

Manuscript title: Deep Learning for Super-Resolution of Mediterranean Sea Surface Temperature Fields

Authored by: Claudia Fanelli, Daniele Ciani, Andrea Pisano, and Bruno Buongiorno Nardelli

Manuscript ID: <https://doi.org/10.5194/egusphere-2024-455>

REVIEWER #1

Reviewer: The article presents a study on improving the resolution of Sea Surface Temperature (SST) fields in the Mediterranean Sea using deep learning models, specifically a dilated convolutional multi-scale learning network. This approach allows for better capture of small scale features and gradients in SST data, overcoming limitations of traditional satellite-based measurements and interpolation methods. The study demonstrates significant improvements in the accuracy and resolution of SST reconstructions, highlighting the potential of deep learning in enhancing oceanographic data analysis and climate research. But the experiment needs some work.

Response: The authors would like to thank the anonymous reviewer for their interest in reading our manuscript. We think that the revised version of the manuscript has been enhanced significantly as a result of the reviewers' feedback. Please find the detailed responses below with the reference to the modifications carried out in the re-submitted files (highlighted in yellow).

Reviewer: Incorporate additional independent datasets for validating the improved SST fields, ensuring the model's robustness across various conditions and regions within the Mediterranean Sea.

Response: We agree with the reviewer that a larger test dataset was needed to provide a more robust assessment. Therefore, we built a new independent dataset using SST fields over one year (2021) and performed a new test. The main results of this extended test are described in new Section 3.2.

Reviewer: Compare the performance of the proposed deep learning model against existing other deep learning super-resolution models, such as GAN series, providing a comprehensive analysis of its advantages and limitations.

Response: Following the reviewer suggestion, we have included in the revised text a comparison between the results obtained with the dADR model and other deep baseline computer vision methods. The outcomes are now presented in Table 1. We are currently working on the application of Generative Adversarial Networks to the problem, but our study is still in a preliminary phase and does still require a substantial effort to be ready for publication. As such, considering also some subtle theoretical aspects related to the application of generative approaches to observational data, we prefer to dedicate an entire new scientific article to discuss/present our findings, as now discussed in lines 321-328 of the revised version of the manuscript.

Reviewer: Conduct a sensitivity analysis to understand the impact of different parameters within the dilated convolutional multi-scale learning network, optimizing the model's performance.

Response: Thanks to the reviewer’s comment, we realized that we did not explain why we had chosen a specific configuration of the network, and we also understood that at least one additional test was needed. In fact, we did not develop a new architecture but exploited the one developed by Buongiorno Nardelli et al. (2022) (<https://doi.org/10.3390/rs14051159>), and also aimed to reduce computational costs. As such, we originally relied on the choices made in that specific work . In the revised paper, however, we now describe one additional test carried out by modifying the network depth, i.e. reducing the number Multiscale Adaptive Residual Blocks (called dADRSR/2 in the revised version of the manuscript) and compare that also with the results obtained by other deep learning methods (i.e., the EDSR and the ADR networks). The assessment of the various network configurations considered are now presented in Table 1.

Reviewer: **Could the article be enriched by including a paragraph discussing how high-resolution SST fields can be incorporated into regional climate models to improve the accuracy of climate projections in the Mediterranean region?**

Response: To our knowledge, using observation-based SST data is not a standard procedure to improve the accuracy of climate projections, which generally do not rely on assimilation strategies. High resolution SST fields are indeed used for the improvement of operational ocean model re-analyses (e.g. within Copernicus Marine Service), which tackle a very different problem with respect to climate projections. Even if there is evidence of the important role of small structures of SST on different processes that are clearly relevant in modulating climate (e.g. atmosphere-ocean interactions, see Renault, et al. “Modulation of the Oceanic Mesoscale Activity by the Mesoscale Thermal Feedback to the Atmosphere”. *J. Phys. Oceanogr.* **53**, 1651–1667, 2023), we feel that introducing a specific discussion on this topic would be distracting and not relevant for the readers.

Manuscript title: Deep Learning for Super-Resolution of Mediterranean Sea Surface Temperature Fields

Authored by: Claudia Fanelli, Daniele Ciani, Andrea Pisano, and Bruno Buongiorno Nardelli

Manuscript ID: <https://doi.org/10.5194/egusphere-2024-455>

REVIEWER #2

Reviewer: This study aims at reconstructing small scales SST from a low resolution SST field and provides a gap free L4 dataset resolving scales up to 5 km using deep learning algorithms. The super Resolution Convolution Network is learning using an ensemble of low resolution and high resolution SST images in the Mediterranean Sea. Results are shown for one snapshot and analyses are shown for cloud-free areas. Compared to the first guess at low resolution, the SST field has indeed been improved but this study needs further investigation and discussion of the results.

Response: The authors would like to thank the anonymous reviewer for their interest in reading our manuscript. We feel that the revised version of the manuscript is significantly enhanced as a result of the reviewer's feedback. Please find the detailed responses below with the reference to the modifications carried out in the re-submitted files (highlighted in yellow).

Reviewer: Comparison with other High resolution L4 products are required as the first guess is at very low resolution and does not reflect what is currently available (such as MUR product from JPL for example). Some SST products perform well when the cloud coverage is small but have a resolved spatial scale that varies a lot with the cloud cover. Comparison of the different SST products on a cloud-free event and on a cloudy day will show the potential of this new SST product much more homogeneous in time.

Response: First of all, thanks to the reviewer's comment, we realized that our objective and the super-resolution workflow within the CNR processing chain could be better introduced. In fact, CNR 1 km grid processing chain is based on a two step algorithm. The first step provides a low resolution field that is needed to obtain robust SST estimates below large cloudy areas, that is successively used as the background field for the application of a more local optimal interpolation algorithm. This background/first guess field is generally quite smooth even in cloud-free areas, and our objective is to improve its effective resolution and the intensity of the related gradients (it is thus an intermediate step in the processing chain). As such, the output of the network is not a new product itself, as clarified in Figure 1 of the revised version of the manuscript.

Then, we agree with the reviewer that it is quite useful and relevant to intercompare our original first-guess and the super-resolved one also with other products. Consequently, we performed a power spectral density analysis on different SST products available over the Mediterranean Sea (see Figure 9 and lines 278-291 of the revised version of the manuscript). The products used are the L4 NRT HR and the L4 NRT UHR provided by CNR for the Copernicus Marine Service, the GLOBAL OCEAN OSTIA product developed by the UK MET OFFICE, the Multi-scale Ultra-high Resolution (MUR) product provided by the NASA-JPL and the super-resolved SST field obtained by the application of the dADR-SR network developed in this work. We focus on cloudy areas and not on cloud-free events as this is where super-resolution is expected to provide an improvement.

Reviewer: From the snapshot with the SST gradient, One wonders if the method enables the reconstruction of smaller scales or is more of a gradient enhancement. It would be helpful to comment on that and illustrate if submesoscales are really generated in this SST reconstruction.

Response: The reviewer is absolutely right: this specific algorithm is mostly enhancing the gradients and it is in fact more conservative in the reconstruction of small scale features. Indeed, other methodologies (those based on Generative Adversarial Networks) have the potential to re-create small scale eddies/filaments which would be compatible with the learned features. However, single-image end-to-end approaches might not be the best choice to obtain accurate and dynamically consistent reconstructions (though surely looking quite “realistic”). Indeed, we are presently carrying out tests with GAN, but considering the time-frame and the amount of work still needed to get to (at least partly) conclusive results, those will necessarily need to be described in a successive paper.

Reviewer: The added value of this SST product is not clear from Table 1 and the spectrum figure (Fig.6). I would recommend adding other SST high resolution L4 products as discussed in point 4. Regarding the SSIM index, it may be interesting to also discuss the details of the decomposition between contrast / structures and luminance if the results prove to be relevant (the contrast and structure are expected to be much more improved than the luminance).

Response: As also explained above, to demonstrate how the super-resolved images reach an enhanced spatial resolution reconstructing SST fields under cloudy conditions, we performed a power spectral density analysis on different SST products available over the Mediterranean Sea. This is now fully documented in Figure 9 and lines 278-291 of the revised version of the manuscript. Given the fact that the SSIM index is not a common metric used in oceanography but it is mainly used in the computer vision field, we feel that is out of the goal of this paper to discuss the details of the single parts of this measure.

Reviewer: Finally, to really describe the effective resolution, the study of the ratio between the spectral content of the reconstructed data and the truth is more relevant than mere spectrum (as detailed in Ballarotta, et al 2019.). It should be included in a further analyses of this SST product to comment on the effective resolution of the SST product.

NB: Ballarota et al 2019: Ballarotta, M., Ubelmann, C., Pujol, M.-I., Taburet, G., Fournier, F., Legeais, J.-F., Faugère, Y., Delepouille, A., Chelton, D., Dibarboure, G., and Picot, N.: On the resolutions of ocean altimetry maps, Ocean Sci., 15, 1091–1109, <https://doi.org/10.5194/os-15-1091-2019>, 2019.

Response: The reviewer is absolutely right.. We now added the computation and the analysis of the PSD ratio as defined in Ballarotta et al. (2019) in lines 273-277 and Figure 8 of the revised version of the manuscript.

Manuscript title: Deep Learning for Super-Resolution of Mediterranean Sea Surface Temperature Fields

Authored by: Claudia Fanelli, Daniele Ciani, Andrea Pisano, and Bruno Buongiorno Nardelli

Manuscript ID: <https://doi.org/10.5194/egusphere-2024-455>

REVIEWER #3

Reviewer: In their work the authors are tackling an important problem in geophysics, namely the reconstruction of high resolution fields of Sea Surface Temperature from partial high-resolution observations. Super-resolution is an ill posed problem, given that identical low-resolution fields can correspond to different high-resolution fields. The authors adapt their previous work on ADT to SST field reconstructions, as well as their well established knowledge of SST field reconstruction, showing improvements over the Mediterranean basin, within the confines of their experiment.

I really appreciated many parts of the article, notably the varying metrics and case studies to evaluate the quality of the reconstructions.

As it stands, I have some major and minor criticisms for the article that, should the authors address, would make for a significant contribution to the community.

Response: The authors would like to thank the anonymous reviewer for their interest in reading our manuscript. We think that the manuscript has been significantly improved by taking into account the reviewer's feedback. Please find the detailed responses below with the reference to the revisions appearing in the re-submitted files (highlighted in yellow).

Reviewer: The validation process is prone to data leakage. 4 days out of a year of data were omitted, but there is no mention of removing some days before or some days after in order to prevent data leakage. The physical reasoning of this is absent. Are the structures that decorelated after one day given the removal of the 200km smoothed field?

In general the L119 statement: “The test dataset is finally selected separating the 15% of the tiles available after 120 the preprocessing, chosen in order to be able to reconstruct the full geographical coverage of four days which are representative of different seasons.” requires clarification. I read it as 4 individual days, one in each season. It could be understood as patches covering the whole area, spread out over each season. I would expect to have more of a cross-validation approach given the one year dataset limitation.

Response: Thank you for pointing out the lack of clarity and previous limitations regarding the choices made to construct the test dataset. Indeed, the initial test was carried out selecting all the tiles covering four individual days (one in each season), as clarified in line 132 of the revised version of the manuscript. However, we do agree about the limitations of that initial test dataset and thus set up a much more extended test including one full year of totally independent data. The new results are presented in the new Section 3.2.

Reviewer: Another major concern is the input. The input, presently, is only the first guess (removing a sliding window), and the information coming from the L3 satellite product is not used as a complimentary input. Why did the authors deprived themselves from potential additional input such as multiple time steps and L3 products? Other works (such as Archambault et al, Martin et al) in SSH fields have training procedures where some of the satellite information is omitted from the target in order to validate the approach.

Response: Thanks to the reviewer's r comment, we realized that we needed to better introduce the objectives and processing steps involved in our workflow . In fact, we are not dealing here with an algorithm designed for gappy-fields interpolation, but rather with an algorithm designed to improve one specific step in our processing chain, namely improving the resolution of the low-resolution gap-free field used as the background for the 1 km optimal interpolation implemented in our chain. As such, unlike the classical interpolation methods, in which it is important and useful to use ground-truth data and multiple time steps as input when available, the goal here is effectively only to learn how to reconstruct a single image at very high resolution starting from a low resolution one. L3 data are indeed e used as input in the final interpolation step, though, and are of course also used as target during the training phase. We have clarified all this in the Introduction of the revised version of the manuscript (lines 87-91).

