SCEWERO

STRENGTHENING THE RESEARCH CAPACITIES
FOR EXTREME WEATHER EVENTS IN ROMANIA

AI for Continental-Scale Super High-Resolution Downscaling

From coarse global climate information to km-scale impact-relevant fields

Isidoros Iakovidis
Euro-Mediterranean Center on Climate Change (CMCC)
Cluj-Napoca, 24 June 2026

UNIVERSITATEA BABEŞ-BOLYAI

cmcc Centro Euro-Mediterraneo sui Cambiamenti Climatici

Universiteit Antwerpen

indeco soft

Funded by the European Union
Grant Agreement 101159497

What this talk is about

One question

How do we turn a coarse climate forecast into local, high-resolution fields that are useful for climate-impact work?

Plan of the talk:

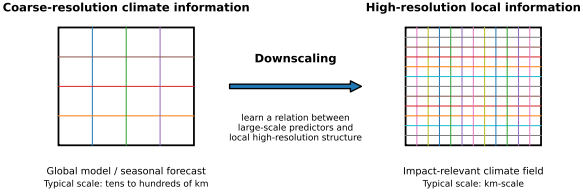
- the resolution gap, and what downscaling means;
- the data: the high-resolution reference (EMO-1) and the ECMWF seasonal forecasts;
- how the model is trained, and how a real forecast is produced;
- how we treat temperature and precipitation differently, and why.

Funded by the European Union
Grant Agreement 101159497

The resolution gap

Two scales

Global models and seasonal forecasts describe the atmosphere at tens to hundreds of kilometres. Decisions are taken at the scale of a city, a valley, a catchment or a coast. Inside one coarse cell there can be a mountain, a valley and a coastline, each with a different climate. (impact)!



Funded by the European Union
Grant Agreement 101159497

What downscaling means

A mapping from coarse to local

We learn a function f that turns coarse fields and fixed local information into a high-resolution field:

$$Y_{\text{high}} = f(X_{\text{coarse}}, S_{\text{local}}).$$

- **predictors** X: the coarse forecast fields;
- **static fields** S: topography and a land–sea mask;
- **target** Y: the high-resolution field we want.

Many fine fields are consistent with one coarse field, so f is a best estimate, not an exact inverse.

Funded by the European Union
Grant Agreement 101159497

Two families, and our choice

Dynamical	Statistical
Run a high-resolution physical model inside the global one. Physically based, but expensive.	Learn the coarse-to-local relation from historical data. Cheap to apply once trained.

In this talk

A statistical method based on a convolutional neural network. A climate field can be represented as a 2-D array, and convolutions learn local spatial patterns by applying shared filters across the domain.

Funded by the European Union
Grant Agreement 101159497

Why a convolutional network

A climate field is an image

Each gridded variable is a 2-D array, so image-based convolutional tools can be used. Convolutional networks are useful here for a few concrete reasons:

- **local filters**: a small filter reads a neighbourhood, matching the local nature of weather;
- **weight sharing**: the same learned filter is used across the grid, reducing the number of parameters;
- **translation equivariance**: a convolutional layer responds similarly to a shifted pattern. Boundaries, padding and location-specific inputs limit this property in the complete model;
- **hierarchy**: stacked layers build from simple local patterns (gradients, fronts) up to larger structures.

Funded by the European Union
Grant Agreement 101159497

The two ingredients

The target (what we want)	The predictors (what we have)
EMO-1: a high-resolution gridded observational reference over Europe. It is the reference target that the model learns to reproduce and that we use for evaluation.	The ECMWF seasonal forecast: coarse atmospheric fields, as an ensemble of members (more on this next).

Plus static fields

Topography and a land–sea mask, which do not change in time, give the model the local geography.

EMO-1: the high-resolution reference

What it is

EMO-1 is a gridded observational dataset covering Europe at 1 arc-minute, about 1.5 km, built from station measurements. It provides daily fields: maximum and minimum temperature, and precipitation.

Its two roles

- during **training**, it is the target the network is taught to reproduce;
- during **evaluation**, it is the observational reference used to assess the prediction.

