providers in the ECA&D project (<https://www.ecad.eu>). Peter Dueben gratefully acknowledges funding from the Royal Society for his University Research Fellowship and the ESiWACE2 project which has received funding from the European Union's Horizon 2020 research and innovation programme under Grant Agreement No. 823988. Peter Watson gratefully acknowledges funding (Grant No. NE/S014713/1) from the Natural Environment Research Council for his Independent Research Fellowship. Yulan He gratefully acknowledges funding (Grant No. EP/V020579/1) from the UK Research and Innovation for her Turing AI Fellowship. We also acknowledge Dr. Sherman Lo's help in processing the datasets.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons

licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Agrawal, S., Barrington, L., Bromberg, C., Burge, J., Gazen, C., & Hickey, J. (2019). Machine learning for precipitation nowcasting from radar images. 1912.12132.
- Burton, C. (2011). How weather patterns have contributed to extreme precipitation in the united kingdom, and links to past flood events.
- Bárdossy, A., & Pegram, G. (2012). Multiscale spatial recorrelation of rcm precipitation to produce unbiased climate change scenarios over large areas and small. *Water Resources Research*, 48(9), 11524. <https://doi.org/10.1029/2011WR011524>.
- Cannon, A. J. (2018). Multivariate quantile mapping bias correction: An N-dimensional probability density function transform for climate model simulations of multiple variables. *Climate Dynamics*, 50(1–2), 31–49. <https://doi.org/10.1007/s00382-017-3580-6>.
- Cornes, R., van der Schrier, G., van den Besselaar, E., & Jones, P. (2019). An ensemble version of the e-obs temperature and precipitation datasets: version 21.0e.
- Ding, D., Zhang, M., Pan, X., Yang, M., & He, X. (2019). Modeling extreme events in time series prediction. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, Association for Computing Machinery, New York, NY, USA, KDD'19* (pp. 1114–1122). <https://doi.org/10.1145/3292500.3330896>.
- Gal, Y., & Ghahramani, Z. (2015). Dropout as a Bayesian approximation: Representing model uncertainty in deep learning. 1506.02142.
- Ghosh, S., Yao, J., & Doshi-Velez, F. (2018). Structured variational learning of Bayesian neural networks with horseshoe priors. 1806.05975.
- Hersbach, H., Bell, B., Berrisford, P., Hirahara, S., Horányi, A., Muñoz-Sabater, J., Nicolas, J., Peubey, C., Radu, R., Schepers, D., Simmons, A., Soci, C., Abdalla, S., Abellan, X., Balsamo, G., Bechtold, P., Biavati, G., Bidlot, J., Bonavita, M., De Chiara, G., Dahlgren, P., Dee, D., Diamantakis, M., Dragani, R., Flemming, J., Forbes, R., Fuentes, M., Geer, A., Haimberger, L., Healy, S., Hogan, R., Hólm, E., Janisková, M., Keeley, S., Laloyaux, P., Lopez, P., Lupu, C., Radnoti, G., de Rosnay, P., Rozum, I., Vamborg, F., Villaume, S., & Thépaut, J. N. (2017). Complete era5: Fifth generation of ecmwf atmospheric reanalyses of the global climate. copernicus climate change service (c3s) data store (cds). <https://doi.org/10.24381/cds.68d2bb30>.
- Hersbach, H., Bell, B., Berrisford, P., Hirahara, S., Horányi, A., & Muñoz-Sabater, J. (2020). The era5 global reanalysis. *Quarterly Journal of the Royal Meteorological Society*, 146(730), 1999–2049. <https://doi.org/10.1002/qj.3803>.
- Hofstra, N., Haylock, M., New, M., & Jones, P. D. (2009). Testing e-obs European high-resolution gridded data set of daily precipitation and surface temperature. *Journal of Geophysical Research: Atmospheres*, 114(D21), 11799. <https://doi.org/10.1029/2009JD011799>.
- Husak, G. J., Michaelsen, J., & Funk, C. (2007). Use of the gamma distribution to represent monthly rainfall in Africa for drought monitoring applications. *International Journal of Climatology*, 27(7), 935–944. <https://doi.org/10.1002/joc.1441>.
- IPCC. (2007). *Fourth Assessment Report: Climate Change 2007: The AR4 Synthesis Report*. Geneva: IPCC.
- Jauhar, S. K., Gamon, M., & Pantel, P. (2018). Neural task representations as weak supervision for model agnostic cross-lingual transfer. 1811.01115.
- Liu, L., Jiang, H., He, P., Chen, W., Liu, X., Gao, J., & Han, J. (2019a). On the variance of the adaptive learning rate and beyond. 1908.03265.
- Liu, Y. T., Li, Y. J., Yang, F. E., Chen, S. F., & Wang, Y. C. F. (2019b). Learning hierarchical self-attention for video summarization. In *2019 IEEE international conference on image processing (ICIP)* (pp. 3377–3381).
- May, W. (2004). Simulation of the variability and extremes of daily rainfall during the Indian summer monsoon for present and future times in a global time-slice experiment. *Climate Dynamics*, 22(2), 183–204. <https://doi.org/10.1007/s00382-003-0373-x>.

