Environmental Modelling and Software 191 (2025) 106509

Contents lists available at ScienceDirect

Environmental Modelling and Software

journal homepage: www.elsevier.com/locate/envsoft

# Improving localized weather predictions for precision agriculture: A Time-Series Mixer approach for hazardous event detection

Marco Zanchi [a,b,c,i,*], Stefano Zapperi [b,d,e], Stefano Bocchi [a], Oxana Drofa [f], Silvio Davolio [f,g], Caterina A.M. La Porta [a,b,h,i]

[a] Department of Environmental Science and Policy, University of Milan, via Celoria 10, 20133 Milano, Italy
[b] Center for Complexity and Biosystems, University of Milan, Via Celoria 16, 20133 Milano, Italy
[c] DNDG srl, Via Po, 57, 10124 Torino, Italy
[d] Department of Physics, University of Milan, Via Celoria 16, 20133 Milano, Italy
[e] CNR - Consiglio Nazionale delle Ricerche, Istituto di Chimica della Materia Condensata e di Tecnologie per l'Energia, Via R. Cozzi 53, 20125 Milano, Italy
[f] Institute of Atmospheric Sciences and Climate, National Research Council of Italy (CNR-ISAC), Via Gobetti 101, Bologna, Italy
[g] Dipartimento di Scienze della Terra "A. Desio", Università di Milano, Via Botticelli 23, Milan, Italy
[h] Fondazione IRCCS Ca' Granda, Ospedale Maggiore Policlinico di Milano, via Francesco Sforza 28, 20122 Milan, Italy
[i] Innovation For Well-Being And Environment (CRC-I-WE), University of Milan, Via Celoria 10, 20133 Milano, Italy

ARTICLE INFO

Dataset link: https://tds.bo.isac.cnr.it/

*Keywords:*
Environmental modeling
Local weather predictions
Artificial intelligence
Microclimate
Precision agriculture

ABSTRACT

Natural environmental systems and human activities are deeply interconnected, especially in agriculture. Despite advancements in agricultural techniques, weather remains a critical factor influencing crop yields and livestock health. Precision agriculture relies on weather predictions to mitigate environmental risks caused by weather. However, numerical weather predictions are generated by global or regional numerical models, lacking the resolution to capture site-specific conditions. Artificial intelligence can address this gap by integrating numerical weather predictions data with local station observations. This study employs the Time-Series Mixer (TSMixer) neural network to forecast temperature, wind speed, relative humidity, and precipitation over a 45-hour horizon. Trained with predictions from the MOLOCH model and data from ARPA stations near six agricultural sites in Northern Italy, TSMixer achieves greater accuracy than the MOLOCH model. Additionally, TSMixer excels in detecting hazardous events for precision agriculture, including frost damage, heat stress, and germination block, highlighting its value for environmental risk management.

## 1. Introduction

Natural environmental systems and human activities are deeply interconnected. These intricate relationships are particularly evident in the domain of agriculture, where climate and weather influence the main management practices developed by humans to optimize production. In particular, mitigate the impacts of hazardous weather events is of pivotal importance for agriculture.

This mitigation can be achieved with weather forecasts. These are produced using numerical weather prediction (NWP) models that integrate data and physical laws to simulate various scenarios of the chaotic evolution of weather (Gneiting and Raftery, 2005). Farmers rely on these forecasts to anticipate conditions that could either benefit or challenge their activities (Guisiano et al., 2020). Weather forecasts can be categorized as either long-term (over months) or short-term (over days) (Kumar et al., 2017). Long-term forecasts provide synoptic information on possible anomalies for upcoming months or seasons, aiding in long-term farming strategies planning and enhancing resilience to climate variability (Hansen, 2005; Calanca et al., 2011). Examples of such weather-sensitive decisions include crop sequence selection, adjustments to the cropping calendar, and crop yield forecasting (Sivakumar, 2006). Short and medium term weather forecasts offer detailed predictions of variables such as temperature, wind speed, humidity, and precipitation over the next few hours or days. These forecasts support the optimization of risk management strategies, such as frost damage prevention, the timing of pesticide or fungicide application, and resource management for irrigation and fertilizer application (Easton et al., 2017).

Precision agriculture leverages data to optimize and enhance farming practices (Tantalaki et al., 2019). In this context, short-term

* Corresponding author at: Department of Environmental Science and Policy, University of Milan, via Celoria 10, 20133 Milano, Italy.
*E-mail address:* marco.zanchi@unimi.it (M. Zanchi).

https://doi.org/10.1016/j.envsoft.2025.106509
Received 10 January 2025; Received in revised form 19 March 2025; Accepted 27 April 2025
Available online 15 May 2025

weather forecasts are particularly synergistic with precision agriculture (Bendre et al., 2015; Pujahari et al., 2022), enabling rapid decision-making for optimizing and protecting agricultural production. However, weather forecasts are typically generated at a scale of about 10 km for global forecasts and around 1 km for regional forecasts, due to computational constraints (Yang et al., 2022). This scale is misaligned with the goals of precision agriculture, which require precise and accurate environmental data at the field level (Hewage et al., 2020). Moreover, local weather conditions are often shaped by terrain morphology and vegetation, factors that are not fully captured by NWP models at these resolutions.

An alternative solution lies in using data-driven methods, such as machine learning algorithms, to predict weather evolution at agricultural sites using local weather data. By training algorithms on data collected from weather stations that capture the conditions experienced at these sites, machine learning can provide valuable future insights. In recent years, deep learning algorithms have made significant strides in numerical weather forecasting at global scale (Olivetti and Messori, 2024a), with global data-driven models like Pangu-Weather (Bi et al., 2023), GraphCast (Lam et al., 2022) and GenCast (Price et al., 2025) achieving performance levels comparable to traditional NWP (Olivetti and Messori, 2024b). The application of these deep learning models is promising; however, their current use is limited to the synoptic scale, which may restrict their effectiveness in the context of precision agriculture. Despite that, different machine learning and deep learning algorithms have been utilized to predict weather patterns using local data and to identify extreme events. For a comprehensive review, see Salcedo-Sanz et al. (2024). Beyond general applications, these algorithms have also been used specifically to forecast weather at agricultural sites. For example, Smith et al. (2009) was among the first to apply AI for temperature prediction. Hewage et al. (2020) employed temporal convolutional neural networks to provide deterministic predictions for weather variables critical to agriculture, with a forecast horizon ranging from single-hour to 24 h intervals. Other examples of deep learning applications for forecasting the evolution of weather variables, which are of fundamental importance for agricultural management, include (Zhou et al., 2023; Marusov et al., 2024; Zanchi et al., 2024; Ma et al., 2024; Zanchi et al., 2025; Tsela et al., 2025; Baumberger et al., 2025). In addition to forecasting, it is essential to identify hazardous agricultural events using AI models (Ghaffarian et al., 2022). For instance, Cadenas et al. (2020) applied a deep learning algorithm to predict frost damages in Spain. Despite the successes of AI, weather is a chaotic system, meaning predictions based solely on past local data may lack reliability and precision beyond a few hours. In addition, local data do not capture larger-scale weather patterns that can influence future climate conditions.

Therefore, there is the need for data-driven tools to integrate local weather data at the field level with large-scale forecasts from NWP. By combining insights from physical laws and broader-scale weather scenarios with real conditions experienced at agricultural sites, data-driven method may offer more accurate and realistic weather forecasts for precision agriculture. While there is extensive literature on using machine learning methods to predict weather data based on past data collected at agricultural sites, the integration of numerical weather prediction (NWP) with local past weather data remains less explored. Shin et al. (2022) uses random forests to provide deterministic wind speed predictions for the next 48 h, combining local weather data with NWP from the Local Data Assimilation and Prediction System over the Korean Peninsula. Rozante et al. (2023) proposes a framework for forecasting frost events in Brazil using an artificial neural network that takes weather forecasts from the Eta model as inputs and assimilated weather data at 5 km resolution. The Neural Forecast model in combination with seamless predictions from the Open Meteo service have been used to predict deterministic temperature evolution over the next 72 h (Perez Garcia et al., 2024b) and Equivalent Temperature Index deterministic evolution over the next 24 h (Perez Garcia et al., 2024a) in a pilot dairy farm located in Emilia-Romagna, Italy. In all the aforementioned studies, the integration of NWP with local weather data through machine learning techniques enhances the accuracy of predictions compared to the NWP alone.

From the current state of the art, three main gaps emerge in the literature: (1) existing studies on machine learning methods that integrate NWP with local data typically focus on a single weather variable and provide only deterministic predictions; (2) existing models that integrate NWP with local data tend to either focus on the prediction of a hazardous event without forecasting the general weather evolution or on forecasting weather evolution without testing the correct prediction of hazardous events; (3) the hazardous event detection using machine learning algorithm trained on NWP and local weather data focuses solely on frost events.

To address these gaps, this research contributes to the current state of the art by pursuing the following objectives:

- Developing a multivariate probabilistic weather forecasting model based on neural networks, leveraging both historical data and numerical weather forecasts.
- Establishing a unified forecasting framework for both general weather conditions and hazardous events.
- Improving hazardous event detection by leveraging probabilistic predictions to identify three major agricultural risks: frost damage to fruit, heat stress in dairy cows, and germination block in rice cultivation.

In particular, we use a Time-Series Mixer (TSMixer) neural network to forecast the probabilistic evolution of temperature, wind speed, relative humidity, and precipitation over the next 45 h for six locations near to agricultural sites in Northern Italy. The neural network takes as input both past weather data from Regional Environment Protection Agency (ARPA) weather stations, located near these agricultural sites, and 45 h weather forecats obtained from the high-resolution MOLOCH NWP model. The multivariate probabilistic predictions have been exploited to detect three different hazardous events for agriculture: frost damage to fruit, heat stress in dairy cows, and germination block in rice cultivation. The TSMixer performances have been compared against the MOLOCH model and a TSMixer trained solely on past weather data from local stations.

By delivering accurate probabilistic forecasts and enhancing hazardous event detection, this research can significantly improve decision-making in precision agriculture, promoting risk management, productivity, and resource optimization. Achieving these innovation is of pivotal importance because agriculture is highly sensitive to weather variability and extreme events, which can significantly impact productivity and profitability. With climate change increasing the frequency and severity of hazardous weather events, developing robust and accurate forecasting methods is crucial for enhancing agricultural resilience and safeguarding food security. The potential impacts and applications of this research include better-informed agricultural practices, reduced economic losses from hazardous events and enhanced resilience of agricultural systems to climate variability.

## 2. Materials and methods

### 2.1. Study design

The workflow design of the present research is illustrated in Fig. 1 and described in the following subsections.