EMO-1 is available only for historical periods with observations, not for a future forecast.

Funded by the European Union
Grant Agreement 101159497

What the ECMWF seasonal forecast does

A coupled model run months ahead

SEAS5 is ECMWF's seasonal forecast system: a coupled atmosphere–ocean–sea-ice model, run forward on a coarse grid (about 36 km) from the first of each month, out to roughly seven months.

What it predicts

At long lead, it provides probabilities for seasonal conditions rather than the exact weather on a particular calendar day. Useful information includes large-scale circulation and temperature and precipitation anomalies. It is run as an *ensemble* of members from slightly perturbed starts, so the spread provides an indication of forecast uncertainty.

Funded by the European Union
Grant Agreement 101159497

ECMWF forecasts: reforecast and real-time

One system, used two ways

SEAS5 can be run for past dates or for the coming season. We use both, for different purposes.

Reforecast (hindcast): 25 members

The same model re-run for past years (we use 1993 on), giving a long, consistent record of past forecasts to pair with EMO-1 observations for training.

Real-time forecast: 51 members

The operational forecast issued for the coming season. This is what we downscale in production, when no observation exists yet.

Both start on the first of each month and run about seven months ahead; the reforecast carries fewer members than the real-time forecast.

Funded by the European Union
Grant Agreement 101159497

The datasets, and what each is for

Dataset	What it is	Role here
EMO-1	1 arc-min observations, Europe	target / observational reference
ECMWF reforecast	25-member past forecasts	training predictors
ECMWF real-time forecast	51-member live forecast	production predictors
ERA5	reanalysis of the past atmosphere	reference to find analogs

Forecast versus reanalysis

A *forecast* evolves forward and can drift from reality. ERA5 combines observations with a model to estimate the past atmosphere and provides a historical reference for pattern comparison.

Funded by the European Union
Grant Agreement 101159497

Which dataset plays which role

Phase	Predictors	Target	Period
Training	ECMWF reforecast (25)	EMO-1	1993–2013
Validation	ECMWF reforecast (25)	EMO-1	2014
Test	ECMWF reforecast (25)	EMO-1	2015
Production	ECMWF forecast (51)	—	new seasons

Why split by whole years

Training fits the model; validation tunes the choices; the test year is judged once, at the end. The split is by whole years because neighbouring days are very similar; mixing them across sets would make the model look better than it is.

Funded by the European Union
Grant Agreement 101159497

What goes into the model

Three kinds of input

- the coarse ECMWF forecast of the **same variable** (for precipitation, the low-resolution forecast precipitation; for temperature, the coarse temperature);
- a set of large-scale ECMWF fields describing the weather pattern: geopotential at several levels, mean sea-level pressure, humidity and temperature at 700 hPa, vorticity at 200 hPa, wind speed;
- two fixed maps: the topography and the land–sea mask.

All the forecast fields come from ECMWF and are stacked as channels (each geopotential level is its own channel). Topography and the mask are fixed and shared by every day.

Funded by the European Union
Grant Agreement 101159497

Temperature and precipitation use different inputs

Temperature: a few fields

The coarse temperature plus geopotential at four levels, about five input channels. The geopotential already sets much of the temperature, so little else is needed.

Precipitation: many fields

The coarse precipitation plus geopotential, pressure, humidity, temperature, vorticity and wind, about ten channels. Rain depends on the whole weather situation, so the model is given more.

Same idea, different richness

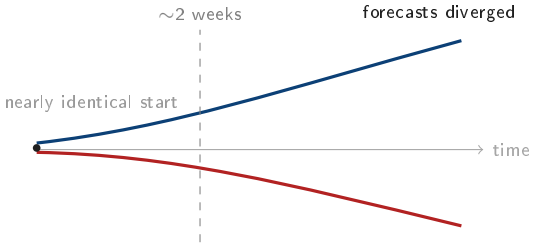
Both models share the architecture and the static inputs; precipitation simply needs a richer description of the atmosphere.