Reviewer: There is no mention of how the total field reconstruction over the whole Mediterranean sea is output. If the image was made by recomposing a sliding window reconstruction that should be mentioned. Given that the network learns filters, one could conceivably apply them on the whole image, though I expect the attention layers to pose an issue.

Response: Following from previous comment/response, we hope to have clarified this point adding a more detailed explanation at the end of Section 2, in lines 208-212 of the revised version of the manuscript.

Reviewer: The choice of architecture, while documented, is not justified. Were the hyper-parameters optimized? Other architectures evaluated? The authors mention a lot of competing methods, but do not compare their architecture to them. (DINEOF, DINCAE, to cite but two) Are the computational and expertise cost justified versus other methods? The results seem to indicate a 0.02°C improvement; is the architecture stable through different initializations?

Response: As clarified in the revised text, we are not dealing here with a novel interpolation algorithm, but rather on a single-image end-to-end algorithm, that is applied to improve our background field used in a two-step optimal interpolation algorithm. As such, comparisons with different interpolation algorithms as DINCAE/DINEOF would be out of scope. On the other hand, we realized that we did not explain why we have chosen a specific configuration of the network, and we also understood that at least one additional test was needed. In fact, we did not develop a new architecture but we exploited the one developed by Buongiorno Nardelli et al. (2022) (<https://doi.org/10.3390/rs14051159>), and also aimed to reduce computational costs. As such, we originally relied on the choices made in that specific work. In the revised paper, however, we now describe one additional test carried out by modifying the network depth, i.e. reducing the number of Multiscale Adaptive Residual Blocks (called dADRSR/2 in the revised version of the manuscript) and

compare that also with the results obtained by applying other deep learning methods (i.e., the EDSR and the ADR networks) The assessment of the various network configurations considered are now shown in Table 1.

Reviewer: The method section (2.2) seems to assume unfamiliarity with neural networks, providing intuitive explanations for basic architectural blocks, but then very quickly skims over important details of the more complicated blocks of the architecture. This part would benefit reducing the initial explanation of activation functions and CNNs (such as the interpretation of lines 59 to 61 which is intuitive but could easily not correspond to the exact explanations provided given the non-linear activations) and expanding on the reasoning of the architectural choices (the adaptive part of the ARB is not discussed, implying the rest of the architecture is non adaptive).

Response: This has now been revised considering the reviewer suggestion. We clarified the adaptive strategy used in the network in lines 72-75 and lines 161-164 of the revised version of the manuscript.

Reviewer: No mention is made to VIT and diffusion-based super resolution techniques that have become state of the art in computer vision. I can understand the daunting nature of these, but should they not be mentioned as potential further steps, at least? The latter is especially significant: the field reconstructions obtained through optimizing RMSE favor smoothness, and often do not represent physically feasible oceanic states. Graphcast for example has been abandoned in favor of Gencast for that very reason. Given that the model is in NRT, and therefore would be used for constraining operational models, it might be interesting to at least think about this. It is even more important given the non-bijective nature of the problem.

Response: We agree with the reviewer that it is worth mentioning also other techniques that have gained a lot of interest lately, which we would like to explore in the future. Following their suggestion, we added a brief paragraph in the conclusion on the topic (lines 321-320 of the revised version of the manuscript).

Reviewer: Fig.1 would benefit from locating with a bounding box or three the patches on the right hand side.

Response: Thank you for your suggestion. We modified Figure 2 of the revised version of the manuscript with an example of a pair of tiles and the corresponding position over the Mediterranean Sea.

Reviewer: 182: $\max(I)$ in denormalized space i.e. K° ? Or in the space where the large field is removed 200km? Is it computed over the patches, or the whole Mediterranean?

Response: All the errors (including the $\max(I)$ used to calculate the PSNR) are computed over the final reconstructed image over the whole Mediterranean Sea, as explained in lines 208-212 of the revised version of the manuscript.

Manuscript title: Deep Learning for Super-Resolution of Mediterranean Sea Surface Temperature Fields

Authored by: Claudia Fanelli, Daniele Ciani, Andrea Pisano, and Bruno Buongiorno Nardelli

Manuscript ID: <https://doi.org/10.5194/egusphere-2024-455>

REVIEWER #4 (Peter Cornillon)

Reviewer: This manuscript explores the use of a deep learning model to enhance the spatial resolution of L4 SST products in regions where missing data, generally associated with cloud cover, results in coarse fields obtained with objective analysis techniques. As I understand it, the model, which the authors have developed, dADR-SR, is trained with high resolution (HR, 1/16 degree) input fields and ultra high resolution (UHR 1/100 degree) target fields. It is then applied to HR fields and shown to reproduce structure at very nearly the same spatial resolution as test (1/100 degree) fields.

Response: The authors would like to thank Peter Cornillon for his interest in reading our manuscript and for his thoughtful suggestions. We think that the manuscript has been significantly improved by taking into account his feedback . Please find the detailed responses below with the reference to the revisions made in the re-submitted files (highlighted in yellow).

Reviewer: The ML model they have developed appears to perform very well for the test dataset they use but I struggled with the manuscript and I believe that it needs a fair amount of editorial work before it is ready for publication. Specifically, after reading the manuscript several times, I think that I sorted out what was done but, I must admit that I am still not sure that I have it right. I've included a figure in which I have tried to show the datasets, which I think they are using, and how these dataset relate to one another. First, there is a set of four datasets, two high resolution (1/16 degree), an L3S and an L4, and two ultra high resolution (1/100 degree) again an L3S and an L4. A subset of these datasets are used to train the dADR-SR algorithm and the HR L4 dataset is then fed into the trained model and the output is compared with an L3S UHR dataset (i.e., one built using the same algorithm as used to build the standard L3S UHR products used to train) constructed with SLSTR data from the Sentinel 3A and 3B satellites. I don't think that the SLSTR data are used in the construction of the standard products but I may have that wrong, well, I may have all of this wrong, for which I apologize. Adding to the confusion is that the authors appear to have changed the terminology they use for the datasets. In the abstract and in the Discussion section the authors refer to low resolution (LR) and high resolution (HR) datasets while in the remainder of the document they refer to high resolution (HR) and ultra high resolution (UHR) datasets. I'm guessing that LR (used in the Introduction and Discussion) is what they later refer to as HR and HR (in the Intro and Discussion) is what they later refer to as UHR.

Response: Thank you for pointing out the lack of clarity of the characterization of the different datasets used and also of the related processing chain steps. To address this point, we have added a more comprehensive description of the methodology used for each

dataset, at in lines 101-114 of the revised version of the manuscript, added the new Figure 1 and removed the acronym “LR” when referring to the first guess background field, due to the clear confusion it was creating throughout the manuscript.

Reviewer: Bottom line: I believe that a bit more description of what goes into the standard datasets that are used to train and later as input to the dADR-SR model, along with a clear description of differences, if any, between the datasets produced as input and/or evaluation for the work undertaken in this study. I also think that a figure showing the relationship of the datasets and processing steps, at a very gross level—sort of like the figure that I have attempted to put together below—would go a long way to making the manuscript easier to follow.

Response: Thank you for your suggestion. We do agree that a much more clear description of the methodology behind the creation of the datasets was necessary. We provided all these details at lines 101-114 and put a sketch describing the workflow as the new Figure 1.

Reviewer: In addition to the general concern outlined above I have made a number of editorial suggestions, which I hope will help to make the manuscript a bit easier to read. These are included in the attached manuscript either as hand-written annotations or as typed comments. Finally, I would like to apologize to the authors for the length of time that I took for this review—I had another manuscript, which I was asked to review, and which had to be completed before addressing this one as well as some family issues.

Response: We really appreciate and do thank Peter Cornilon for the attention paid to revise this manuscript. We made all corrections following his suggestions, as highlighted in yellow throughout the revised version of the manuscript.

Deep Learning for Super-Resolution of Mediterranean Sea Surface Temperature Fields

Claudia Fanelli^{1*}, Daniele Ciani², Andrea Pisano², and Bruno Buongiorno Nardelli¹

¹Consiglio Nazionale delle Ricerche, Istituto di Scienze Marine (CNR-ISMAR), Calata Porta di Massa, 80133, Naples, Italy

²Consiglio Nazionale delle Ricerche, Istituto di Scienze Marine (CNR-ISMAR), Via del Fosso del Cavaliere 100, 00133, Rome, Italy

Correspondence: Claudia Fanelli (claudia.fanelli@cnr.it)

Abstract.

Sea surface temperature (SST) is one of the essential variables of the Earth climate system. Being at the air-sea interface, SST modulates heat fluxes in and out of the ocean, provide insight into several upper/interior ocean dynamical processes, and it is a fundamental indicator of climate variability potentially impacting the health of marine ecosystems. Its accurate estimation and regular monitoring from space is therefore crucial. However, even if satellite infrared/microwave measurements provide much better coverage than what is achievable from in situ platforms, they cannot sense the sea surface under cloudy/rainy conditions. Large gaps are present even in merged multi-sensor satellite products and different statistical strategies have thus been proposed to obtain gap-free (L4) images, mostly based on Optimal Interpolation algorithms. These techniques, however, filter out the signals below the space-time decorrelation scales considered, significantly smoothing most of the small mesoscale and submesoscale features. Here, deep learning models, originally designed for single image Super Resolution (SR), are applied to enhance the effective resolution of SST products and the accuracy of SST gradients. SR schemes include a set of computer vision techniques leveraging Convolutional Neural Networks to retrieve high-resolution data from low-resolution images. A dilated convolutional multi-scale learning network, which includes an adaptive residual strategy and implements a channel attention mechanism, is used to reconstruct features in SST data at $1/100^\circ$ spatial resolution starting from $1/16^\circ$ data over the Mediterranean Sea. The application of this technique shows an improvement in the high resolution reconstruction, capturing small scale features and providing a root-mean-squared-difference improvement of 0.02°C with respect to the L3 ground-truth data.

1 Introduction

Investigating ocean dynamics and climate variability requires accurate, regular and systematic observations of the Sea Surface Temperature (SST). SST plays indeed a key role in air-sea interaction and upper ocean circulation processes (Warner et al., 1990; Deser et al., 2010; Chang and Cornillon, 2015), it is used to track climate variability and change (Jha et al., 2014; Pisano et al., 2020), and it is at the base of various chemical and biological processes (MacKenzie and Schiedek, 2007; Dong et al., 2022a). SST and the estimate of its gradients have also been proven to be a powerful tool to assess and investigate mesoscale and submesoscale variability (e.g., Bowen et al., 2002; Isern-Fontanet et al., 2006; González-Haro and Isern-Fontanet, 2014; Rio

25 et al., 2016; Castro et al., 2017; Ciani et al., 2020). Therefore, the availability of high resolution SST fields is crucial, since they serve as the primary data source for many scientific and operational applications. However, their reliability is hindered by the limitations of infrared (IR) and microwave-based (MW) measurements (Minnett et al., 2019). In fact, thermal IR instruments are able to provide SST images at kilometer to sub-kilometer scale resolution, although their application is limited by cloud cover, aerosols radiation emission/absorption and scattering. Conversely, SST retrieval in the microwave is hampered only due
30 to sunglint, rain, **radio frequency interference** or proximity to land, but the lower spatial resolution ($\simeq 25$ km **nominal resolution**) achievable with present platforms represents a significant disadvantage. Higher resolution MW SSTs ($\simeq 15$ km resolution) will only be available after the launch of the Copernicus Imaging Microwave Radiometer (Pearson et al., 2019), expected during 2029. Consequently, SST fields at high resolution are generally affected by several data voids. For this reason, a few statistical techniques have been developed to obtain gap-free SST images, mostly based on optimal interpolation (OI) (Bretherton et al.,
35 1976). However, as a result of the temporal and spatial averaging applied during the interpolation, the effective resolution of the interpolated products can be significantly coarser than the nominal grid resolution, rarely getting down to less than a few tens of kilometres (Chin et al., 2017; Ciani et al., 2020; Yang et al., 2021). As such, providing interpolated data increases the accessibility of sea surface temperature fields for a wide community of users, but this improvement comes with a trade-off, as statistical interpolation leads to a strong smoothing of small scale ocean features.