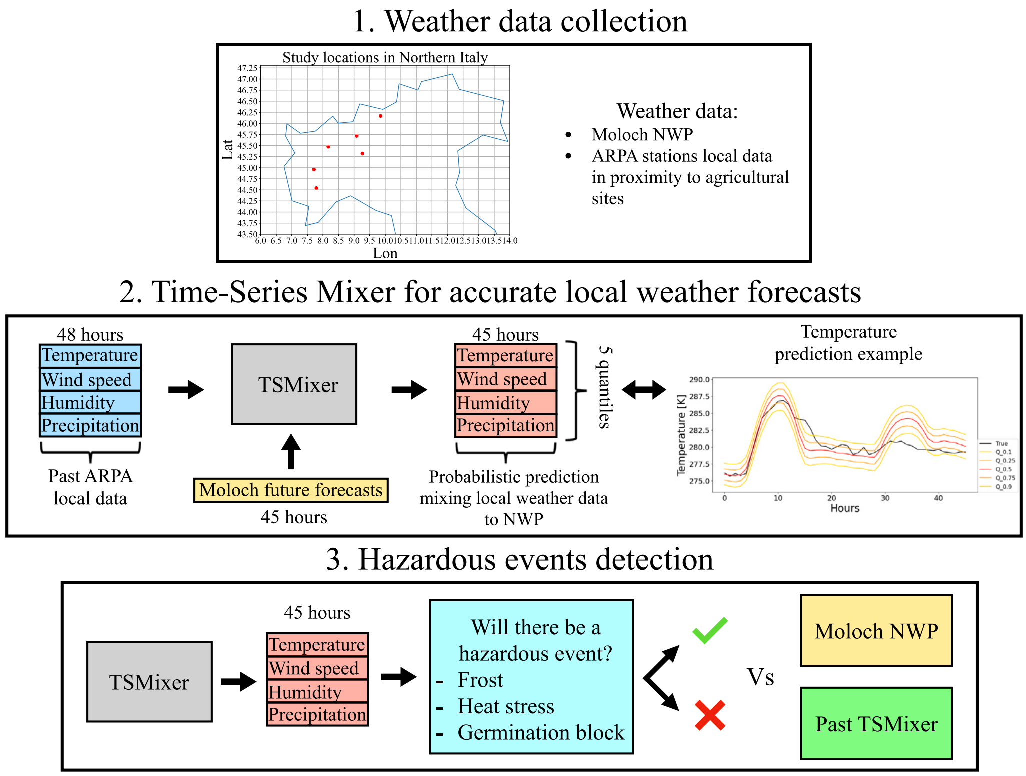


**Fig. 1.** Study design scheme. The first block displays the locations of the six ARPA weather stations in Northern Italy. The second block outlines the design of the Time-Series Mixer neural network used for probabilistic weather predictions. The third block illustrates the protocol for detecting hazardous events.

#### *2.1.1. Weather data collection*

We collected temperature, wind speed, relative humidity, and precipitation data from six ARPA weather stations located near agricultural sites in Northern Italy. Two stations are situated close to apple orchards, two near dairy farms, and two near rice fields. These ARPA weather stations serve as proxies for the actual weather conditions experienced by the agricultural sites. Additionally, we gathered weather forecasts from a NWP model, the MOLOCH model, for each grid cell containing these weather stations (and equivalently, the agricultural sites). The locations of the stations are shown in the first block of Fig. 1.

#### *2.1.2. Time-Series Mixer for precise local weather forecasts*

We train a TSMixer to generate accurate weather forecasts for the agricultural sites. The TSMixer uses as inputs the past 48 h of temperature, wind speed, relative humidity, and precipitation data collected by ARPA stations, combined with weather forecasts from the MOLOCH model for the next 45 h. The output consists of probabilistic predictions for temperature, wind speed, relative humidity, and precipitation over the following 45 h. An example of the temperature quantile predictions is shown in the second block of Fig. 1. By combining past local weather data with weather from the MOLOCH model, the TSMixer effectively calibrates general weather forecasts to reflect the microclimate characteristics of the agricultural sites. To ensure the TSMixer adds value beyond existing numerical forecasts, its performance is compared with that of the MOLOCH model and to another TSMixer trained only with past weather data collected by the ARPA station (this model is named Past TSMixer).

#### *2.1.3. Hazardous event detection*

We test the TSMixer forecasting capability in detecting hazardous events relevant to agriculture. Specifically, we identify three different events, each associated with a particular agricultural site: (1) frost damage to apple orchards, (2) heat stress for dairy cows, and (3) germination block for rice cultivation. We then compare the TSMixer probabilistic predictions with those of the MOLOCH model and the Past TSMixer to assess whether our approach adds value for prevention strategies.

### *2.2. Study area: Lombardy and Piedmont*

The climates of Lombardy and Piedmont, where the study areas are located, are classified as humid subtropical ("Cfa" according to the Köppen and Geiger classification Beck et al., 2018), characterized by warm, wet summers and cold winters. However, these regions display significant climate variations from the Köppen model due to local factors such as elevation, proximity to large water bodies, and metropolitan influences. In Lombardy, the annual average temperature is approximately 13.1 °C, with an average annual rainfall of around 853 mm (Zanchi et al., 2022). In Piedmont, the annual average temperature is about 11.4 °C, with an average annual rainfall of approximately 612 mm.

### *2.3. Background on hazardous events*

Frost occurrence, defined as temperatures dropping to 0 °C or lower measured at a height between 1.25 and 2.0 m in a meteorological shelter (Hogg, 1971), can significantly impact agricultural crops, especially in sensitive stages of development, such as apple buds and blossoms. The intensity and duration of frost are influenced by local factors like soil type and terrain orientation, and damage to crops can vary accordingly. While techniques such as frost protection sprinkling are effective to minimize damage caused by frost (Unterberger et al., 2018), it remains crucial to anticipate frost events to mitigate potential damage.

When dairy cows face high temperatures and humidity conditions, they experience heat stress. The high temperature cause cows to rely heavily on evaporative cooling to regulate their body temperature.

**Table 1**
ARPA station locations. The table describes the ARPA station locations.

| Station location | Coordinates Wgs84 (Lat,Lon) | Height (m) | Study case |
|---|---|---|---|
| Massazza, Biella | (45.474699, 8.17086) | 226 | Rice |
| Landriano, Pavia | (45.320567, 9.26714) | 88 | Rice |
| Fojanini, Sondrio | (46.167553, 9.85094) | 307 | Fruits |
| Minoprio, Como | (45.719011, 9.085488) | 310 | Fruits |
| Fossano, Cuneo | (44.540635, 7.788169) | 403 | Livestock farm |
| Bauducchi, Torino | (44.961901, 7.709583) | 226 | Livestock farm |

However, high humidity interferes with the process, making it difficult for the cows to effectively release body heat. As a result, their core temperature rises. The primary and most recognized outcome of heat stress in dairy cows is a reduction in both milk production and feed intake, leading to significant financial losses for dairy farmers (West, 2003). To counteract the impact of high temperatures, efficient air circulation in the barn, along with the use of sprinklers or misters in resting areas, can enhance heat dissipation and help alleviate the effects of heat stress.

Low temperatures significantly limit rice growth and yield, with germination being one of the most sensitive developmental stages to cold stress (Teixeira et al., 2021). Under low-temperature stress, rice seeds experience imbalances in intracellular oxygen metabolism, leading to damage to the cell membrane system and potentially resulting in a germination block (Yang et al., 2019). In Northern Italy, 12 °C is considered the threshold below which germination block occurs.

### 2.4. Weather data from ARPA stations

The local weather data are provided by the ARPA. Specifically, data have been collected from six weather stations located in Lombardy and Piedmont regions, Italy. The dataset includes hourly measurements of temperature (°C), wind speed (m/s), relative humidity (%), and precipitation (mm) from January 1, 2019, to June 30, 2024. The six locations have been selected based on the following criteria: (1) they have continuous weather data over the study period with minimal missing values; (2) they are situated near agricultural fields or livestock farms, providing practical case studies for agricultural applications. The handling of missing values will be addressed in the following sections. The 6 locations are resumed by Table 1.

The station in Massazza, in the province of Biella, is surrounded by rice fields cultivating varieties such as Carnaroli and Canovese rice. The station in Landriano, in the province of Pavia, is also notable for its proximity to areas known for producing Carnaroli rice. The Fojanini station, Sondrio, is located in the Valtellina region, renowned for apple production, and is near Fondazione Fojanini, which conducts biological research on apple orchards. The station in Minoprio, in the province of Como, is close to an experimental site involved in a regional project on apple orchard development, "Valorizzazione della frutticoltura lombarda". The station near Fossano, in the province of Cuneo, is situated 14 km from a commercial farm. The station near Bauducchi, in the province of Torino, is located 12 km from a commercial farm. Given the plain terrain characterizing these two locations, we assume that microclimatic variables do not vary significantly across this distance. According to Kottek et al. (2006), the ARPA weather stations located in Massazza, Landriano, Minoprio, Fossano, and Bauducchi are classified as Cfa (humid subtropical climate) under the Köppen and Geiger system. In contrast, the Fojanini station is classified as Dfc (subarctic climate).

### 2.5. Weather data from MOLOCH model

The MOLOCH model (Malguzzi et al., 2006) generates forecasts with fine spatial resolution. It solves the nonhydrostatic, fully compressible atmospheric equations on a 1.25 km grid with 60 vertical levels, utilizing a hybrid height-based coordinate system. The model produces 2-day forecasts over a 1154 × 1154 grid covering Italy and adjacent seas, with runs starting daily at 03:00 UTC. See Davolio et al. (2020) for more details on model description and operational implementation. The MOLOCH variables used in this study are listed in Table 2.

**Table 2**
MOLOCH data. The table describes the MOLOCH data.

| Variable | Unit of measure | Symbol |
|---|---|---|
| Mean sea level pressure | Pa | Pmsl |
| Hourly total precipitation | mm | Pt |
| Total cloud cover | % | Cc |
| Downward short-wave radiation flux | W m$^{-2}$ | Rsw |
| Hourly total snowfall | m | Sn |
| Water equivalent of accumulated snow depth | kg m$^{-2}$ | Wsd |
| Lapse rate | K m$^{-1}$ | Lr |
| Soil temperature | K | Ts |
| 2-m temperature | K | T |
| 10-m u wind component | m/s | Wu |
| 10-m v wind component | m/s | Wv |
| 2-m relative humidity | % | Hr |
| 2-m dew point temperature | K | Tdew |

Since wind speed components are computed at a height of 10 m, we apply a logarithmic downscaling to adjust the MOLOCH wind speed to 2 m, ensuring comparability with ARPA station measurements. First, we calculate the wind speed magnitude as follows:

$$W_s^{10\ \mathrm{m}} = \sqrt{W_u^2 + W_v^2} \tag{1}$$

Then, we derive the wind speed at 2 m using a logarithmic downscaling approach (Monteith and Unsworth, 2013):

$$W_s^{2\ \mathrm{m}} = W_s^{10\ \mathrm{m}} \frac{\log(\frac{2}{z_0})}{\log(\frac{10}{z_0})} \tag{2}$$

where $z_0$ is the roughness height, estimated at 0.03 for open fields.

### 2.6. Time-Series Mixer: architecture and application to weather forecasts

#### 2.6.1. Time-Series Mixer architecture

The Time-Series Mixer (TSMixer) (Chen et al., 2023) is a neural network architecture tailored for forecasting the evolution of multivariate time series. It employs mixing operations across both the temporal and feature dimensions to effectively extract relevant information. The model consists of multi-layer perceptrons (MLPs) organized along the temporal and feature axes. This structure facilitates time-mixing and feature-mixing processes, enabling TSMixer to capture both temporal dependencies and relationships between features. Its performance demonstrated strong competitiveness when compared to state-of-the-art models on the M-competition (Makridakis et al., 2022) datasets, including those based on transformer architectures (Vaswani et al., 2017). TSMixer takes as input a set of observations of length $I$ composed by $C_i$ time series and predicts the future evolution of a set of $C_o$ features for a temporal horizon of length $O$. It also possible to add a set of future covariates to reinforce the prediction. The future covariates are represented by a multivariate time series of $C_f$ features of length $O$.

