Deep learning-based downscaling of seasonal forecasts over the Iberian Peninsula

Carlos Gómez-Gonzalez, Lluís Palma Garcia, Llorenç Lledó,
Raül Marcos, Nube Gonzalez-Reviriego, Giulia Carella and Albert Soret

carlos.gomez@bsc.es

EGU 2021

March 30, 2021

Objectives

- Statistical downscaling aims at learning empirical links between the large-scale and local-scale climate, i.e., a mapping from a low-resolution gridded variable to a higher-resolution grid
- In this study, we aim at improving the coarse spatial resolution of seasonal forecasts
- Concretely, we downscale the SEAS5 ECMWF seasonal forecast of temperature over the north-east of the Iberian Peninsula (Catalunya) using deep convolutional networks in supervised and generative adversarial training frameworks
- We rely on temperature observations from two reanalyses (UERRA, ERA5) to learn how to take SEAS5 from its native 1º grid to a 0.05º resolution (20x upscaling factor)

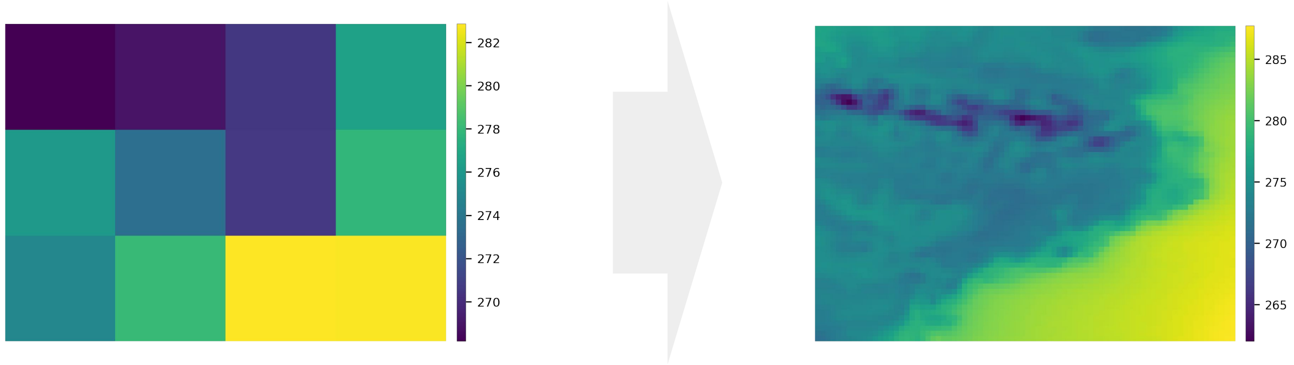

Data and methods

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

Data

SEAS5 tas 1º (1981-2018)

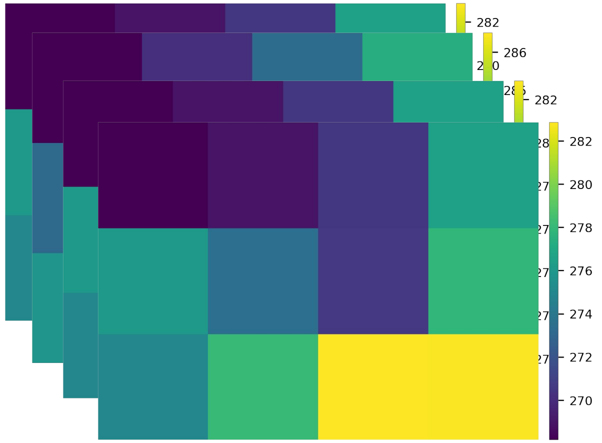

UERRA tas 0.05º (1979-2018)

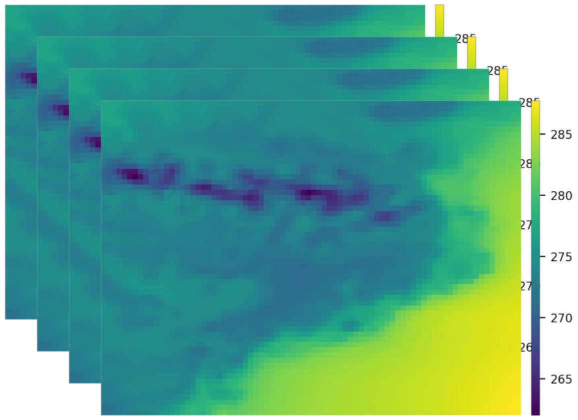

ERA5 tas 0.25º (1979-2018)