Funded by the European Union
Grant Agreement 101159497

Why the exact day cannot be predicted

Chaos, after Lorenz

Small differences in the atmospheric initial state grow with lead time. Beyond roughly one to two weeks, daily evolutions can diverge substantially, so a forecast issued months ahead cannot determine one specific day's weather.



At seasonal range, useful skill is probabilistic and concerns anomalies and seasonal statistics.

Funded by the European Union
Grant Agreement 101159497

Why we use analogs

Possible trajectories, not an exact daily sequence

Seasonal ensemble members provide dynamically consistent possible trajectories. Beyond the weather-predictability range, their daily sequences sample a forecast distribution rather than the sequence that will occur.

What the analog gives instead

An analogue method searches historical reanalysis for circulation states similar to a forecast state. The matches supply historically plausible spatial structures.

Funded by the European Union
Grant Agreement 101159497

How the analogs are built

A pattern search in the historical record

For each forecast day and ensemble member, every large-scale SEAS5 field is compared with historical ERA5 candidates:

- the **Teweles–Wobus score** emphasises spatial-gradient similarity rather than directly matching the absolute field level;
- candidates lie within about 15 days of the date, so the season matches;
- fields and levels may select different analogue dates.

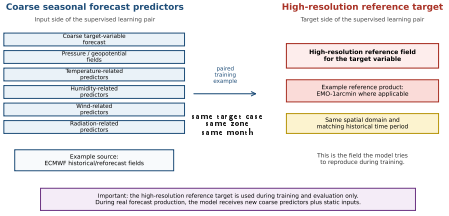
Historical ERA5 remains searchable in production; only the future verifying state is unavailable. Validation and test cases need a training-only archive. Per-field matches are not jointly physically consistent.

Funded by the European Union
Grant Agreement 101159497

A supervised training pair

One example

For each day in the historical period we form a pair: the coarse predictors and the static fields on the input side, the EMO-1 reference on the output side. The network learns from many such pairs.



Funded by the European Union
Grant Agreement 101159497

Scaling the data

Why and how

Variables have different units and ranges, and training is more stable when they share a comparable scale. We compute the scaling for **every grid cell separately**, using only the **training years**, then reuse it unchanged for validation, test and real forecasts. The parameters are saved, so a prediction can be returned to physical units.

Per-cell scaling matters: a coastal cell and a mountain cell have very different climates, and a single global scaling would distort both. Sea cells, which have no value, are handled with the mask.

Funded by the European Union
Grant Agreement 101159497

The standardisation options

Scaler	Forward (per cell)	Inverse
min–max	$\tilde{y} = \dfrac{y - y_{\min}}{y_{\max} - y_{\min}}$	$y = \tilde{y}\,(y_{\max} - y_{\min}) + y_{\min}$
standard	$\tilde{y} = \dfrac{y - \mu}{\sigma_y}$	$y = \tilde{y}\,\sigma + \mu$
abs–max	$\tilde{y} = \dfrac{y}{\max\lvert y\rvert}$	$y = \tilde{y}\ \max\lvert y\rvert$
robust	$\tilde{y} = \dfrac{y - \mathrm{med}}{\mathrm{IQR}}$	$y = \tilde{y}\,\mathrm{IQR} + \mathrm{med}$

$\mathrm{IQR} = Q_3 - Q_1$, replaced by 1 below 0.4 to avoid dividing by a near-zero spread. Standard scaling suits roughly symmetric variables; robust scaling resists outliers; abs–max keeps the sign and the zero. Every option is invertible, so the prediction can return to physical units.

Funded by the European Union
Grant Agreement 101159497

Temperature and precipitation are scaled differently

Temperature: min–max	Precipitation: abs–max
Maps each cell's training range to $[0, 1]$. This is a numerical transform, not a physical bound.	Maps an observed zero to zero during scaling. It does not guarantee a zero prediction.

The reason

These are numerical choices. Precipitation is non-negative and contains many zeros, while temperature is continuous. New values can still fall outside the training range.