Fig. 2 shows the TSMixer architecture. It primarily consists of the Mixer layer (second block in Fig. 2), which can be stacked multiple times, and the temporal projection (third block in Fig. 2). The Mixer layer plays a key role in capturing temporal dynamics and leveraging cross-variate information across feature dimensions. It is composed of two main components: the time-mixing multilayer perceptron (MLP) and the feature-mixing MLP. The time-mixing MLPs are designed to extract temporal patterns within the time series. They consist of a single fully connected layer followed by an activation function and dropout (Zeng et al., 2023). The input is transposed to enable the fully connected layer to operate along the temporal dimension. The

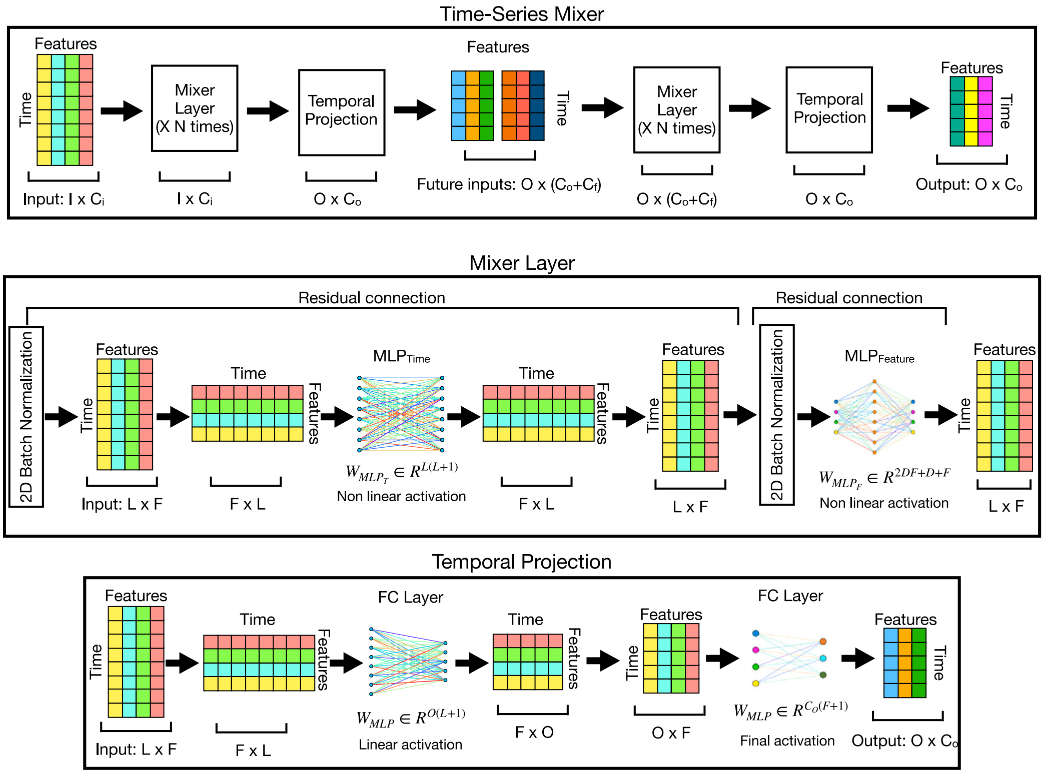


**Fig. 2.** Time-Series Mixer architecture. The first block: general structure of a Time-Series Mixer which receives as input an historical instance of $I$ time-steps for $C_i$ features and predict the evolution of $C_o$ features for $O$ time-steps, using $C_f$ future covariates. The second block: Mixer layer which is mainly composed by the time-mixing multilayer perceptron and the feature-mixing multilayer perceptron. The third block: temporal projection operation.
*Source:* Modified from Chen et al. (2023).

fully connected layer weights are shared across the features dimension. The feature-mixing MLPs are applied across each time step to capture covariate information between features. To enhance the model ability to learn complex inter-feature relationships, a two-layer MLP configuration is adopted. This module receives the transposed output from the time-mixing MLP and applies two fully connected layers along the feature dimension, with shared weights across time steps. The temporal projection consists of a fully connected layer that operates on the time axis, enabling the model to learn temporal patterns and map the input time series from its original length $I$ to the target forecast length $O$. In addition to the original fully connected layer, an additional layer has been incorporated to adjust the number of output features predicted by TSMixer.

Initially, the TSMixer receives as input the historical multivariate time series of dimensions $I \times C_i$, which are processed through Mixer layers and a Temporal Projection layer, resulting in an output of dimensions $O \times C_o$. Next, the $C_f$ future covariates of length $O$ are appended to this output, creating a tensor of dimensions $O \times (C_o + C_f)$. This tensor is then processed by another set of Mixer layers and a Temporal Projection layer, producing a final output of dimensions $O \times C_o$. This scheme is depicted in the first block of Fig. 2.

### *2.6.2. Time-Series Mixer application for probabilistic local weather forecasts*

In the present study, the TSMixer produces probabilistic local weather predictions as follows. Each day at 3:00 UTC, the model uses as historical inputs the local weather data from the past 48 h, sourced from an ARPA station, including temperature, wind speed, relative humidity, and precipitation data. These data are processed by 6 Mixer layers with 64 neurons each, using a ReLU activation function and a dropout rate of 0.33, resulting in an output of 4 features, each of length 45. Ideally, these values represent the 45 h future evolution predictions for temperature, wind speed, relative humidity, and precipitation, based solely on local past data. These four arrays are then combined with the 13 arrays, each of 45 h, provided by the MOLOCH model, as described in Table 2. The multivariate time series, now comprising 17 features of length 45 h, is processed by 8 additional Mixer layers with 128 neurons, a GeLU activation function, and a dropout rate of 0.13. The TSMixer outputs the probabilistic future evolution of local temperature, wind speed, relative humidity, and precipitation over the next 45 h. By probabilistic evolution, we refer to predictions of the following quantiles for each hour: 0.1, 0.25, 0.5, 0.75, and 0.9. Predictions examples are shown in Fig. 3. The TSMixer hyperparameters have been optimized using the Hyperopt library (Bergstra et al., 2013), with a Tree-structured Parzen Estimator (TPE) for the search space.

To generate probabilistic predictions, the TSMixer is trained using a custom quantile loss function:

$$L_Q = L_Q^s + L_Q^p \tag{3}$$

Here, $L_Q^s$ represents a standard quantile loss function (Wen et al., 2017), which is used to evaluate prediction errors for temperature, wind speed, and relative humidity. In contrast, $L_Q^p$ is a zero-inflated quantile loss function employed to assess precipitation errors. This choice reflects the statistical properties of hourly precipitation, which is highly skewed towards zero. In order to counteract the negative effects of predicting a variable highly skewed to 0, we have predicted the logarithm of precipitation by computing:

$$log_{prec} = \log(1 + prec) \tag{4}$$

Each predictions has been then transformed back to the precipitation in mm units (therefore the results on precipitation are always represented in mm units). The choice of using the log precipitation is further discussed in the Supplementary Materials.

The loss functions are defined as follows.

$$L_Q^s = \frac{1}{3*5*45}\sum_{f=1}^{3}\sum_{q=1}^{5}\sum_{h=1}^{45} L_q(y_f(h), \hat{y}_f^q(h)) \tag{5}$$

where

$$L_q(y, \hat{y}) = q(y - \hat{y}^q)_+ + (1-q)(\hat{y}^q - y)_+ \tag{6}$$

where $y_f$ and $\hat{y}_f$ are the true and predicted values for the output feature $f$, which can be temperature, wind speed, relative humidity, $q$ is the quantile and $h$ is the hour in the forecasting horizon. The $(x)_+$ operation for a generic value $x$ is defined as $\max(0, x)$.

$$L_Q^P = \frac{1}{5*45}\sum_{q=1}^{5}\sum_{h=1}^{45}(L_q(y_p(h), \hat{y}_p^q(h)))P + \alpha(1-P)(L_q(y_p(h), \hat{y}_p^q(h))) \tag{7}$$

where

$$P = \begin{cases} 1 & \text{if } y_p(h) \ \geq 0, \\ 0 & \text{if } y_p(h) = 0 \end{cases} \tag{8}$$

and $\alpha$ is a penalization term equal to 0.225, used to penalize the errors for the more frequent zero precipitation. Since the value of $\alpha$ directly influences the magnitude of the loss function, we manually calibrated it. After optimizing the other hyperparameters, we tested different $\alpha$ values (ranging from 0 to 1) and selected the highest value that did not cause a collapse of the quantile predictions for precipitation. Specifically, we avoided values that led to the median, 0.75, and 0.90 quantile predictions collapsing to zero, a phenomenon driven by the highly skewed nature of hourly precipitation distribution. In addition, since relative humidity is physically constrained between 0% and 100%, and the ReLU activation function does not impose any upper bounds, we manually clip predicted humidity values to a maximum of 100% to prevent unphysical predictions. Further details regarding this choice are provided in the Supplementary Materials.

The TSMixer is trained on weather data from January 1, 2019, to December 31, 2021, validated on data from January 1, 2022, to December 31, 2022, and tested on data from January 1, 2023, to June 30, 2024. The train, validation and test data division follows the guidelines outlined in Schultz et al. (2021). Since the original data measured by the ARPA stations contains some missing values, we removed instances with missing data, resulting in a total of 718 discarded instances. All the data have been normalized according to a min–max scaler. The scaler parameters have been derived from the training set to avoid data leakage. The training has been performed using single precision floating point (FP32). The gradient descent process in this study employs the AdamW optimizer (Loshchilov and Hutter, 2017) with an initial learning rate set to 0.001 and a weight decay set to 0.004 (the value of the weight decay has been chosen with the hyperparameter optimization). The TSMixer has been trained for 2000 epochs on the training set, with an early stopping patience of 500 epochs based on validation set performance. The final model weights are those that minimize the loss function on the validation set. After training, the TSMixer is tested on the test set, and all reported results are derived solely from this model application.

In addition to the TSMixer model described above, which is trained on past covariates from ARPA weather stations and future covariates from the MOLOCH model, another TSMixer model is trained using the same procedure but only on past covariates. The hyperparameters of the model have been optimized using the TPE algorithm as was done for the TSMixer model. Their values are presented in the Supplementary Materials. This model serves to demonstrate the importance of including future covariates in the predictions. This model will be addressed as Past TSMixer.

*2.6.3. Time-Series Mixer interpretation*

To address the black-box nature of neural networks, we implement the method proposed in Zanchi et al. (2023) to interpret the influence of input variables on the TSMixer model output. This technique involves systematically replacing each input variable with random values, uniformly extracted within its minimum and maximum range, while keeping the other inputs unchanged. The TSMixer is then tested on this randomized dataset, and the mean absolute error respect to the median prediction (quantile 0.5) is calculated. These results are compared to the MAE values obtained from the original, non-randomized test set. To account for statistical variability, the randomization process is repeated 300 times. The mean absolute percentage error of the differences in MAE between the randomized and original tests is computed to derive the final values and it is described by the following equation:

$$R_{mape}^{i,f,h} = \sum_{n=1}^{300} \frac{MAE_r^{n,i,f,h} - MAE_s^{f,h}}{MAE_s^{f,h}} \tag{9}$$

where $f$ is the output feature, $h$ is the prediction hour, $i$ is the perturbed input feature, $MAE_r^{i,f}$ is the mean absolute error for the output feature $f$ at hour $h$ obtained by replacing the input feature $i$ with random samples at the $n{-}esime$ sampling and $MAE_s^f$ is the mean absolute error for the output feature $f$ at hour $h$ on the non perturbed data. A high mean absolute percentage error for a given input variable indicates its significant role in the prediction. This process reveals the importance of each input variable by analyzing how MAE changes when that variable is randomized.