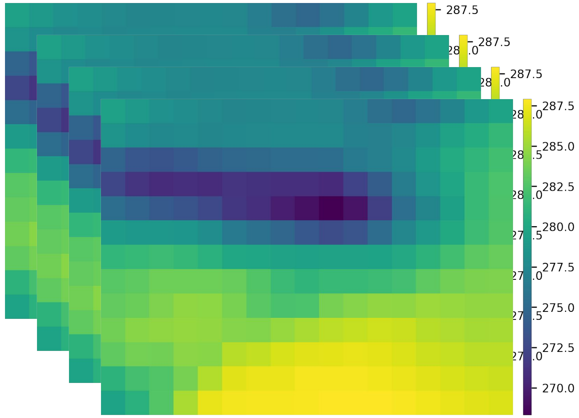

Static fields: elevation and land-ocean mask

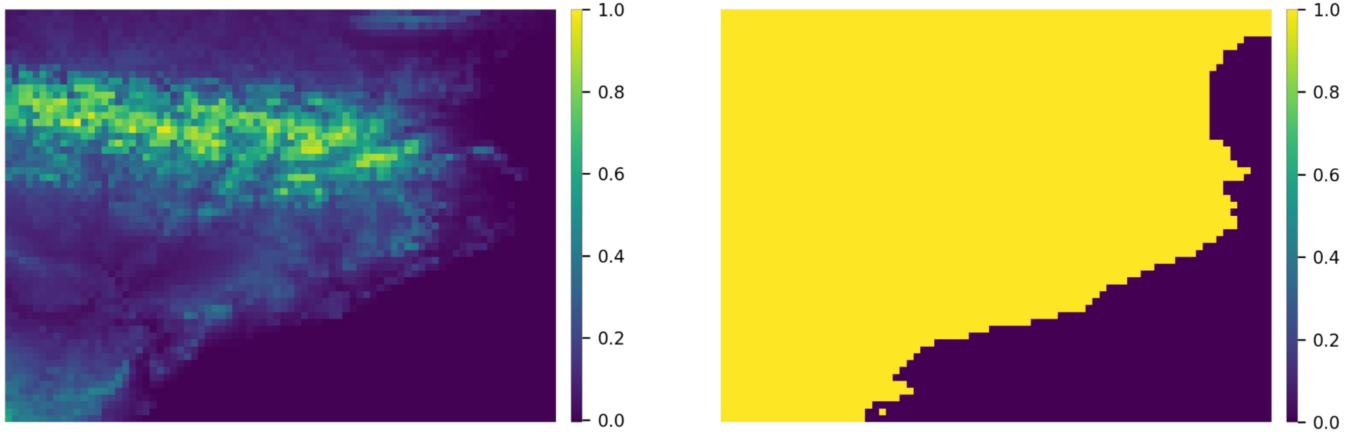

Other ERA5 variables: humidity, geopotential, wind, etc

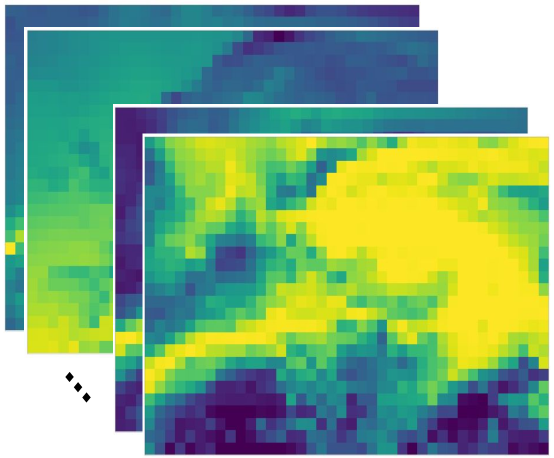

Experimental setup

Historical observations

training dataset (32 years)

Train / hypertune downscaling methods

field RMSE

holdout dataset (8 years)

Verify downscaling methods on unseen observations

- field RMSE
- spatial Correlation

quality of the downscaling method on daily mean fields

Perfect Prognosis: The models are trained with obs, but used later with seasonal forecasts

Seasonal forecasts

hindcast @daily scale

Downscale SEAS5 predictions:

- Nov start dates for DJF
- 1981-2017
- member by member

hindcast @seasonal scale

Verify SEAS5 downscaled

- EnsCorr
- RPSS

quality of the downscaled SEAS5 DJF forecasts issued in November

Baseline methods

- The **conservative** interpolation is used as a “no-downscaling” baseline:
    1. take the data from the closest coarse-scale grid point
    2. compare it directly to fine-scale observations

- The **analogs** method has been used in the Earth Sciences for long to downscale field predictions:
    1. look in the historical archive for the 15 days with most similar large-scale temperature conditions (according to RMSE).
    2. The high-resolution fields of those 15 analog days are then averaged and used as high-resolution forecast

  From a data science perspective, this corresponds to a K-Nearest Neighbors regression with k=15 and euclidean distance, where each coarse-scale grid point value is a feature, and the high-resolution observations are the (multi-output) predictands. The k=15 optimal parameter was found by cross-validation in the training dataset.

DL-based super resolution for statistical spatial downscaling

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

- Several super-resolution approaches have been proposed in the field of computer vision
- These ideas have inspired deep learning-based downscaling methods in climate science, e.g., Vandal et al. 2017, Leinonen et al. 2020, Stengel et al. 2020

