

RESEARCH ARTICLE OPEN ACCESS

Time-Aware UNet and Super-Resolution Deep Residual Networks for Spatial Downscaling

Mika Sipilä¹  | Sabrina Maggio² | Sandra De Iaco^{2,3,4}  | Klaus Nordhausen⁵ | Monica Palma^{2,3}  | Sara Taskinen¹

¹Department of Mathematics and Statistics, University of Jyväskylä, Jyväskylä, Finland | ²Department of Economic Sciences, University of Salento, Lecce, Italy | ³National Centre for HPC, Big Data and Quantum Computing, Bologna, Italy | ⁴National Biodiversity Future Center, Palermo, Italy | ⁵Department of Mathematics and Statistics, University of Helsinki, Helsinki, Finland

Correspondence: Klaus Nordhausen (klaus.nordhausen@helsinki.fi)

Received: 18 February 2026 | **Revised:** 1 July 2026 | **Accepted:** 4 August 2026

Keywords: coarse satellite data | ozone | radial basis function | temporal encoding

ABSTRACT

Satellite observations of atmospheric pollutants are often available only at coarse spatial resolution, which limits their use in local-scale environmental analysis. Spatial downscaling methods aim to transform such data into high-resolution fields. In this work, two widely used deep learning architectures—the super-resolution deep residual network (SRDRN) and the encoder–decoder-based UNet—for spatial downscaling, are extended with a lightweight temporal module that encodes observation time using either sinusoidal or radial basis function representations and integrates temporal features with spatial information. The proposed time-aware extensions are evaluated in a case study on ozone downscaling over Italy. Results show that, with only a minor increase in computational cost, incorporating temporal information significantly improves downscaling accuracy and accelerates model convergence.

1 | Introduction

Air pollution continues to be one of the most critical challenges in the current era, with profound consequences for human health, ecosystems, and the climate system. Among the major air pollutants, ground-level ozone (O_3) has received particular attention due to its harmful effects. Unlike stratospheric ozone, which protects the Earth from ultraviolet radiation, tropospheric ozone is a secondary air pollutant produced through photochemical processes from primary pollutants such as nitrogen oxides (NO , NO_2) and volatile organic compounds (VOCs) under sunlight and stagnant meteorological conditions (Zhang et al. 2019). Elevated tropospheric ozone levels have been linked to several harmful issues such as respiratory and cardiovascular diseases, premature mortality, and reduced agricultural and forest productivity (Nuvolone et al. 2018; Emberson 2020). Consequently, accurate monitoring

and prediction of ozone concentrations are of high societal and scientific importance.

Spatial downscaling is a paramount method for transforming coarse-resolution data into the fine-scale information required for effective environmental management to meet the growing demand for detailed environmental information at local and regional scales, especially in fields such as air quality, climate change, and the use of natural resources (Li et al. 2024).

Recently, Deep Learning (DL) as a core component within Artificial Intelligence (AI) has greatly improved the ability to estimate and downscale ground-level O_3 concentrations. Machine Learning (ML) methods, such as Light Gradient Boosting Machine (LightGBM), Random Forest (RF), Extreme Random Forest (ERF), and eXtreme Gradient Boosting (XGBoost),

This is an open access article under the terms of the [Creative Commons Attribution](https://creativecommons.org/licenses/by/4.0/) License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

© 2026 The Author(s). *Environmetrics* published by John Wiley & Sons Ltd.

among others, have been applied by many authors to capture complex non-linear relationships between diverse input features and O_3 concentrations (Zhu et al. 2022; Zeng et al. 2023; Shams et al. 2024; Cheng et al. 2025). However, a common limitation of traditional approaches is their simplified handling of spatio-temporal information, which may not allow them to effectively exploit all the valuable space and time features in environmental data (Zeng et al. 2023). A new generation of modeling techniques is being developed that incorporates advanced spatio-temporal techniques in order to tackle these shortcomings.

For instance, Zeng et al. (2023) introduced MixNet, a novel spatio-temporal hybrid network built upon a UNet and Long Short-Term Memory (LSTM) architecture. MixNet processed features in a multi-dimensional image format to estimate daily O_3 concentrations in the Yangtze River Delta region, demonstrating superior accuracy as compared to other ML/DL models. Similarly, Cheng et al. (2025) applied a comprehensive set of 15 AI models, including both tree-based ML models and DL models, by integrating multidimensional spatio-temporal information to enhance the prediction of surface O_3 in China. Their best model (represented by the 4D-STET model based on Extremely Randomized Trees) exhibited outstanding capability in detecting diurnal variations of surface O_3 concentrations. In addition, Generalized Additive Models with dynamic coefficients have been proposed by Jalilian et al. (2025, 2026) as flexible frameworks for spatiotemporal ozone prediction and for obtaining harmonized data across Italy.

Beyond these prediction frameworks, neural network architectures have also been extensively explored for direct spatial downscaling of environmental data. Jha et al. (2025) applied Super-Resolution Convolutional Neural Networks (SRCNN) and residual networks to downscale temperature data. In addition, Rastogi et al. (2025) used SRCNNs to emulate dynamical downscaling for daily precipitation over the conterminous United States. Furthermore, Wang et al. (2022) proposed a mutual information-guided method to spatially downscale surface O_3 concentrations from satellite data. Besides downscaling, another fundamental aspect for ensuring the reliability and accuracy of model output is bias correction. Menapace et al. (2025) provided an in-depth review of bias correction methods for climate model outputs by classifying them into univariate and multivariate approaches. These methods are essential for enhancing the accuracy of hydrological inputs and reducing uncertainties in climate change projections. Current research has also shown the efficacy of bias correction for O_3 prediction. For example, Shams et al. (2024) used a CNN-based framework to refine hourly forecasts in South Korea, while Gouldsbrough et al. (2024) employed an ML-based methodology to downscale and correct the chemical transport model output for O_3 in the UK, resulting in both enhanced accuracy and reduced biases.

Despite the above-mentioned advances documented in the literature, spatial resolution remains a critical factor for many practical applications. Satellite-derived ozone data provide broad spatial coverage but typically at coarser resolutions, which are insufficient for local-scale exposure assessments and policy decisions. Spatial downscaling methods aim to bridge this gap by enhancing coarse-resolution satellite estimates into high-resolution fields. However, most existing approaches treat downscaling as a purely

spatial task, neglecting the temporal structure of ozone variability. Since ozone levels follow a complex spatio-temporal process exhibiting strong cyclical and long-term dynamics (Donzelli and Suarez-Varela 2024), incorporating temporal information into downscaling models is crucial, especially when the goal is to generalize to future periods beyond the training data.

A straightforward way to introduce temporal information into neural network models is through sinusoidal transformations, which have been widely adopted in sequence modeling and positional encodings to capture periodicity in time (Vaswani et al. 2017). While sinusoidal transformations are effective in capturing global cyclical patterns, they might be too simple and restrictive for capturing more complex and localized spatio-temporal processes. To address this issue, time-aware spatial downscaling methods are proposed through radial basis function (RBF) transformations in time to capture temporal patterns at multiple scales without introducing substantial model complexity.

In this paper, two widely used spatial downscaling methods, Super-Resolution Deep Residual Network (SRDRN) (Wang et al. 2021) and UNet (Ronneberger et al. 2015), are extended by adding a temporal RBF module to include the temporal information in modeling the satellite images.

Three variants of each model are comprehensively compared on an ozone spatial downscaling task over Italy: (i) pure spatial baselines of SRDRN and UNet, (ii) sinusoidal time encoding SRDRN and UNet, and (iii) SRDRN and UNet with temporal RBF layers. This study focuses on evaluating the robustness of these models to generalize to future test periods, where the ability to handle temporal shifts is crucial. In addition, all proposed spatial downscaling methods along with their baseline variants are implemented in the R package *SpatialDownscaling*, which is published alongside this paper.