*2.6.4. Hazardous event detection with Time-Series Mixer*

This research aims to evaluate the effectiveness of TSMixer in detecting events that pose risks to precision agriculture. Three specific hazardous events have been identified: frost for apple orchards, heat stress for dairy cow farms, and germination block for rice cultivation. Frost events occur when temperatures drop below 0 °C. Heat stress events are detected using the equivalent temperature index (ETI) (Yan et al., 2021), which is defined as:

$$\begin{aligned} ETI &= 27.88 - 0.456T + 0.010754T^2 - 0.4905RH + 0.00088RH^2 \\ &\quad + 1.15W_s - 0.12644W_s^2 + 0.019876TRH - 0.046313TW_s \end{aligned} \tag{10}$$

where $T$ is the temperature (°C), $RH$ is the relative humidity (%) and $W_s$ is the wind speed (m/s). When the ETI exceeds a value of 30, a heat stress condition is identified (Yan et al., 2021). The germination block events occur when temperatures fall below 12 °C during the months of March, April, or May, which is the rice germination period. For each event, we selected two weather stations located near apple orchards, cattle farms, and rice fields, as summarized in Table 1.

We propose the following strategy to detect the occurrence of hazardous events using TSMixer probabilistic predictions. The approach varies between detecting frost and germination block versus heat stress, so we describe them separately.

For frost and germination block detection, we examine whether the 0.10, 0.25, and 0.5 quantiles from the probabilistic temperature predictions for the next 45 h fall below 0 °C (for frost) or 12 °C (for germination block) at least once. If this condition is met, a value of 1 is assigned to indicate the prediction of the hazardous event; otherwise, it is assigned a value of 0. This classification is then compared to the actual temperature measurements recorded by the ARPA weather stations. We have decided to compare classifications across three quantiles to balance the trade-off between the cost and reliability of implementing countermeasures based on probabilistic predictions.

For heat stress detection, we compute the future ETI scenario over the next 45 h using the 0.5, 0.75, and 0.9 quantile predictions for temperature and humidity, along with the 0.5, 0.25, and 0.1 quantile predictions for wind speed. As indicated by Eq. (10), ETI is maximized when high temperature and humidity coincide with low wind speed. By combining the highest quantile predictions for temperature and humidity with the lowest quantile predictions for wind speed, we aim to characterize the worst-case scenario for heat stress events. We define this combination of quantiles in heat stress detection as $q1,5$ and $q2,4$, where the first index refers to the wind speed quantile and the second to the temperature and relative humidity quantile. A value of 1 is assigned for each hour in which the ETI exceeds 30 (0 otherwise), resulting in 45 classifications across each prediction window.

We apply different classification strategies for frost and germination block detection versus heat stress detection due to the nature of the respective countermeasures. Frost and germination block countermeasures involve high-cost, irreversible actions such as tree shielding, anti-frost spraying, or flooding rice fields. Therefore, the detection of at least one hazardous event will trigger countermeasures which affect all the other time steps after it. In contrast, heat stress countermeasures, like activating ventilation systems or water showers for dairy cows, are more flexible and can be readily switched on and off at hourly frequency. To evaluate the innovation and value offered by the TSMixer, we compare its accuracy in detecting agricultural hazards with that of the MOLOCH model and the Past TSMixer.

### 2.7. Evaluation metrics

#### 2.7.1. Absolute error with respect to the median prediction

The mean absolute error has been used to evaluate the accuracy of the median predictions (the predictions originating from the quantile 0.5):

$$MAEMP_f = \frac{\sum_{i=1}^{N} \left\langle |y_f - \hat{y}_f^{0.5}| \right\rangle_H}{N} \tag{11}$$

where $\hat{y}_f^{0.5}$ is a prediction of the weather feature $f$ for the quantile 0.5, $y_f$ represents the respectively real values and $N$ is the number of samples in the test set. The operation $\langle\rangle_H$ indicates the average over the prediction horizon of 45 h. The distribution of the absolute error has also been computed by collecting all the values in the sum over the test set samples. In addition to the mean absolute error, we have defined the absolute error profile, which shows the average absolute error of the test set over the prediction horizon hour:

$$AEP_f = \frac{\sum_{i=1}^{N} |y_f - \hat{y}_f^{0.5}|}{N} \tag{12}$$

where $\hat{y}_f^{0.5}$ is a prediction of the weather feature $f$ for the quantile 0.5, $y_f$ represents the respectively real values and $N$ is the number of samples in the test set. The $AEP_f$ for a weather feature $f$ is a numerical array of length 45.

#### 2.7.2. Validity

A validity measure has been exploited to evaluate the reliability of the quantile prediction. The validity is defined as follows. For a given quantile band $(a, b)$, a prediction sampled from a quantile distribution should fall within this band in the $(b-a)\%$ of cases. Therefore for each data sample in the test set we have computed how many times the real values fall within the respectively predicted quantile bands. The counting has been performed for each hour of the prediction horizon. The validity measure for a quantile band $(a, b)$ for the weather feature $f$ can be defined as the average of this counts over all the predictions:

$$V_f(a,b) = \frac{\sum_{i=1}^{N} \sum_{h=1}^{45} v(y_f(h), \hat{y}_f^a(h), \hat{y}_f^b(h))}{45N} \tag{13}$$

where $y_f(h)$ represents the real weather values at hour $h$, $\hat{y}_f^b(h)$ is the prediction of quantile $b$ at hour $h$ and $v(x,z,k)$ is defined as:

$$v(x,z,k) = \begin{cases} 1 & z < x < k, \\ 0 & \text{else} \end{cases} \tag{14}$$

If the $Validity(a, b)$ value is near to $(b-a)\%$ the probabilistic predictions works. The validity measure helps assessing whether TSMixer can produce a reliable probabilistic forecast. We have applied the validity measures over the quantile bands (0.1, 0.9) and (0.25, 0.75).

#### 2.7.3. Efficiency

An efficiency metric has been utilized to assess the model confidence in generating probabilistic predictions. This metric evaluates the width of the quantile band. For a given quantile band $(a, b)$, for a weather feature $f$ the efficiency is defined as:

$$E_f(a,b) = \frac{\sum_{i=1}^{N} \sum_{h=1}^{45} |\hat{y}_f^a(h) - \hat{y}_f^b(h)|}{45N} \tag{15}$$

where $\hat{y}_f^a(h)$ is the prediction of quantile $a$ (same for $b$) at hour $h$ and $N$ is the number of samples in the test set. A lower efficiency value (i.e., a smaller average thickness) indicates that the model produces a tighter band, which may suggest a more confident prediction. However, if the bands are too narrow, it may indicate that the model underestimates the uncertainty.

#### 2.7.4. Classification metrics

The TSMixer model used in the present study is primarily designed for probabilistic regression tasks. However, as outlined in previous sections, this approach can also be leveraged to classify if a hazardous event will happen in the predicted future horizon. For instance, given a predicted temperature quantile over the next 45 h, we can check for the occurrence of frost (if the quantile falls below 0 °C) by assigning a value of 1 for a frost event or 0 for no frost event. This method allows for the application of various classification metrics, which are discussed below.

Accuracy is the proportion of correctly classified instances (both true positives and true negatives) out of all the instances. It is calculated as follows:

$$\text{Accuracy} = \frac{TP + TN}{TP + TN + FP + FN} \tag{16}$$

where:

- $TP$ = True Positives (correctly predicted positive instances)
- $TN$ = True Negatives (correctly predicted negative instances)
- $FP$ = False Positives (incorrectly predicted as positive)
- $FN$ = False Negatives (incorrectly predicted as negative)

The true positive rate, also known as sensitivity or recall, measures the proportion of actual positives that are correctly identified. It is given by:

$$\text{TPR} = \frac{TP}{TP + FN} \tag{17}$$

The false positive rate is the proportion of negative instances that are incorrectly classified as positive. It is given by:

$$\text{FPR} = \frac{FP}{FP + TN} \tag{18}$$

The true negative rate, also known as specificity, measures the proportion of actual negatives that are correctly identified:

$$\text{TNR} = \frac{TN}{TN + FP} \tag{19}$$

The Receiver Operating Characteristic (ROC) curve is a graphical representation of a classifier performance across different threshold values. The curve plots the true positive rate (TPR) against the false positive rate (FPR):

- X-axis: False Positive Rate (FPR)
- Y-axis: True Positive Rate (TPR)

The ROC curve allows visualization of the trade-off between sensitivity and specificity.

The area under the ROC curve (AUC) quantifies the overall ability of the model to distinguish between classes. A perfect model has an AUC of 1, whereas a random model has an AUC of 0.5:

$$\text{AUC} = \int_0^1 \text{TPR}(t)\, d\text{FPR}(t) \tag{20}$$

where $t$ represents the threshold used for classification.

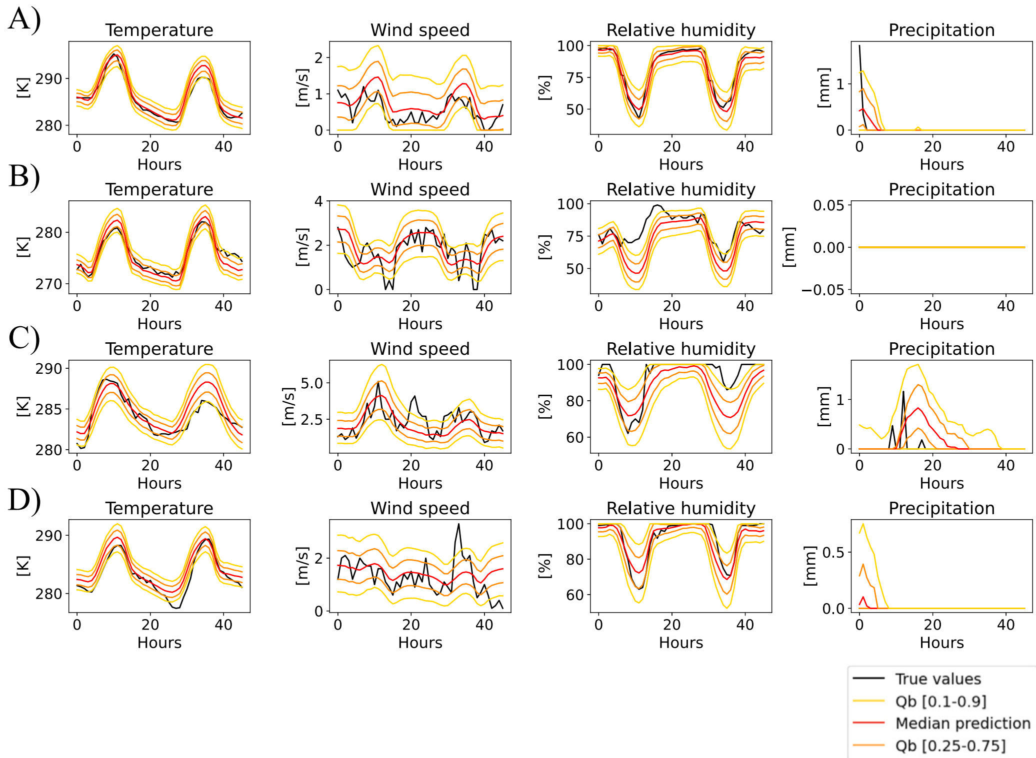


**Fig. 3.** Time-Series Mixer prediction examples. The figure shows four examples of Time-Series Mixer predictions, with each row corresponding to a different 45 h forecast (X-axis across all plots). The black curve represents the actual values measured by the ARPA stations. The red line represents the median of the probabilistic prediction, while the yellow and orange bands show the quantile ranges: [0.1–0.9] and [0.25–0.75], respectively. Temperature is given in Kelvin, wind speed in meters per second, relative humidity in percentage, and precipitation in millimeters.