40 In this context, we investigate here the potential of applying deep learning models to improve the effective resolution of a gap-free SST field, providing those small scale features even when direct measurements are missing. We exploit techniques generally used in the field of computer vision, which have proven **to be** very successful especially for processing gridded data, managing large-scale datasets while controlling the computational efficiency. In the field of oceanography, the research community started only recently to explore the applicability of machine learning (ML) methods to ocean remote sensing images
45 (Dong et al., 2022b). The applications range from oil spill (Singha et al., 2013) to eddy detection (Lguensat et al., 2018; Duo et al., 2019) and parametrization (Bolton and Zanna, 2019), to marine algae species discrimination (Balado et al., 2021; Cui et al., 2022), to forecasting of ocean variables (Deo and Naidu, 1998; Ham et al., 2019) and estimation of meteorological parameters (Krasnopolsky et al., 2013; Zanna et al., 2019). Moreover, good results have been achieved from applying Neural Networks (NNs) to space-time interpolation and short-term forecasting issues with satellite altimetry data (Fablet et al.,
50 2021), to high-performance description of turbulence processes (Mohan et al., 2020; Zanna and Bolton, 2020) and to SST reconstruction (Meng et al., 2021; Lloyd et al., 2021).

Among image processing techniques, impressive **results** have been obtained **with** Convolutional Neural Networks (CNNs) due to **their** high ability to extract the most important information from two-dimensional spatial fields. Recently, the application of CNN architectures in the process of reconstructing high-resolution images from low-resolution ones, the so-called single
55 image Super Resolution (SR) problem, has attracted much attention in a wide range of scientific challenges. The idea is to implement a network that directly learns the end-to-end mapping between low and high resolution images. One of the simplest attempts made by Dong et al. (2015) was the construction of a network for image restoration composed of three 2D convolutional layers with different kernel sizes. The three layers might be seen as the three conceptual phases of this NN algorithm: a first extraction of overlapping tiles from the input images representing them into features maps; the non-linear

60 mapping of these maps onto one high-dimensional vector of high-resolution feature maps; the final reconstruction aggregating
the above representations to generate the final high-resolution image. This final image is expected to be similar to the ground
truth one. Despite the simplicity of the architecture, the SRCNN developed by Dong et al. (2015) achieved excellent **re-**
sults with respect to more traditional methods and it has already been applied to reconstruct satellite-derived SST data, with
promising results (Ducournau and Fablet, 2016). Building on Dong's work, several more complicated structures have been
65 developed to tackle the Super Resolution problem. In subsequent years, a few **attempts** to develop residual networks have shown
the convergence improvement of deeper architectures, mainly given by the introduction of skip-connections and recursive
convolutions (He et al., 2016; Kim et al., 2016a, b). Similarly, a step forward has been made by Lim et al. (2017) with the
development of an Enhanced Deep Residual Network for Super Resolution (EDSR) which makes use of residual blocks with
constant scaling layers after the last convolutional layer, in order to stabilize the training even in presence of a large number of
70 filters. This modification **led** to significantly better accuracy using much deeper networks, while controlling the computational
cost of the training phase. A further step has been proposed by Liu et al. (2019) with the Adaptive Residual Blocks (ARBs),
which replace the constant factor with adaptive residual factors, increasing the adaptability of the network. **More specifically, in**
the ARB, feature responses (i.e. the filter output channels) are re-calibrated on a channel-wise basis using a so-called Squeeze
and Excitation module, before being combined with the block input. This process is in fact able to enhance the network's
75 **capability to capture intricate relationships among the learned feature channels.** Recently, to further push these networks to
efficiently handle different spatial scales in a multichannel input (each channel including a different variable with characteristic
feature scales), Buongiorno Nardelli et al. (2022) introduced dilated convolutional multi-scale learning modules in the network
developed by Liu et al. (2019), expanding the network receptive fields while still controlling its computational cost. In that
work, the deep learning network was designed to super-resolve absolute dynamic topography (ADT) learning from both low
80 resolution ADT and high resolution SST data, through an observing system simulation experiment (namely simulating all
observations through an ocean general circulation numerical model).

Our aim is to exploit the ability of the dilated SR network to increase gap-free SST effective resolution, directly training our
network on remote sensing SST data for both the input and target datasets. We make use of the data produced by the Italian
National Research Council - Institute of Marine Sciences (CNR-ISMAR), within the Copernicus Marine Service, consisting
85 of merged multi-sensor (L3S) and gap-free (L4) Sea Surface Temperature products over the Mediterranean Sea at high (HR,
nominal $1/16^\circ$ resolution) and ultra-high (UHR, nominal $1/100^\circ$ resolution) spatial resolution (Buongiorno Nardelli et al.,
2013). Considering that our UHR interpolation accounts for space-time decorrelation scales of 5 km and 1 day, in presence
of valid UHR L3 observations the resulting L4 data can be considered as submesoscale resolving (Kurkin et al., 2020). UHR
OI processing makes use of upsized HR L4 **fields as an initial guess**. This guess is left unchanged in the absence of UHR L3
90 data and L4 effective resolution is thus by definition lower than $1/100^\circ$ there. Therefore, our goal is to enhance the effective
resolution of the **gaps-free** upsized $1/16^\circ$ background field.

2 Materials and Methods

2.1 Training and test datasets

When dealing with deep learning methods it is important to construct the training and the test datasets, to ensure a sufficient generalization capability and, more specifically, to avoid under and over-fitting problems. Under-fitting occurs when the model fails to achieve a suitably low error on the training set, while over-fitting occurs when the gap between the training error and test error becomes excessively wide (Goodfellow et al., 2016). Remote sensing data are a very suitable resource to prevent the occurrence of these problems, due to the wide availability of large-scale gridded datasets which are complex enough to encapsulate an extensive variety of features.

The suite of products considered for this project provides the foundation SST (i.e., the temperature free of diurnal variability) over the Mediterranean Sea from 2008 to present (<https://doi.org/10.48670/moi-00172>) at Near Real Time (NRT). These data are built from level 2 (L2) infrared measurements (i.e., data in satellite native grid/swath coordinates) processed through a series of steps divided into different modules as detailed by Buongiorno Nardelli et al. (2013) and sketched in Figure 1.

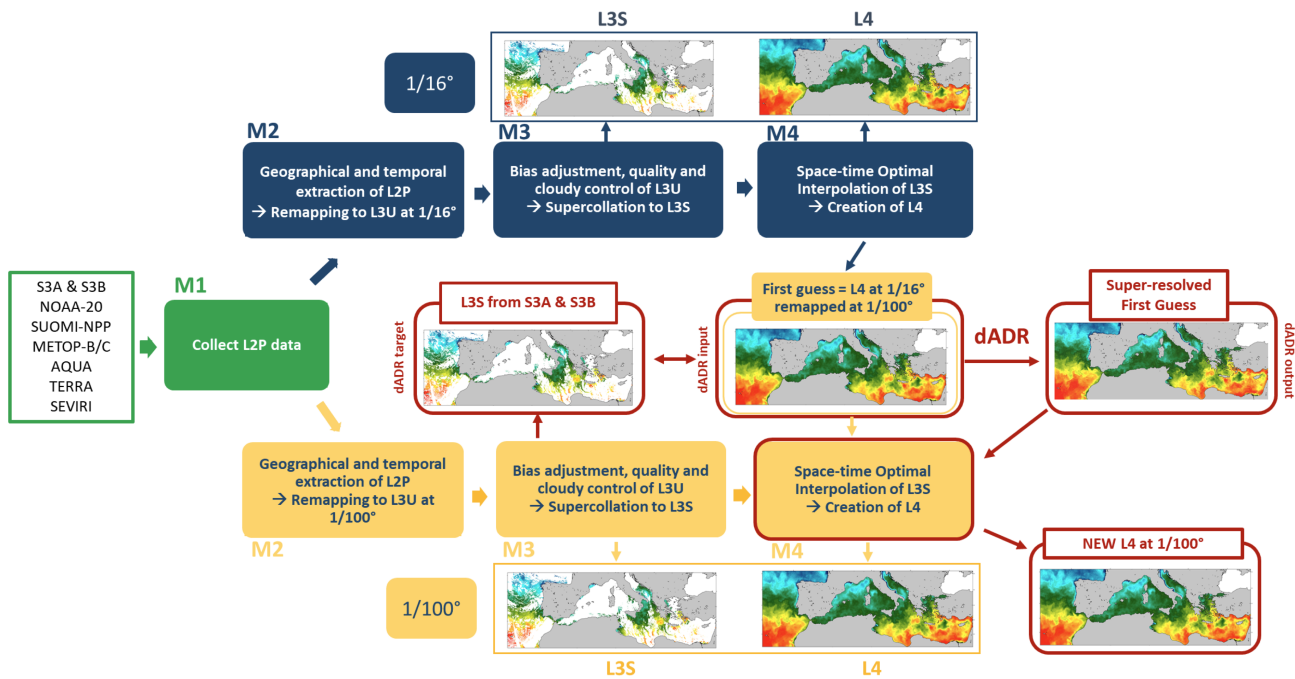


Figure 1. Sketch of the processing chains for the production of L3S and L4 SST fields at $1/16^\circ$ and $1/100^\circ$ spatial resolution over the Mediterranean Sea. Green boxes represent shared modules, blue boxes contain the information and the data for the HR product, yellow boxes refer to the UHR processing chain and red boxes highlight the application of the super-resolution technique.

105 The L2 measurements are obtained from several instruments on board both geostationary and polar orbiting satellites (including Sentinel-3A and Sentinel-3B, NOAA-20, SUOMI NPP, Metop-B, Metop-C, AQUA, TERRA and SEVIRI). After a first shared module for the upstream data collection (M1), two separated processing chains are used to obtain the SST products at $1/16^\circ$ and $1/100^\circ$ spatial resolution. Both include a second module for the geographical and temporal extraction of the L2 data and the remapping onto the corresponding regular grid (M2), and a third module for the bias correction and quality control which lead to the merging of all the data from the different sensors (M3). Therefore, the L3S products at both resolution are
 110 created. The final gap-free fields are provided by applying a space-time Optimal Interpolation algorithm to the L3 data for both HR and UHR products (M4). However, the two OI schemes use different initial guess in absence of the observed data. While the HR processing chain makes use of a climatological background field to produce the L4 SST image at $1/16^\circ$ resolution, in the case of the finer resolution product the initial guess used by the OI algorithm is the L4 HR SST field, preliminary upsized onto a $1/100^\circ$ regular grid (through a thin plate spline). As a consequence, due to the small decorrelation scales assumed in
 115 the UHR interpolation, small scale features are correctly represented only when valid UHR L3 observations are present close to the interpolation point within a short temporal window. For this reason, in the absence of these observations, the final UHR product will have an effective resolution equal to or coarser than $1/16^\circ$.

Therefore, we train the network to improve the satellite SST effective resolution introducing realistic small scale features in the interpolated and upsized gap-free L4 HR SST images (dADR input in Fig. 1). The target data are derived from a ground-
 120 truth super-collated UHR SST dataset (L3S, namely merged multi-sensor data) specifically built for this purpose. The dataset is obtained by applying the first three modules of the UHR SST processing chain to acquisitions only from the Sea and Land Surface Temperature Radiometer (SLSTR) on board of the Sentinel 3A and 3B satellites, due to their high radiometric accuracy and km-scale resolving capabilities (Coppo et al., 2020).

Both input and target datasets are mapped on a regular grid at $1/100^\circ$ spatial resolution over the Mediterranean Sea for the
 125 year 2020. Since our goal is to retrieve small scale features, a first moving average high pass filter (with a kernel radius of 200 km) is applied to remove the large scale dynamics in both input and target images (Kurkin et al., 2020). The data are then selected considering overlapping patches of dimensions 100×100 km, chosen extracting all the tiles containing at least 95% of valid pixels. SST values are then transformed into anomalies to avoid seasonal variability and scaled between -1 and 1 through a classical min-max normalization technique, i.e.:

$$130 \quad SST_{norm} = (b - a) \frac{SST - \min(SST)}{\max(SST) - \min(SST)} + a, \quad (1)$$

where $a = -1$ and $b = 1$. The test dataset is finally selected setting aside the 15% of the tiles available after the preprocessing, using all the tiles extracted from four individual days which are representative of the different seasons. These data are separated in a fully independent dataset which is composed by 18576 tiles (374Mb). The training dataset finally consists of 94110 pairs of tiles (6.4Gb), of which one part (properly separated from the original training dataset at the beginning of the training phase,
 135 following a 85:15 ratio between the two) is used by the network in the validation phase. In Figure 2 an example of the images used to construct the two datasets is shown. The top left panel depicts the gap-free first guess map on 4 January 2020 and below

the corresponding L3 image derived from merged Sentinel 3A and 3B data. On the right, we provide an example of the a pair of tiles extracted from the related low and high resolution SST, fed as input and target to train the network, respectively.