The paper is organized as follows. Section 2 explores other dedicated contributions on spatial downscaling, covering traditional statistical approaches and modern deep learning methods. Section 3 details the architecture of the proposed time-aware SRDRN and UNet extensions, and Section 4 provides details about the implementation of the method in R via the R package *SpatialDownscaling*. Moreover, Section 5 presents the case study of ozone downscaling in Italy, with the specific model parameters and the corresponding performance evaluation. Finally, some concluding remarks are reported in Section 6.

2 | Related Works and Deep Learning Contributions

Spatial downscaling has been of interest for several decades, motivated by the need to obtain high-resolution fields from coarse-gridded input data. Traditional statistical approaches, such as bias correction and spatial disaggregation (BCSD), have been widely used to bridge the resolution gap by applying statistical relationships between coarse predictors and fine-scale target fields (Wood et al. 2004; Maurer et al. 2010). While effective, these methods are limited in their ability to capture nonlinear and heterogeneous relationships.

With advances in deep learning and increasing computational power, many new approaches have been introduced beyond the early statistical methods. Convolutional neural networks (CNNs, Raitoharju 2022) in particular have become a cornerstone of modern spatial downscaling due to their ability to learn hierarchical spatial features directly from gridded data.

Modern deep learning methods for downscaling include super-resolution CNNs (Dong et al. 2014), encoder–decoder based networks, such as UNet (Ronneberger et al. 2015), and generative adversarial network (GAN) based super-resolution methods (Ledig et al. 2017). A key challenge in early deep architectures was that information from the input tended to vanish as it passed through many layers, leading to a loss of fine-scale details. The introduction of residual connections (He et al. 2016) alleviated this problem by allowing information to bypass intermediate layers and be reintroduced later in the network, enabling much deeper and more effective structures. Among the widely used architectures for spatial downscaling are the UNet, an encoder–decoder model with extensive residual connections originally proposed for biomedical segmentation (Ronneberger et al. 2015), and SRDRNs (Lim et al. 2017; Wang et al. 2021), which use residual blocks and skip connections to enhance coarse inputs into high-resolution outputs.

Although a multitude of spatial downscaling methods have been developed, most standard approaches neglect the temporal dependence structure inherent in geophysical and environmental data. To address this, temporal extensions have been explored. To account for temporal dependence in the model, approaches such as convolutional long-short-term memory (LSTM) extensions (Chou et al. 2021; Harilal et al. 2021) and models incorporating temporal attention mechanisms (Adewoyin et al. 2021) have been investigated. While effective, these methods are computationally heavy and might be challenging to use in practice. A lightweight alternative for incorporating temporal information is to represent the observation time with a positional encoding. Sinusoidal positional encodings are widely used in deep learning, for example, in transformers and attention mechanisms (Vaswani et al. 2017) and in seasonal modeling to capture cyclical components of time series. A more flexible approach is to use RBFs (Buhmann 2000), which can capture temporal dynamics across multiple scales, from short-term fluctuations to long-term structures. RBF-based encodings have been successfully applied in the deep learning community, including spatio-temporal models such as deep kriging (Nag et al. 2023) and identifiable variational autoencoders (Sipilä et al. 2025a, 2025b). More specifically, the implementation of RBFs for temporal encoding has been supported by their demonstrated capacity in capturing non-stationary spatiotemporal dependencies within deep latent variable models. Based on these findings, Sipilä et al. (2026b) have recently proposed some enhancements to explicitly integrate satellite-derived auxiliary covariates and handle missing observations for spatiotemporal ozone modeling. In contrast, most existing spatial downscaling methods have either neglected temporal information entirely or have relied on heavy spatio-temporal modules such as convolutional LSTMs or attention mechanisms, which substantially increase model complexity. In this work, a simpler alternative is proposed by extending SRDRN and UNet models with either sinusoidal time encodings or RBF-based temporal embeddings. This approach introduces only minimal additional complexity

while effectively incorporating temporal information into spatial downscaling.

3 | Time-Aware SRDRN and UNet

Both SRDRN and UNet are well-established convolutional neural network architectures for image-to-image translation tasks, including spatial downscaling tasks. While in many cases the coarse-resolution input alone provides sufficient information to estimate the fine-scale representation, in more complex spatio-temporal processes the appropriate downscaling transformation may itself vary depending on the time of year or day—for instance, when seasonal variation in radiation or atmospheric dynamics alters the spatial structure of the fine-scale field even for similar coarse-scale values. In this work, the SRDRN and UNet methods have been broadened by incorporating a temporal module into the models, enabling them to learn the spatio-temporal structures of the data, and to allow the downscaling transformation to differ based on the temporal position of the observation. A similar temporal module based on RBFs or sinusoidal encoding is used in both of the model architectures. In this section, the RBFs and the temporal encoding module are first introduced, then the time-aware SRDRN and UNet are presented.

3.1 | Temporal Encoding Module

To encode the temporal information, that is, the time point of the observation, into a more representative form, either sinusoidal or RBF transformations are considered. Let $t \in [1, \dots, T]$ be the time point of the observation $\mathbf{x}_t$ that contains the gridded spatial data at the time point t . The sinusoidal encoding transforms the time point t of the observation as follows:

$$g_1(t; c) = \cos\left(\frac{2\pi t}{c}\right), \quad (1)$$

$$g_2(t; c) = \sin\left(\frac{2\pi t}{c}\right), \quad (2)$$

where c is a cyclical period. This effectively allows the model to capture the global and smooth cyclical pattern present in the data.

In the RBF transformation, several node points are selected across the temporal domain at different resolution levels. Let o_{ji} be the i th selected node point at resolution level G_j . Then, the Gaussian radial basis function $v(t; \zeta, o_{ij})$ is defined as:

$$v(t; \zeta_j, o_{ij}) = \exp\left(\frac{-|t - o_{ij}|^2}{2\zeta_j}\right), \quad (3)$$

where the parameter ζ_j is a scaling parameter defined based on the selected resolution level G_j as $\zeta_j = \frac{T}{G_j+2}$. The node points for a resolution level G_j are formed by evenly selecting the G_j node points across the temporal domain. Therefore, by selecting multiple different resolution levels, large-scale and small-scale radial basis functions are obtained. Total number of radial basis functions for a resolution levels $G = (G_1, \dots, G_M)$ are $G_1 + G_2 + \dots + G_M$.

After the initial temporal encoding using either sinusoidal or RBF transformations, the encodings are fed into a feed-forward network (Svozil et al. 1997), followed by a stack of convolutional layers. The feed-forward network outputs $W \cdot H$ units, where the parameters W and H are the dimensions of the image to which the temporal information is concatenated. The output of the feed-forward network is reshaped into a $W \times H$ representation that embeds the temporal information. This 2D representation is subsequently processed by the convolutional layers, which further refine the features by capturing spatial dependencies within the temporal patterns. In this architecture, the feed-forward network performs the initial feature extraction, while the convolutional layers model higher-order structure by exploiting the spatial organization of the reshaped temporal features. The overall architecture is independent of the temporal encoding strategy, and either sinusoidal or RBF encodings can be used as input. The choice between these alternatives depends on the complexity of the temporal patterns to be represented.

The sinusoidal positional encoding is a simpler approach with fewer parameters than RBF positional encoding, and is a natural starting point when the temporal structure is expected to be smooth and periodic. The RBF encoding is more flexible and better suited when temporal dynamics are more complex, and the downscaling transformation is expected to vary nonlinearly over time. However, by introducing more parameters, RBF encoding carries a higher risk of overfitting. The choice of encoding approach, as well as the resolution levels G for RBF encoding, should be selected by validation on held-out data.