## 3. Results and discussion

### 3.1. Time-Series Mixer predictions

The following section presents the results of the TSMixer predictions on the test set. We first showcase four prediction examples to illustrate how individual forecasts are generated. Then, we analyze the overall performance of the neural network across all test samples.

Fig. 3 presents four different examples of TSMixer probabilistic predictions (each row corresponds to one example). These examples demonstrate that the median prediction can serve as a deterministic forecast, allowing comparison with actual values (and with MOLOCH NWP). The probabilistic ranges provide additional insights, which are useful for risk assessment. Narrower bands indicate greater model confidence in the median prediction. From these examples, we observe that the temperature and relative humidity have much narrower quantile bands compared to wind speed. For wind speed, the median follows the general trend, while the quantile bands capture the fluctuations around this trend. Precipitation represents a special case. In the examples, the lower quantiles ($[0.1, 0.25]$) are consistently zero, while the median and higher quantiles reach non-zero values. This behavior is explained by the statistical properties of hourly precipitation, which is highly skewed towards zero. As a result, the lower quantiles are compressed towards zero (since precipitation cannot be negative). The custom loss function, defined in Eq. (7), enables non-zero values for the higher quantiles, providing a useful proxy for the observed precipitation. Due to the differences in prediction behavior, we analyze the overall performance of the TSMixer model in predicting temperature, wind speed, and humidity separately from precipitation.

#### 3.1.1. Time-Series Mixer predictions: temperature, wind speed and relative humidity

Since the TSMixer prediction includes both deterministic values (the median prediction) and two probabilistic intervals based on the four other quantiles used in the loss function, we evaluate these two aspects separately. The mean absolute error with respect to the median prediction is calculated on the test set for temperature, wind speed, and relative humidity, using Eq. (11). The resulting values are: 1.44 K for temperature, $0.61\,\frac{\text{m}}{\text{s}}$ for wind speed, and 6.86% for relative humidity. The error values for temperature and relative humidity are comparable to the typical accuracy of weather observation sensors, while the wind speed mean absolute error is comparable with that obtained in Shin et al. (2022). The empirical cumulative distribution function of absolute errors with respect to the TSMixer median (shown in red) is compared to those of the MOLOCH model (light blue) and the Past TSMixer model (green) in Fig. 4A, for each variable. In all cases, the TSMixer model exhibits smaller error distributions compared to the other two models. Additionally, the Past TSMixer model outperforms the MOLOCH model, except for the relative humidity where the MOLOCH model has lighter tails. We then compute the absolute error profile with respect to the median over the 45 h prediction horizon, as defined in Eq. (12). The error profiles are displayed in Fig. 4B. The TSMixer model (red curve) consistently outperforms both the Past TSMixer (green curve) and the MOLOCH model (light blue curve) at each hour for all predicted variables. It is interesting to note that both AI models exhibit a bias around 1 p.m., where higher temperatures may lead to more unstable weather conditions, resulting in larger absolute errors. In contrast, the MOLOCH model shows a minimum error around these hours for temperature. It is also noteworthy that the Past TSMixer shows an increasing trend over the prediction horizon for both temperature and relative humidity. In the first few hours, the error is comparable or even lower to that of the TSMixer, but it gradually increases, eventually surpassing the MOLOCH curve. This can be explained by the chaotic nature of weather system, where the information from past observations loses relevance after a short period. In contrast, the MOLOCH model incorporates the dynamic future evolution of weather, governed by physical laws, and therefore do not exhibit an increasing trend in absolute error. The TSMixer model

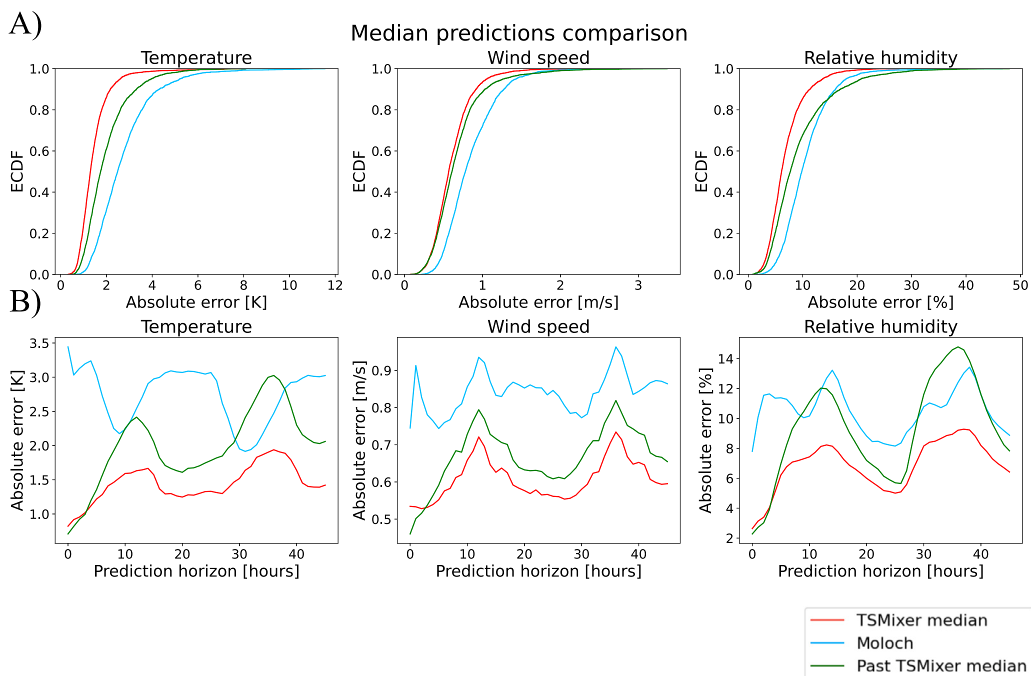


**Fig. 4.** Time-Series Mixer absolute error comparisons: temperature, wind speed and relative humidity. (A) The figure compares the empirical cumulative distribution function (ECDF) for three models: Time-Series Mixer (red), Past Time-Series Mixer (green), and the MOLOCH model (light blue), across the three weather variables. The $X$-axis represents the absolute error, while the $Y$-axis shows the ECDF. (B) The figure shows the average absolute error profile over the 45 h prediction horizon for the three models. The $X$-axis represents the prediction horizon (from 3 a.m. to 12 p.m. the following day), and the $Y$-axis represents the absolute error averaged over the test set.

**Table 3**
Time-Series Mixer validity. The table describes the validity metrics results.

| Model | Temperature | Wind speed | Relative humidity |
|---|---|---|---|
| TSMixer q1,5 | 0.76 | 0.80 | 0.77 |
| TSMixer q2,4 | 0.46 | 0.50 | 0.46 |
| Past TSMixer q1,5 | 0.78 | 0.81 | 0.77 |
| Past TSMixer q2,4 | 0.48 | 0.51 | 0.47 |

**Table 4**
Time-Series Mixer efficiency. The table describes the efficiency metrics results.

| Model | Temperature | Wind speed | Relative humidity |
|---|---|---|---|
| TSMixer q1,5 | 3.93 | 1.83 | 19.07 |
| TSMixer q2,4 | 2.06 | 0.95 | 9.92 |
| Past TSMixer q1,5 | 5.89 | 2.03 | 26.42 |
| Past TSMixer q2,4 | 2.94 | 1.01 | 12.37 |

shows a mixed behavior: in the first hours of the prediction horizon the absolute error has an increasing trend which then saturates after few hours.

To evaluate the TSMixer ability in generating accurate probabilistic forecasts, we use the validity and efficiency metrics, defined in Eqs. (13) and (15). Validity measures the neural network ability to capture the quantiles of the true distribution, while efficiency assesses the neural network confidence in its quantile predictions. The validity results for the TSMixer and Past TSMixer quantile predictions are presented in Table 3.

Both neural networks show good validity, with values approaching the expected 0.80 for the $[0.1, 0.9]$ quantile band and 0.50 for the $[0.25, 0.75]$ quantile band. This indicates that the neural networks are capable of capturing the statistical variations of the true processes and providing reliable risk intervals for future predictions.

The efficiency results for the quantiles predictions of the TSMixer and Past TSMixer are shown in Table 4.

The TSMixer produces smaller risk intervals in all cases, indicating greater confidence when incorporating NWP as future covariates. These results, combined with validity metrics, show that both models exhibit comparable validity, with the Past TSMixer performing slightly better by a few percentage points. However, the efficiency intervals of the Past TSMixer are significantly larger than those of the TSMixer model. The larger risk intervals produced by the Past TSMixer increase the likelihood of encompassing the true values, making its predictions appear more valid. However, the TSMixer achieves the same level of validity while maintaining narrower risk intervals, demonstrating higher confidence in its predictions. This highlights the importance of incorporating future covariates in local weather forecasting.

### 3.1.2. Time-Series Mixer predictions: precipitation

Next, we analyze the precipitation predictions. We confirm that the TSMixer generally predicts zero values for the first two quantiles. A similar behavior is observed in the Past TSMixer model. However, the use of an innovative custom quantile loss function defined in Eq. (7), which penalizes zero predictions, combined with the prediction of log precipitation, has enabled TSMixer to predict non-zero values for the median and fourth and fifth quantiles. Despite these improvements, predicting local hourly precipitation remains an open challenge in the literature. Fig. 5 presents the results of the precipitation predictions on the test set. Fig. 5A compares the cumulative distribution of the logarithmic absolute error among the TSMixer model, the MOLOCH model, and the Past TSMixer model. A logarithmic scale is used to better visualize the distribution profiles, which would otherwise be concentrated near zero. As shown, the median predictions of the TSMixer model are slightly more accurate than those of the MOLOCH and Past TSMixer models. The Past TSMixer exhibits higher errors at smaller error scales. Fig. 5B compares the errors of the three models in predicting cumulative precipitation over the entire prediction horizon (i.e., total rainfall over the next two days). In this case, the MOLOCH model is slightly more accurate than the TSMixer model, while the Past TSMixer produces larger errors. Additionally, we perform a classification task to determine whether precipitation will occur in each of the next 45 h based on the models predictions. We then compute the ROC curve for this classification, along with the AUC, as defined in Eq. (20). The