Model families

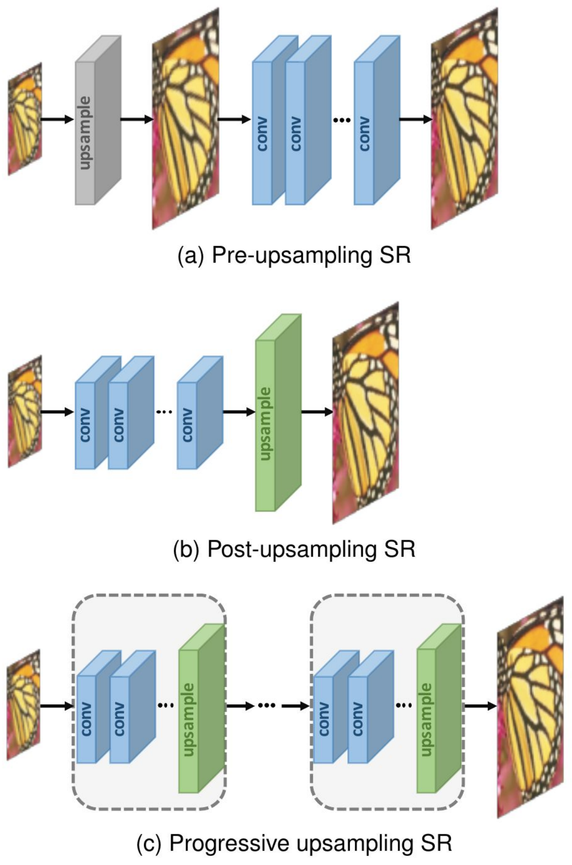

Upscaling methods

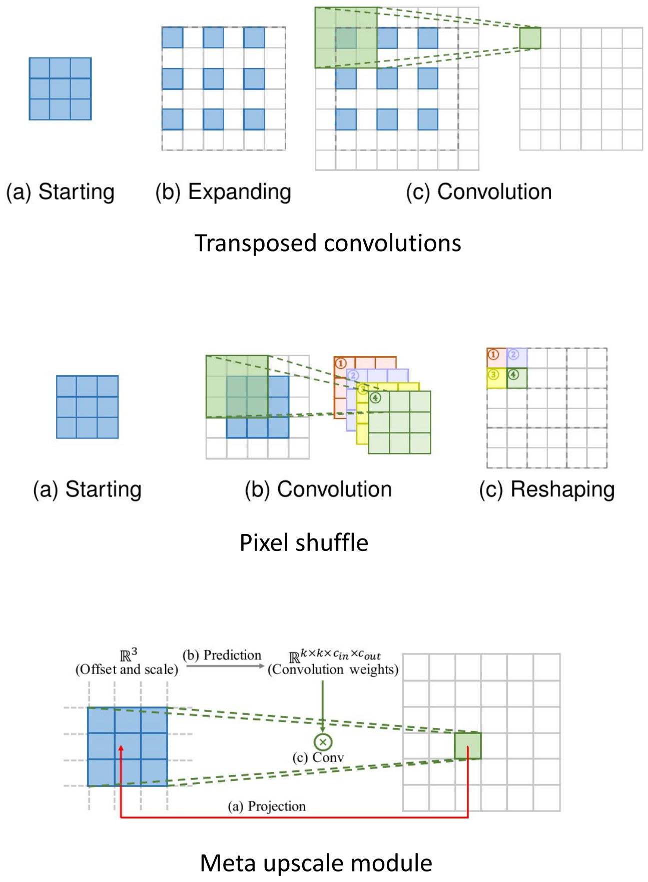

Model architectures

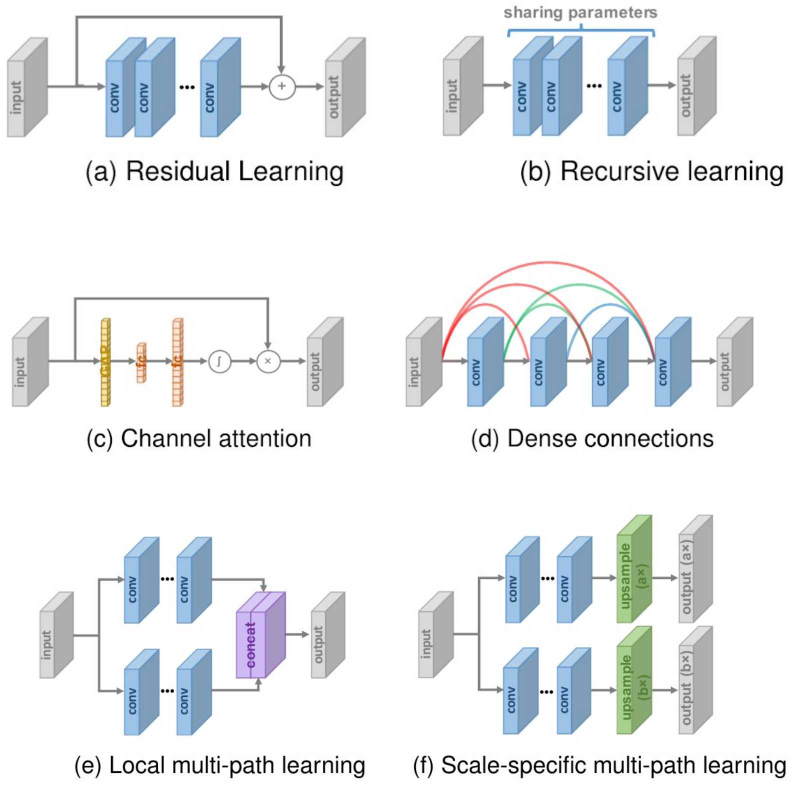

Taken from Wang, et al. 2020

DL-based methods

- Several deep neural network architectures were implemented and trained to learn the mapping between the coarse and fine spatial resolutions
- We compare:
  - single (20x) and progressive (4x + 5x) models
  - adversarial training, as conditional generative adversarial networks (CGANs), and non-adversarial training
  - various scaling methods (here we showcase two, one for pre and one in post-upscaling)
- Deep residual networks (He et al. 2015) are used as a backbone model
  - w/o batch norm (Lim et al. 2017)
  - with a channel attention mechanism (Woo et al. 2018)