3.2 | Time-Aware SRDRN

The SRDRN is designed to reconstruct high-resolution fields from coarse-resolution inputs by leveraging CNNs, residual learning, and upsampling layers. The method utilizes residual blocks (He et al. 2016) to allow information from previous layers to bypass some intermediate layers, enabling deeper models without losing information from the earlier states. A single residual block consists of two convolutional layers with rectified linear unit (ReLU) activation in between, followed by element-wise summation of the output and input of the layer, which is called the residual connection. Mathematically, the l th residual block is defined as follows:

$$h_{l+1} = h_l + F(h_l, W_l), \quad (4)$$

where h_l is the input for the first convolutional layer and W_l contains the parameters for the l th residual block. The function F consists of the first convolutional layer, ReLU activation, and the second convolutional layer.

Another key component in the architecture of SRDRN is the upsampling block, which increases the spatial resolution of the feature maps by gradually transforming low-resolution representations into higher-resolution outputs. Rather than relying on fixed interpolation schemes, the upsampling block learns how to distribute information across spatial dimensions during training, which leads to more accurate and sharper reconstructions.

The upsampling operation in SRDRN is implemented using the sub-pixel convolution procedure (Shi et al. 2016). Instead of explicitly interpolating the feature maps using, for example, bilinear interpolation, a convolution is applied that outputs $C \cdot r^2$ channels, where C is the number of desired output features and r is the spatial upscaling factor, typically $r = 2$. The resulting array of shape $(H, W, C \cdot r^2)$ is then rearranged by a pixel shuffle operator into an array of shape $(H \cdot r, W \cdot r, C)$. This operation effectively distributes the channel information into higher-resolution spatial dimensions. Finally, a ReLU activation is applied after upsampling, enabling the block to learn non-linear feature transformations. The whole upsampling procedure allows the network to learn how to upsample the features instead of relying on naive interpolation or transposed convolution. Multiple low-factor (e.g., $r = 2$) upsampling blocks can be used in a row instead of one with a higher upsampling factor to allow the network to perform the upsampling more accurately.

The standard SRDRN is primarily composed of stacked residual blocks followed by upsampling blocks, which allow the network to reconstruct high-resolution outputs from coarse inputs. However, this baseline architecture does not incorporate temporal information. To address this, a lightweight temporal module, that captures temporal patterns and can be efficiently integrated into the network, is introduced. In the time-aware SRDRN, the input first passes through an initial convolutional layer, the output of which is stored for a global residual connection. The feature maps propagate through a sequence of residual blocks, each containing its own local residual connections to preserve spatial information from earlier layers. After the residual block sequence, an intermediate convolutional layer is applied, and the stored global residual is added to the output. The temporal features extracted by the temporal module are then concatenated with the obtained spatial features, allowing the network to efficiently model the spatio-temporal interactions. Then, the upsampling blocks progressively increase the spatial resolution, and the final convolutional layer produces the high-resolution output field. This integration of residual blocks, temporal fusion, and upsampling blocks forms the complete time-aware SRDRN architecture illustrated in Figure 1.

The parameters of the network are trained by minimizing the mean squared error (MSE) between the model output and the true high-resolution field. Prior to training, the training data are normalized to have zero mean and unit standard deviation. The main hyperparameters include the number of residual blocks n_L , the number of upsampling blocks n_B and their scale parameters r_b , $b = 1, \dots, B$, the number of channels in each convolutional layer, and the structure of the temporal module (number of dense layers, units, and temporal encoding procedure). For RBF encoding, the resolution levels G to form the radial basis functions have to be chosen.

In practice, a relatively small number of residual blocks (e.g., $n_L = 2 - 5$) is typically sufficient to balance model depth and computational efficiency, while the number of upsampling blocks and their scaling factors are chosen to match the desired overall magnification (e.g., two blocks with $r_b = 2$ for $4\times$ upscaling). The number of convolutional channels is usually kept constant (e.g., 64 or 128) across layers to maintain compatibility




FIGURE 1 | Schematic overview of time-aware SRDRN architecture. The residual blocks are presented within dashed red lines, the upsampling blocks within dashed blue lines, and the temporal module within dashed green lines.

between the residual connections. In the initial feed-forward stack of the temporal module, a relatively small number of hidden layers (e.g., 1–5) with a gradually increasing number of units is sufficient, as temporal encoding already provides an expressive representation. Likewise, only a few subsequent convolutional layers (e.g., 1–3) are typically required, since their role is not to build deep hierarchical structure but simply to refine the 2D embedding and capture local spatial interactions within a single time point. The proportion of the temporal information fed into the upsampling path can be increased by increasing the number of filters in the convolutional layers. 4–32 filters in the final convolution layers is a recommended starting point, depending on the influence of the temporal part in the application of interest.

In the upscaling path, the number of filters in the convolutional layers typically varies between 8 and 128, so that the number of filters decreases when the resolution of the features increases.

For RBF encoding, multiple resolution levels should be used to capture both large-scale temporal patterns and small-scale details. The network is typically optimized using Adam optimizer (Kingma and Ba 2015) with an initial learning rate between 10^{-4} and 10^{-3} , a batch size of 16–64. The number of training epochs depends on the size of the dataset, and the convergence should be monitored by using separate validation data.

3.3 | Time-Aware UNet

The UNet architecture is an encoder–decoder network originally designed for image segmentation (Ronneberger et al. 2015; Adewoyin et al. 2021). It employs skip connections between the encoder and decoder paths to preserve spatial information across different scales of the field. In the baseline UNet, the input passes through a series of downsampling convolutional blocks, which gradually reduce the spatial resolution while increasing the number of feature channels, followed by a symmetric upsampling path that reconstructs the high-resolution output.

A key feature of the UNet architecture is the use of skip connections that link corresponding layers in the encoder and decoder paths. At each stage of the decoder, the feature maps from the encoder at the same spatial resolution are concatenated with the upsampled decoder features. This allows the network to preserve high-resolution spatial information that might otherwise be lost during downsampling, improving the reconstruction of fine-scale structures. Unlike the residual connections in SRDRN, which use element-wise summation, the UNet skip connections use concatenation, providing the decoder with access to the full set of encoder features.

To efficiently incorporate temporal information, this architecture has been extended by adding a temporal module similarly to




FIGURE 2 | A schematic representation of time-aware UNet. The temporal module is presented within the dashed green line, the encoder blocks within the dashed red line, and the decoder blocks within the dashed blue line.

that of SRDRN. The temporal features, obtained from the temporal module, are concatenated with the feature maps at the beginning of the decoder, where the spatial resolution is low. This placement allows the temporal information to be propagated through the upsampling path, enabling the network to modulate both coarse and fine-scale spatial features and effectively learn spatio-temporal interactions across the entire high-resolution output.

The complete time-aware UNet is constructed as follows. An initial bilinear upsampling layer is first applied to bring the coarse-resolution input to the desired target resolution. This ensures that the encoder and decoder paths operate at the same spatial scale, allowing skip connections to effectively transfer information between the corresponding encoder and decoder layers. After this, n_K initial feature extraction convolutional layers with ReLU activation are used to obtain rich, higher-dimensional, low-level representation. The features are then passed through a sequence of n_M encoder blocks, each consisting of two convolutional layers with ReLU activation and a max pooling operation. Each encoder block reduces the spatial resolution by a factor of two while increasing the number of feature channels. The outputs of the encoder blocks are stored for the skip connections.