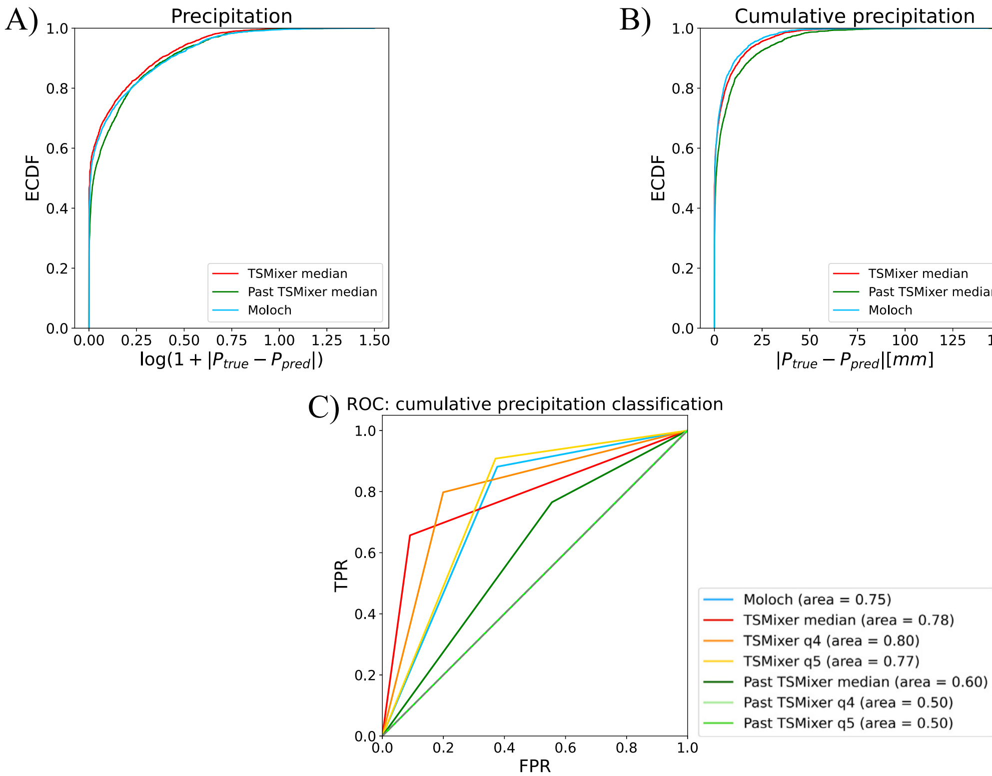


**Fig. 5.** Time-Series Mixer predictions: precipitation. (A) The figure shows the empirical cumulative distribution function (Y-axis) of the logarithm of the absolute error (X-axis) for precipitation predictions. It compares the Time-Series Mixer median prediction (red curve), the MOLOCH predictions (light blue curve), and the Past Time-Series Mixer median prediction (green curve). (B) The figure shows the empirical cumulative distribution function (Y-axis) of the absolute error (X-axis) for the cumulative precipitation predictions over the prediction horizon. It compares the Time-Series Mixer median prediction (red curve), the MOLOCH predictions (light blue curve), and the Past Time-Series Mixer median prediction (green curve). (C) The plot compares the ROC curves for predicting whether there is precipitation for each hour, using the Time-Series Mixer median, fourth and fifth quantiles (red, orange and yellow curves), the MOLOCH predictions (light blue curve), and the Past Time-Series Mixer median, fourth and fifth quantiles prediction (green, pale green and light green curves). The $X$-axis represents the false positive rate, and the $Y$-axis represents the true positive rate. The area under the curve (AUC) is displayed in the legend.

results are shown in Fig. 5C. All TSMixer quantile bands outperform the MOLOCH model, achieving higher AUC values, which indicate a better balance between true positive and false positive rates. Notably, the MOLOCH model has a false positive rate (FPR) comparable to the fifth quantile of the TSMixer model, suggesting that it tends to overpredict precipitation. This may explain the discrepancy between Fig. 5A and B, where the MOLOCH model cumulative precipitation predictions outperform the TSMixer model median predictions. The Past TSMixer model performs similarly to a random classifier for the fourth and fifth quantiles, while its median prediction exhibits a very high false positive rate. Overall, the MOLOCH and TSMixer models demonstrate comparable performance, whereas the Past TSMixer model performs significantly worse. This result is expected, as predicting future precipitation solely from past data patterns is inherently challenging. The MOLOCH model, specifically designed for convective phenomena, naturally excels in precipitation forecasting; however, at such small scales it is far from guaranteed that convection can be correctly predicted in time and space, given the consistently significant uncertainty. The TSMixer model is able to approximate these forecasts but does not provide additional predictive value. The statistical properties of hourly precipitation present significant challenges for neural networks, making accurate predictions for localized fields particularly difficult.

For precipitation, the model faces clear limitations due to the complex physical nature of precipitation and its challenging statistical properties on hourly time scales. An improvement in precipitation predictions could be achieved by incorporating additional factors such as the spatial characteristics of the terrain and vegetation presence, which may influence precipitation patterns.

### 3.2. Time-Series Mixer interpretation

Given the inherently black-box nature of the TSMixer algorithm, it is crucial to interpret how the neural network produces predictions, or at least identify which inputs are most influential. To achieve this, we apply an input perturbation strategy, as described in the Methods Section 2.6.3.

The results of the input perturbation analysis for each hour of the prediction horizon are shown in Fig. 6 for temperature, wind speed, and relative humidity predictions. We have excluded precipitation from this analysis, as the TSMixer model does not outperform the MOLOCH model in that case, making it less relevant for interpretation. Fig. 6A illustrates the input importance for the median temperature prediction. Perturbing the past temperature measured by the ARPA weather station results in mean absolute errors 5 times higher during the initial hours of the prediction horizon (the errors have been defined in Eq. (9)). Among the future variables, air temperature is the most influential, followed by wind speed, dew point temperature, and soil temperature. Fig. 6B shows the input importance for the median wind speed prediction. The most significant past variable is wind speed, especially in the first hours of the prediction horizon. The most relevant future covariates are precipitation and the U and V components of wind speed. Fig. 6C highlights the input importance for the relative humidity prediction. Past relative humidity is the most important input, particularly in the early hours of the prediction horizon. Among future covariates, downward short-wave radiation flux, temperature, and wind speed are the most influential. This analysis reveals that past input variables related to the predicted variable play a crucial role in the prediction process, especially during the first hours of the prediction horizon

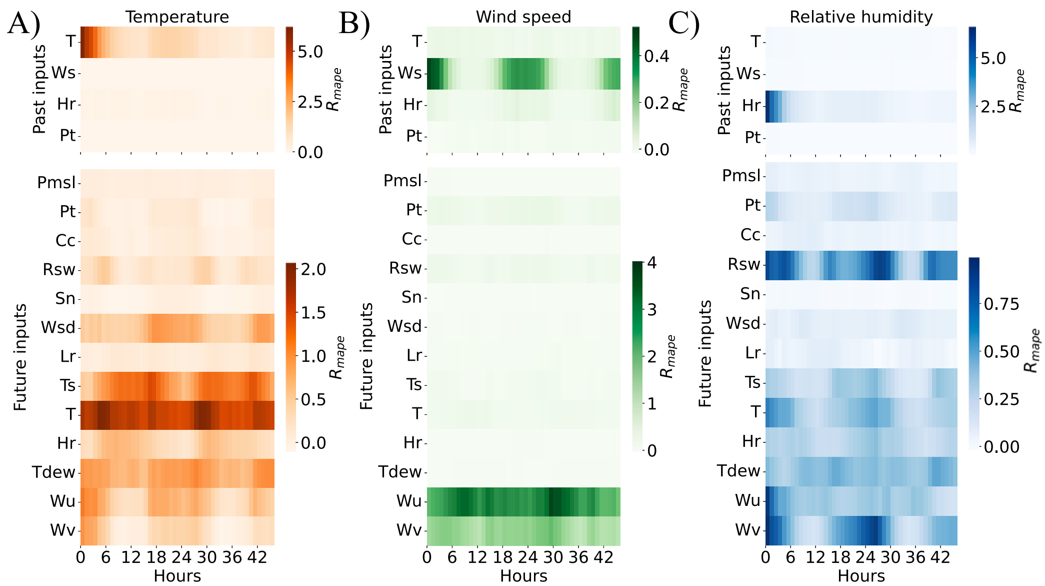


**Fig. 6.** Time-Series Mixer input importance. The figure illustrates the importance of past and future inputs for Time-Series Mixer predictions of temperature (A), wind speed (B), and relative humidity (C). The $Y$-axis lists the input variables, with the upper plot displaying the importance of past inputs and the bottom plot showing the importance of future inputs. The $X$-axis represents the prediction horizon in hours. The colorbar indicates the mean absolute percentage error resulting from input perturbation, defined in Eq. (9).

for temperature and relative humidity. This finding aligns with the results in Fig. 4B, where the TSMixer and Past TSMixer show similar performance in the early hours, when past information still plays a significant role compared to the chaotic nature of weather systems. Additionally, future covariates related to the predicted variable are important. For temperature predictions, MOLOCH air temperature, soil temperature, and dew point are particularly influential, while for wind speed predictions, the U and V components of wind speed and precipitation are key variables. The case of relative humidity is particularly interesting, as MOLOCH relative humidity appears to have little influence on local relative humidity predictions. Instead, the most important future variables are downward short-wave radiation flux and temperature and wind speed. This behavior can be explained by the fact that local humidity is heavily affected by local vegetation and terrain, factors not fully captured by the MOLOCH model, while temperature and short-wave radiation may help adjust predictions for relative humidity.

### 3.3. Prediction of hazardous events for agriculture

Since we have confirmed that the TSMixer neural network, trained with local weather data and NWP, outperforms both the MOLOCH model and the Past TSMixer, and that its predictions are interpretable, we now aim to assess its practical value in the context of precision agriculture. Specifically, we want to ensure that its application is not merely technical but offers real value. To this end, we evaluate TSMixer ability to predict hazardous events for agriculture. Each hazardous event is associated with specific weather conditions, as defined in the Methods section. We verify that the TSMixer predictions accurately capture these events. To assess TSMixer performance, we conduct a classification analysis. A value of 1 is assigned when the hazardous event is correctly detected (i.e., TSMixer predicts it and it actually occurs), and 0 otherwise. We then apply classification metrics, as defined in the Methods section, to compare the performance of TSMixer with that of the MOLOCH model and the Past TSMixer. For each hazardous event, we select two weather stations located near agricultural sites susceptible to the event. In the following paragraphs, we describe the results obtained for each event separately.

Fig. 7A presents the ROC curve for detecting frost events within the next 45 h. Based on the AUC metric, the TSMixer model outperforms

**Table 5**
Frost event classification. The table describes the frost event classification metrics.

| Model | TP | FP | TN | FN | Sens. | Spec. |
|---|---|---|---|---|---|---|
| MOLOCH | 205 | 188 | 663 | 6 | 0.97 | 0.78 |
| TSMixer median | 167 | 17 | 834 | 44 | 0.79 | 0.98 |
| TSMixer q1 | 205 | 109 | 742 | 6 | 0.97 | 0.87 |
| TSMixer q2 | 199 | 57 | 794 | 12 | 0.94 | 0.93 |
| Past TSMixer median | 135 | 7 | 844 | 76 | 0.64 | 0.99 |
| Past TSMixer q1 | 193 | 83 | 768 | 18 | 0.91 | 0.90 |
| Past TSMixer q2 | 166 | 46 | 805 | 45 | 0.79 | 0.95 |

both the Past TSMixer and the MOLOCH model. From a cost optimization perspective, the second quantile prediction of TSMixer offers a strong balance between true positive rate (TPR) and false positive rate (FPR) compared to the other models. However, as shown in Table 5, the first quantile prediction of TSMixer and the MOLOCH achieve the highest sensitivity. The MOLOCH model presents an higher FPR with respect to the first quantile prediction of TSMixer. Therefore, the last model should be taken into consideration, if the cost of implementing frost protection measures in orchards is considered lower than the potential damage caused by frost misclassification.