- Murphy, J., Harris, G., Sexton, D., Kendon, E., Bett, P., & Clark, R. (2018). Ukcpl8 land projections: Science report. Tech. rep., Met Office.
- Panjwani, S., Naresh Kumar, S., & Ahuja, L. (2021). Bias correction of gcm data using quantile mapping technique. In S. D. Purohit, D. SinghJat, R. C. Poonia, S. Kumar, & S. Hiranwal (Eds.), *Proceedings of international conference on communication and computational technologies* (pp. 617–621). Singapore: Springer.
- Rasp, S., Pritchard, M. S., & Gentile, P. (2018). Deep learning to represent subgrid processes in climate models. *Proceedings of the National Academy of Sciences*, 115(39), 9684–9689. <https://doi.org/10.1073/pnas.1810286115>.
- Rasp, S., Dueben, P. D., Scher, S., Weyn, J. A., Mouatadid, S., & Thuerey, N. (2020). Weatherbench: A benchmark dataset for data-driven weather forecasting. 2002.00469.
- Ronneberger, O., Fischer, P., & Brox, T. (2015). U-net: Convolutional networks for biomedical image segmentation. 1505.04597.
- Shaw, P., Uszkoreit, J., & Vaswani, A. (2018). Self-attention with relative position representations. In *Proceedings of the 2018 Conference of the North American chapter of the association for computational linguistics: Human language technologies (Short Papers)* (Vol. 2, pp. 464–468). <https://doi.org/10.18653/v1/N18-2074>.
- Shi, X., Chen, Z., & Wang, H. (2015). Convolutional lstm network: A machine learning approach for precipitation nowcasting. In *Advances in neural information processing systems* (Vol. 28, pp. 802–810). Curran Associates, Inc.
- Shi, X., Gao, Z., Lausen, L., Wang, H., Yeung, D. Y., Wong, W. k., & Woo, W. c. (2017). Deep learning for precipitation nowcasting: A benchmark and a new model. In *Advances in neural information processing systems* (pp. 5617–5627).
- Shukla, P., Skea, J., Buendia, E. C., Masson-Delmotte, V., & Pörtner, H. (2019). Ipcc special report on climate change, desertification, land degradation, sustainable land management, food security, and greenhouse gas fluxes in terrestrial ecosystems.
- Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., & Salakhutdinov, R. (2014). Dropout: A simple way to prevent neural networks from overfitting. *Journal of Machine Learning Research*, 15(56), 1929–1958.
- Stern, R., & Coe, R. (1984). A model fitting analysis of daily rainfall data. *Journal of the Royal Statistical Society Series A (General)*, 147(1), 1–34.
- Tai, K. S., Socher, R., & Manning, C. D. (2015). Improved semantic representations from tree-structured long short-term memory networks. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, Association for Computational Linguistics (pp. 1556–1566). Beijing, China, <https://doi.org/10.3115/v1/P15-1150>.
- Trebing, K., & Mehrkanoon, S. (2020). Smaat-unet: Precipitation nowcasting using a small attention-unet architecture. 2007.04417.
- Vandal, T., Kodra, E., Dy, J., Ganguly, S., Nemani, R., & Ganguly, A. R. (2018). Quantifying uncertainty in discrete-continuous and skewed data with bayesian deep learning. In *Proceedings of the 24th ACM SIGKDD international conference on knowledge discovery and data mining*. <https://doi.org/10.1145/3219819.3219996>.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., et al. (2017). Attention is all you need. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, & R. Garnett (Eds.), *Advances in neural information processing systems* (Vol. 30, pp. 5998–6008). New York: Curran Associates Inc.
- Wallemacq, P., & Herden, C. (2015). The Human Cost Of Weather Related Disasters. CRED, UNISDR.
- Woo, W. C., & Wong, W. K. (2017). Operational application of optical flow techniques to radar-based rainfall nowcasting. *Atmosphere*, 8(3), 48. <https://doi.org/10.3390/atmos8030048>.
- Wood, A. W., Maurer, E. P., Kumar, A., & Lettenmaier, D. P. (2002). Long-range experimental hydrologic forecasting for the eastern united states. *Journal of Geophysical Research: Atmospheres*, 107(D20), ACL6-1-ACL6-15. <https://doi.org/10.1029/2001JD000659>.
- Zhao, Y., Shenm Y., & Yao, J. (2019). Recurrent neural network for text classification with hierarchical multiscale dense connections. In *Proceedings of the twenty-eighth international joint conference on artificial intelligence, IJCAI-19, international joint conferences on artificial intelligence organization* (pp. 5450–5456). <https://doi.org/10.24963/ijcai.2019/757>.

Authors and Affiliations

Rilwan A. Adewoyin^{1,2} · Peter Dueben³ · Peter Watson⁴ · Yulan He¹ · Ritabrata Dutta⁵

¹ Department of Computer Science, University of Warwick, Coventry, UK

² Department of Computer Science and Engineering, Southern University of Science and Technology, Shenzhen, China

³ Earth System Modelling Section, The European Centre for Medium-Range Weather Forecasts, Reading, UK

⁴ School of Geographical Sciences, University of Bristol, Bristol, UK

⁵ Department of Statistics, University of Warwick, Coventry, UK