Ocean Science

Deep learning for the super resolution of Mediterranean sea surface temperature fields

Claudia Fanelli[1], **Daniele Ciani**[2], **Andrea Pisano**[2], and **Bruno Buongiorno Nardelli**[1]

[1]Consiglio Nazionale delle Ricerche, Istituto di Scienze Marine (CNR-ISMAR), Calata Porta di Massa, 80133 Naples, Italy
[2]Consiglio 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's climate system. Being at the air–sea interface, SST modulates heat fluxes in and out of the ocean, provides insight into several upper and interior ocean dynamical processes, and 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 and rainy conditions. Large gaps are present even in merged multi-sensor satellite products, and different statistical strategies, mostly based on optimal interpolation (OI) algorithms, have thus been proposed to obtain gap-free (L4) images. 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 $1/100°$ spatial resolution starting from $1/16°$ 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 indeed plays 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 for assessing and investigating mesoscale and submesoscale variability (e.g., Bowen et al., 2002; Isern-Fontanet et al., 2006; González-Haro and Isern-Fontanet, 2014; Rio 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 a 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 by sunglint, rain, radio frequency interference or proximity to land, but the lower spa-

tial resolution achievable with present platforms represents a significant disadvantage (for instance, the footprint of each sample of the NASA Advanced Microwave Scanning Radiometer for Earth Observing System is an ellipse of approximately 45 km × 64 km in size). 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), which is expected to be released during 2029. Consequently, SST fields at high resolution are generally affected by several data voids. For this reason, a few statistical techniques, mostly based on optimal interpolation (OI), have been developed to obtain gap-free SST images (Bretherton et al., 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 kilometers (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.

In this context, here we investigate 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 and managing large-scale datasets while controlling the computational efficiency. In the field of oceanography, the research community only recently started to explore the applicability of machine learning (ML) methods to ocean remote sensing images (Dong et al., 2022b). The applications include oil spill (Singha et al., 2013) and eddy detection (Lguensat et al., 2018; Duo et al., 2019) and parametrization (Bolton and Zanna, 2019), marine algae species discrimination (Balado et al., 2021; Cui et al., 2022), 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 by applying neural networks (NNs) to space–time interpolation and short-term forecasting issues with satellite altimetry data (Fablet et al., 2021), high-performance descriptions of turbulence processes (Mohan et al., 2020; Zanna and Bolton, 2020), and 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 level of 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-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 two-dimensional 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 converting them to feature maps and the non-linear 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 Super Resolution Convolutional Neural Network (SRCNN) developed by Dong et al. (2015) achieved excellent results 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 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, which was 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 the presence of a large number of 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., 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 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) by learning from both low-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 by directly training our network on remote sensing SST data for both input and target datasets. We make use of the data produced by the Institute of Marine Sciences of the Italian National Research Council (CNR-ISMAR) within the Copernicus Marine Service, which consists of merged multi-sensor (L3S) and gap-free (L4) sea surface temperature products over the Mediterranean Sea at a high resolution (HR; nominally 1/16°) and ultra-high resolution (UHR; nominally 1/100°) (Buongiorno Nardelli et al., 2013). Considering that our UHR interpolation accounts for space–time decorrelation scales of 5 km and 1 d, in the presence of valid L3 UHR observations, the resulting L4 data can be considered to be submesoscale-resolving (Kurkin et al., 2020). UHR OI processing makes use of upsized L4 HR fields as an initial guess. This guess is left unchanged in the absence of L3 UHR data and the L4 effective resolution is thus, by definition, lower than 1/100° there. Therefore, our goal is to enhance the effective resolution of the gap-free upsized 1/16° 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 test datasets; ensure a sufficient generalization capability; and, more specifically, avoid under- and over-fitting problems. Under-fitting occurs when the model fails to achieve a suitably low error in 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 used 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 the 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 Fig. 1.