Fig. 7B shows the ROC curve for detecting rice germination block events within the next 45 h. Based on the AUC metric, the TSMixer and Past TSMixer median predictions outperform the MOLOCH model. The first and second quantiles of TSMixer perform better than the Past TSMixer model, primarily due to a lower FPR. Table 6 provides additional classification metrics for germination block detection. Overall, both the TSMixer and Past TSMixer median offer the best combination of sensitivity and specificity compared to the other predictions. However, TSMixer exhibits a higher number of false negatives than the Past TSMixer, and the MOLOCH model, which on the contrary, have higher false positive rates. These results suggest that for purely protective purposes, the MOLOCH model may suffice to prevent germination blockages. However, from a cost-optimization perspective, AI models like TSMixer can reduce the number of false positives, thus lowering resource use for unnecessary preventive measures.

Fig. 7C presents the ROC curve for detecting heat stress events over the next 45 h (one event per hour). Based on the AUC metric, the TSMixer model outperforms both the Past TSMixer and the MOLOCH model, achieving AUC values of 0.95, 0.97, and 0.95 for the median,

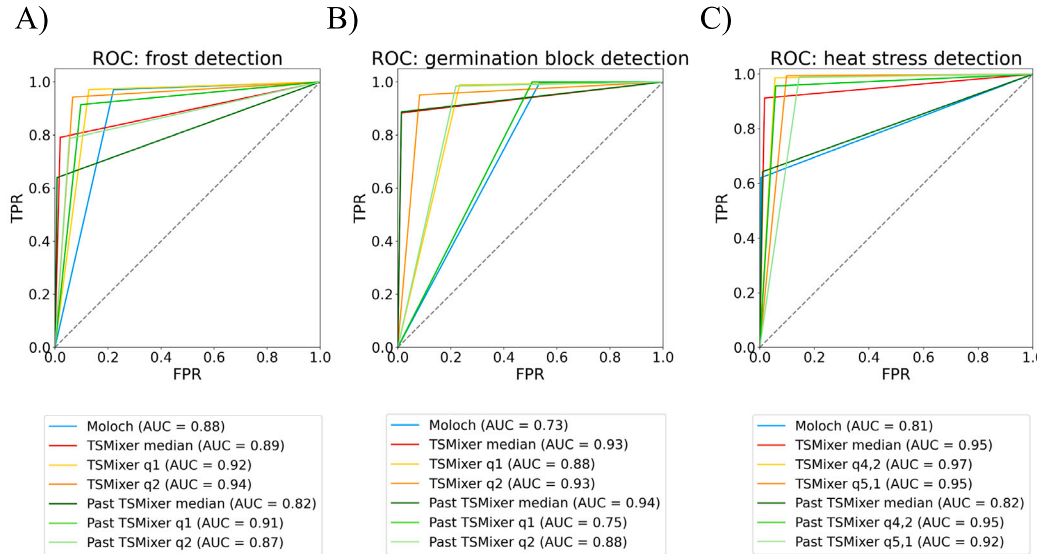


**Fig. 7.** Hazardous event classification with Time-Series Mixer. The figure presents the ROC curves for hazardous event classification: (A) frost, (B) germination block, and (C) heat stress. The following predictions are compared: MOLOCH model (light blue), Time-Series Mixer median (red), Time-Series Mixer first quantile (yellow), Time-Series Mixer second quantile (orange), Past Time-Series Mixer median (green), Past Time-Series Mixer first quantile (lime green), and Past Time-Series Mixer second quantile (pale green). The legend also includes the AUC values for each model.

**Table 6**
Germination block event classification. The table describes the germination block event classification metrics.

| Model | TP | FP | TN | FN | Sens. | Spec. |
|---|---|---|---|---|---|---|
| MOLOCH | 265 | 39 | 34 | 1 | 0.99 | 0.47 |
| TSMixer median | 235 | 1 | 72 | 31 | 0.88 | 0.99 |
| TSMixer q1 | 263 | 17 | 56 | 3 | 0.99 | 0.77 |
| TSMixer q2 | 253 | 6 | 67 | 13 | 0.95 | 0.92 |
| Past TSMixer median | 236 | 1 | 72 | 30 | 0.89 | 0.99 |
| Past TSMixer q1 | 266 | 37 | 36 | 0 | 1.00 | 0.49 |
| Past TSMixer q2 | 262 | 16 | 57 | 4 | 0.98 | 0.78 |

**Table 7**
Heat stress event classification. The table describes the heat stress event classification metrics.

| Model | TP | FP | TN | FN | Sens. | Spec. |
|---|---|---|---|---|---|---|
| MOLOCH | 881 | 188 | 42 048 | 537 | 0.62 | 0.99 |
| TSMixer median | 1294 | 736 | 41 500 | 124 | 0.91 | 0.98 |
| TSMixer q4,2 | 1399 | 2352 | 39 884 | 19 | 0.99 | 0.94 |
| TSMixer q5,1 | 1410 | 4137 | 38 099 | 8 | 0.99 | 0.90 |
| Past TSMixer median | 912 | 436 | 41 800 | 506 | 0.64 | 0.64 |
| Past TSMixer q4,2 | 1357 | 2459 | 39 777 | 61 | 0.96 | 0.94 |
| Past TSMixer q5,1 | 1403 | 6084 | 36 152 | 15 | 0.99 | 0.86 |

first-fifth, and second-fourth quantile predictions, respectively. Table 7 provides additional classification metrics for heat stress detection. Also for these metrics the TSMixer shows better performances with respect to the MOLOCH and Past TSMixer models. The TSMixer first-fifth quantile model shows less false negatives than the second-fourth quantile model, whilst the TSMixer second-fourth quantile model shows less false positives than the first-fifth quantile model. However, in the case of detecting individual hours of heat stress—which does not incur significant costs in terms of prevention or milk production loss—a well-balanced model like the second-fourth quantile predictions from TSMixer may be the optimal choice.

The classification analysis shows that the TSMixer model achieves high accuracy in detecting hazardous events, also optimizing cost-effective counteracting strategies. Overall, the TSMixer model outperforms both the MOLOCH and Past TSMixer models. In addition, we have evaluated the TSMixer capabilities in detecting consecutive hours with hazardous event in the Supplementary materials, showing its great capabilities. From a risk prevention perspective, the MOLOCH model proves reliable for frost and germination blockage events, showing very few false negatives. In contrast, MOLOCH predictions are less reliable for heat stress prevention, likely because the heat stress indicator combines several weather variables and is more sensitive to local field conditions. In addition, the MOLOCH model exhibits a cold bias at higher temperatures, potentially impairing its ability to accurately detect heat stress events. From a cost-optimization standpoint, the AI models outperform the MOLOCH model, positioning them as promising tools for providing reliable weather predictions tailored to local agricultural sites.

## 4. Conclusions

We develop an AI-based framework using TSMixer to produce multivariate weather forecasts. Trained on past local data and MOLOCH model predictions, TSMixer predicts the 45 h probabilistic evolution of temperature, wind speed, humidity, and precipitation for six agricultural sites in Northern Italy. The model outperforms both MOLOCH and a past-data-only TSMixer for temperature, wind speed, and humidity, whilst the hourly precipitation predictions are comparable between the MOLOCH and TSMixer model. Input analysis confirms that TSMixer focuses on key physical variables. Additionally, TSMixer surpasses both models in detecting agricultural hazards, including frost, germination block, and heat stress.

The primary limitation of our approach lies in hourly precipitation predictions. The statistical properties of precipitation, which are highly skewed towards zero, make accurate neural network predictions challenging. Improvements could be achieved by incorporating additional factors such as terrain characteristics and vegetation presence, which likely influence precipitation patterns. A further potential solution could involve interpolating the MOLOCH grid point data to smooth variations in weather variables. Another limitation is the relatively short period covered by the test set (from January 1, 2023, to June 30, 2024). While this period provides a reasonable sample of hazardous events, it does not span a long time frame. This constraint arises from the need for substantial data for training and validation. Future studies could extend the testing period or augment the data by incorporating additional weather stations.

Despite the challenges in hourly precipitation prediction, our study addresses key gaps in the literature by introducing probabilistic, multivariate weather forecasting that combines local data and numerical predictions using AI. These forecasts are effective for detecting hazardous events related to temperature, humidity, and wind speed. By providing reliable probabilistic forecasts and enhancing the early detection of hazardous events, our study substantially supports decision-making processes in precision agriculture. This approach facilitates improved risk management, boosts agricultural productivity, and optimizes resource utilization. These innovations are especially critical given agriculture high sensitivity to weather variability and extreme events, both of which strongly influence farm productivity and economic returns. As climate change continues to increase the occurrence and intensity of extreme weather conditions, robust and precise forecasting methods become indispensable. Additionally, the lightweight nature of TSMixer makes it well-suited for on-field implementation. By transmitting NWP data for the agricultural site to IoT sensors and weather stations via Internet connections, TSMixer can run directly on these devices, enabling real-time monitoring of local microclimate conditions in precision agriculture.

## Software and data availability section

- Software name: LocalWeatherPredictionsDL
- Developer: Marco Zanchi
- Contact information: marco.zanchi97@gmail.com
- First year available: 2025
- Software required: Google Colab
- Program language: Python
- Source code at https://github.com/dndg/LocalWeatherPredictionsDL
- Documentation: Detailed documentation for application installation, testing can be found at https://github.com/dndg/LocalWeatherPredictionsDL
- Data availability: MOLOCH data are available in grib2 format upon request from the server https://tds.bo.isac.cnr.it/. Please send your request to dinamica@isac.cnr.it or to the corresponding author.

## CRediT authorship contribution statement

**Marco Zanchi:** Writing – review & editing, Writing – original draft, Visualization, Validation, Software, Methodology, Investigation, Formal analysis, Data curation, Conceptualization. **Stefano Zapperi:** Writing – review & editing, Supervision, Investigation, Conceptualization. **Stefano Bocchi:** Writing – review & editing, Supervision, Investigation, Conceptualization. **Oxana Drofa:** Writing – review & editing, Investigation, Data curation, Conceptualization. **Silvio Davolio:** Writing – review & editing, Writing – original draft, Supervision, Investigation, Data curation, Conceptualization. **Caterina A.M. La Porta:** Writing – review & editing, Supervision, Investigation, Conceptualization.

## Declaration of Generative AI and AI-assisted technologies in the writing process

During the preparation of this work the authors used ChatGPT in order to improve the text readability. After using this tool/service, the authors reviewed and edited the content as needed and take full responsibility for the content of the publication.

## Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

## Acknowledgments

The work presented in this paper was made possible in part through the financial support provided by DNDG Srl through the doctoral scholarship of Marco Zanchi. Forecast data are provided by the MOLOCH model, at CNR-ISAC. The MOLOCH model is part of an operational prediction chain implemented as part of an agreement with the National Department of Civil Protection (https://www.isac.cnr.it/dinamica/projects/forecasts/).

## Appendix A. Supplementary data

Supplementary material related to this article can be found online at https://doi.org/10.1016/j.envsoft.2025.106509.

## Data availability

MOLOCH data are available in grib2 format upon request from the server https://tds.bo.isac.cnr.it/. Please send your request to dinamica@isac.cnr.itor to the corresponding author.

## References

Baumberger, M., Haas, B., Tewes, W., Risse, B., Meyer, N., Meyer, H., 2025. Gated recurrent units for modelling time series of soil temperature and moisture: An assessment of performance and process reflectivity. Environ. Model. Softw. 183, 106245.