At the bottleneck, two convolutional layers with ReLU activation are applied to obtain the final encoded spatial features, after which an optional dropout layer is applied, and the features obtained from the temporal module are concatenated. The concatenated spatial and temporal features are fed into the decoder. The decoder is symmetric to the encoder and therefore consists of n_M decoding blocks, each containing a bilinear upsampling layer followed by two convolutional layers with ReLU activation.

The skip connections are applied by concatenating the features from the encoder blocks with the corresponding decoder features at each spatial scale. Finally, a convolutional output layer is applied to generate the high-resolution output field. A schematic overview of the time-aware UNet is provided in Figure 2.

The UNet parameters are trained by minimizing the MSE between the estimated high-resolution output and the true high-resolution field. The main hyperparameters include the number of encoder and decoder blocks n_M , the number of initial feature extraction layers n_K , the number of channels within each convolution layer, the temporal encoding process, and the number of hidden layers within it. In addition, general hyperparameters such as the optimizer, learning rate, batch size, and the number of epochs must be specified.

In practice, UNet architectures for spatial downscaling typically employ two to four encoder-decoder blocks ($n_M = 2-4$), each halving and doubling the spatial resolution in the encoder and decoder paths, respectively. The number of initial feature extraction layers n_K is usually one or two, allowing the network to learn low-level spatial representations before entering the encoder path. The number of channels is often doubled at each deeper encoder level and symmetrically reduced in the decoder, thereby facilitating the transfer of information through the skip connections. The temporal module and the remaining hyperparameters are selected analogously to the time-aware SRDRN.

The preceding sections have described the proposed time-aware extensions of SRDRN and UNet. It is also worthwhile to briefly contrast these deep learning architectures with more traditional probabilistic approaches for spatial downscaling, since both classes of methods aim to exploit spatial dependence but differ

substantially in their modeling philosophy and computational characteristics.

Probabilistic approaches based on Markov Random Fields (MRFs) describe spatial dependence through an explicit probabilistic framework in which each location is linked to neighboring locations according to a predefined conditional dependence structure (Besag 1974; Li 2009). This framework offers several attractive features, including model interpretability, uncertainty quantification, and the possibility of incorporating prior physical knowledge into the model specification (Rue and Held 2005; Cressie and Wikle 2011). However, as spatial resolution increases and the number of locations grows, parameter estimation and probabilistic inference may become computationally demanding, often requiring iterative optimization or sampling procedures whose computational cost increases substantially with domain size and model complexity (Rue and Held 2005).

Deep learning approaches adopt a fundamentally different strategy. Rather than explicitly specifying spatial dependence, convolutional neural networks learn multiscale spatial representations directly from data through hierarchical feature extraction (LeCun et al. 2015; Goodfellow et al. 2016). Architectures such as SRDRN and UNet have demonstrated excellent performance in image-to-image prediction problems and have been successfully adapted to environmental and remote-sensing applications. This data-driven paradigm is particularly attractive for satellite-based spatial downscaling, where large volumes of observations and highly nonlinear atmospheric processes make explicit probabilistic modeling increasingly difficult (Reichstein et al. 2019).

Furthermore, deep neural networks naturally exploit modern parallel computing architectures, enabling efficient training and rapid inference over large spatial domains (LeCun et al. 2015; Reichstein et al. 2019). The proposed time-aware extensions further enrich these models by incorporating temporal information into the learned spatial representations, allowing the downscaling transformation itself to adapt to temporal variability while maintaining computational efficiency.

Consequently, the choice between probabilistic and deep learning approaches ultimately depends on the objectives of the application. MRF-based models remain attractive when probabilistic interpretation, uncertainty quantification, and interpretable spatial dependence are primary goals. In contrast, for operational spatial downscaling involving large satellite datasets, complex nonlinear spatio-temporal relationships, and high-resolution prediction, architectures such as the proposed time-aware SRDRN and UNet provide a scalable and computationally efficient alternative with strong predictive performance (Reichstein et al. 2019; Vandal et al. 2017; Baño-Medina et al. 2020).

4 | R Package for Spatial Downscaling

As part of this work, all considered spatial downscaling methods have been implemented in the R package `SpatialDownscaling` (Sipilä et al. 2026a). The main functions of the package are `srdrn` and `unet`, which implement the SRDRN and UNet

architectures as described in Sections 3.2 and 3.3, respectively. Both functions support optional temporal conditioning through a shared temporal module architecture. The methods are built mainly on top of the packages `keras3` (Kalinowski et al. 2024) and `tensorflow` (Allaire and Tang 2024).

Both `srdrn` and `unet` follow a consistent interface, accepting coarse-resolution input data and fine-resolution target data as multi-dimensional arrays with dimensions $[x, y, \text{time}]$. The spatial dimensions x and y represent the grid coordinates, while the time dimension corresponds to different temporal snapshots. Optional validation data can be provided with the parameters `val_coarse_data` and `val_fine_data`. The most important parameters controlling the network architectures correspond directly to the mathematical formulations presented in Sections 3.2 and 3.3.

Temporal conditioning is activated by providing the `time_points` parameter, which is a vector of time indices corresponding to each temporal snapshot in the data. When temporal information is included, the temporal module parameters control how temporal patterns are encoded. The `cyclical_period` parameter specifies the period for cyclical temporal patterns (e.g., 365 for daily data with annual cycles), while `temporal_basis` controls the number and placement of radial basis functions for temporal encoding (default: `c(9, 17, 37)`). The `temporal_layers` parameter defines the hidden layer sizes in the initial temporal feed-forward network (default: `c(32, 64, 128)`) and the parameters `temporal_cnn_filters` (default: `c(8, 16)`) and `temporal_cnn_kernel_sizes` (default: `list(c(3, 3), c(3, 3))`) define the number of temporal CNN layers and their filters and kernel sizes, respectively. Alternatively, users can set `cos_sin_transform` to `TRUE` to use sinusoidal transformation instead of radial basis functions for cyclical time representation. If the validation data are provided, the user should also provide the validation time points (`val_time_points`) if the temporal modeling is included.

For the SRDRN architecture, the key structural parameter is `num_residual_blocks`, which controls the number of residual blocks n_R in the network (default: 3). The number of upsampling blocks is automatically calculated as $\lfloor \log_2(r) \rfloor$, where r is the spatial upscaling factor determined from the ratio of fine to coarse resolution dimensions. The parameter `num_res_block_filters` defines the number of filters in the convolutional layers of the residual blocks (default: 64), while the vector `upsampling_filters` defines the number of filters used in each upsampling block, with defaults taken sequentially from `c(64, 32, 16, 8, 4, 2)` based on the number of upsampling blocks needed.

For the UNet architecture, `initial_filters` specifies the number of filters in the n_K initial feature extraction layers (default is 16), while `filters` defines the number of filters in each of the n_M encoder/decoder blocks (default: `c(32, 64, 128)`). The `kernel_sizes` parameter controls the convolutional kernel sizes for each encoder/decoder block.

Both functions share common training parameters including `learning_rate` for the Adam optimizer (default: 0.001),

epochs for the number of training iterations (default is 10), and `batch_size` for mini-batch size of the stochastic gradient descent (default: 32). Loss function, optimizer and nonlinear activation functions can be modified with parameters `loss` (default: “mse”), `optimizer` (default: “adam”) and `activation` (default: “relu”), respectively. The options for the parameters are inherited from the `keras3` package.

The typical workflow for spatial downscaling using the package consists of three main steps:

1. Prepare coarse and fine resolution training data as arrays with dimensions [x, y, time].
2. Train a model by calling `srdn` or `unet` with the training data and desired parameters.
3. Generate high-resolution predictions for new coarse-resolution inputs using the `predict` method.