Funded by the European Union
Grant Agreement 101159497

The network, in one picture

Three inputs, one output

$$X_{\text{coarse predictors}} + \text{topography} + \text{land–sea mask} \longrightarrow \widehat{Y}.$$

- an **encoder** compresses the coarse predictors into features;
- **attention** lets the network focus on the useful channels and regions;
- a **decoder** rebuilds the field at high resolution, guided by topography;
- the **land–sea mask** excludes sea cells from the loss and evaluation; masked output cells are stored as missing values.

It is an encoder–decoder of the kind used for image super-resolution.

Funded by the European Union
Grant Agreement 101159497

Same network family, different per variable

	Temperature	Precipitation
Input fields	∼5	∼10
Scaling	min–max	abs–max
Output activation	swish	ReLU
Network size	smaller	wider and deeper

- **ReLU for precipitation**: enforces a non-negative scaled output;
- **swish for temperature**: a smooth modelling choice. Negative Celsius values are recovered after inverse scaling;
- precipitation is intermittent and spatially patchy, which can justify greater model capacity.

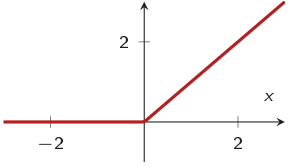
Funded by the European Union
Grant Agreement 101159497

The two output activations

Different numerical constraints

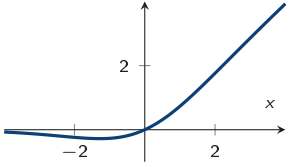
ReLU enforces non-negative scaled precipitation. Swish is a smooth modelling choice for temperature, but the sign of a scaled temperature is not its sign in physical units.

Precipitation: ReLU



$f(x) = \max(0, x)$: clips negatives to zero

Temperature: swish



$f(x) = x\,\sigma(x)$: smooth, allows small negatives

Funded by the European Union
Grant Agreement 101159497

How the network is trained

Fit, then keep the best

The network minimises mean squared error against the EMO-1 reference using the Adam optimiser. Training stops at the epoch with the best validation-year performance.

One thing to remember

Mean squared error targets the conditional mean. It can smooth uncertain fine-scale structure and underestimate some extremes, although this is not guaranteed in every case.

What a trained model is

Not only a weights file

For each variable, zone and month, the saved package contains everything needed to reproduce a forecast:

- the trained network weights;
- the scaling parameters, to return to physical units;
- the topography and the land–sea mask;
- the expected order of the predictor fields, and the output grid.

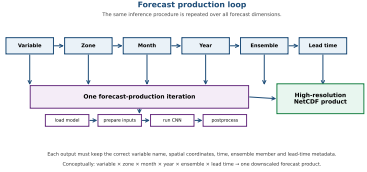
Funded by the European Union
Grant Agreement 101159497

Producing a real forecast

From a new seasonal forecast to a local field

When a real-time forecast arrives, no future observation exists, so we apply the trained model.

1. prepare the predictors with the same fields, order, scaling and analogue procedure used in training;
2. run the network and return the output to physical units;
3. restore the sea cells with the mask and write a NetCDF file with full metadata.



Funded by the European Union
Grant Agreement 101159497

The scale of production

One model per combination

The pipeline is repeated across several dimensions:

$$\text{variable} \times \text{zone} \times \text{month} \times \text{year} \times \text{ensemble} \times \text{lead time}.$$

- Europe is split into **zones** that are physically coherent and fit in memory;
- there is one model per **calendar month**, since seasons behave differently;
- the ensemble members (51 in the real-time forecast) carry the forecast uncertainty into the high-resolution product.

Funded by the European Union
Grant Agreement 101159497

Quality control before the data is used

Treat the output as a climate dataset

- were all expected files produced, with the right dimensions and coordinates?
- are units and variable names correct, and missing values handled?
- are values physically plausible, for instance no negative rainfall?
- are the ensemble members and lead times preserved?

Why it matters

These fields feed impact indicators such as heat, drought and wind risk, where a small error can grow when it passes through a threshold.