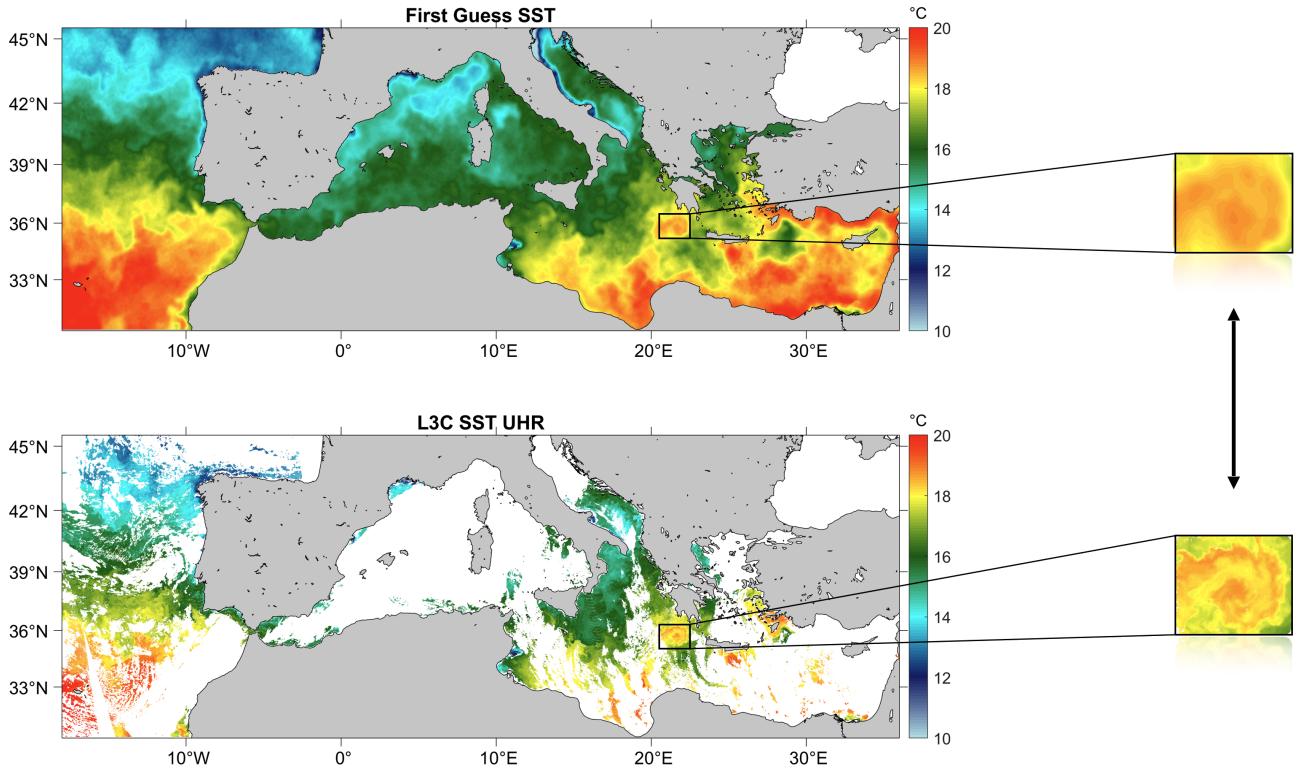


Figure 2. On the left the SST First Guess map used to extract the input tiles (top) and the SST L3C target image (bottom) on 4th January 2020. On the right an example of an extracted tile used for the training and test datasets.

2.2 Super Resolution Convolutional Neural Network

140 In deep learning, Super Resolution algorithms are example-based methods, which generate exemplar patches from the input
 image. As mentioned above, the application of Convolutional Neural Networks to the Super Resolution problem is based on
 networks that directly learn an end-to-end mapping between low and high resolution images. These networks consist of a series
 of interconnected layers, which make use of the convolution operator simulating the connectivity between neurons observed
 in the organization of the animal visual cortex. Formally, the output Y of each layer i is a function of a transformation of the
 145 previous layer output X , i.e.

$$Y = F(W_i * X + B_i), \quad (2)$$

where F is the non-linear activation function (in this case the Rectified Linear Unit or ReLU), W_i the weights, B_i the biases and $*$ the convolution operator. The array of weights, generally called filters or kernels, is able to detect a specific type of feature in the input (generating what is called a features map) and might include for instance Gaussian-like filters or edge detectors along different directions, or any other kind of filter learned during the training. Having its own functionality, each layer will contain different structures. The main idea is that the network learns the weights of the filters and updates the parameters of the system through an optimization process based on minimizing the error between the output and the data from the validation set. This simple yet powerful idea can be augmented as much as desired, especially exploiting the potential of deep networks. The CNN implemented here was originally developed by Buongiorno Nardelli et al. (2022) and called dilated Adaptive Deep Residual Network for Super-Resolution (dADR-SR). In the dADR-SR network (schematized in Fig. 3) the low resolution input dataset is initially fed to three parallel dilated convolutional layers with the same number of filters (equal to 3×3) but increasing dilation factor (1, 3 and 5, respectively), which allows to have a larger receptive field, extracting information at different scales without increasing the number of parameters. After this first stage, the data pass through a sequence of twelve Multiscale Adaptive Residual Blocks (M-ARB), each one including two sets of parallel dilated convolutional layers (with 120 and 10 filters, respectively) and a Squeeze-and-Excitation (SE) module able to improve channel interdependencies at almost no computational cost (Hu et al., 2018), before being summed up to produce the final high resolution output. The main characteristics of an adaptive strategy involves replacing the fixed scaling of learned features with an adaptive scaling mechanism, carried out by the SE block (which is a kind of attention mechanism). It captures the global importance of each channel by initially squeezing the feature maps to a single numeric value (therefore obtaining a vector of size equal to the number of channels) and, finally, feeding this output to a two-layer "bottleneck" network which will produce new features maps, scaling each channel based on its importance. The detailed discussion of the architecture of the network may be found in (Buongiorno Nardelli et al., 2022). The training algorithm follows an early stopping rule which terminates the iterations as soon as the validation loss function increases for a previously chosen number of epochs (defined by the *patience* parameter which is set here equal to 20). An adaptive learning rate (initialized at $l_r = 10^{-4}$) is given by the Adam optimizer (Kingma and Ba, 2014), where the hyperparameters are set following the values found in most of the recent literature (Lim et al., 2017; Liu et al., 2019; Buongiorno Nardelli et al., 2022): the numerical stability constant is $\varepsilon = 10^{-8}$, the exponential decay rates for the first and the second moment estimates are set to $\beta_1 = 0.9$ and $\beta_2 = 0.999$, respectively. Instead of using the classical dropout regularization technique, a DropBlock strategy is implemented where contiguous regions of a feature map are dropped together, which has been shown to increase the accuracy of the network for convolutional layers (Ghiasi et al., 2018). To evaluate the accuracy, the mean-squared error is used as reference in the loss function. The dADR-SR training model finally uses almost 1.6 M trainable parameters. All codes are written in Python using the deep learning framework *Keras* and the training was performed on a single NVIDIA T4 GPU in almost 4 days.

2.3 Evaluation of model performances

Three different error measures are calculated to evaluate the network reconstruction performance between the ground-truth image x and the network output y .




Figure 3. Schematic of the dilated Adaptive Deep Residual Network for Super-Resolution developed by Buongiorno Nardelli et al. (2022). In the figure: i) Conv stands for convolutional; ii) M-ARB stands for Multiscale Adaptive Residual Block; iii) SE stands for Squeeze and Excitation module; iv) Add indicates aggregations of the outputs from the networks blocks.

Firstly, we consider the classical Root Mean Squared Error (RMSE) given by:

$$RMSE(x, y) = \sqrt{\frac{\sum_i^N (x_i - y_i)^2}{N}}, \quad (3)$$

where x_i and y_i are the pixel i of the images x and y , respectively, and N is the total number of pixels. This measure can be useful to evaluate the accuracy of a reconstructed value pixel by pixel, but it can be misleading in the assessment of the ocean state reconstruction, depending on the objectives of the application considered. For instance, if the NN is able to reproduce an ocean structure in a position which is slightly misplaced with respect to the ground-truth measurement, the RMSE will be high indicating a poor reconstruction, even worse than if it would have entirely missed the structure. This issue is often referred to as the double penalty issue, since point-matching measures will penalize the misplacement twice (where the structure should actually be and where it is **incorrectly** predicted). However, in some cases it is possible that capturing an ocean phenomenon, even if in a slightly wrong position, is better than missing it completely. For this reason, we also consider two additional measures usually considered in image processing.

The most commonly used measure for reconstructed image quality is the Peak Signal to Noise Ratio (PSNR), representing the ratio between the maximum possible pixel value of the image I and the power of distorting noise that affects the quality of its representation, usually represented by the RMSE itself:

$$PSNR(x, y) = 20 \log_{10} \left(\frac{\max(I)}{RMSE(x, y)} \right). \quad (4)$$

The PSNR can be seen as an approximation to human perception of reconstruction quality, where the higher the value the better the quality of the image.

The third error measure is the structural similarity index measure (SSIM) proposed by Wang et al. (2004), widely used for measuring image quality and especially the similarity between two images. The concept is based on the idea that while PSNR estimates perceived errors to quantify image degradation, SSIM captures changes in perceived structural information variation. That is, if a reconstructed image is altered with a different type of degradation (for instance, mean-shifted, blurred or with a salt-pepper impulsive noise effect), while the MSE will come out nearly identical for all the cases, the SSIM will capture the different perceptual quality, being a weighted combination of luminance, contrast and structure measurements. The mathematical formulation of the SSIM between two images x and y is given by:

$$SSIM(x, y) = \frac{(2\mu_x\mu_y + c_1)(2\sigma_{xy} + c_2)}{(\mu_x^2 + \mu_y^2 + c_1)(\sigma_x^2 + \sigma_y^2 + c_2)}, \quad (5)$$

where μ_x and μ_y are the mean values of x and y , respectively, σ_{xy} is the cross-correlation of x and y , σ_x^2 and σ_y^2 are the variance of x and y , respectively, and c_1 and c_2 are the regularization constants for the luminance, contrast, and structural terms.

All these errors are computed for the final reconstructed maps over the whole Mediterranean Sea. The combination of the super-resolved tiles is made considering a linear combination of the values obtained for the same pixel, weighted according to their position within the tile. That is, the coefficient of each value decreases as its distance becomes larger from the central pixel of the tile. Finally, the large scale field initially removed to isolate the small scale features is added back to obtain the original SST field.

3 Results and discussion

Our aim is to verify whether the dADR-SR network, trained by means of satellite-derived observations, is able to improve the effective resolution of the SST fields in the areas where our interpolation technique removes most of the spatial variability associated with mesoscale and submesoscale processes. Due to computational costs, a preliminary study based on a restricted test dataset was performed. The performances of the dADR network in comparison with other deep learning algorithms and a sensitivity analysis on the impact of different choices for the architecture is presented in Sec. 3.1. Finally, an additional test on an independent dataset built on one year long series of SST fields was carried out (Sec. 3.2).

3.1 Preliminary validation study

Figure 4 shows the comparison of the result obtained on one of the SST fields included in the test dataset which corresponds to the SST daily map of 1 August 2020. The reconstruction of the dADR-SR network at $1/100^\circ$ spatial resolution over the Mediterranean Sea is shown in the top panel. In order to visually evaluate the reconstruction of the network, the corresponding L3S SST merged field observed by Sentinel 3A and 3B (central panel) and the First Guess map used in the optimal interpolation algorithm (bottom panel) are shown below the network output. Evident from this figure is that the SST features estimated by the CNN appear much sharper than the ones approximated by the low resolution map, showing promising capabilities to effectively reconstruct dynamical features.