For example, to train a time-aware SRDRN model with coarse training data of dimensions $32 \times 32 \times 1000$ and fine training data of dimensions $128 \times 128 \times 1000$ (corresponding to a $4\times$ upscaling factor), one would call `srdn` with the input and target data arrays, provide the time points vector (e.g., 1 to 1000), specify a cyclical period of 365 for annual cycles, and set the desired number of residual blocks and training epochs. The resulting model object can then be used with the `predict` function to generate downscaled predictions for new coarse-resolution data by providing both the new coarse data array and corresponding time points. The same workflow applies to UNet training using the `unet` function, with appropriate architectural parameters such as the filter sizes for encoder/decoder blocks.

Both methods handle missing values automatically by masking them during training and prediction. The methods use `keras3` and `tensorflow` backends for efficient computation. Model objects returned by the training functions contain the trained neural network, normalization parameters, and metadata required for prediction, allowing trained models to be saved and reused later for inference.

In addition to SRDRN and UNet methods, the package contains an implementation for the statistical baseline model BCSD. The method is called similarly to UNet and SRDRN methods by giving coarse and fine training datasets. The key parameters include `method` for specifying the interpolation approach (default: bilinear), `n_quantiles` for controlling the number of quantile levels used in the quantile mapping (default: 100), and `reference_period` for restricting the quantile calculation to a specific temporal subset of the training data. The `extrapolate` parameter (default: TRUE) determines whether values outside the training range are extrapolated using constant shift corrections, while `normalize` controls data standardization (default: TRUE). Unlike the neural network approaches, BCSD does not require iterative training and instead creates a quantile mapping function directly from the training data, making it computationally efficient but usually less flexible in capturing complex spatial patterns. The `predict` method for BCSD objects applies the learned quantile mapping followed by spatial interpolation to generate high-resolution fields.

5 | Case Study

Tropospheric ozone is a key air pollutant with strong spatio-temporal variability, influencing both human health and ecosystem functioning. High-resolution ozone fields are essential for accurately assessing local air quality, but the satellite products are often provided only at coarse resolution levels. In this study, the focus is on spatial downscaling of daily gridded satellite-derived ozone concentration data over Italy. Two well-known image-to-image architectures, SRDRN and UNet, are compared with their proposed time-aware variants. For the time-aware variants, both sinusoidal and RBF encodings are considered. In addition, the BCSD method is included as a statistical baseline for comparison.

The analysis is performed in R 4.4.1 (R Core Team 2023) and all implemented methods are available in the R package `SpatialDownscaling` described in Section 4. The reproducible code to replicate all results in the case study is provided on GitHub.¹

5.1 | Analyzed Data

This study focuses on the spatial downscaling of gridded tropospheric O_3 concentrations ($\mu g/m^3$) at ground level, over the entire Italian territory. High-resolution ozone fields are essential for accurately assessing local air quality, particularly in regions characterized by challenging landscapes and intricate meteorological dynamics. The target area includes the whole national territory, encompassing diverse geographical features (from high Alpine and Apennine systems to coastlines and the extensive plains). Italy represents a key economic and demographic hub in Southern Europe, whose importance is intrinsically linked to its environmental complexity. The most persistent pollution issues of the territory are driven by several factors, notably the distinct geographic basin effect experienced in large flat areas such as the Po Valley (Pianura Padana), as well as in other densely populated or industrialized areas distributed across the peninsula and coastal zones. This configuration, bordered by the Alps to the north and the Apennines to the south, critically traps pollutants and restricts atmospheric mixing. Consequently, in those areas, O_3 concentrations often exceed established health and environmental thresholds, especially during the summer. In particular, elevated contamination in several Italian regions is due to a confluence of conditions: high levels of precursor pollutants (such as nitrogen oxides from vehicle traffic and industrial operations), intense solar radiation, and region-specific weather patterns.

A visual overview of the study area and an example of ozone distributions at different spatial resolutions are provided in Figure 3a–c.

The gridded O_3 concentrations at ground level were sourced from the Copernicus Atmosphere Monitoring Service (AQ-CAMS). The daily data span six years, from 2018 to 2023, yielding a total of 2191 records. The input dataset consisted of low-resolution O_3 data (0.4° on a 30×30 grid), obtained by averaging the original high-resolution fields (0.1° on a 120×120 grid) over non-overlapping 4×4 spatial blocks.




FIGURE 3 | The map of the study area is marked with a red square (a). The satellite-derived ozone levels are provided on 1/7/2022 at fine resolution (b) and at coarse resolution (c).

The objective of the study is to reconstruct the high-resolution ozone fields from their coarse counterparts over time.

The performance of the proposed approach is evaluated using this satellite gridded tropospheric ozone dataset. The compared methods include variants of SRDRN and UNet, both with and without positional encoding. For model evaluation, the last 200 time points are reserved as test data.

Figure 3a–c show the study area and examples of satellite-derived ozone concentration maps at fine and coarse resolutions for July 1, 2022. The figure highlights the strong contrast between fine and coarse-scale representations: fine-resolution data capture local gradients and hotspots (particularly in urbanized or industrial regions and over challenging landscapes), while coarse-resolution fields smooth out these spatial variations. This difference visually demonstrates the importance of this downscaling approach for recovering small-scale features crucial to accurate air quality assessment across Italy.

5.2 | Model Parameters

In all SRDRN variants, $n_L = 3$ encoder blocks and $n_B = 2$ upsampling blocks with upscale factors $r_1, r_2 = 2$ are used. The convolutional layers of the encoder blocks have 64 filters, while the upsampling blocks have 64 and 32 filters, respectively. In the temporal modules of the time-aware UNet and SRDRN models, the initial feed-forward network consists of layers with 32, 64, 128, 256, and 900 hidden units, respectively. The last hidden layer has 900 hidden units in order to reshape the output into a 30×30 array, which is the size of the coarse-resolution input. Then, two consecutive convolutional layers with 8 and 16 filters, respectively, are applied to produce a distinctive spatial representation of the temporal information.

For RBF variants, the radial basis functions are constructed based on the seasonal time period $t = 1, \dots, 365$ instead of the absolute

time point to allow the model to generalize better to future values. The resolution levels $G = (9, 17, 37)$ are used to create the seasonal RBFs.

In each UNet variant, one initial convolutional layer with 16 filters and two encoding blocks with filter sizes 32 and 64 are used. The bottleneck convolutional layers have 128 filters each, and the decoder is symmetric to the encoder, having 64 and 32 filters in the two decoding blocks.

The number of training epochs for SRDRN and UNet was determined based on an initial training where a validation accuracy was monitored. In the validation step, the models were trained on 1500 first time points of the training dataset and evaluated based on the remaining 490 time points. The convergence and good generalization to unseen data were ensured by selecting the number of epochs at which the validation accuracy stabilizes. The convergence curves are provided in Figure 4. As the validation accuracy for SRDRN methods steadily decreases and stabilizes only right before the final training epoch, the final models are trained for a total of 80 epochs. Because the UNet models converge more rapidly and stabilize at around epoch 40, the final models are trained for a total of 40 epochs. Moreover, based on the validation losses per training epoch, the time-aware variants of UNet and SRDRN achieve lower validation loss and converge significantly faster than their baseline models. The learning rate was set to 0.0005 for all models in the validation step and in final training.