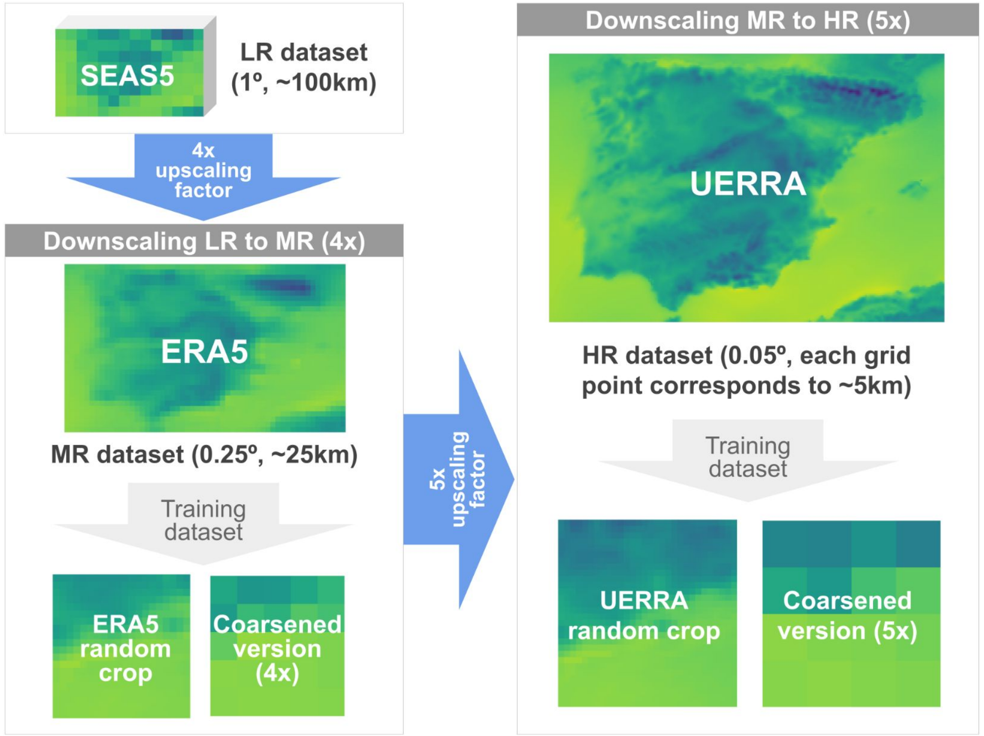

(a) Progressive downscaling

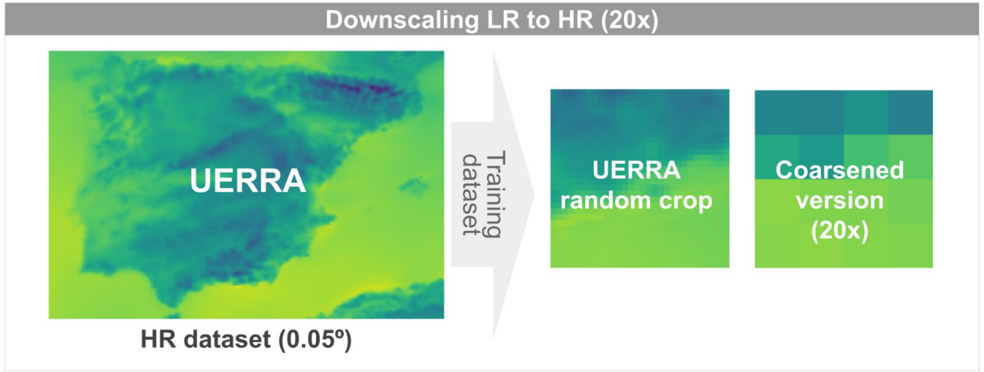

(b) Downscaling with a single model

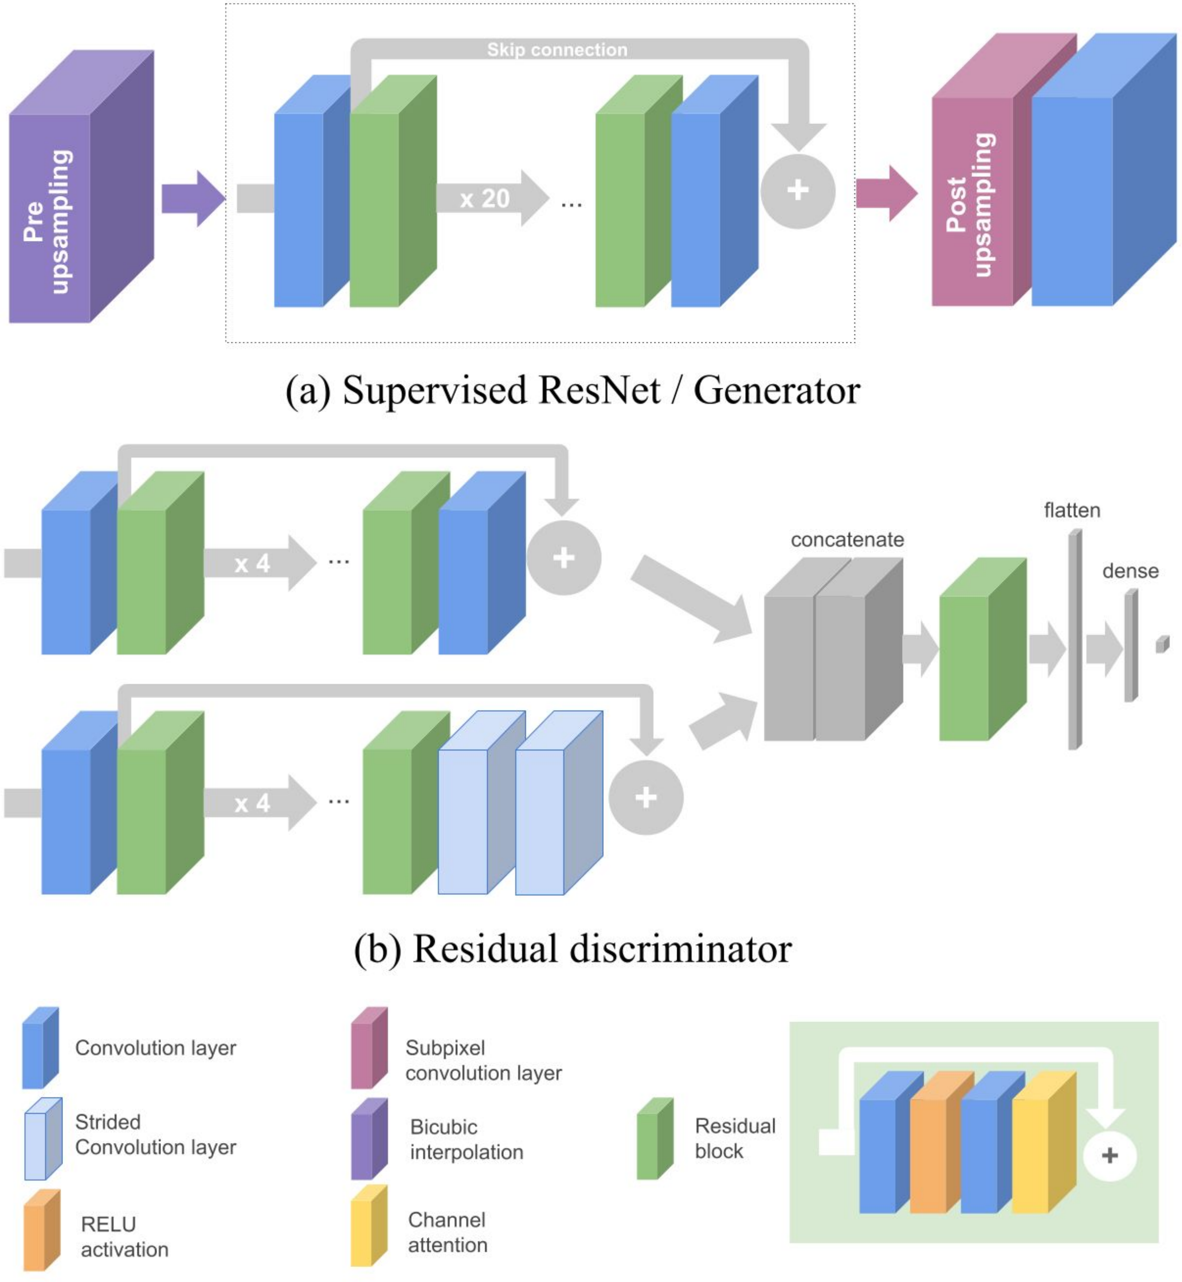

(a) Supervised ResNet / Generator

(b) Residual discriminator

- We have different architectures, depending on the stage at which the upscaling is performed
- For panel *(a)*, only one of the blocks “pre-upsampling” and “post-upsampling” is active for a given model
  - ResNet-INT uses pre-upsampling
  - ResNet-SPC uses post-upsampling
- The purely supervised models are trained with a mean absolute error loss (MAE)
- Our CGANs feature either the ResNet-SPC or ResNet-INT as generators, shown in panel *(a)*. The architecture of the discriminator is shown in panel *(b).* The CGAN models are trained as standard GANs (Isola et al. 2016)
- The networks are implemented in Tensorflow/Keras and trained on the POWER-CTE cluster of the Barcelona Supercomputing Center