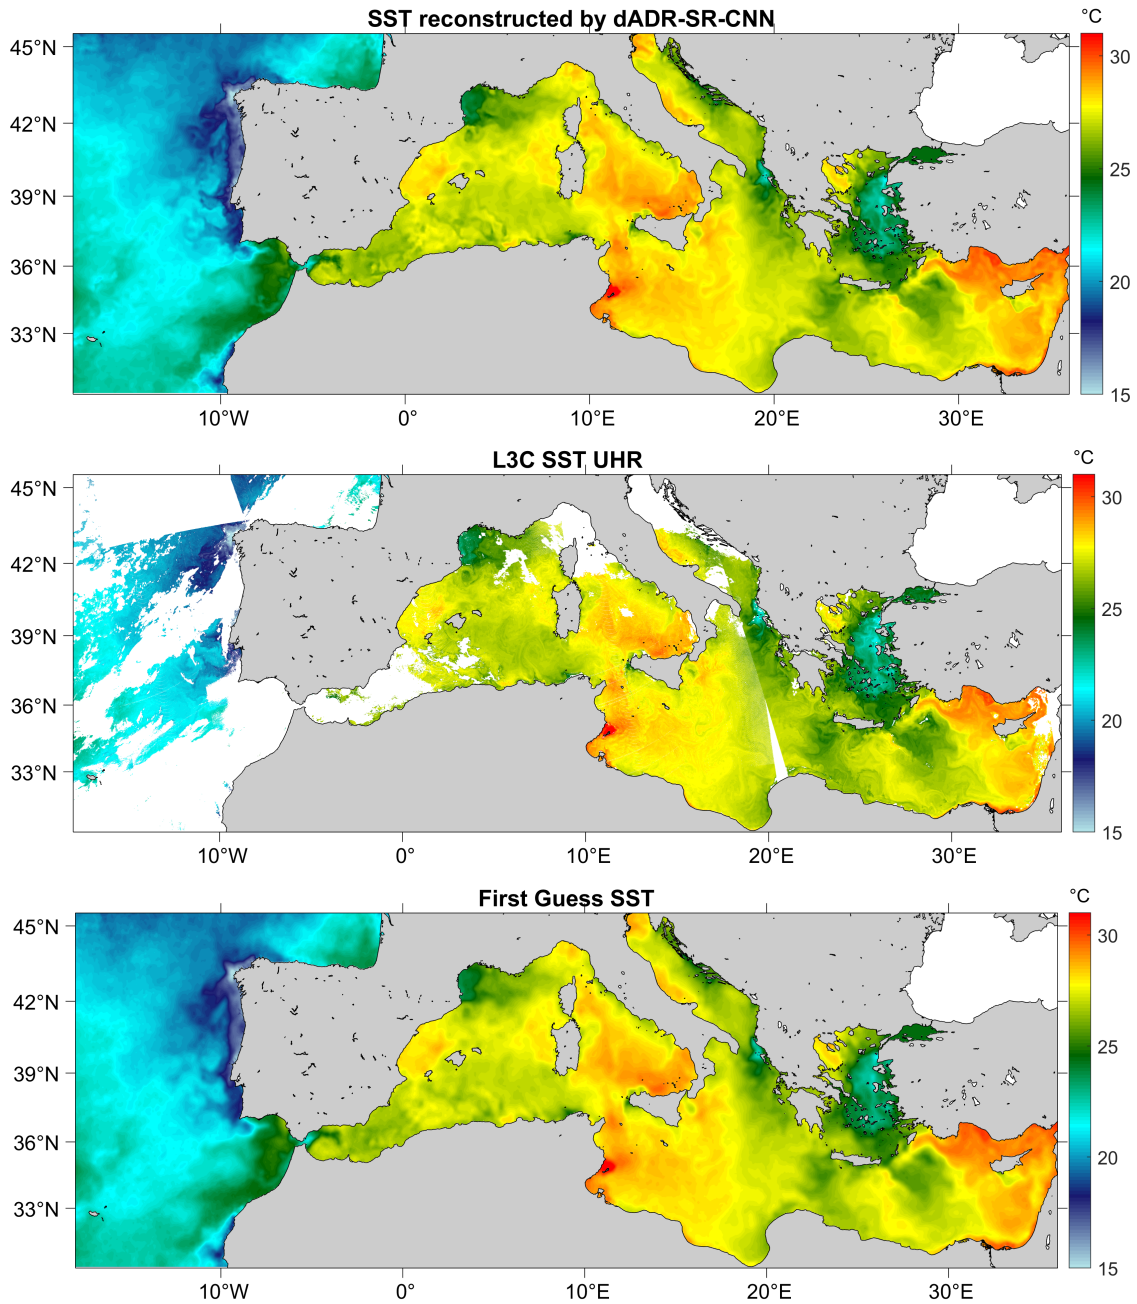


Figure 4. Comparison of the SST fields of 1st August 2020 provided by the reconstruction of the dADR-SR-CNN (top panel), the collated L3 data measured by Sentinel 3A and 3B (central panel) and the Optimal Interpolated First Guess (bottom panel).

To highlight the ability of the CNN to more accurately capture small scale features with respect to the statistical algorithm, we show in Figure 5 three smaller panels corresponding to zoomed-in regions of the fields shown in Fig. 4 delimited by the

230 coordinates [30, 40]^oN and [19, 36]^oE, with the correspondent SST spatial gradients (calculated using the Sobel operator)
 for the dADR-SR reconstruction, L3S ground-truth data and the first guess approximation (from top to bottom, respectively).
 This particularly structure-rich area is an excellent example to demonstrate the ability of the network to capture dynamical
 processes which are quite clear in the high resolution SLSTR data (shown in the central panels). In fact, SST fronts are
 strongly connected with the surface dynamics and they are generally associated with energetic motions at the mesoscale and
 235 submesoscale. While the CNN is able to capture at least the most energetic structures (as shown in the top panels of Fig. 5),
 the optimal interpolation algorithm produces an extremely smooth field, where all the small scale features have been filtered
 out even when high resolution data are present.

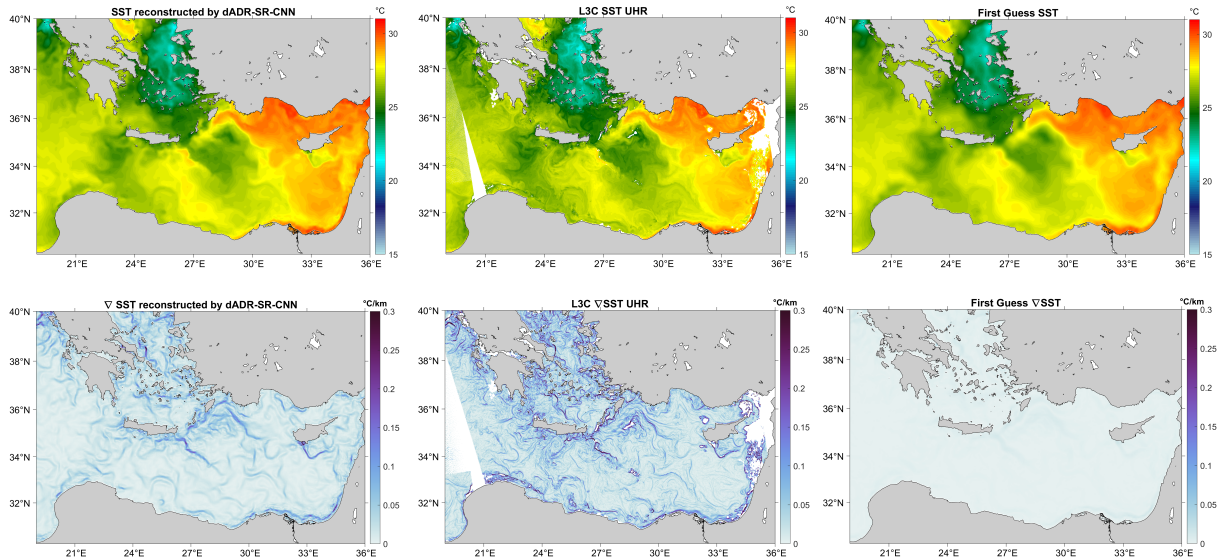


Figure 5. Comparison of the SST (on the left) and SST gradients (on the right) provided by the reconstruction of the dADR-SR-CNN (top panels), the collated L3 data measured by Sentinel 3A and 3B (central panels) and the Optimal Interpolated First Guess (bottom panels) in the selected region delimited by the coordinates [30, 40]^oN in latitude and [19, 36]^oE in longitude.

This visual analysis is quantitatively confirmed by the maps in Figure 6, displaying the difference between the error made by the low resolution approximation and the dADR-SR model with respect to the original L3 image, averaged on $1^{\circ} \times 1^{\circ}$
 240 boxes. Here, red indicates an improvement of the CNN with respect to the low resolution image and blue a degradation. A clear predominance of red boxes is found both in the SST and the SST gradients error maps.

In Table 1, the comparison is summarized quantitatively, with the network reconstruction presenting a RMSE = 0.31°C (against the 0.33°C obtained from the OI approximation), a mean PSNR equal to 37.9 and a SSIM equal to 0.54, both larger than the low resolution result. To highlight the quality of the dADR reconstruction we performed the same test using other
 245 deep learning super-resolution models: the Enhanced Deep Residual network for Super-Resolution (EDSR) proposed by (Lim et al., 2017), the Adaptive Deep Residual Network for Super-Resolution (ADR) developed by (Liu et al., 2019) and the dADR

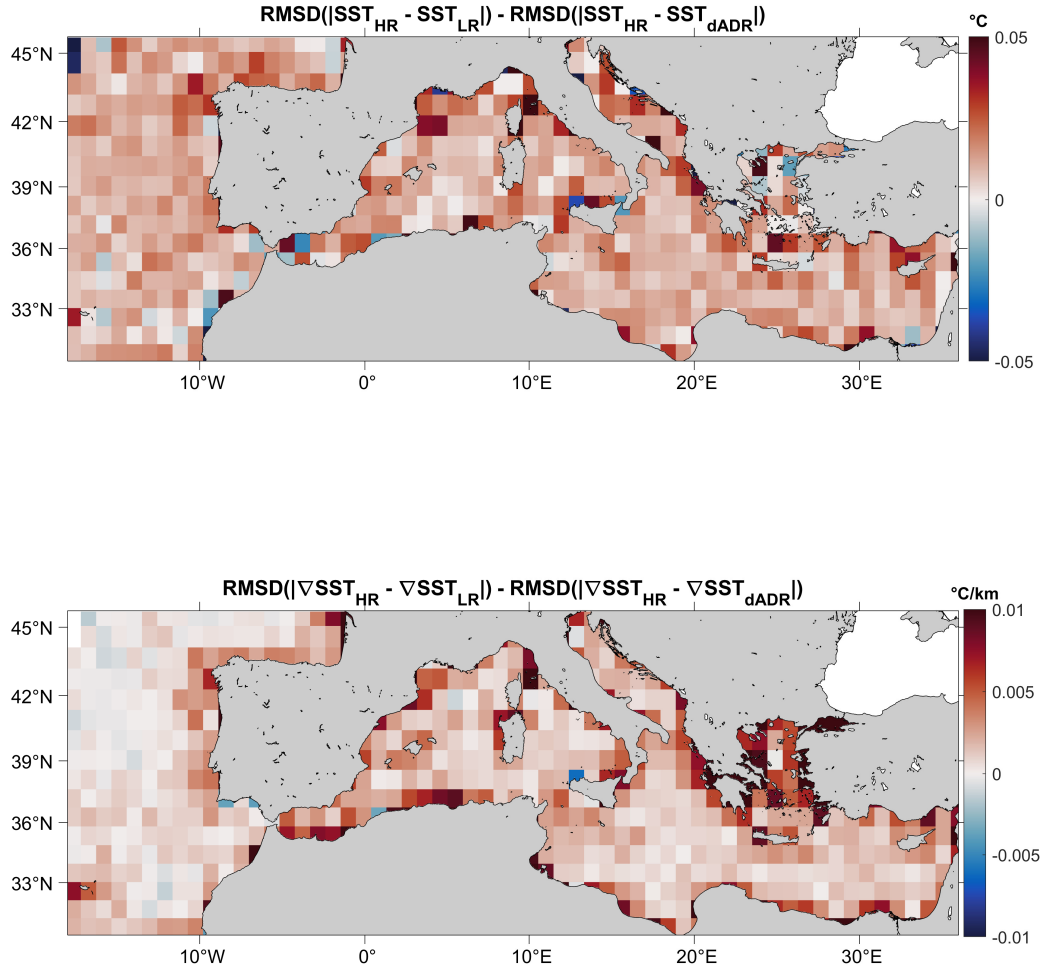


Figure 6. Comparison of the performance of the SST (top) and ∇SST (bottom) dADR-SR reconstruction and the L4 First Guess with respect to the L3 data measured by Sentinel 3A and 3B satellites. Red positive values show an improvement of the network reconstruction with respect to the optimal interpolated First Guess.

proposed by (Buongiorno Nardelli et al., 2022) setting the number of M-ARB equal to six instead of 12 (called dADRSR/2 hereinafter). We recall that while the RMSE should be low to ensure a good approximation, for the other two quantities high values indicate an improvement. The dADR output achieves the best value for all the evaluation methods presented.

