Deep learning for high-resolution climate projections: a Latent Diffusion Model emulating dynamical downscaling of precipitation and temperature

Elena Tomasi (1), Giacomo Tomezzoli (1), Gabriele Franch (1), Sandro Calmanti (2), Marco Cristoforetti (1)
(1) Data Science for Industry and Physics, Fondazione Bruno Kessler, Trento - Italy, (2) ENEA, Rome - Italy

Motivation

Traditional ways to produce high-res data:

- Dynamical models → physically consistent but expensive to run
- Statistical models → 'cheap' to run but less reliable

Can DL learn to mimic dynamical models and provide reliable, fast, and less resource-demanding high- resolution climate projections?

Methodology and Reference Datasets

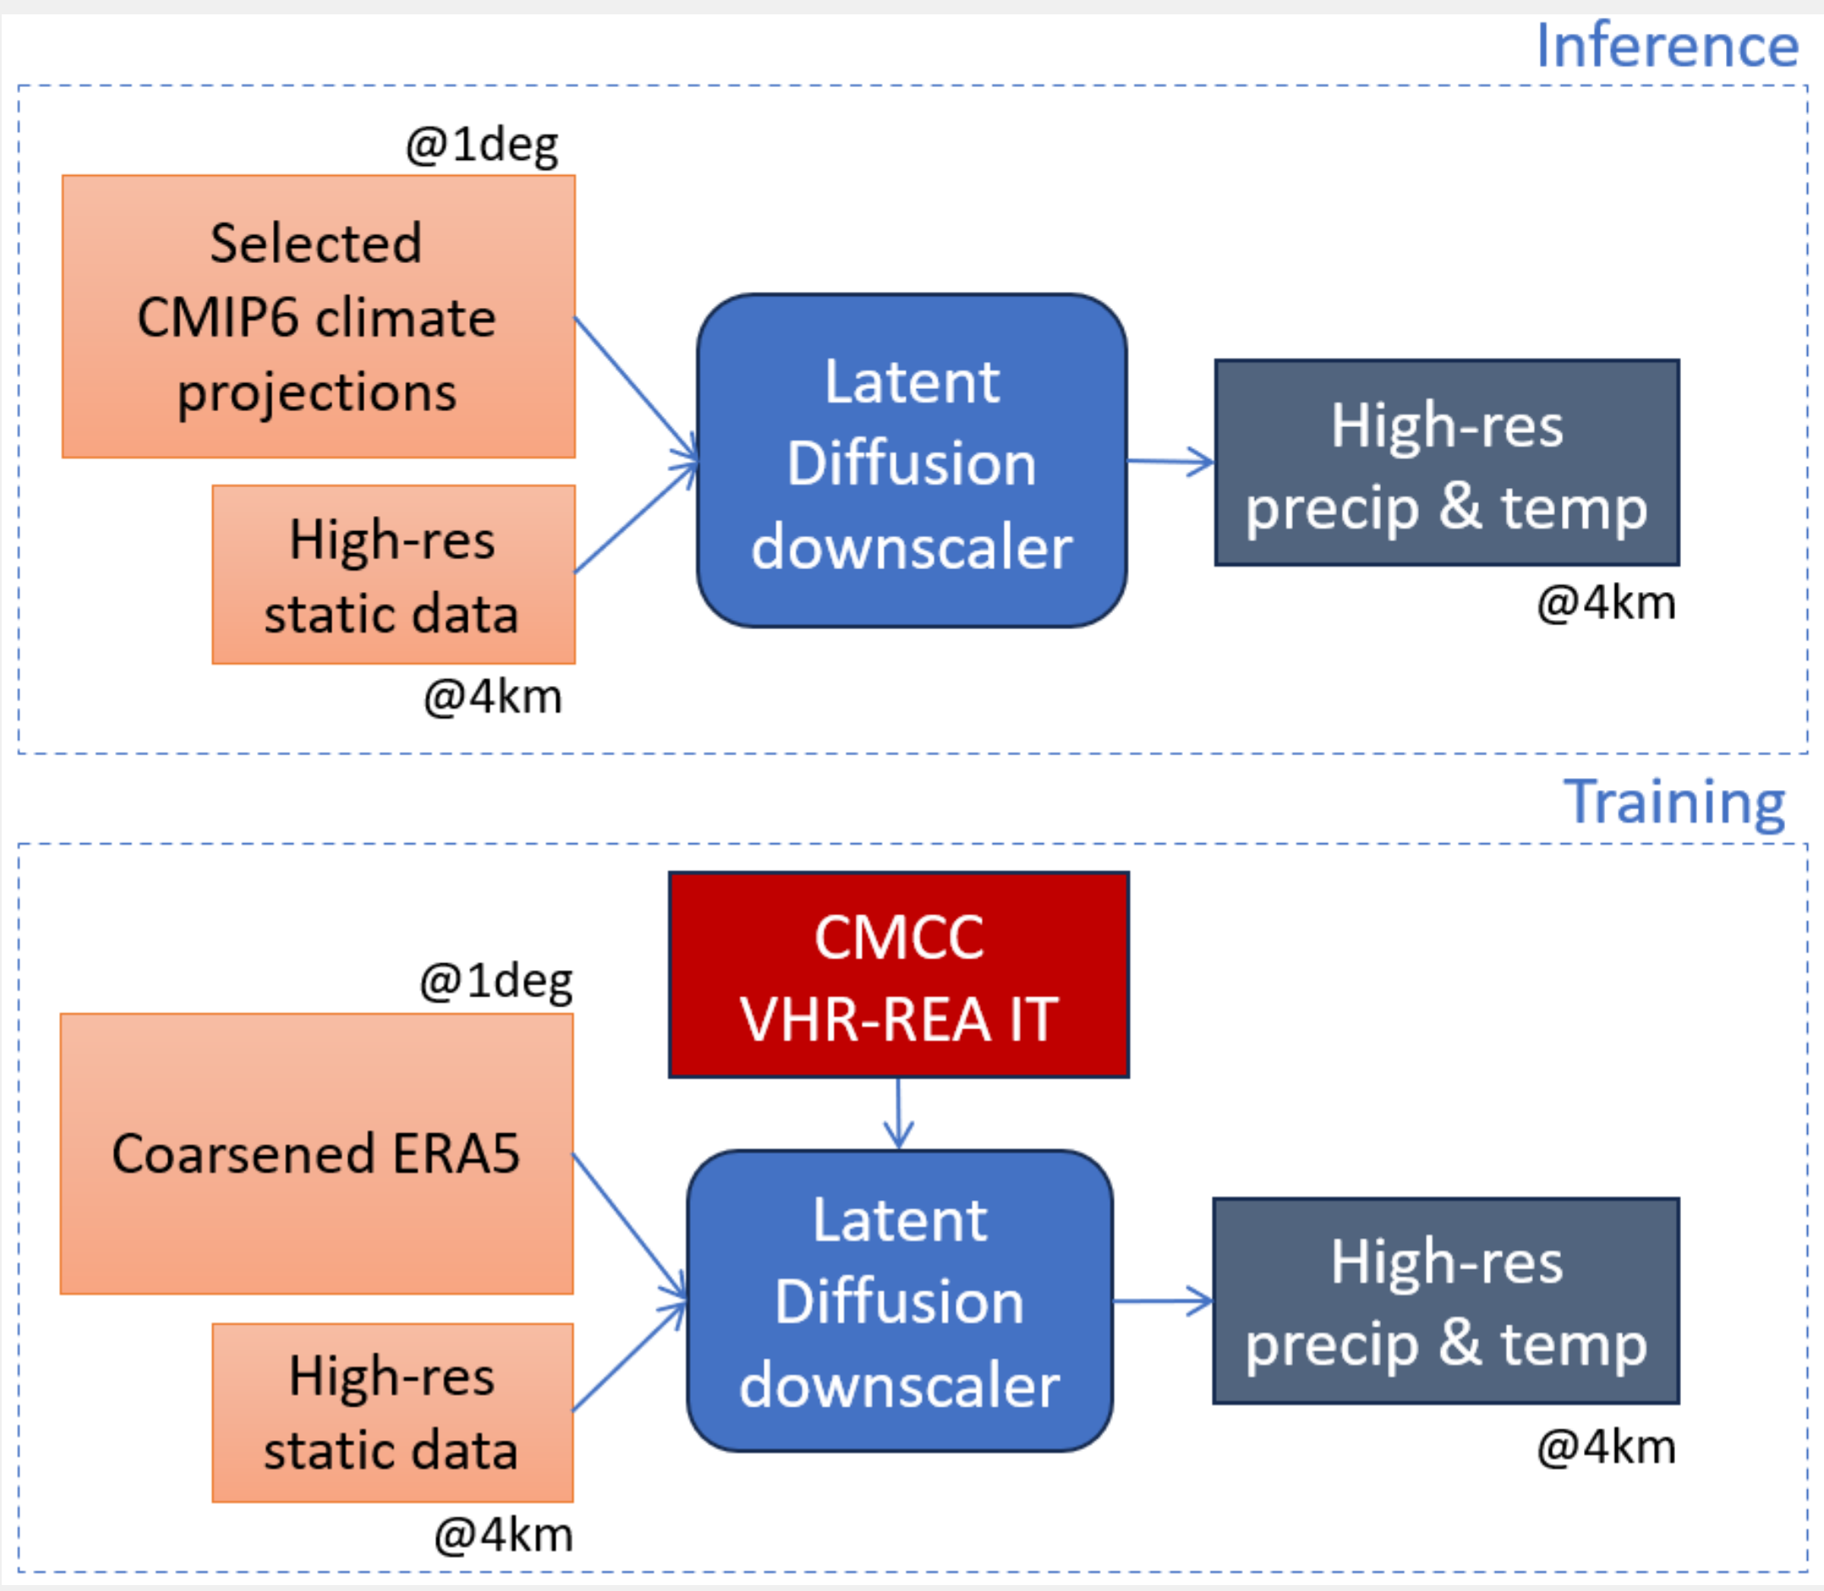

Input low-res data: CMIP6 GCMs/ERA5[1]

- Resolution: ~1deg → ~64km @Italy
- Time resolution: 6h-inst, 6h-avg, 3h-avg
- 9 predictors @different heights → 19 input channels
- For CMIP6: historical + 4 emission scenarios

Target high-res data: VHR-REA_IT[2]

- Resolution: ~2.2km over Italy
- Time resolution: 6h
- 3 target variables (t_{min}, t_{max}, precip)
- Dynamical downscaling of ERA5 with COSMO-CLM

42 years for training

[1] Hersbach et al., https://doi.org/10.1002/qj.3803.

[2] Raffa et al., https://doi.org/10.3390/data6080088

Deep Generative Models

Latent Diffusion architecture: LDM_res

Stable diffusion derived architecture. The network allows resolution independent inference.

300M trainable parameters
20 denoising steps in inference

Three main components:

- **Conv VAE**: projection to/from latent space of target variables
- **CONDITIONER**: embedding of low-res and high-res static data
- **DENOISER**: latent space diffusion process
  - 2D UNET architecture: ResNET+AFNO+Conv blocks