Results and validation

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

Holdout verification

RMSE and correlation measured using the holdout UERRA dataset as a reference

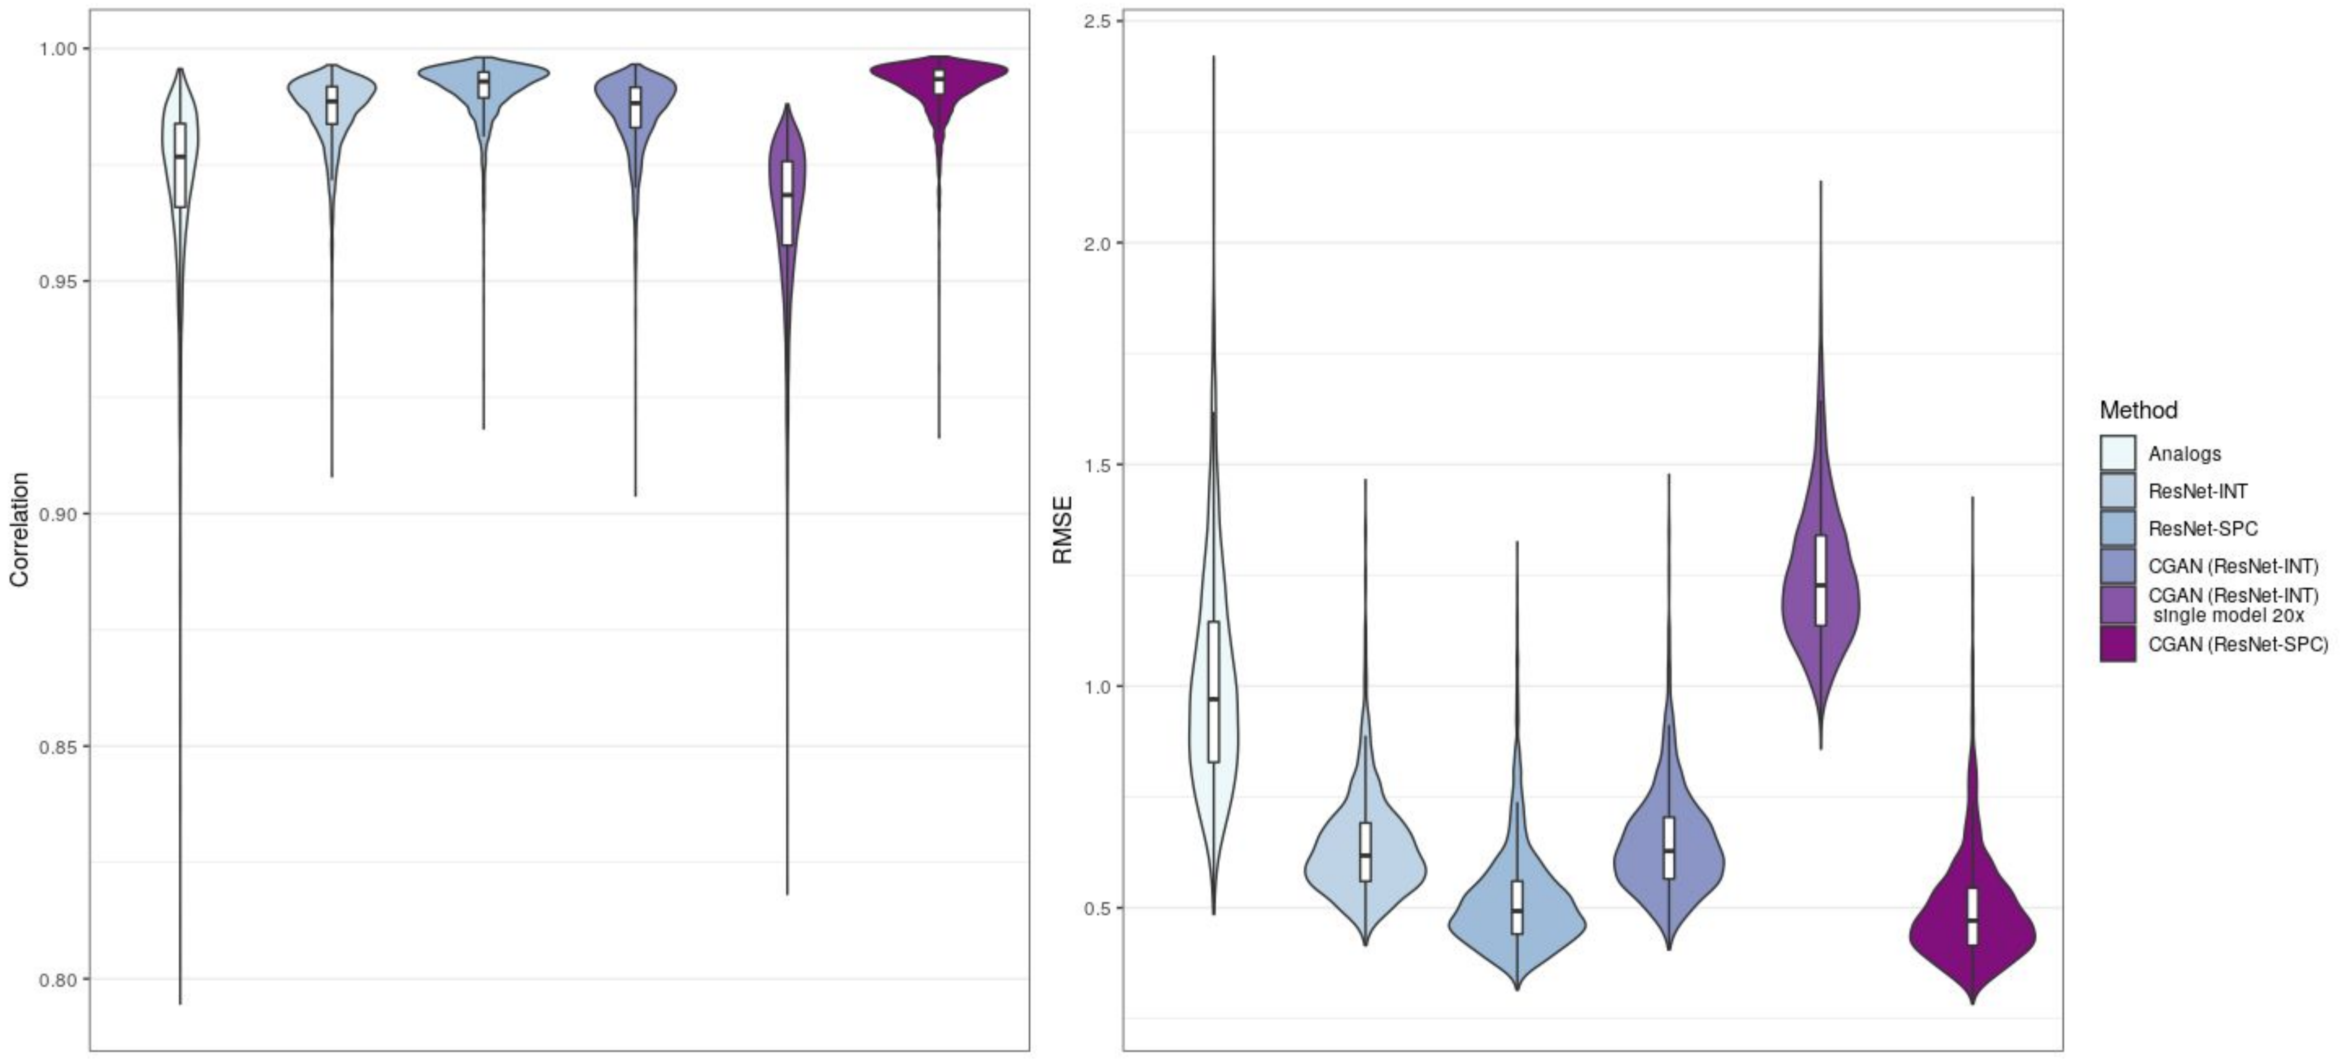

Visual comparison of SEAS5 downscaled products

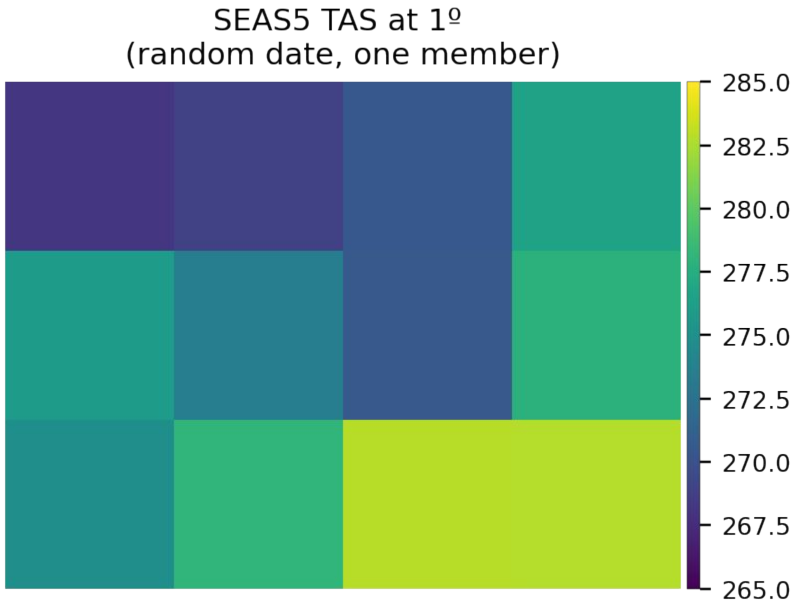

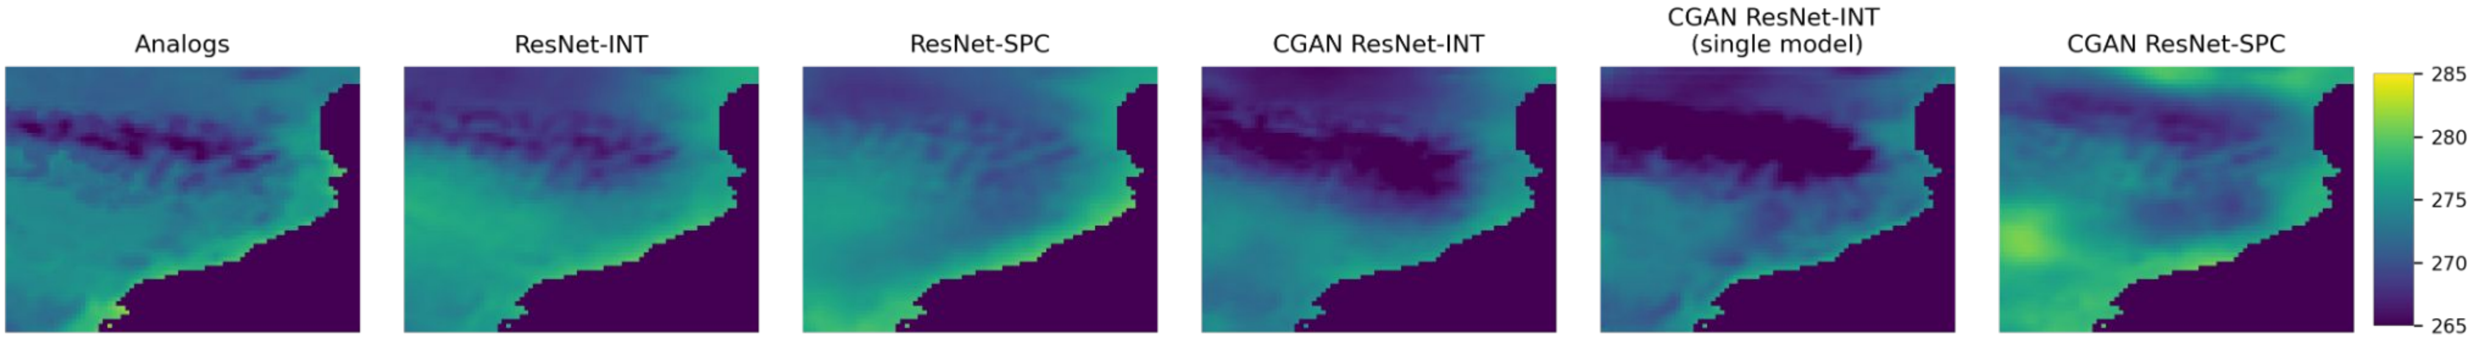

Downscaled SEAS5 verification

Start dates: Nov
Lead time: 1 month
Valid period: DJF

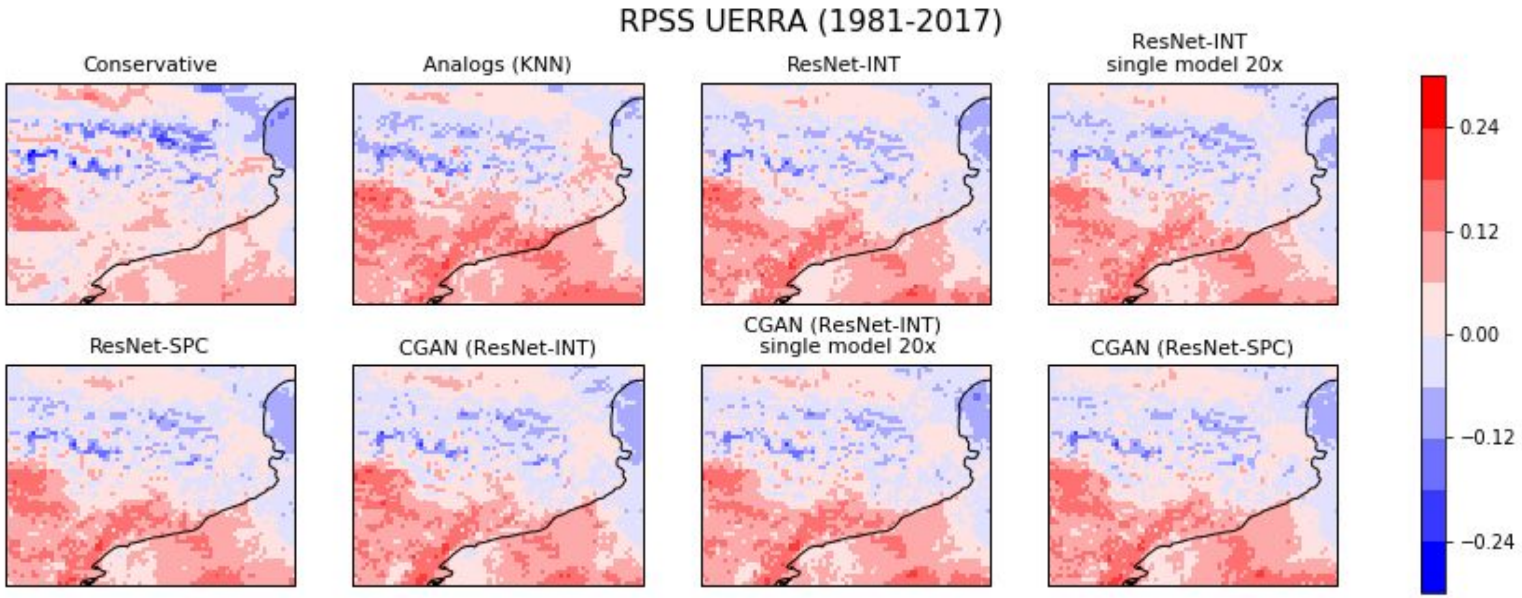

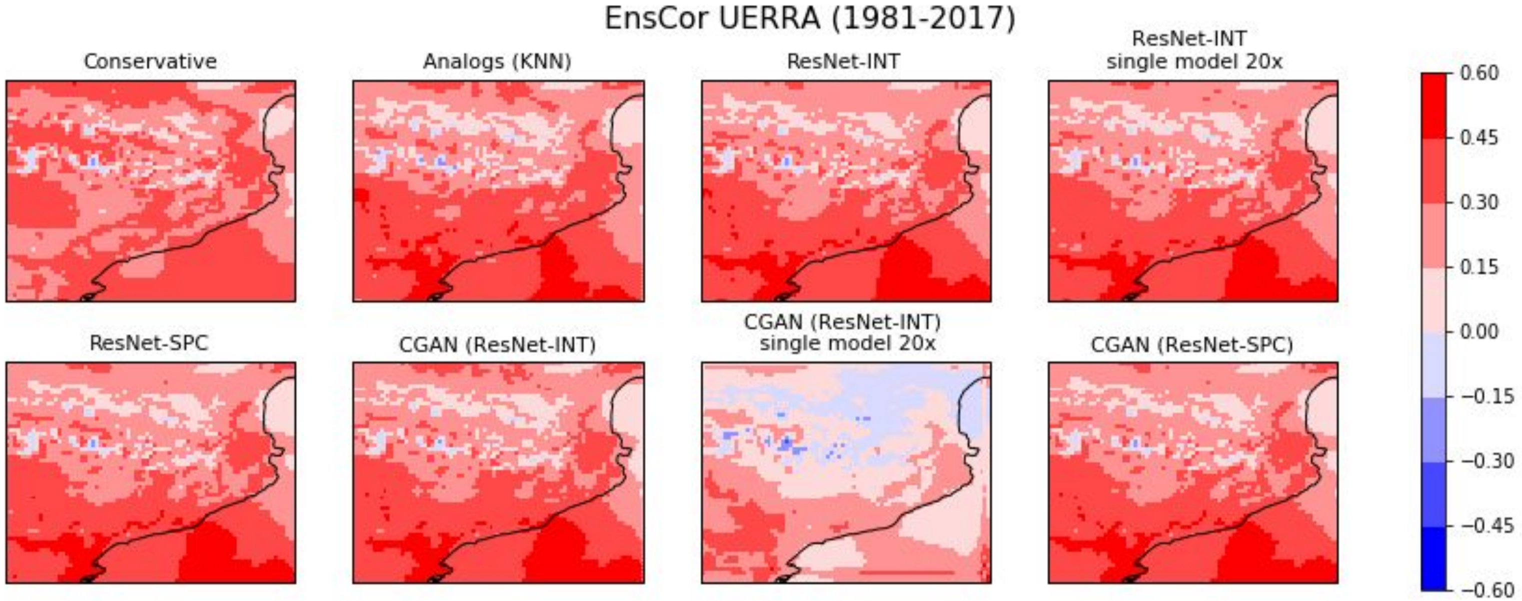

