U-Net algorithm details

Building a machine learning climate emulator

Harvard Earth and Planetary Science 226, Eli Tziperman

The model first adds seven zero rows and six zero columns so that the input maps can be reduced evenly through four encoder levels. At each level, residual convolutional blocks halve the grid resolution while increasing the number of feature channels. A bottleneck processes the coarsest representation, after which four decoder levels use transposed (up) convolutions to restore the padded grid resolution. U-Net skip connections pass same-resolution features from the encoder to the decoder, preserving fine-scale spatial information during reconstruction. Finally, the padding is removed to produce next-month maps of SST, D20, and zonal wind stress.

1. State and model input

1.1 State, normalization, and calendar

Let $c \in \{1,2,3\}$ denote SST, D20, and zonal wind stress; let $M_{ij} \in \{0,1\}$ be the fixed Pacific mask on the 41×170 grid, $\Omega = \{(i,j) : M_{ij} = 1\}$, and $N = 5810$. From the first $8400 = 700 \times 12$ months, one scalar mean and population standard deviation are stored per field,

$$\mu_c = \frac{1}{8400N} \sum_{t=0}^{8399} \sum_{(i,j)\in\Omega} F_{tcij}, \qquad \sigma_c = \left[\frac{1}{8400N} \sum_{t=0}^{8399} \sum_{(i,j)\in\Omega} (F_{tcij} - \mu_c)^2 \right]^{1/2}. \tag{1}$$

At month t, define $S_{tcij} = (F_{tcij} - \mu_c)/\sigma_c$ on Ω and $S_{tcij} = 0$ outside Ω. Let $m_t \in \{0, \ldots, 11\}$ be the zero-based calendar month. Tensor dimensions below omit the batch axis. During deterministic inference, batch members are evaluated independently using the same learned parameters.

1.2 Six-channel grid input and padding

The six supplied channels are

$$X_t[:,i,j] = \big(S_{t1ij}, S_{t2ij}, S_{t3ij}, \sin(2\pi m_t/12), \cos(2\pi m_t/12), M_{ij}\big)^\top \in \mathbb{R}^6. \tag{2}$$

Seven zero rows are appended in latitude and six zero columns in longitude (on the high-index sides), making both dimensions divisible by 2^4 and giving $X^{(0)} \in \mathbb{R}^{6\times 48\times 176}$.

2. Network building blocks

2.1 Shared convolution and activation

No dropout is used. Use a 2-D convolution with zero bias, a 3×3 kernel, padding one, and stride $s \in \{1,2\}$,

$$[\mathcal{C}^s_W(A)]_{oij} = \sum_{r=1}^{C_{\text{in}}} \sum_{a=-1}^{1} \sum_{b=-1}^{1} W_{or,a+1,b+1}\, \bar{A}_{r,si+a,sj+b}, \tag{3}$$

where $\bar{A}$ is assumed zero outside its grid. Each convolutional output channel is batch-normalized with its stored running statistics and then rectified:

$$\mathcal{A}^s_W(A)_{oij} = \mathrm{ReLU}\left(\gamma_o \frac{\mathcal{C}^s_W(A)_{oij} - \widehat{\mu}_o}{\sqrt{\widehat{v}_o + 10^{-5}}} + \beta_o \right), \qquad \mathrm{ReLU}(u) = \max(0,u). \tag{4}$$

2.2 Encoder and bottleneck residual blocks

The encoder blocks decrease the grid resolution while increasing the number of feature channels. The bottleneck keeps the grid resolution fixed while increasing the number of feature channels. Here p is the number of input feature channels, and q the number of output feature channels. The residual block used by the encoder and bottleneck is the explicit map

$$\mathcal{D}^s_{p\to q}(A) = \mathcal{A}^1_{W_2}(\mathcal{A}^s_{W_1}(A)) + \mathcal{P}(A). \tag{5}$$

$\mathcal{P}(A)$ is the output of the in-block residual skip connection, and the $+$ sign denotes its element-by-element addition to the output of the main path. Here $W_1 : q \times p \times 3 \times 3$ and $W_2 : q \times q \times 3 \times 3$. The in-block residual skip connection $\mathcal{P}$ therefore makes A match the main path: it reduces the grid and changes the channel count when $s = 2$, changes only the channel count when $s = 1$ and $p \neq q$, and leaves A unchanged when $s = 1$ and $p = q$:

$$[\mathcal{P}(A)]_{oij} = \begin{cases} [\mathcal{C}^2_{W_\mathcal{P}}(A)]_{oij}, & s = 2, & 3\times 3 \text{ stride-2 convolution,} \\ \sum_{r=1}^{p} (W_\mathcal{P})_{or} A_{rij}, & s = 1 \text{ and } p \neq q, & 1\times 1 \text{ channel convolution,} \\ A_{oij}, & s = 1 \text{ and } p = q, & \text{input unchanged.} \end{cases} \tag{6}$$

where $W_\mathcal{P} \in \mathbb{R}^{q\times p\times 3\times 3}$ in the first case and $W_\mathcal{P} \in \mathbb{R}^{q\times p}$ in the second. There is no activation after the residual sum.