Leveraging a residual approach

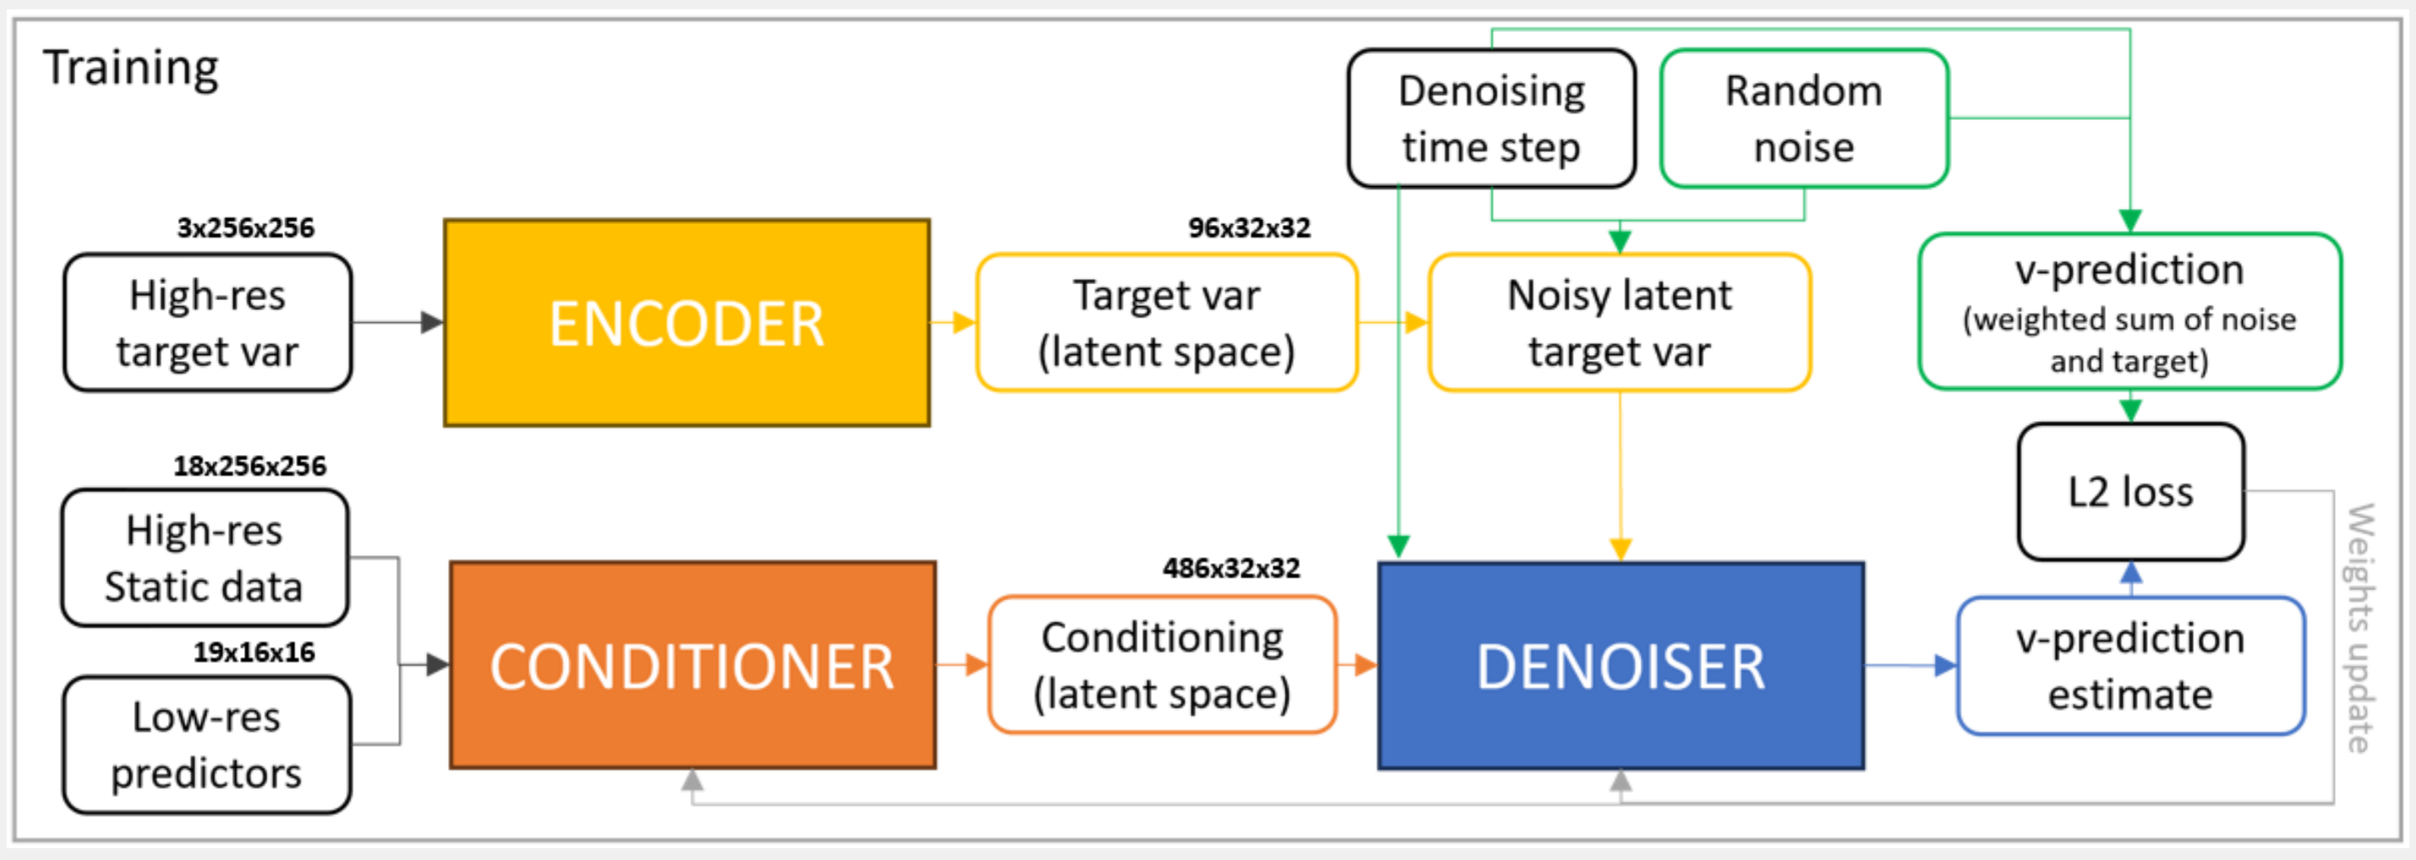

Tomasi et al 2025

Results on the test set

Snapshots of downscaled data

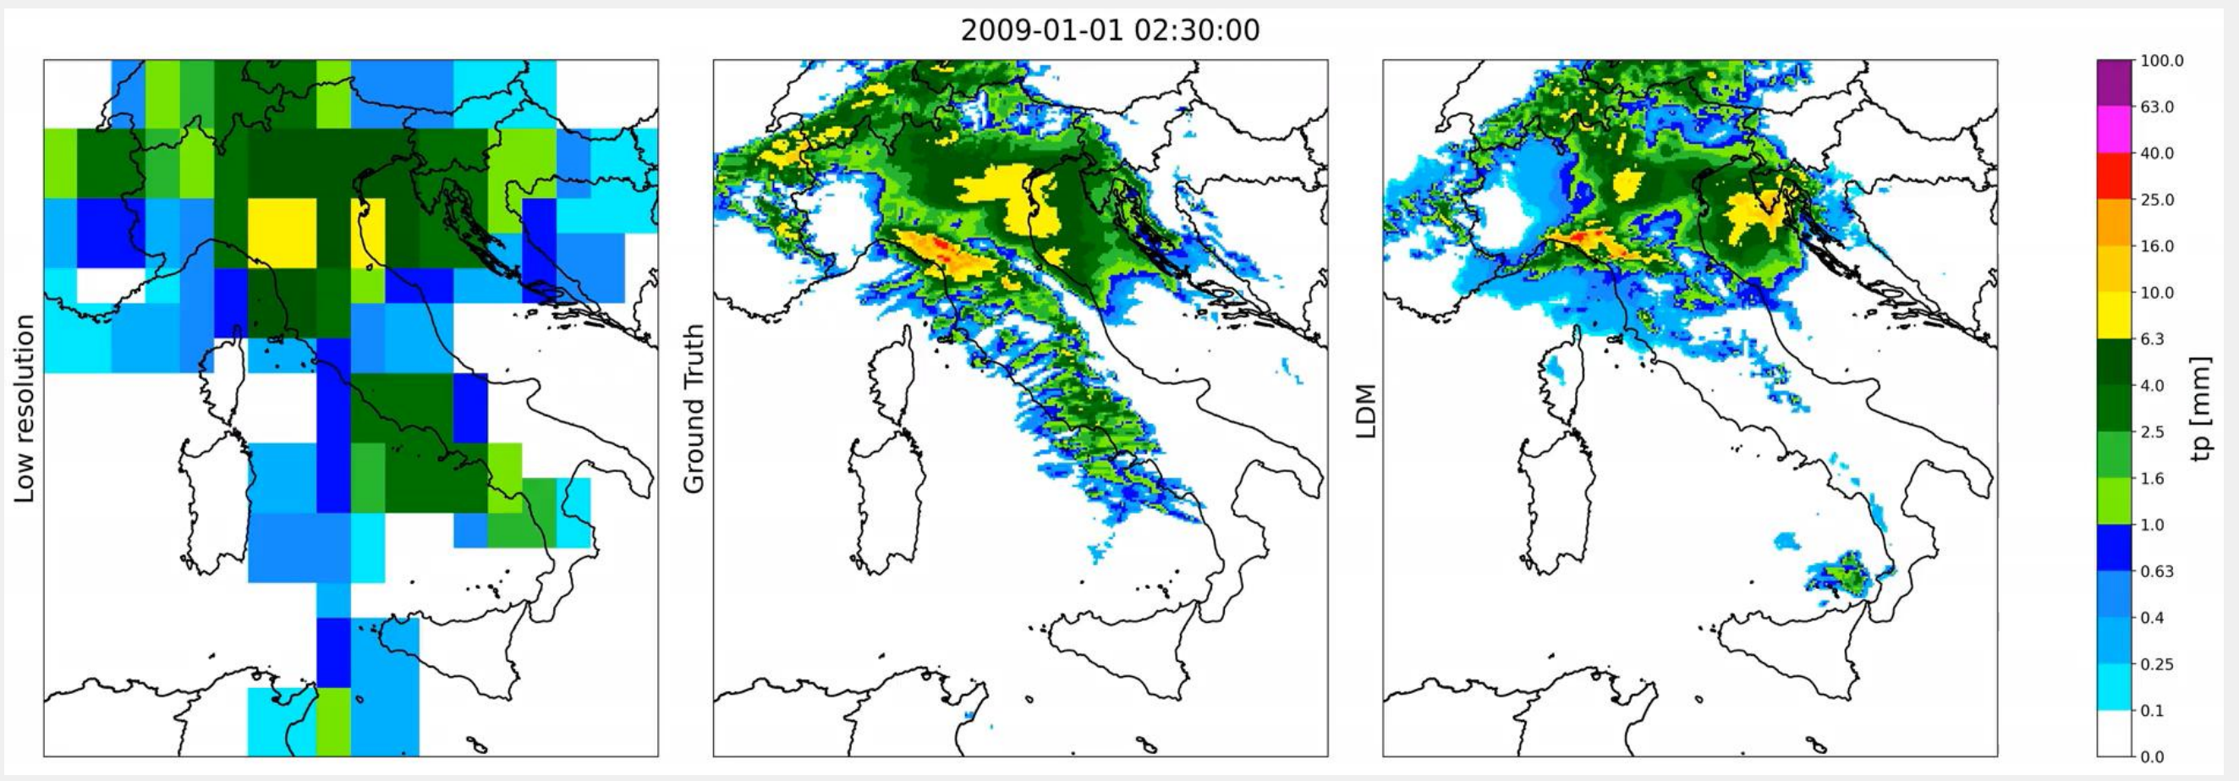

Fig 1: Snapshots of precipitation data: from the left, ERA5 low-res data (NOT used as input), target VHR-REA-IT high-res data, predicted high-res data with the Latent Diffusion Model.