Table 1. Error estimations of the SST given by the First Guess map, the EDSR network (Lim et al., 2017), the ADR reconstruction (Liu et al., 2019), the dADRSR built with half of the M-ARB (called dADRSR/2) and the dADRSR output with respect to the L3S ground-truth: the RMSE given by Equations (3) and the corresponding confidence interval calculated by the bootstrapping procedure, the Peak-Signal to Noise Ratio obtained by (4) and the Structural Similarity Index Measure given by (5).

Model	RMSE ($^{\circ}C$)	PSNR	SSIM
Low Resolution	$0.33 \pm 7 \times 10^{-5}$	37.5	0.53
EDSR	$0.32 \pm 6 \times 10^{-5}$	37.7	0.54
ADR	$0.32 \pm 6 \times 10^{-5}$	37.9	0.54
dADRSR/2	$0.32 \pm 8 \times 10^{-5}$	37.7	0.54
dADRSR	$0.31 \pm 7 \times 10^{-5}$	37.9	0.54

250 To analyse the effectiveness of the high resolution reconstruction, we compare the Power Spectral Density (PSD) of the network output with the L3 product and the First Guess map over three selected zones. We compute the PSD via Fast Fourier Transform (FFT) with a Blackman–Harris window over the three areas corresponding to the boxes with labels *a*, *b* and *c* in Figure 7. The zones are chosen in order to have the maximum number of valid pixels available in the L3 observations and represent different dynamical regimes:

- 255 a - A region over the Sea of Sardinia of low spatial variability.
- b - A region over the Ionian Sea with an important SST variability between the eastern and the western part of the area.
- c - A region over the Levantine Sea characterized by small scale structures.

The three central panels of Figure 7 show the PSD, presented as a function of the wavenumber, of the SST reconstructed by the dADR-SR (in yellow), the first guess map (in red) and the high resolution observations (in blue) over the three zones delimited by the black rectangles over the L3 SST field on the top panel. In all the cases note that the PSD follows the same behavior for wavenumbers smaller than 1 deg^{-1} ; i.e., for scales larger than 100 km. This means that for such scales the SST fields reconstructed by both the CNN and the OI algorithm characterize well large mesoscale features. Conversely, for regions *a* and *b*, both the Low Resolution and the network reconstruction exhibit a significant PSD decrease for wavenumbers higher than 1 deg^{-1} , but starting from approximately 10 deg^{-1} , the first guess spectrum separates from the CNN spectrum (and consequently increases the distance from the HR one), indicating a poorer reconstruction of spatial features below 10 km. The CNN spectrum, on the other hand, shows that the machine learning algorithm is able to capture those small scale features in all the cases. An analogous behaviour is found for the spectra of the SST gradients over the same three regions (three bottom panels). The abrupt decreases are probably due to artifacts introduced by the re-gridding, as already discussed by Liberti et al. (2023), leading to the lack of physical meaning for the spectra from the "bumps" onwards (i.e. for the highest wavenumbers).

265

270 Overall, we can conclude that the CNN spectra tracks the observed spectra from 1 deg^{-1} to 10 deg^{-1} quite well while the OI

spectra are more energetic in this region. On the other hand, for wave numbers greater than 10 deg^{-1} the CNN spectra are more energetic than the observed spectra but still more representative of the spectra on this portion of the curve than the OI spectra.

275 To quantify the differences between the effective resolution of the products, we calculated the ratio between the spectral content of the mapping error and the spectral content of the L3S observation of the SST for the dADR-SR reconstructed fields and the first guess map (see 8) over the three zones delimited by black rectangles on the L3 SST field on 1 August 2020 presented in the bottom panel of Figure 7. The effective resolution of the products, based on the intersection between this PSD ratio and the ratio equal to 0.5 as defined by (Ballarotta et al., 2019), shows the ability of the dADR network to resolve smaller scales.

280 In order to demonstrate the effective spatial resolution enhancement of the super-resolved images under cloudy conditions, we performed a power spectral density analysis on different SST products available over the Mediterranean Sea (Figure 9). The products used are the L4 NRT HR at nominal $1/16^\circ$ and the L4 NRT UHR at nominal $1/100^\circ$ spatial resolution provided by CNR for the Copernicus Marine Service (<https://doi.org/10.48670/moi-00172>), the GLOBAL OCEAN OSTIA product developed by the UK MET OFFICE at 0.05° (<https://doi.org/10.48670/moi-00165>), the Multi-scale Ultra-high Resolution (MUR) product provided by the NASA-JPL at 0.01° (Chin et al., 2017) and the super-resolved SST field obtained by the application of the dADR-SR network developed in this work at 0.01° . The PSD are computed along the three transects (a, b and c in the bottom panel of Figure 9), chosen in three areas affected (at different levels) by cloud coverage. In all the cases the slope of the spectra is very similar for the three products at higher resolution (namely the UHR, the MUR and the dADR approximations), while we can observe lower values for the HR and OSTIA products (which also do not reach wavenumbers larger than 0.1 deg^{-1} by construction). The PSD analysis for section (a) shows how the dADR approximation is not affected by excessive noise as the other products, presenting PDS values higher than the MUR spectrum and the UHR product (especially at small scales). For sections (b) and (c), the green line representing the dADR spectrum stands above all the other lines for almost the whole wavenumber range.

285

290

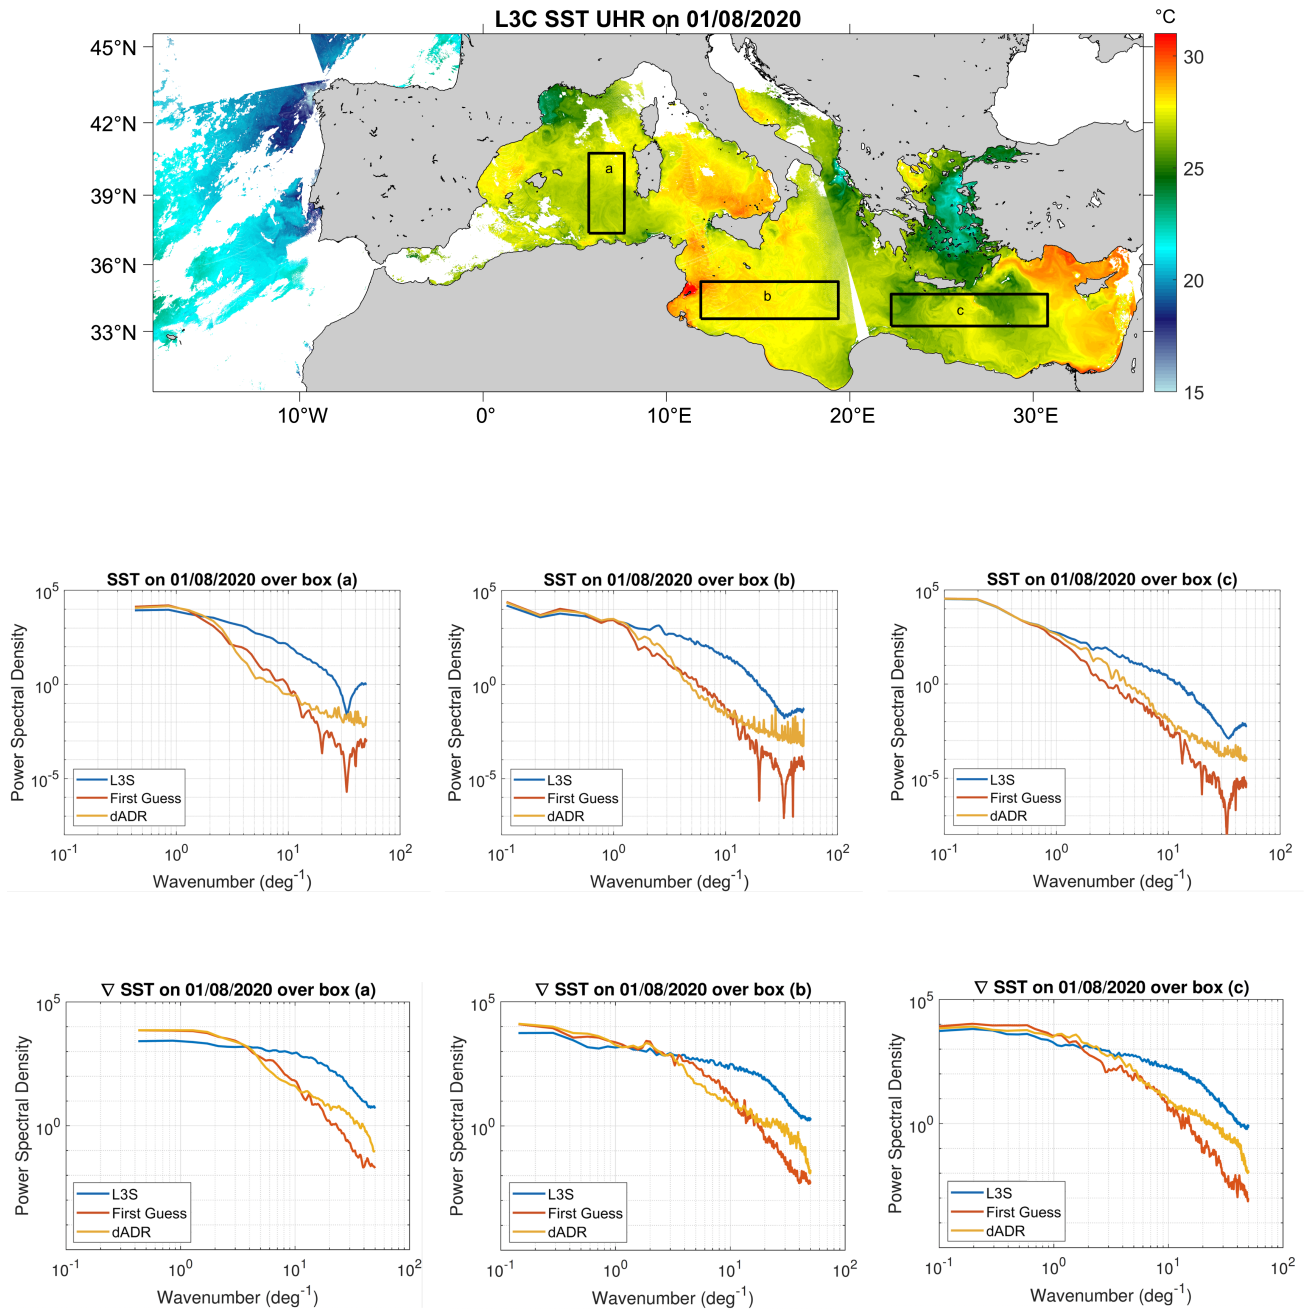


Figure 7. The PSD of the SST (central panels) and SST gradients (bottom panels) reconstructed by the dADR-SR (in yellow), the first guess map (in red) and the high resolution observations (in blue) over the three zones delimited by black rectangles on the L3 SST field on 1 August 2020 in the top panel.




Figure 8. The ratio between the spectral content of the mapping error and the spectral content of the L3S observation of the SST for the dADR-SR reconstructed fields (in yellow) and the first guess map (in red) over the three zones delimited by black rectangles on the L3 SST field on 1 August 2020 presented in the bottom panel of Figure 7. The black line represents the ratio equal to 0.5, the dashed black line and the dotted black line represent the intersection between PSD ratio of the first guess and the dADR, respectively, showing the effective resolution of the products.