5.3 | Performance Metrics

The performance of the different downscaling methods is evaluated using mean absolute error (MAE), root mean squared error (RMSE), and Kling-Gupta efficiency (KGE) developed by Gupta et al. (2009). MAE and RMSE are standard error measures that quantify the average magnitude of prediction errors, with lower values indicating better performance. KGE provides a measure




FIGURE 4 | The comparison of the validation losses against the number of training epochs for the baseline UNet (a) and SRDRN (b) models along with their time-aware extensions. The dashed lines represent the exact validation losses, while the solid lines are smoothed using locally weighted regression (LOESS) with a smoothing span of 0.3.

that decomposes the model performance into three components: correlation, bias, and variability. It is defined as

$$\text{KGE} = 1 - \sqrt{(r - 1)^2 + (\alpha - 1)^2 + (\beta - 1)^2}, \quad (5)$$

where r is the linear correlation coefficient between the predictions and the true values, $\alpha = \sigma_{\hat{x}}/\sigma_x$ is the variability ratio with $\sigma_{\hat{x}}$ and σ_x denoting the standard deviations of predicted and observed values, respectively, and $\beta = \mu_{\hat{x}}/\mu_x$ is the bias ratio with $\mu_{\hat{x}}$ and μ_x denoting their means. KGE ranges from negative infinity to 1, where the optimal value of 1 indicates a perfect fit for the model. The value greater than -0.41 indicates performance of mean flow (Knoben et al. 2019), meaning that the true variables would be estimated by their mean.

5.4 | Results

After initial validation steps, the different models are trained by using the whole training dataset. The performance metrics, presented in Table 1, are calculated for the test dataset.

All SRDRN and UNet variants substantially outperformed BCSD. Among the baseline methods without a temporal module, SRDRN reached lower errors and higher KGE than UNet. Inclusion of the temporal information consistently improved the performance for both architectures. For SRDRN, the sinusoidal positional encoding yielded the best results, whereas in UNet, RBF encoding was slightly better. The relative MAE improvements of the time-aware methods compared to the baseline methods were 11.5% for sinusoidal SRDRN, 9.7% for RBF SRDRN, 15.7% for sinusoidal UNet, and 18.2% for RBF UNet. While both time-aware variants outperformed their baseline counterparts, the differences between sinusoidal and RBF encoding were small in this application. Although the time-aware UNet with RBF encoding yielded the highest performance across all metrics, these marginal differences suggest that the simple periodic structure captured by sinusoidal encoding may be sufficient to represent the temporal variation in the downscaling transformation for ozone.

From a computational perspective, the temporal module adds negligible overhead. It consists of a simple feed-forward network

TABLE 1 | Performance metrics (MAE, RMSE, and KGE) for different methods in spatial downscaling of ozone in Italy.

Method	MAE	RMSE	KGE
BCSD	1.453	2.165	0.9875
SRDRN (baseline)	0.765	1.076	0.9924
SRDRN (sinusoidal)	0.677	0.965	0.9982
SRDRN (RBF)	0.691	0.973	0.9932
UNet (baseline)	0.782	1.080	0.9809
UNet (sinusoidal)	0.659	0.950	0.9925
UNet (RBF)	0.640	0.929	0.9970

together with an optional stack of lightweight convolutional layers, while the computationally demanding operations remain in the main SRDRN and UNet modules. Consequently, the training times per epoch for baseline and time-aware variants were nearly identical, and no practical differences in runtime were observed. However, because of the faster training convergence of the time-aware models, they can be trained with lower computational resources than the baseline models.

The performance of the methods is further compared by calculating average MAE spatial maps over the test time period. The spatial patterns are presented in Figure 5 for SRDRN and Figure 6 for UNet. While the maps show only slightly lower errors for time-aware models in Southern and Central Italy, more pronounced differences are observed across Northern Italy. For all methods, the lowest errors occur over the sea, whereas the highest errors are concentrated in Northern Italy. These elevated errors are likely associated with the particularly challenging atmospheric and emission conditions that characterize this region. Several interacting factors in Northern Italy, particularly across the Piedmont, Lombardy, Veneto, and Emilia-Romagna regions (the Po Valley Basin), can make ozone forecasting less accurate. Topographical features that restrict ventilation, together with meteorological conditions characterized by persistent thermal inversions and sea-breeze circulations from the Adriatic and Ligurian seas (Pivato et al. 2023; Capozzi et al. 2026), make this area the major air pollution hotspot in Italy, with




FIGURE 5 | The mean absolute prediction errors over the test set for the SRDRN variants. The first column presents the errors for the baseline SRDRN, the second for time-aware SRDRN with RBF encoding, and the third for time-aware SRDRN with sinusoidal encoding.




FIGURE 6 | The mean absolute prediction errors over the test set for the UNet variants. The first column presents the errors for the baseline UNet, the second for time-aware UNet with RBF encoding, and the third for time-aware UNet with sinusoidal encoding.

excessive accumulation of ozone and its precursors. The region is also characterized by high population density, heavy industry, and intensive agriculture, resulting in substantial emissions of ozone precursors, especially nitrogen oxides and VOCs. Under such complex local conditions, models may struggle to accurately reproduce small-scale dynamics, leading to higher prediction errors.

6 | Conclusions and Discussion

In this study, time-aware extensions of SRDRN and UNet architectures were implemented by incorporating a temporal module that extracts information from positionally encoded time points. The resulting models were evaluated for the spatial downscaling of ozone concentrations in Italy.

This downscaling process is critical because satellite observations of O_3 are typically available only at coarse resolution (e.g., 0.4° resolution), which is insufficient for local-scale exposure assessments and policy decisions. The specific case study focused on reconstructing high-resolution fields (0.1° resolution) across the Italian territory, a vast area characterized by high O_3 levels due to high precursor pollutants, intense solar radiation, and a distinctive geographical configuration.

All deep learning models considered substantially outperformed the statistical baseline BCSD. Across all metrics, the time-aware variants consistently outperformed their baseline counterparts in terms of prediction accuracy, while introducing only marginal computational overhead. The time-aware extensions also provided considerably faster training convergence compared to the baseline methods.

The spatial patterns of the prediction errors were further studied by spatial maps of MAEs averaged over the test time period for the different methods. Although the differences are modest, the average MAEs of the time-aware variants were clearly lower, especially across Northern Italy.

In conclusion, the performance improvements of the time-aware variants highlight the effectiveness and potential of including a lightweight temporal module with positional encoding to incorporate temporal information into convolutional architectures for spatial downscaling. Given the faster convergence and similar training time per epoch, the proposed models can be trained successfully with lower computation time than the baseline models without a temporal module.

While both encoding strategies performed well in this study, future work should explore their performance across different

datasets and under varying levels of downscaling complexity using simulation studies. In particular, for more challenging scenarios such as higher spatial scaling factors or more complex environmental phenomena, the flexibility of RBF-based encoding may offer an advantage.

Although the motivation of this work was to develop a computationally light yet effective way to incorporate temporal information into common spatial downscaling deep-learning architectures, it will be important to benchmark these time-aware methods against more resource-intensive temporal models, such as recurrent or attention-based networks, which can explicitly leverage temporal dependencies and past observations. The temporal module proposed in this paper encodes the position of an observation within the seasonal cycle instead of conditioning on the previous observations explicitly. This design choice is well-suited for variables such as ozone, which are governed by smooth and recurring seasonal and diurnal cycles that systematically shape the fine-scale spatial structure. However, in more complex settings where the downscaling transformation differs substantially between time points sharing both similar coarse-resolution inputs and similar positional encodings, as may occur in highly dynamic processes such as precipitation, including only the positional encoding may be sub-optimal. In such settings, explicitly incorporating recent past observations through a sequence-based architecture might be beneficial.

Finally, while the proposed framework was evaluated for a single variable and geographical domain, the architecture itself generalizes to any gridded environmental variable, any spatial region, and allows varying downscaling factors without structural modifications. However, the fitted model weights are specific to the variable, region, and downscaling factor for which the model was trained. This is consistent with the broader spatial downscaling literature and does not represent a limitation specific to the proposed approach. Future work should assess the framework's performance across other environmental variables and across diverse spatial regions to establish its external validity more comprehensively. Transfer learning strategies, in which a model pretrained on one domain is fine-tuned on another, may also be effective in reducing training data requirements and computational cost when adapting the models to new variables and spatial areas.