Verification metrics

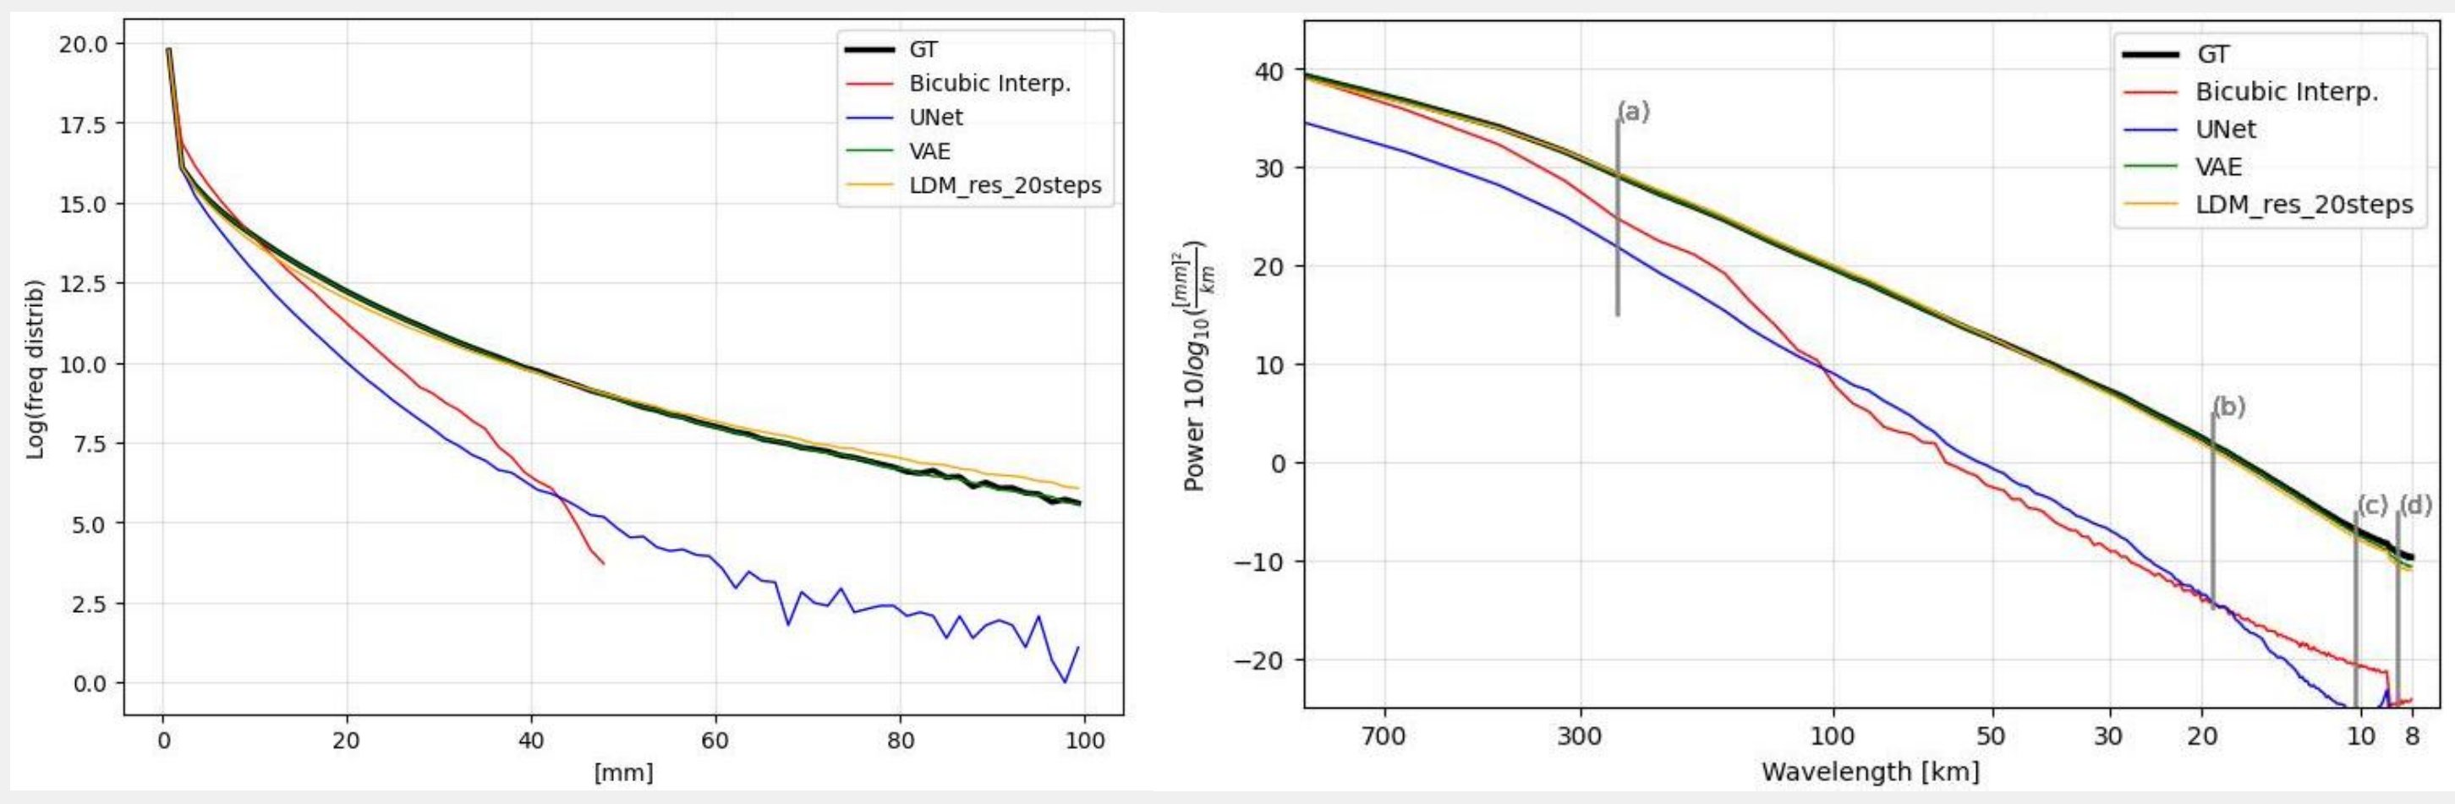

Fig 2: Verification metrics for the precipitation: on the left, frequency distribution; on the right, radially averaged power spectra density