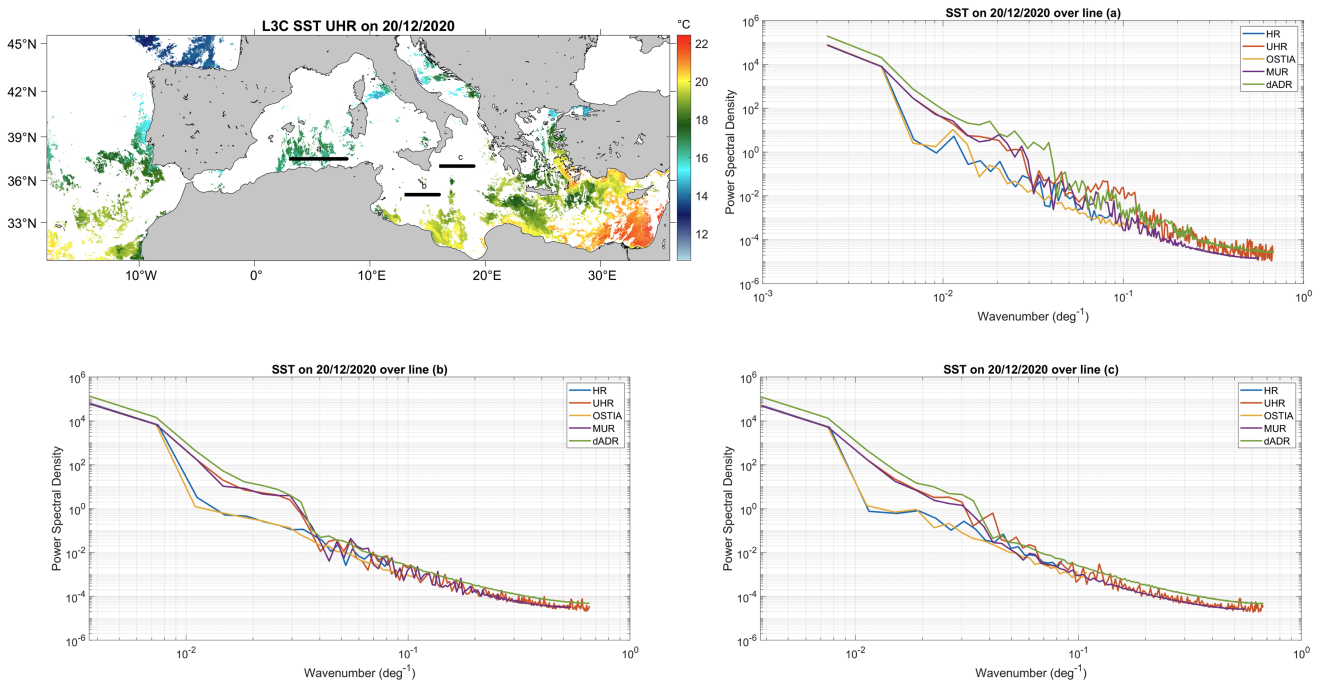


Figure 9. The Power Spectral Density profiles (bottom panels) of different SST products under cloudy conditions calculated along the three transects (black lines) on the L3 SST field on 20 December 2020 in the top panel: in blue the MED L4 NRT HR product at $1/16^\circ$ (Copernicus), in red the MED L4 NRT UHR at $1/100^\circ$ (Copernicus), in yellow the GLOBAL OCEAN OSTIA product at 0.05° (Copernicus), in purple the Multi-scale Ultra-high Resolution (MUR) product at 0.01° (NASA-JPL) and in green the dADR reconstruction developed in this work.

3.2 An extensive test dataset

Once established the potential ability of the network, we performed an additional validation test with the aim to strengthen the robustness of the statistics. Therefore, we built a new totally independent dataset exploiting the L4 SST fields used as initial guess for the OI algorithm which produces gap-free SST maps over the Mediterranean Sea for the year 2021 at $0.01^\circ \times 0.01^\circ$ spatial resolution. Before to be fed to the dADR network in order to obtain the super resolved SST fields, this dataset was preprocessed following the same steps described in Sec. 2.1. The comparison of the evolution of the daily, basin-scale RMSE for the input L4 SSTs and the dADR reconstruction along the whole year confirms the improvement of the network output with respect to the interpolated fields (Fig. 10). The mean values of the RMSE for the 2021 are in line with the ones found for the previous test (i.e., 0.31°C for the dADR output against the 0.33°C of the low resolution approximation).




Figure 10. Comparison of the daily RMSE time series of the SST obtained via dADR-SR reconstruction (red) and the L4 First Guess (blue) with respect to the L3 data measured by Sentinel 3A and 3B satellites during the year 2021. The jump at Julian day 134 is due to the absence of L3 observation for that day which made impossible the comparison with the input and the output of the network.

The analogous of Fig. 6, showing the differences (averaged on $1^\circ \times 1^\circ$ boxes) between the errors of the two methods approximating the L3 observed data, presents an overall improvement of the output of the dADR network with respect to the interpolated maps. The red positive values are found almost everywhere for both SST and SST gradients fields.

4 Conclusions

The advance obtained by the application of machine learning-based techniques for the improvement of the effective resolution of remote sensing observations have recently opened a new way to approach satellite-derived data processing. The great advantages provided by making high resolution gap-free images available for a wide range of scientific users are severely limited by the number of valid L3 observations. In the case of sea surface temperature measurements, infrared data are commonly contaminated by cloud cover, reducing the quality of the L4 data that can be obtained via statistical interpolation techniques. The

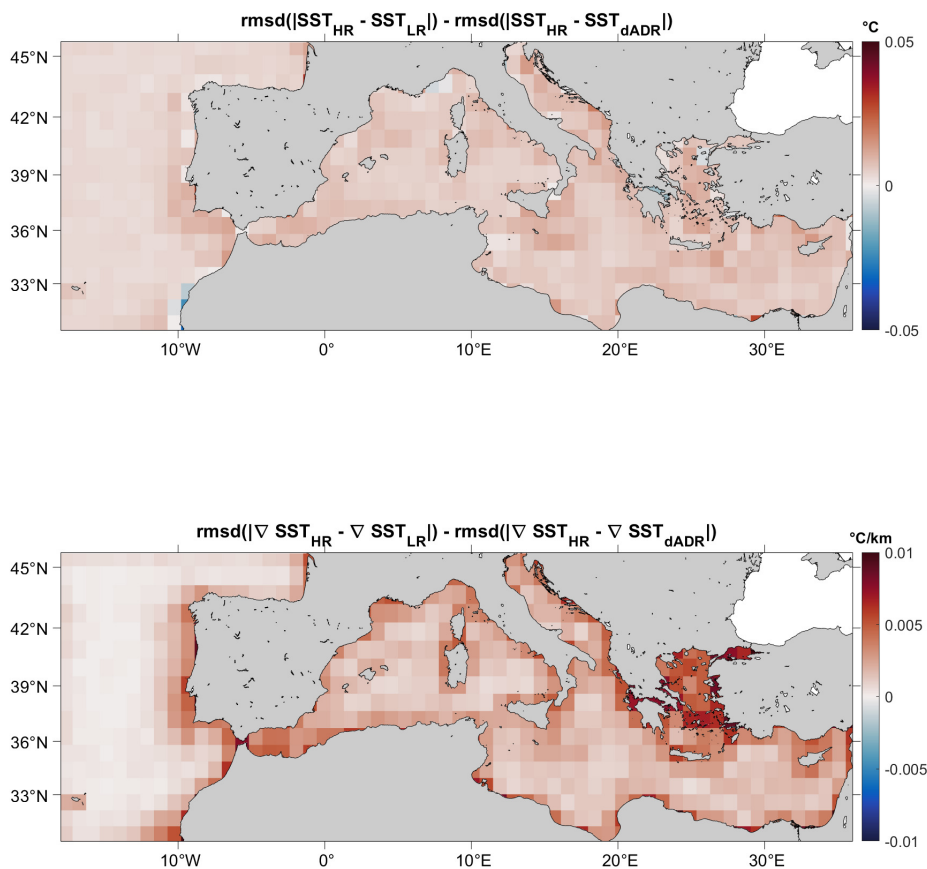


Figure 11. Comparison of the performance of the SST (top) and ∇SST (bottom) dADR-SR reconstruction and the L4 First Guess with respect to the L3 data measured by Sentinel 3A and 3B satellites during the year 2021. Red positive values show an improvement of the network reconstruction with respect to the optimal interpolated First Guess.

310 machine learning approach used here exploits progress made in the field of computer vision for extrapolating high resolution features even when a direct measurement is missing. Learning directly from ground-truth data, and taking advantage of both dilated convolution and attention mechanisms, the deep neural network employed here proved able to reproduce small scale signals generally smoothed out by Optimal Interpolation algorithms. The strong variability of the SST in the Mediterranean Sea allowed us to obtain excellent results even considering just one year of data during the training phase. However, it would
 315 be important to investigate whether using longer time series may help to improve the network ability to reconstruct SST fields, as well as to rely on more robust statistics. Moreover, given the inhomogeneity of the spatial error distribution related to the

interpolation technique, it would be interesting to expand the present investigation in order to take into account the OI error field as an additional predictor and to consider the contribution of the error of the SST gradients in the loss function. Another aspect that would deserve further investigation concerns the applicability of the dADR-SR network to different sea/ocean areas, even though a fine-tuning of the model would probably be needed.

In the future, we also plan to study other super-resolution techniques recently became very popular in the field of computer vision. Firstly, we are currently investigating the possibility to improve the reconstruction of small scale features in SST fields via other successful generative AI, such as GANs or diffusion models. The former exploits the outcomes of a minmax game between a generator of reconstructed images and a discriminator which tries to distinguish the real image from the output of the other network; the latter builds super-resolved fields initially introducing noise into the initial signal to then reverse this process until convergences to the desired distribution. However, this kind of networks applied to remote sensing data, while have been proven to be able to reconstruct very realistic small scale structures, seems to fail to optimize a point-match evaluation of a reconstructed remote-sensed SST. Secondly, we would like to explore the usage of Vision Transformers (ViT) for understanding and reproducing high-level structures by understanding contextual relationships between the patches of an image.

The results achieved here, however, may already benefit a wide range of applications. Super-resolved SST fields would facilitate the challenging task of 2D/3D ocean dynamics reconstruction in synergy with other variables (e.g., Buongiorno Nardelli et al., 2022; Fablet et al., 2023) or the monitoring of ocean fronts in areas of particular interest (e.g., areas affected by vertical exchange and upwelling regions). To enhance the effective resolution of SST data and especially SST gradients may also benefit data assimilation in forecast modelling, given their proven sensibility to small structures of sea surface temperature (Woollings et al., 2010; Maloney and Chelton, 2006). We also plan to validate our results exploiting the high resolution SST data derived by the CNN reconstruction within the operational SST chain in the framework of the Copernicus Marine Service.

Author contributions. Conceptualization, C.F., A.P. and B.B.N.; methodology, C.F., D.C., A.P. and B.B.N.; investigation, C.F., D.C., A.P. and B.B.N.; visualization, C.F.; writing—original draft, C.F.; writing—review and editing, C.F., D.C., A.P. and B.B.N. All authors have read and agreed to the published version of the manuscript.

Financial support. This work has been supported by the Copernicus Marine Service, funded through contract agreement no. 21001L03-COP-TAC SST-2300– Lot 3: Provision of Sea Surface Temperature Observation Products (SST-TAC).

Competing interests. The authors declare no conflict of interest.