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

Conclusions and final thoughts

- These promising results are part of a work in progress (paper in preparation)
- A python package with all the implementations and a clear common API is planned for release
- The DL-based downscaling techniques presented here are efficient at generating high-resolution gridded fields
- While the training time varies depending on the model (from 30 min to 12 hours on a single V100 GPU), the inference takes only a few seconds and can be applied to arbitrary domain sizes
- In terms of RMSE and correlation, the conditional GAN with a ResNet-SPC generator outperforms the other approaches
- In terms of RPSS and ensemble correlation, all the models behave similarly (with the exception of a single CGAN ResNet-INT model)
    - It is important to notice that even a simple method such as analogs-KNN provides good results
    - improving the skill and correcting the bias of the seasonal prediction require fine-tuned loss functions and training procedures. This is something we are currently exploring

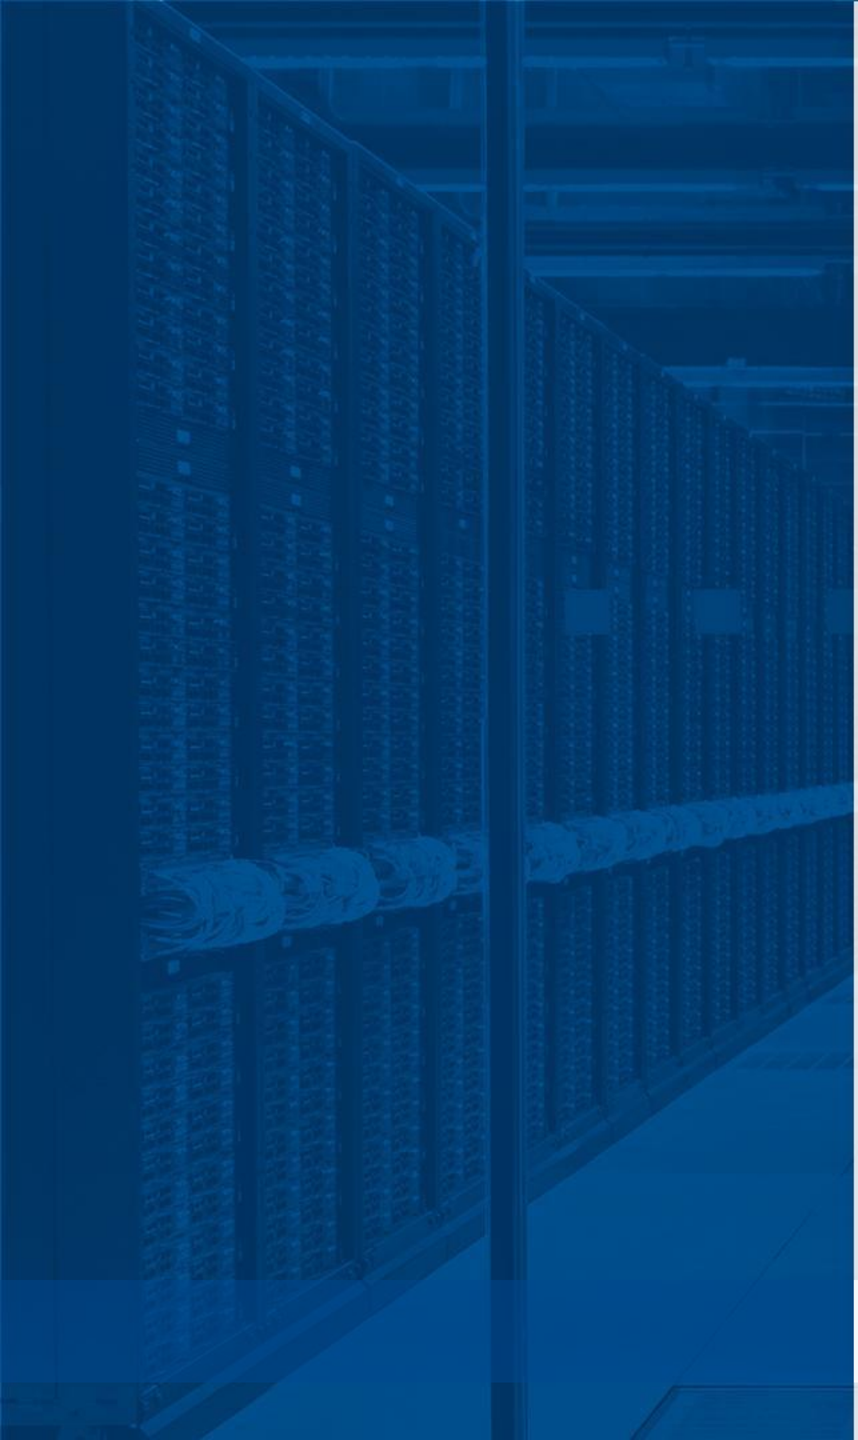

Barcelona
Supercomputing
Center
Centro Nacional de Supercomputación

Gracias

Acknowledgements:

The research leading to these results has received funding from the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement H2020-MSCA-COFUND-2016-754433 and the EU H2020 Framework Programme under grant agreements n° GA 823988 (ESiWACE-2), GA 869575 (FOCUS-Africa) and GA 869565 (VitiGEOSS).

carlos.gomez@bsc.es