Verification of applicability to future climate – ongoing work

Methodology

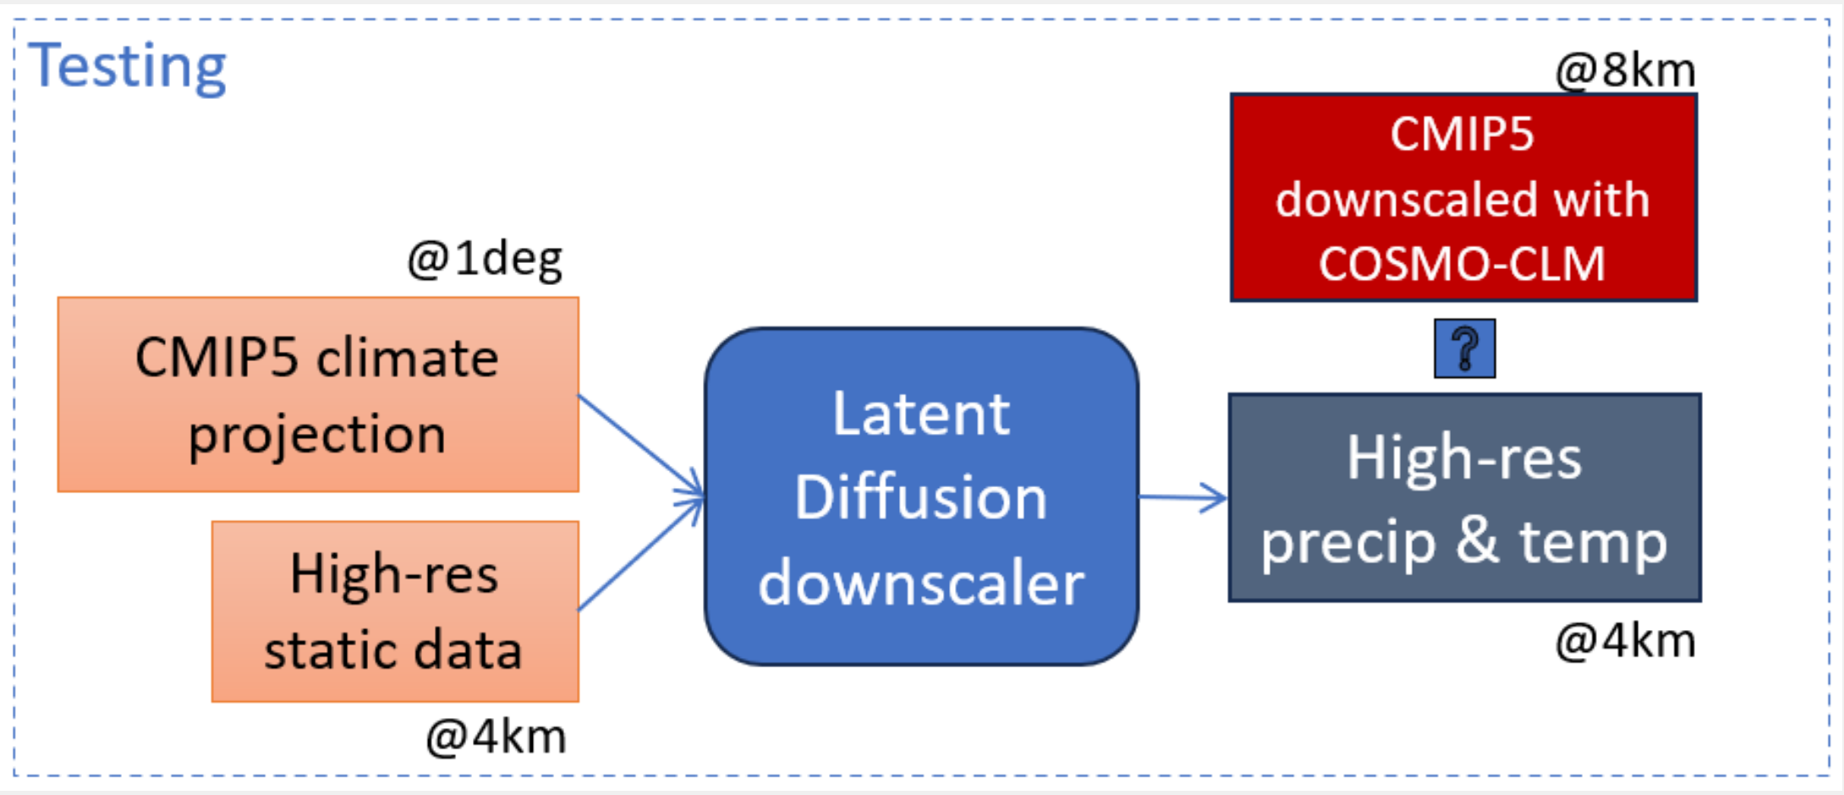

Reference high-res data: COSMO-8km[3]

- Dynamical downscaling of CMIP5(CMCC-CM historical and rcp4.5) with COSMO-CLM
- Resolution: ~8km over Italy
- Time resolution: daily
- 3 target variables (t_{min}, t_{max}, precip)

Test run:
historical = **1985-2005**,
RCP4.5 = **2098-2099**

Resampling of CMIP5 downscaled with LDM_res:

- Temporally: daily t_{min}= min(6-h t_{min}), daily t_{max}= max(6-h t_{max}), daily precip=cum(6-h precip)
- Spatially: average from 4 to 8 km