Beck, H.E., Zimmermann, N.E., McVicar, T.R., Vergopolan, N., Berg, A., Wood, E.F., 2018. Present and future Köppen–Geiger climate classification maps at 1-km resolution. Sci. Data 5 (1), 1–12.

Bendre, M., Thool, R., Thool, V., 2015. Big data in precision agriculture: Weather forecasting for future farming. In: 2015 1st International Conference on Next Generation Computing Technologies. NGCT, IEEE, pp. 744–750.

Bergstra, J., Yamins, D., Cox, D., 2013. Making a science of model search: Hyperparameter optimization in hundreds of dimensions for vision architectures. In: International Conference on Machine Learning. PMLR, pp. 115–123.

Bi, K., Xie, L., Zhang, H., Chen, X., Gu, X., Tian, Q., 2023. Accurate medium-range global weather forecasting with 3D neural networks. Nature 619, 1–6. http://dx.doi.org/10.1038/s41586-023-06185-3.

Cadenas, J.M., Garrido, M., Martínez-España, R., Guillén-Navarro, M.A., 2020. Making decisions for frost prediction in agricultural crops in a soft computing framework. Comput. Electron. Agric. 175, 105587.

Calanca, P., Bolius, D., Weigel, A.P., Liniger, M.A., 2011. Application of long-range weather forecasts to agricultural decision problems in Europe. J. Agric. Sci. 149 (1), 15–22.

Chen, S.-A., Li, C.-L., Yoder, N., Arik, S.O., Pfister, T., 2023. Tsmixer: An all-mlp architecture for time series forecasting. arXiv preprint arXiv:2303.06053.

Davolio, S., Malguzzi, P., Drofa, O., Mastrangelo, D., Buzzi, A., 2020. The piedmont flood of november 1994: A testbed of forecasting capabilities of the CNR-ISAC meteorological model suite. Bull. Atmos. Sci. Technol. 1, 263–282.

Easton, Z.M., Kleinman, P.J., Buda, A.R., Goering, D., Emberston, N., Reed, S., Drohan, P.J., Walter, M.T., Guinan, P., Lory, J.A., et al., 2017. Short-term forecasting tools for agricultural nutrient management. J. Environ. Qual. 46 (6), 1257–1269.

Ghaffarian, S., van der Voort, M., Valente, J., Tekinerdogan, B., de Mey, Y., 2022. Machine learning-based farm risk management: A systematic mapping review. Comput. Electron. Agric. 192, 106631.

Gneiting, T., Raftery, A.E., 2005. Weather forecasting with ensemble methods. Science 310 (5746), 248–249.

Guisiano, J., Chiky, R., Orensanz, J., Ahvar, S., 2020. Multilayered ensemble learning for short-term forecasting in agro-climatology. In: Proceedings of the 2020 6th International Conference on Computer and Technology Applications. pp. 7–11.

Hansen, J.W., 2005. Integrating seasonal climate prediction and agricultural models for insights into agricultural practice. Phil. Trans. R. Soc. B 360 (1463), 2037–2047.

Hewage, P., Behera, A., Trovati, M., Pereira, E., Ghahremani, M., Palmieri, F., Liu, Y., 2020. Temporal convolutional neural (TCN) network for an effective weather forecasting using time-series data from the local weather station. Soft Comput. 24, 16453–16482.

Hogg, W., 1971. Spring frosts. Agriculture 78 (1), 28.

Kottek, M., Grieser, J., Beck, C., Rudolf, B., Rubel, F., 2006. World map of the Köppen–Geiger climate classification updated.

Kumar, V., Khan, S., et al., 2017. Importance of weather prediction for sustainable agriculture in Bihar, India. Arch. Agric. Environ. Sci. 2 (2), 105–108.

Lam, R., Sanchez-Gonzalez, A., Willson, M., Wirnsberger, P., Fortunato, M., Alet, F., Ravuri, S., Ewalds, T., Eaton-Rosen, Z., Hu, W., et al., 2022. GraphCast: Learning skillful medium-range global weather forecasting. arXiv preprint arXiv:2212.12794.

Loshchilov, I., Hutter, F., 2017. Decoupled weight decay regularization. arXiv preprint arXiv:1711.05101.
Ma, H.-J., Jin, X.-B., Li, Z.-M., Bai, Y.-T., 2024. Fuzzy adaptive-normalized deep encoder-decoder network: Medium and long-term predictor of temperature and humidity in smart greenhouses. Comput. Electron. Agric. 226, 109480.
Makridakis, S., Spiliotis, E., Assimakopoulos, V., 2022. M5 accuracy competition: Results, findings, and conclusions. Int. J. Forecast. 38 (4), 1346–1364.
Malguzzi, P., Grossi, G., Buzzi, A., Ranzi, R., Buizza, R., 2006. The 1966 "century" flood in Italy: A meteorological and hydrological revisitation. J. Geophys. Res.: Atmos. 111 (D24).
Marusov, A., Grabar, V., Maximov, Y., Sotiriadi, N., Bulkin, A., Zaytsev, A., 2024. Long-term drought prediction using deep neural networks based on geospatial weather data. Environ. Model. Softw. 179, 106127.
Monteith, J., Unsworth, M., 2013. Principles of Environmental Physics: Plants, Animals, and the Atmosphere. Academic Press.
Olivetti, L., Messori, G., 2024a. Advances and prospects of deep learning for medium-range extreme weather forecasting. Geosci. Model. Dev. 17 (6), 2347–2358.
Olivetti, L., Messori, G., 2024b. Do data-driven models beat numerical models in forecasting weather extremes? A comparison of IFS HRES, Pangu-weather and GraphCast. EGUsphere 2024, 1–35.
Perez Garcia, C.A., Benni, S., Tassinari, P., Torreggiani, D., Bovo, M., 2024a. Predicting equivalent temperature index in dairy cows with the NeuralProphet model. In: 11th European Conference on Precision Livestock Farming. pp. 1421–1428.
Perez Garcia, C.A., Bovo, M., Torreggiani, D., Tassinari, P., Benni, S., 2024b. Indoor temperature forecasting in livestock buildings: A data-driven approach. Agriculture 14 (2), 316.
Price, I., Sanchez-Gonzalez, A., Alet, F., Andersson, T.R., El-Kadi, A., Masters, D., Ewalds, T., Stott, J., Mohamed, S., Battaglia, P., et al., 2025. Probabilistic weather forecasting with machine learning. Nature 637 (8044), 84–90.
Pujahari, R.M., Yadav, S.P., Khan, R., 2022. Intelligent farming system through weather forecast support and crop production. In: Application of Machine Learning in Agriculture. Elsevier, pp. 113–130.
Rozante, J.R., Ramirez, E., Ramirez, D., Rozante, G., 2023. Improved frost forecast using machine learning methods. Artif. Intell. Geosci..
Salcedo-Sanz, S., Pérez-Aracil, J., Ascenso, G., Del Ser, J., Casillas-Pérez, D., Kadow, C., Fister, D., Barriopedro, D., García-Herrera, R., Giuliani, M., et al., 2024. Analysis, characterization, prediction, and attribution of extreme atmospheric events with machine learning and deep learning techniques: a review. Theor. Appl. Climatol. 155 (1), 1–44.
Schultz, M.G., Betancourt, C., Gong, B., Kleinert, F., Langguth, M., Leufen, L.H., Mozaffari, A., Stadtler, S., 2021. Can deep learning beat numerical weather prediction? Phil. Trans. R. Soc. A 379 (2194), 20200097.
Shin, J.-Y., Min, B., Kim, K.R., 2022. High-resolution wind speed forecast system coupling numerical weather prediction and machine learning for agricultural studies—a case study from South Korea. Int. J. Biometeorol. 66 (7), 1429–1443.
Sivakumar, M.V., 2006. Climate prediction and agriculture: current status and future challenges. Clim. Res. 33 (1), 3–17.
Smith, B.A., Hoogenboom, G., McClendon, R.W., 2009. Artificial neural networks for automated year-round temperature prediction. Comput. Electron. Agric. 68 (1), 52–61.
Tantalaki, N., Souravlas, S., Roumeliotis, M., 2019. Data-driven decision making in precision agriculture: The rise of big data in agricultural systems. J. Agric. Food Inf. 20 (4), 344–380.
Teixeira, S.B., Pires, S.N., Ávila, G.E., Silva, B.E.P., Schmitz, V.N., Deuner, C., da Silva Armesto, R., da Silva Moura, D., Deuner, S., 2021. Application of vigor indexes to evaluate the cold tolerance in rice seeds germination conditioned in plant extract. Sci. Rep. 11 (1), 11038.
Tsela, R., Maladaki, S., Kolios, S., 2025. An integrated, automated and modular approach for real-time weather monitoring of surface meteorological variables and short-range forecasting using machine learning. Environ. Model. Softw. 183, 106203.
Unterberger, C., Brunner, L., Nabernegg, S., Steininger, K.W., Steiner, A.K., Stabentheiner, E., Monschein, S., Truhetz, H., 2018. Spring frost risk for regional apple production under a warmer climate. PloS One 13 (7), e0200201.
Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, Ł., Polosukhin, I., 2017. Attention is all you need. Adv. Neural Inf. Process. Syst. 30.
Wen, R., Torkkola, K., Narayanaswamy, B., Madeka, D., 2017. A multi-horizon quantile recurrent forecaster. arXiv preprint arXiv:1711.11053.
West, J.W., 2003. Effects of heat-stress on production in dairy cattle. J. Dairy Sci. 86 (6), 2131–2144.
Yan, G., Li, H., Shi, Z., 2021. Evaluation of thermal indices as the indicators of heat stress in dairy cows in a temperate climate. Animals 11 (8), 2459.
Yang, M., Yang, J., Su, L., Sun, K., Li, D., Liu, Y., Wang, H., Chen, Z., Guo, T., 2019. Metabolic profile analysis and identification of key metabolites during rice seed germination under low-temperature stress. Plant Sci. 289, 110282.
Yang, J., Yu, M., Liu, Q., Li, Y., Duffy, D.Q., Yang, C., 2022. A high spatiotemporal resolution framework for urban temperature prediction using IoT data. Comput. Geosci. 159, 104991.
Zanchi, M., La Porta, C.A., Forte, C., Zapperi, S., Ozella, L., 2025. Influence of microclimatic conditions on dairy production in an automatic milking system: Trends and time-series mixer predictions. Comput. Electron. Agric. 229, 109730.
Zanchi, M., Zapperi, S., La Porta, C.A., 2023. Harnessing deep learning to forecast local microclimate using global climate data. Sci. Rep. 13 (1), 21062.
Zanchi, M., Zapperi, S., La Porta, C.A., 2024. Optimized placement of sensor networks by machine learning for microclimate evaluation. Comput. Electron. Agric. 225, 109305.
Zanchi, M., Zapperi, S., Zanotti, C., Rotiroti, M., Bonomi, T., Gomarasca, S., Bocchi, S., La Porta, C.A., 2022. A pipeline for monitoring water pollution: The example of heavy metals in lombardy waters. Heliyon 8 (12), e12435.
Zeng, A., Chen, M., Zhang, L., Xu, Q., 2023. Are transformers effective for time series forecasting? In: Proceedings of the AAAI Conference on Artificial Intelligence. Vol. 37, pp. 11121–11128.
Zhou, I., Lipman, J., Abolhasan, M., Shariati, N., 2023. Intelligent spatial interpolation-based frost prediction methodology using artificial neural networks with limited local data. Environ. Model. Softw. 165, 105724.