Funded by the European Union
Grant Agreement 101159497

How we check the model

On the test year, against EMO-1

We compare the downscaled fields with the observations using simple, standard scores:

- average error (bias), typical error size (mean absolute error, root mean square error);
- how much of the variability is captured (R^2);
- maps of the error, to see whether it is systematic in particular places.

Funded by the European Union
Grant agreement 101159497

Trade-offs: dynamical versus statistical

Dynamical (regional model)

- \+ physically consistent, all variables together;
- \+ may extrapolate more plausibly because physical equations constrain it;
- − computational cost usually limits ensemble size;
- − model biases and structural errors remain.

Statistical (learned)

- \+ cheap and fast once trained;
- \+ easy to run many members and zones;
- − needs good training data;
- − assumes the past relation holds.

Why statistical here

At seasonal range we must downscale many members, months and zones; a learned model makes that affordable.

Funded by the European Union
Grant Agreement 101159497

Trade-offs: among statistical methods

Method	Strengths	Limitations
Regression	simple and interpretable	limited by the chosen functional form
Analogue methods	non-parametric and intuitive	depend on archive coverage and distance metric
Trees and kernels	nonlinear and flexible	spatial structure must be encoded explicitly
Deep learning (CNN)	learns shared spatial features	data intensive and less interpretable

Why a CNN

CNNs encode local spatial structure directly through shared filters. Other methods can also model spatial fields when they use suitable features, kernels or multivariate formulations.

Funded by the European Union
Grant Agreement 101159497

Trade-offs: our CNN approach

Strengths

+ learns local patterns from the fields;
+ one pass gives a full high-res map;
+ fast, so many members are affordable;
+ takes many predictors and static maps as channels.

Limitations

− needs substantial training data;
− less interpretable, a black box;
− the squared-error loss smooths extremes;
− weak outside the training range.

We accept these because the gains in speed and spatial skill fit a seasonal, many-member production setting.

Funded by the European Union
Grant Agreement 101159497

What to keep in mind

- the model assumes the past relation still holds for new forecasts and a changing climate;
- historical reanalysis remains available for analogue searches, but the future verifying state does not;
- any difference between training and production preprocessing introduces a distribution shift;
- squared-error training can smooth the sharpest extremes;
- predictions become less reliable outside the training distribution; this alone does not quantify uncertainty.

Ways to reduce these risks

Use blocked validation, checks for unfamiliar inputs, ensemble calibration and losses that better represent extremes.

Funded by the European Union
Grant Agreement 101159497

Take-home messages

- Downscaling turns a coarse forecast into a km-scale field; the model learns it from past data.
- The data is the heart: EMO-1 is the high-resolution target, the ECMWF reforecast trains the model, and the real-time forecast is what we downscale in production.
- Temperature and precipitation are handled differently, and the differences follow the physics of each variable.
- The output is a climate dataset and must be checked as one, not only as a network output.

Funded by the European Union
Grant Agreement 101159497

Closing

Thank you very much for your attention.

Mulțumesc foarte mult!

Funded by the European Union
Grant Agreement 101159497

References I

Maraun, D., Wetterhall, F., Ireson, A. M., et al. (2010). Precipitation downscaling under climate change: recent developments to bridge the gap between dynamical models and the end user. *Reviews of Geophysics*, 48, RG3003.

Maraun, D., and Widmann, M. (2018). *Statistical Downscaling and Bias Correction for Climate Research*. Cambridge University Press.

Lorenz, E. N. (1969). Atmospheric predictability as revealed by naturally occurring analogues. *Journal of the Atmospheric Sciences*, 26(4), 636–646.

Vandal, T., Kodra, E., Ganguly, S., et al. (2017). DeepSD: generating high-resolution climate change projections through single-image super-resolution. *Proc. 23rd ACM SIGKDD*, 1663–1672.

Wang, F., Tian, D., Lowe, L., Kalin, L., and Lehrter, J. (2021). Deep learning for daily precipitation and temperature downscaling. *Water Resources Research*, 57(4), e2020WR029308.

Funded by the European Union
Grant Agreement 101159497