[3] Bucchignani et al., https://doi.org/10.1002/joc.4379

Verification metrics

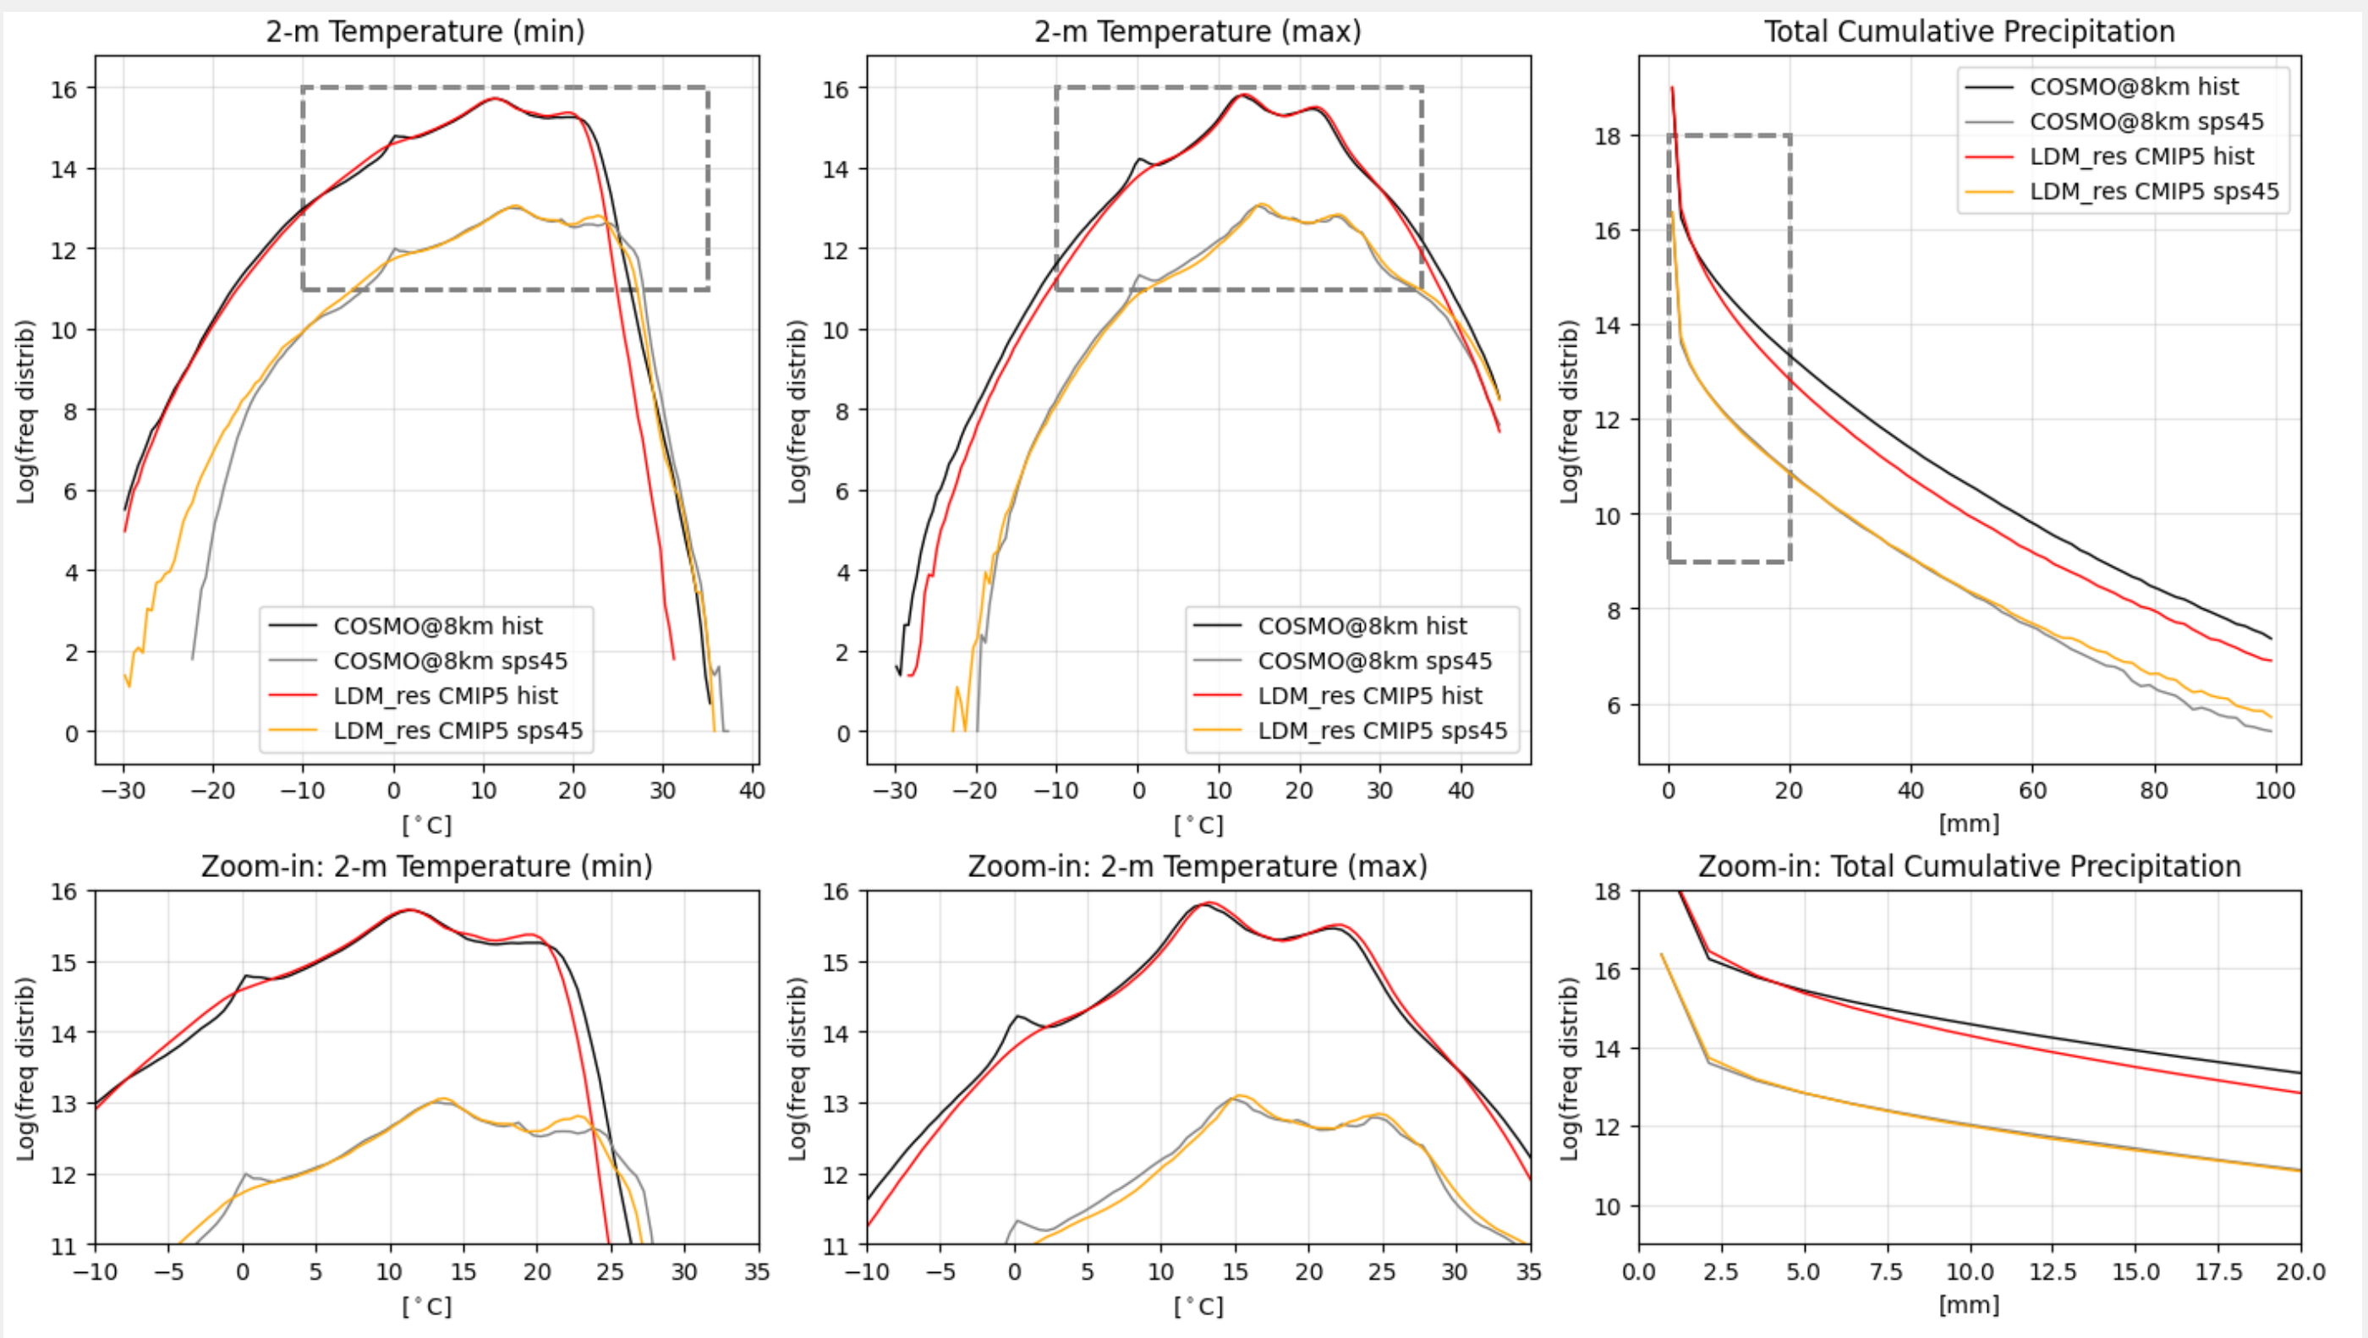

- LDM_res is able to reconstruct COSMO@8km results both on the historical period and the future scenario
- Deviations from expected behavior of LDM_res:
  - future minima of t_{min}
  - historical precipitation

Fig 3: Verification metrics - frequency distributions: from the left, t_{min}, t_{max} and precipitation