Acknowledgments

The authors are grateful to the Editor and the anonymous referees for their useful suggestions which contribute to improve the paper.

This research was partially supported by the European Union-NextGenerationEU with the Cascade Open Calls published by ALMA MATER STUDIUM – University of Bologna, inside the Project GRINS funded by PNRR – Mission 4, Component 2, Investment 1.3 “Partnership extended to Universities, Research Centers, Firms and research projects funding”, D.D. 341 of 15/03/2022, “ECOST-DATA, Exploring Spatio-Temporal Environmental Conditions: Harmonized Databases and Analytical Techniques”, CUP: J33C22002910001.

The work of Klaus Nordhausen was supported by the Research Council of Finland (363261). The work of Mika Sipilä was supported by the Vilho, Yrjö and Kalle Väisälä foundation. The work of Sara Taskinen was supported by the Research Council of Finland (356484).

Open access publishing facilitated by Helsingin yliopisto, as part of the Wiley - FinELib agreement.

Funding

This work was supported by the European Union-NextGenerationEU with the Cascade Open Calls published by ALMA MATER STUDIO – University of Bologna, inside the Project GRINS funded by PNRR – Mission 4, Component 2, Investment 1.3 “Partnership extended to Universities, Research Centers, Firms and research projects funding”, D.D. 341 of 15/03/2022, “ECOST-DATA, Exploring Spatio-Temporal Environmental Conditions: Harmonized Databases and Analytical Techniques”, CUP: J33C22002910001. The work of Mika Sipilä was supported by the Vilho, Yrjö and Kalle Väisälä foundation. Klaus Nordhausen and Sara Taskinen were supported by the Research Council of Finland (363261, 356484).

Conflicts of Interest

The authors declare no conflicts of interest.

Data Availability Statement

The data that support the findings of this study are openly available in GRINS AQCLIM: The GRINS datasets on Air Quality and Climate at: <https://doi.org/10.5281/zenodo.17103414>, reference number zenodo.17103414.

Endnotes

¹<https://github.com/mikasip/OzoneSpatialDownscaling>.

References

- Adewoyin, R. A., P. Dueben, P. Watson, Y. He, and R. Dutta. 2021. “TRU-NET: A Deep Learning Approach to High Resolution Prediction of Rainfall.” *Machine Learning* 110, no. 8: 2035–2062.
- Allaire, J., and Y. Tang. 2024. “Tensorflow: R Interface to ‘Tensorflow’ (Computer software manual).” <https://CRAN.R-project.org/package=tensorflow>. R package version 2.16.0.
- Baño-Medina, J., R. Manzanar, and J. M. Gutiérrez. 2020. “Configuration and Intercomparison of Deep Learning Neural Models for Statistical Downscaling.” *Geoscientific Model Development* 13, no. 4: 2109–2124. <https://doi.org/10.5194/gmd-13-2109-2020>.
- Besag, J. 1974. “Spatial Interaction and the Statistical Analysis of Lattice Systems.” *Journal of the Royal Statistical Society. Series B, Methodological* 36, no. 2: 192–236.
- Buhmann, M. D. 2000. “Radial Basis Functions.” *Acta Numerica* 9: 1–38.
- Capozzi, V., A. Fucci, G. Budillon, et al. 2026. “Identification of Atmospheric Circulation Schemes That Promote Summer Hail Events in the Northern Italy.” *Atmospheric Research* 338: 108990.
- Cheng, F., Z. Li, Z. Yang, et al. 2025. “First Retrieval of 24-Hourly 1-Km-Resolution Gapless Surface Ozone (O_3) From Space in China Using Artificial Intelligence: Diurnal Variations and Implications for Air Quality and Phytotoxicity.” *Remote Sensing of Environment* 316: 114482.
- Chou, C., J. Park, and E. Chou. 2021. “Generating High-Resolution Climate Change Projections Using Super-Resolution Convolutional LSTM Neural Networks.” In *2021 13th International Conference on Advanced Computational Intelligence (ICACI)*, 293–298. IEEE.
- Cressie, N., and C. K. Wikle. 2011. *Statistics for Spatio-Temporal Data*. John Wiley & Sons. <https://doi.org/10.1002/9781118180112>.
- Dong, C., C. C. Loy, K. He, and X. Tang. 2014. “Learning a Deep Convolutional Network for Image Super-Resolution.” In *Computer Vision – ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6–12, 2014, Proceedings, Part IV*, edited by D. Fleet, T. Pajdla, B. Schiele, and