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° and 1/100° 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). The output of each of these processing paths is an L3S product at the associated spatial resolution. 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 guesses in the 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° resolution, in the case of the finer-resolution product, the initial guess used by the OI algorithm is the L4 HR SST field, preliminarily upsized onto a 1/100° regular grid (through a thin plate spline). As a consequence, due to the small decorrelation scales assumed in 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 that is equal to or coarser than 1/16°.

We train the network to improve the satellite SST effective resolution by introducing realistic small-scale features in the interpolated and upsized gap-free L4 HR SST images (dilated Adaptive Deep Residual Network for Super Resolution, dADRSR, input in Fig. 1). The target data are derived from a ground-truth super-collated L3S UHR SST dataset (namely, merged multi-sensor data) that was specifically built for this purpose. The dataset (hereinafter called L3S Sentinel) 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 Sentinel-3B satellites due to their high radiometric accuracy and kilometer-scale resolving capabilities (Coppo et al., 2020).

Both input and target datasets are mapped on a regular grid at 1/100° spatial resolution over the Mediterranean Sea for the 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). Then, all 100 km × 100 km regions (herein referred to as tiles) consisting of at least 95 % valid values were extracted from the filtered fields. Tiles were allowed to overlap at most one other tile by 50 %. 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.,

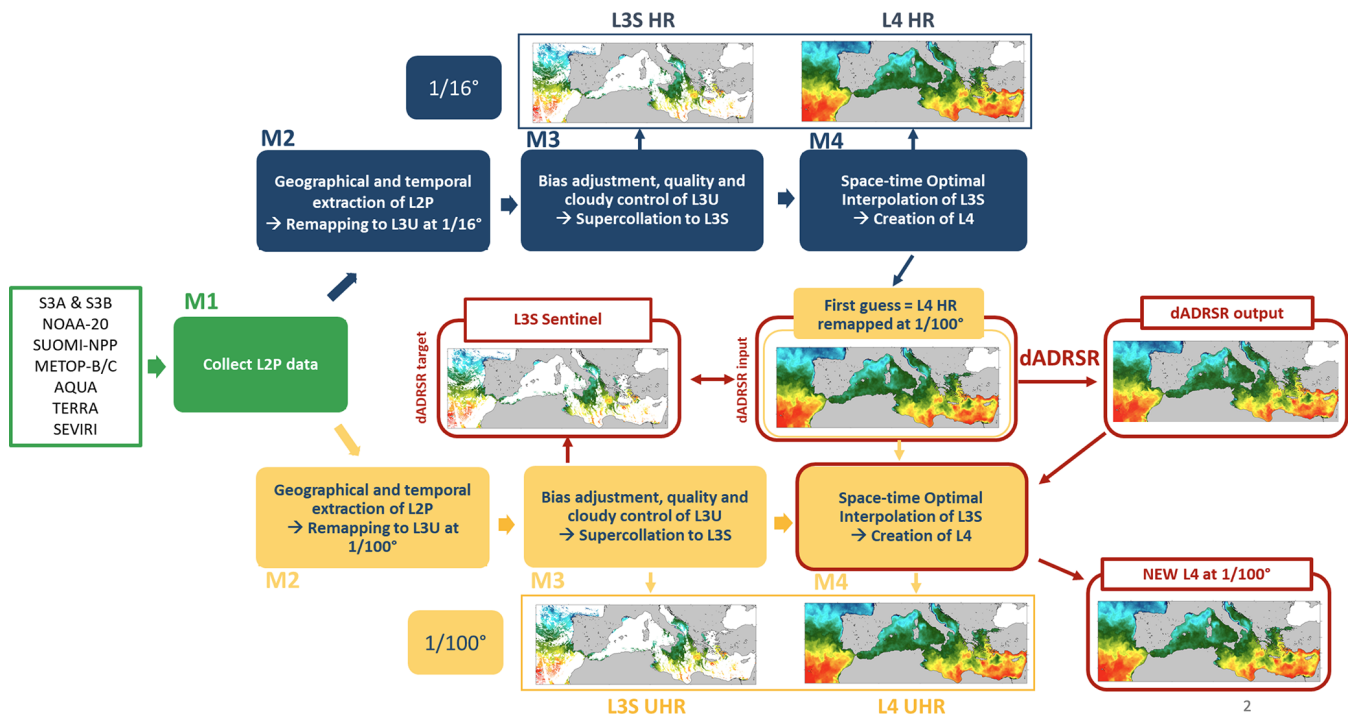$$SST_{norm} = 2 \left(\frac{SST - \min(SST)}{\max(SST) - \min(SST)} \right) - 1. \tag{1}$$