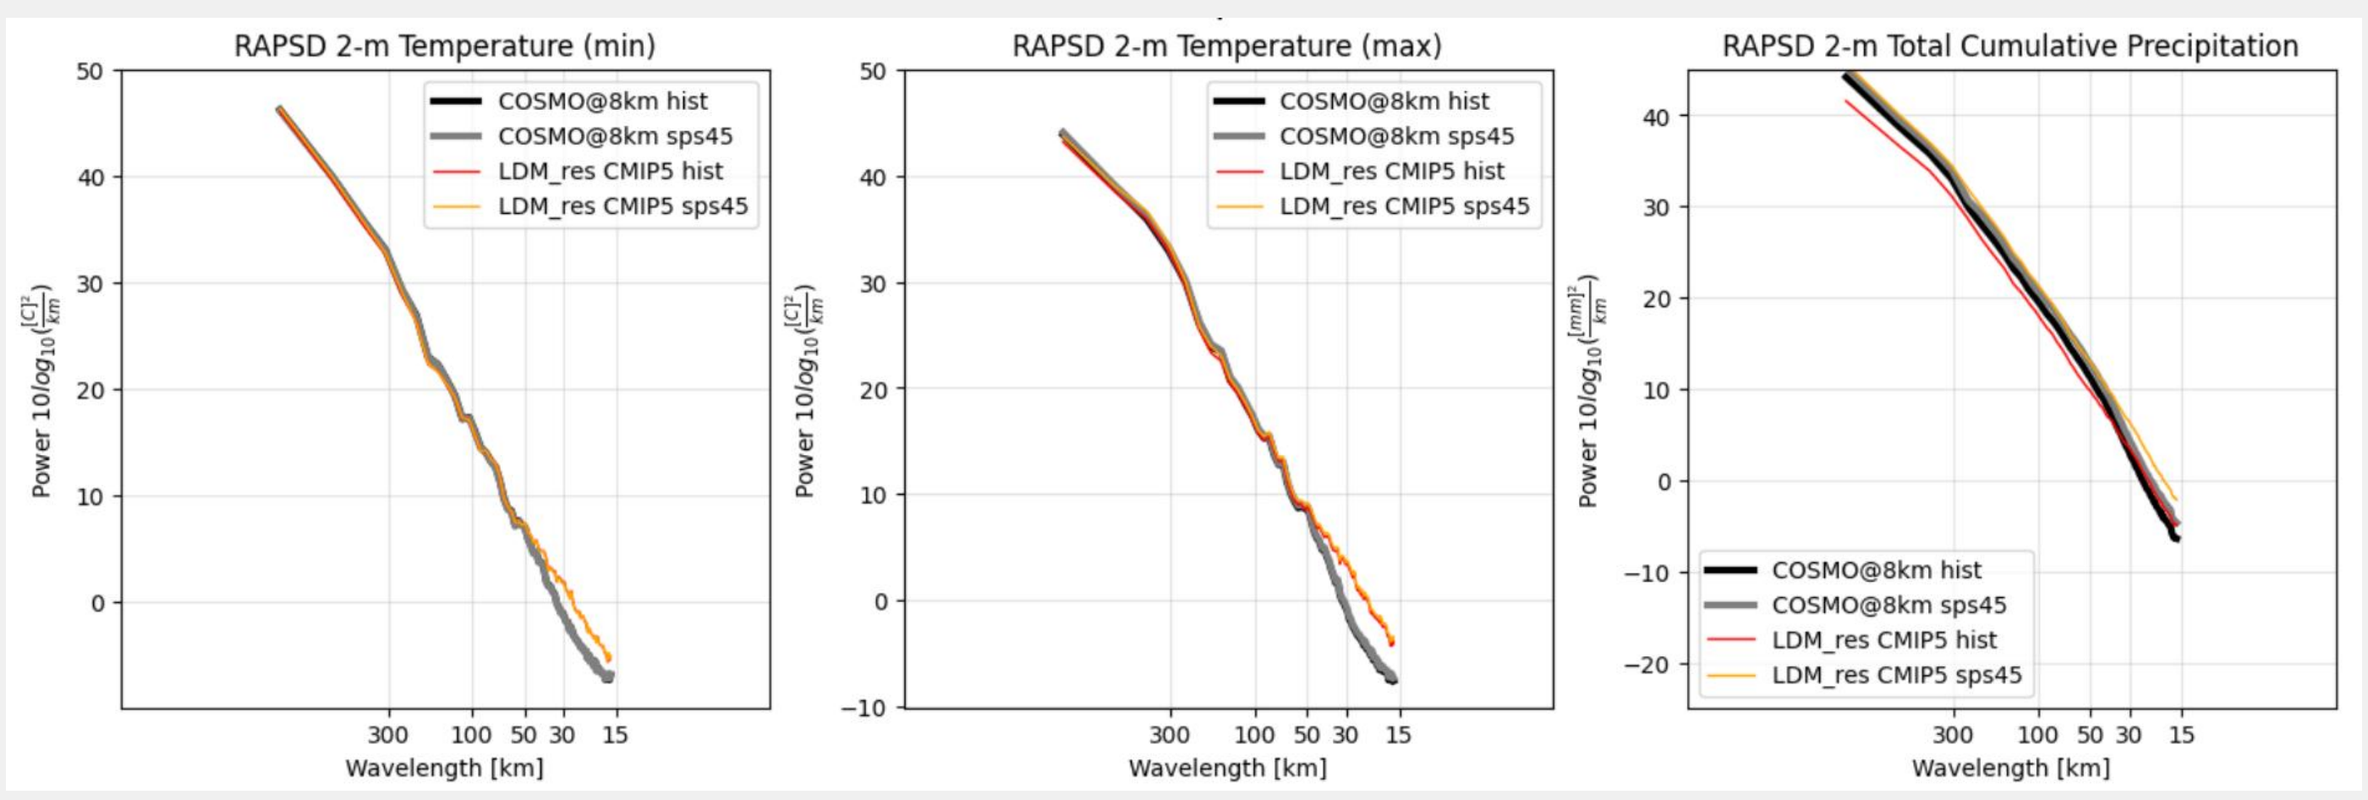

- RAPSD distributions do not change from historical to future scenario: LDM_res provides a consistent reconstruction (8km vs 4km)

Fig 3: Verification metrics – Radially Averaged Power Spectra Density: from the left, t_{min}, t_{max} and precipitation

Preliminary conclusions

- LDM_res shows great performance also applied for precipitation and bigger (25x) downscaling factors in terms of frequency distributions, energy spectra, spatial errors (not shown) due to: the diffusion process and the residual approach
- Future climate reconstruction is reasonable against its numerical counterpart, but further analyses are needed and are currently undergoing

Future Work

- Evaluate performance over target extreme events and perform a sensitivity analysis to the predictors
- Explore the ensemble potential of the diffusion model

contacts: Elena Tomasi - eltomasi@fbk.eu | Gabriele Franch - franch@fbk.eu

ESA-ECMWF WORKSHOP
Machine Learning for Earth System Observation and Prediction