2.3 Decoder upsampling blocks

The goal of each transposed (up) convolution is to increase the grid resolution of each output channel. The bias-free 3×3 transposed (up) convolution (stride 2, padding 1, output padding 1) is, with invalid or noninteger source indices omitted,

$$[\mathcal{T}_W(A)]_{o,u,v} = \sum_r \sum_{a=0}^{2} \sum_{b=0}^{2} W_{roab}\, A_{r,(u+1-a)/2,(v+1-b)/2}. \tag{7}$$

It is followed by the same stored-statistics batch normalization and ReLU; call the result $\mathcal{A}_W^{\top}(A)$. At decoder levels 4–2 the complete up map is

$$\mathcal{U}_{p\to q}(A) = Z + \mathcal{A}_V^1(Z), \qquad Z = \mathcal{A}_W^{\top}(A), \tag{8}$$

where $V : q \times q \times 3 \times 3$. At the final level only, the refinement is linear: $\mathcal{U}^{\text{out}}_{p\to 3}(A) = Z + \mathcal{C}_V^1(Z)$ (no batch norm or ReLU on $\mathcal{C}_V^1$).

3. Layer-by-layer evaluation

3.1 Network evaluation and output

The encoder, bottleneck, U-Net skip connections, and decoder are used as follows. The instantiated network is exactly the following composition. Brackets $[A; B]$ denote U-Net skip connections: they concatenate same-resolution encoder and decoder features channel-wise. Each listed $\mathcal{D}$ contains the two convolutions and the in-block residual skip connection defined in Section 2.

part	operation	output size	goal
encoder	$E_1 = \mathcal{D}^2_{6\to 16}(X^{(0)})$	$16 \times 24 \times 88$	reduce grid resolution, increase # of channels
encoder	$E_2 = \mathcal{D}^2_{16\to 32}(E_1)$	$32 \times 12 \times 44$	reduce grid resolution, increase # of channels
encoder	$E_3 = \mathcal{D}^2_{32\to 64}(E_2)$	$64 \times 6 \times 22$	reduce grid resolution, increase # of channels
encoder	$E_4 = \mathcal{D}^2_{64\to 96}(E_3)$	$96 \times 3 \times 11$	reduce grid resolution, increase # of channels
bottleneck	$B = \mathcal{D}^1_{96\to 176}(E_4)$	$176 \times 3 \times 11$	same grid; more channels
decoder	$G_4 = \mathcal{U}_{272\to 64}([E_4; B])$	$64 \times 6 \times 22$	combine with E_4; double grid
decoder	$G_3 = \mathcal{U}_{128\to 32}([E_3; G_4])$	$32 \times 12 \times 44$	combine with E_3; double grid
decoder	$G_2 = \mathcal{U}_{64\to 16}([E_2; G_3])$	$16 \times 24 \times 88$	combine with E_2; double grid
decoder	$Y_t^{\text{pad}} = \mathcal{U}^{\text{out}}_{32\to 3}([E_1; G_2])$	$3 \times 48 \times 176$	combine with E_1; restore padded grid
output	$Y_t = Y_t^{\text{pad}}[:, 0{:}41, 0{:}170]$	$3 \times 41 \times 170$	remove padding; recover original grid

(9)

All ordinary and transposed (up) convolutions have no bias. Ordinary 3×3 convolutions use padding 1; all four downsamplers have stride 2. The in-block residual skip connection in the bottleneck uses a 1×1 convolution because it changes $96 \to 176$ channels without changing resolution. The in-block residual skip connections in the four downsampling blocks use 3×3 convolutions because they also change resolution. Decoder transposed (up) convolutions use kernel 3 and stride 2.

3.2 Dimensions and parameter accounting

stage	E_1	E_2	E_3	E_4	B	G_4	G_3	G_2	output
output channels	16	32	64	96	176	64	32	16	3
parameters	4,096	18,560	73,984	193,920	448,448	193,792	46,208	11,584	951

(10)

The total is 991,543 trainable parameters. Batch-normalization scale and offset are included in these counts; stored running means and variances are not trainable parameters.

4. Deterministic monthly recurrence

Starting from a standardized state S_t, construct the model input described in Section 1 and evaluate the network to obtain $Y_t = f_\theta(S_t, m_t; M)$, with no output activation. During a deterministic free run, impose the ocean mask and advance the calendar:

$$S_{t+1,cij} = M_{ij} Y_{t,cij}, \qquad m_{t+1} = (m_t + 1) \bmod 12. \tag{11}$$

Rebuild the model input from the new state and calendar month and repeat the complete network. Saved ocean output is converted to physical units field by field:

$$\widehat{F}_{t+1,cij} = \sigma_c S_{t+1,cij} + \mu_c \quad ((i,j) \in \Omega). \tag{12}$$

Land is written as missing (NaN). The only state carried between forecast months is the three-channel standardized prediction; no hidden features are carried forward. The calendar advances analytically, while the mask and grid remain fixed.

5. Architecture at a glance