- T. Tuytelaars, 184–199. Lecture Notes in Computer Science, Vol. 8692. Springer.
- Donzelli, G., and M. M. Suarez-Varela. 2024. "Tropospheric Ozone: A Critical Review of the Literature on Emissions, Exposure, and Health Effects." *Atmosphere* 15, no. 7: 779.
- Emberson, L. 2020. "Effects of Ozone on Agriculture, Forests and Grasslands." *Philosophical Transactions of the Royal Society A* 378, no. 2183: 20190327.
- Goodfellow, I., Y. Bengio, and A. Courville. 2016. *Deep Learning*. MIT Press.
- Gouldsbrough, L., R. Hossaini, E. Eastoe, P. J. Young, and M. Vieno. 2024. "A Machine Learning Approach to Downscale EMEP4UK: Analysis of UK Ozone Variability and Trends." *Atmospheric Chemistry and Physics* 24: 31633196.
- Gupta, H. V., H. Kling, K. K. Yilmaz, and G. F. Martinez. 2009. "Decomposition of the Mean Squared Error and NSE Performance Criteria: Implications for Improving Hydrological Modelling." *Journal of Hydrology* 377, no. 1–2: 80–91.
- Harilal, N., M. Singh, and U. Bhatia. 2021. "Augmented Convolutional LSTMs for Generation of High-Resolution Climate Change Projections." *IEEE Access* 9: 25208–25218.
- He, K., X. Zhang, S. Ren, and J. Sun. 2016. "Deep Residual Learning for Image Recognition." In *Proceedings of the Ieee Conference on Computer Vision and Pattern Recognition*, 770–778. IEEE.
- Jalilian, A., C. Cappello, M. Palma, and S. De Iaco. 2025. "Ozone Predictions Through a Generalized Additive Model." *Techniques. ECoST-DATA 2025. Springer Proceedings in Mathematics & Statistics*, edited by S. De Iaco, M. Palma and D. Posa, 95–103, Vol. 531. Springer. https://doi.org/10.1007/978-3-032-17526-7_9.
- Jalilian, A., C. Cappello, M. Palma, and S. De Iaco. 2026. "Generalized Additive Model With Dynamic Coefficients for Spatiotemporal Ozone Predictions." *Environmetrics* 37, no. 3: e70076. <https://doi.org/10.1002/env.70076>.
- Jha, S. K., V. Gupta, P. J. Sharma, A. Mishra, and S. Joshi. 2025. "Deep Learning Super-Resolution for Temperature Data Downscaling: A Comprehensive Study Using Residual Networks." *Frontiers in Climate* 7: 1572428.
- Kalinowski, T., J. Allaire, and F. Chollet. 2024. "Keras3: R Interface to 'Keras' (Computer software manual)." <https://CRAN.R-project.org/package=keras3>. R package version 1.1.0.
- Kingma, D. P., and J. Ba. 2015. "Adam: A Method for Stochastic Optimization." In *Proceedings of the 3rd International Conference on Learning Representations (ICLR)*.
- Knoben, W. J., J. E. Freer, and R. A. Woods. 2019. "Inherent Benchmark or Not? Comparing Nash–Sutcliffe and Kling–Gupta Efficiency Scores." *Hydrology and Earth System Sciences* 23, no. 10: 4323–4331.
- LeCun, Y., Y. Bengio, and G. Hinton. 2015. "Deep Learning." *Nature* 521, no. 7553: 436–444. <https://doi.org/10.1038/nature14539>.
- Ledig, C., L. Theis, F. Huszár, et al. 2017. "Photo-Realistic Single Image Super-Resolution Using a Generative Adversarial Network." In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 4681–4690. IEEE.
- Li, S., C. Xu, M. Su, et al. 2024. "Downscaling of Environmental Indicators: A Review." *Science of the Total Environment* 916: 170251.
- Li, S. Z. 2009. *Markov Random Field Modeling in Image Analysis*. 3rd ed. Springer. <https://doi.org/10.1007/978-1-84882-425-6>.
- Lim, B., S. Son, H. Kim, S. Nah, and K. Mu Lee. 2017. "Enhanced Deep Residual Networks for Single Image Super-Resolution." In *Proceedings of the Ieee Conference on Computer Vision and Pattern Recognition Workshops*, 136–144. IEEE.
- Maurer, E. P., H. G. Hidalgo, T. Das, M. D. Dettinger, and D. R. Cayan. 2010. "The Utility of Daily Large-Scale Climate Data in the Assessment of Climate Change Impacts on Daily Streamflow in California." *Hydrology and Earth System Sciences* 14, no. 6: 1125–1138.
- Menapace, A., P. Dhawan, D. Dalla Torre, et al. 2025. "Review of Bias Correction Methods for Climate Model Outputs in Hydrology." *Journal of Hydrology* 660: 133213.
- Nag, P., Y. Sun, and B. J. Reich. 2023. "Spatio-Temporal Deepkriging for Interpolation and Probabilistic Forecasting." *Spatial Statistics* 57: 100773.
- Nuvolone, D., D. Petri, and F. Voller. 2018. "The Effects of Ozone on Human Health." *Environmental Science and Pollution Research* 25: 8074–8088.
- Pivato, A., L. Pegoraro, M. Masiol, et al. 2023. "Long Time Series Analysis of Air Quality Data in the Veneto Region (Northern Italy) to Support Environmental Policies." *Atmospheric Environment* 298: 119610.
- R Core Team. 2023. "R: A Language and Environment for Statistical Computing (Computer Software Manual). Vienna, Austria." <https://www.R-project.org/>.
- Raitoharju, J. 2022. "Convolutional Neural Networks." In *Deep Learning for Robot Perception and Cognition*, 35–69. Elsevier.
- Rastogi, D., H. Niu, L. Passarella, et al. 2025. "Complementing Dynamical Downscaling With Super-Resolution Convolutional Neural Networks." *Geophysical Research Letters* 52: e2024GL111828.
- Reichstein, M., G. Camps-Valls, B. Stevens, et al. 2019. "Deep Learning and Process Understanding for Data-Driven Earth System Science." *Nature* 566, no. 7743: 195–204. <https://doi.org/10.1038/s41586-019-0912-1>.
- Ronneberger, O., P. Fischer, and T. Brox. 2015. "U-Net: Convolutional Networks for Biomedical Image Segmentation." In *International Conference on Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015*, edited by N. Navab, J. Hornegger, W. M. Wells III, and A. F. Frangi, 234–241. Lecture Notes in Computer Science, Vol. 9351. Springer.
- Rue, H., and L. Held. 2005. *Gaussian Markov Random Fields: Theory and Applications*. Chapman & Hall/CRC.
- Shams, S. R., Y. Choi, D. Singh, M. Ghahremanloo, M. Momeni, and J. Park. 2024. "Innovative Approaches for Accurate Ozone Prediction and Health Risk Analysis in South Korea: The Combined Effectiveness of Deep Learning and AirQ+." *Science of the Total Environment* 946: 174158.
- Shi, W., J. Caballero, F. Huszár, et al. 2016. "Real-Time Single Image and Video Super-Resolution Using an Efficient Sub-Pixel Convolutional Neural Network." In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 1874–1883. IEEE.
- Sipilä, M., C. Cappello, S. De Iaco, K. Nordhausen, and S. Taskinen. 2025a. "Modelling Multivariate Spatio-Temporal Data With Identifiable Variational Autoencoders." *Neural Networks* 181: 106774.
- Sipilä, M., K. Nordhausen, and S. Taskinen. 2025b. "Identifiable Autoregressive Variational Autoencoders for Nonlinear and Nonstationary Spatio-Temporal Blind Source Separation." In *Machine Learning and Knowledge Discovery in Databases. Research Track: European Conference, ECML PKDD 2025, Porto, Portugal, September 15–19, 2025, Proceedings, Part VII*, 362–380. Springer-Verlag.
- Sipilä, M., C. Cappello, S. De Iaco, K. Nordhausen, and S. Taskinen. 2026a. "Spatialdownscaling: Methods for Spatial Downscaling Using Deep Learning (Computer Software Manual)." <https://CRAN.R-project.org/package=SpatialDownscaling>. R package version 0.1.2.
- Sipilä, M., S. De Iaco, C. Cappello, K. Nordhausen, M. Palma, and S. Taskinen. 2026b. "Enhancing Identifiable Variational Autoencoder in the Presence of Missing Values and Auxiliary Covariates for Spatio-Temporal Ozone Modeling." *Environmetrics*: e70121. <https://doi.org/10.1002/env.70121>.

- Svozil, D., V. Kvasnicka, and J. Pospichal. 1997. "Introduction to Multi-Layer Feed-Forward Neural Networks." *Chemometrics and Intelligent Laboratory Systems* 39, no. 1: 43–62.
- Vandal, T., E. Kodra, S. Ganguly, A. Michaelis, R. Nemani, and A. R. Ganguly. 2017. "DeepSd: Generating High Resolution Climate Change Projections Through Single Image Super-Resolution." In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 1663–1672. ACM. <https://doi.org/10.1145/3097983.3098004>.
- Vaswani, A., N. Shazeer, N. Parmar, et al. 2017. "Attention Is All You Need." In *Advances in Neural Information Processing Systems 30 (NeurIPS 2017)*, 5998–6008. Curran Associates, Inc.
- Wang, F., D. Tian, L. Lowe, L. Kalin, and J. Lehrter. 2021. "Deep Learning for Daily Precipitation and Temperature Downscaling." *Water Resources Research* 57, no. 4: e2020WR029308.
- Wang, X., Y. Xue, C. Jin, Y. Sun, and N. Li. 2022. "Spatial Downscaling of Surface Ozone Concentration Calculation From Remotely Sensed Data Based on Mutual Information." *Frontiers in Environmental Science* 10: 925979.
- Wood, A. W., L. R. Leung, V. Sridhar, and D. P. Lettenmaier. 2004. "Hydrologic Implications of Dynamical and Statistical Approaches to Downscaling Climate Model Outputs." *Climatic Change* 62, no. 1: 189–216.
- Zeng, Q., Y. Wang, J. Tao, et al. 2023. "Estimation of Ground-Level O₃ Concentration in the Yangtze River Delta Region Based on a High-Performance Spatiotemporal Model MixNet." *Science of the Total Environment* 896: 165061.
- Zhang, J., Y. Wei, and Z. Fang. 2019. "Ozone Pollution: A Major Health Hazard Worldwide." *Frontiers in Immunology* 10: 2518.
- Zhu, S., J. Xu, C. Yu, et al. 2022. "Learning Surface Ozone From Satellite Columns (LESO): A Regional Daily Estimation Framework for Surface Ozone Monitoring in China." *IEEE Transactions on Geoscience and Remote Sensing* 60: 1–11.