References

- 345 Balado, J., Olabarria, C., Martínez-Sánchez, J., Rodríguez-Pérez, J. R., and Pedro, A.: Semantic segmentation of major macroalgae in coastal environments using high-resolution ground imagery and deep learning, *International Journal of Remote Sensing*, 42, 1785–1800, 2021.
- Ballarotta, M., Ubelmann, C., Pujol, M.-I., Taburet, G., Fournier, F., Legeais, J.-F., Faugère, Y., Delepouille, A., Chelton, D., Dibarboure, G., et al.: On the resolutions of ocean altimetry maps, *Ocean science*, 15, 1091–1109, 2019.
- Bolton, T. and Zanna, L.: Applications of deep learning to ocean data inference and subgrid parameterization, *Journal of Advances in*
350 *Modeling Earth Systems*, 11, 376–399, 2019.
- Bowen, M. M., Emery, W. J., Wilkin, J. L., Tildesley, P. C., Barton, I. J., and Knewton, R.: Extracting multiyear surface currents from sequential thermal imagery using the maximum cross-correlation technique, *Journal of Atmospheric and Oceanic Technology*, 19, 1665–1676, 2002.
- Bretherton, F. P., Davis, R. E., and Fandry, C.: A technique for objective analysis and design of oceanographic experiments applied to
355 *MODE-73*, in: *Deep Sea Research and Oceanographic Abstracts*, vol. 23, pp. 559–582, Elsevier, 1976.
- Buongiorno Nardelli, B., Tronconi, C., Pisano, A., and Santoleri, R.: High and Ultra-High resolution processing of satellite Sea Surface Temperature data over Southern European Seas in the framework of MyOcean project, *Remote Sensing of Environment*, 129, 1–16, 2013.
- Buongiorno Nardelli, B., Cavaliere, D., Charles, E., and Ciani, D.: Super-resolving ocean dynamics from space with computer vision algorithms, *Remote Sensing*, 14, 1159, 2022.
- 360 Castro, S. L., Emery, W. J., Wick, G. A., and Tandy, W.: Submesoscale sea surface temperature variability from UAV and satellite measurements, *Remote Sensing*, 9, 1089, 2017.
- Chang, Y. and Cornillon, P.: A comparison of satellite-derived sea surface temperature fronts using two edge detection algorithms, *Deep Sea Research Part II: Topical Studies in Oceanography*, 119, 40–47, 2015.
- Chin, T. M., Vazquez-Cuervo, J., and Armstrong, E. M.: A multi-scale high-resolution analysis of global sea surface temperature, *Remote*
365 *sensing of environment*, 200, 154–169, 2017.
- Ciani, D., Rio, M.-H., Nardelli, B. B., Etienne, H., and Santoleri, R.: Improving the altimeter-derived surface currents using sea surface temperature (SST) data: A sensitivity study to SST products, *Remote Sensing*, 12, 1601, 2020.
- Coppo, P., Brandani, F., Faraci, M., Sarti, F., Dami, M., Chiarantini, L., Ponticelli, B., Giunti, L., Fossati, E., and Cosi, M.: Leonardo spaceborne infrared payloads for Earth observation: SLSTRs for Copernicus Sentinel 3 and PRISMA hyperspectral camera for PRISMA
370 satellite, *Applied Optics*, 59, 6888–6901, 2020.
- Cui, B., Zhang, H., Jing, W., Liu, H., and Cui, J.: SRSe-net: Super-resolution-based semantic segmentation network for green tide extraction, *Remote Sensing*, 14, 710, 2022.
- Deo, M. and Naidu, C. S.: Real time wave forecasting using neural networks, *Ocean engineering*, 26, 191–203, 1998.
- Deser, C., Alexander, M. A., Xie, S.-P., and Phillips, A. S.: Sea surface temperature variability: Patterns and mechanisms, *Annual review of*
375 *marine science*, 2, 115–143, 2010.
- Dong, C., Loy, C. C., He, K., and Tang, X.: Image super-resolution using deep convolutional networks, *IEEE transactions on pattern analysis and machine intelligence*, 38, 295–307, 2015.
- Dong, C., Liu, L., Nencioli, F., Bethel, B. J., Liu, Y., Xu, G., Ma, J., Ji, J., Sun, W., Shan, H., et al.: The near-global ocean mesoscale eddy atmospheric-oceanic-biological interaction observational dataset, *Scientific Data*, 9, 436, 2022a.

- 380 Dong, C., Xu, G., Han, G., Bethel, B. J., Xie, W., and Zhou, S.: Recent developments in artificial intelligence in oceanography, *Ocean-Land-Atmosphere Research*, 2022, 2022b.
- Ducournau, A. and Fablet, R.: Deep learning for ocean remote sensing: an application of convolutional neural networks for super-resolution on satellite-derived SST data, in: 2016 9th IAPR Workshop on Pattern Recognition in Remote Sensing (PRRS), pp. 1–6, IEEE, 2016.
- Duo, Z., Wang, W., and Wang, H.: Oceanic mesoscale eddy detection method based on deep learning, *Remote Sensing*, 11, 1921, 2019.
- 385 Fablet, R., Amar, M., Febvre, Q., Beauchamp, M., and Chapron, B.: End-to-end physics-informed representation learning for satellite ocean remote sensing data: Applications to satellite altimetry and sea surface currents, *ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, 3, 295–302, 2021.
- Fablet, R., Febvre, Q., and Chapron, B.: Multimodal 4DVarNets for the reconstruction of sea surface dynamics from SST-SSH synergies, *IEEE Transactions on Geoscience and Remote Sensing*, 2023.
- 390 Ghiasi, G., Lin, T.-Y., and Le, Q. V.: Dropblock: A regularization method for convolutional networks, *Advances in neural information processing systems*, 31, 2018.
- González-Haro, C. and Isern-Fontanet, J.: Global ocean current reconstruction from altimetric and microwave SST measurements, *Journal of Geophysical Research: Oceans*, 119, 3378–3391, 2014.
- Goodfellow, I., Bengio, Y., and Courville, A.: *Deep learning*, MIT press, 2016.
- 395 Ham, Y.-G., Kim, J.-H., and Luo, J.-J.: Deep learning for multi-year ENSO forecasts, *Nature*, 573, 568–572, 2019.
- He, K., Zhang, X., Ren, S., and Sun, J.: Deep residual learning for image recognition, in: *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778, 2016.
- Hu, J., Shen, L., and Sun, G.: Squeeze-and-excitation networks, in: *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 7132–7141, 2018.
- 400 Isern-Fontanet, J., Chapron, B., Lapeyre, G., and Klein, P.: Potential use of microwave sea surface temperatures for the estimation of ocean currents, *Geophysical research letters*, 33, 2006.
- Jha, B., Hu, Z.-Z., and Kumar, A.: SST and ENSO variability and change simulated in historical experiments of CMIP5 models, *Climate dynamics*, 42, 2113–2124, 2014.
- Kim, J., Lee, J. K., and Lee, K. M.: Accurate image super-resolution using very deep convolutional networks, in: *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 1646–1654, 2016a.
- 405 Kim, J., Lee, J. K., and Lee, K. M.: Deeply-recursive convolutional network for image super-resolution, in: *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 1637–1645, 2016b.
- Kingma, D. P. and Ba, J.: Adam: A method for stochastic optimization, *arXiv preprint arXiv:1412.6980*, 2014.
- Krasnopolsky, V. M., Fox-Rabinovitz, M. S., and Belochitski, A. A.: Using ensemble of neural networks to learn stochastic convection parameterizations for climate and numerical weather prediction models from data simulated by a cloud resolving model, *Advances in Artificial Neural Systems*, 2013, 5–5, 2013.
- 410 Kurkin, A., Kurkina, O., Rybin, A., and Talipova, T.: Comparative analysis of the first baroclinic Rossby radius in the Baltic, Black, Okhotsk, and Mediterranean seas, *Russian Journal of Earth Sciences*, 20, 8, 2020.
- Lguensat, R., Sun, M., Fablet, R., Tandeo, P., Mason, E., and Chen, G.: EddyNet: A deep neural network for pixel-wise classification of oceanic eddies, in: *IGARSS 2018-2018 IEEE International Geoscience and Remote Sensing Symposium*, pp. 1764–1767, IEEE, 2018.
- 415 Liberti, G. L., Sabatini, M., Wethey, D. S., and Ciani, D.: A Multi-Pixel Split-Window Approach to Sea Surface Temperature Retrieval from Thermal Imagers with Relatively High Radiometric Noise: Preliminary Studies, *Remote Sensing*, 15, 2453, 2023.

- Lim, B., Son, S., Kim, H., Nah, S., and Mu Lee, K.: Enhanced deep residual networks for single image super-resolution, in: Proceedings of the IEEE conference on computer vision and pattern recognition workshops, pp. 136–144, 2017.
- 420 Liu, S., Gang, R., Li, C., and Song, R.: Adaptive deep residual network for single image super-resolution, *Computational Visual Media*, 5, 391–401, 2019.
- Lloyd, D. T., Abela, A., Farrugia, R. A., Galea, A., and Valentino, G.: Optically enhanced super-resolution of sea surface temperature using deep learning, *IEEE Transactions on Geoscience and Remote Sensing*, 60, 1–14, 2021.
- MacKenzie, B. R. and Schiedek, D.: Long-term sea surface temperature baselines—time series, spatial covariation and implications for biological processes, *Journal of Marine Systems*, 68, 405–420, 2007.
- 425 Maloney, E. D. and Chelton, D. B.: An assessment of the sea surface temperature influence on surface wind stress in numerical weather prediction and climate models, *Journal of climate*, 19, 2743–2762, 2006.
- Meng, Y., Rigall, E., Chen, X., Gao, F., Dong, J., and Chen, S.: Physics-guided generative adversarial networks for sea subsurface temperature prediction, *IEEE transactions on neural networks and learning systems*, 2021.
- 430 Minnett, P., Alvera-Azcárate, A., Chin, T., Corlett, G., Gentemann, C., Karagali, I., Li, X., Marsouin, A., Marullo, S., Maturi, E., et al.: Half a century of satellite remote sensing of sea-surface temperature, *Remote Sensing of Environment*, 233, 111 366, 2019.
- Mohan, A. T., Lubbers, N., Livescu, D., and Chertkov, M.: Embedding hard physical constraints in neural network coarse-graining of 3D turbulence, *arXiv preprint arXiv:2002.00021*, 2020.
- Pearson, K., Good, S., Merchant, C. J., Prigent, C., Embury, O., and Donlon, C.: Sea surface temperature in global analyses: Gains from the Copernicus Imaging Microwave Radiometer, *Remote Sensing*, 11, 2362, 2019.
- 435 Pisano, A., Marullo, S., Artale, V., Falcini, F., Yang, C., Leonelli, F. E., Santoleri, R., and Buongiorno Nardelli, B.: New evidence of Mediterranean climate change and variability from sea surface temperature observations, *Remote Sensing*, 12, 132, 2020.
- Rio, M.-H., Santoleri, R., Bourdalle-Badie, R., Griffa, A., Piterbarg, L., and Taburet, G.: Improving the altimeter-derived surface currents using high-resolution sea surface temperature data: a feasibility study based on model outputs, *Journal of Atmospheric and Oceanic Technology*, 33, 2769–2784, 2016.
- 440 Singha, S., Bellerby, T. J., and Trieschmann, O.: Satellite oil spill detection using artificial neural networks, *IEEE Journal of selected topics in applied earth observations and remote sensing*, 6, 2355–2363, 2013.
- Wang, Z., Bovik, A. C., Sheikh, H. R., and Simoncelli, E. P.: Image quality assessment: from error visibility to structural similarity, *IEEE transactions on image processing*, 13, 600–612, 2004.
- 445 Warner, T. T., Lakhtakia, M. N., Doyle, J. D., and Pearson, R. A.: Marine atmospheric boundary layer circulations forced by Gulf Stream sea surface temperature gradients, *Monthly Weather Review*, 118, 309–323, 1990.
- Woollings, T., Hoskins, B., Blackburn, M., Hassell, D., and Hodges, K.: Storm track sensitivity to sea surface temperature resolution in a regional atmosphere model, *Climate dynamics*, 35, 341–353, 2010.
- Yang, C., Leonelli, F. E., Marullo, S., Artale, V., Beggs, H., Buongiorno Nardelli, B., Chin, T. M., De Toma, V., Good, S., Huang, B., et al.: Sea surface temperature intercomparison in the framework of the Copernicus Climate Change Service (C3S), *Journal of Climate*, 34, 5257–5283, 2021.
- 450 Zanna, L. and Bolton, T.: Data-driven equation discovery of ocean mesoscale closures, *Geophysical Research Letters*, 47, e2020GL088 376, 2020.
- Zanna, L., Brankart, J., Huber, M., Leroux, S., Penduff, T., and Williams, P.: Uncertainty and scale interactions in ocean ensembles: From seasonal forecasts to multidecadal climate predictions, *Quarterly Journal of the Royal Meteorological Society*, 145, 160–175, 2019.
- 455