Figure 1. Sketch of the processing chains for the production of L3S and L4 SST fields at 1/16° and 1/100° spatial resolutions 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.

The test dataset is finally selected, setting aside 4 individual days which are representative of the different seasons and composing the 15 % of the total number of the tiles available after the preprocessing. These data are separated in a fully independent dataset (i.e., data never used in the training phase), which is composed of 18 576 overlapping tiles. The training dataset finally consists of 94 110 pairs of tiles extracted from the remaining 362 d, of which one part (properly separated from the original training dataset at the beginning of the training phase following an 85 : 15 ratio between the two) is used by the network in the validation phase. In Fig. 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 (i.e., L4 HR upsized at 1/100° spatial resolution) on 4 January 2020 and below the corresponding L3S Sentinel image derived from merged Sentinel-3A and Sentinel-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 the input and target to train the network, respectively.

2.2 Super-resolution convolutional neural network

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 previous layer output X; that is, the following applies:

$$Y = F(W_i * X + B_i), \tag{2}$$

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 feature map) and might include, for instance, Gaussian-like filters or edge detectors along different direction or any other kind of filter learned during the training. Due to each of the layers having its own functionality, they 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 when exploiting the potential of deep networks. The CNN implemented here was originally developed by Buongiorno Nardelli et al. (2022) and
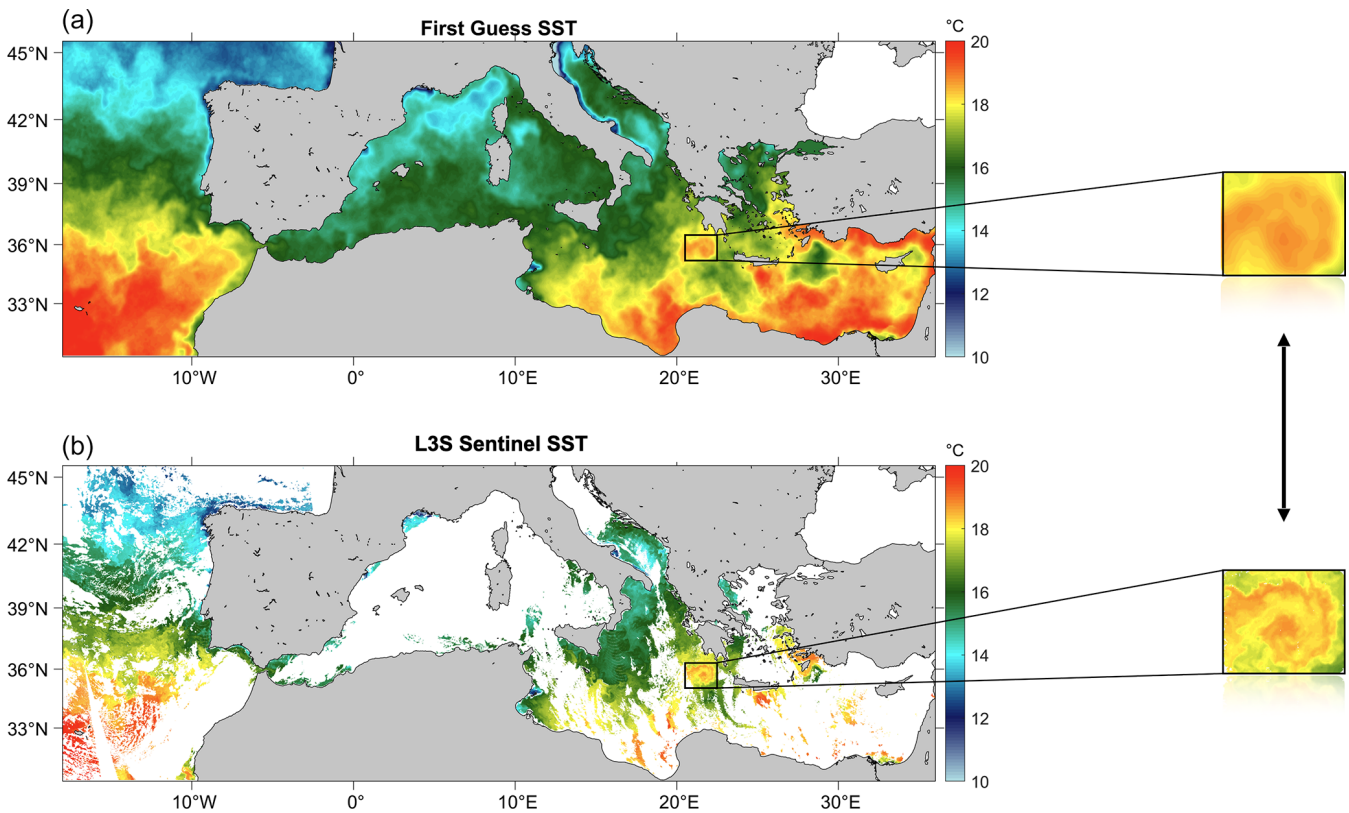
Figure 2. On the left are the SST first-guess maps used to extract the input tiles **(a)** and the SST L3S Sentinel target image **(b)** on 4 January 2020. On the right are examples of extracted tiles used for the training and test datasets.
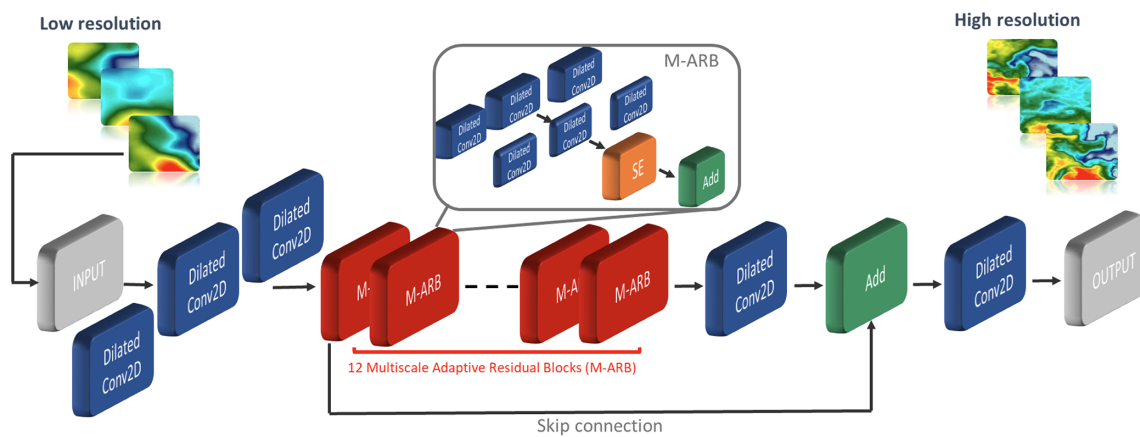

Figure 3. Schematic of the dilated adaptive deep residual network for super resolution developed by Buongiorno Nardelli et al. (2022). Conv stands for convolutional, M-ARB for multi-scale adaptive residual block, SE for squeeze-and-excitation module, and Add indicates aggregations of the outputs from the networks blocks.

called dilated Adaptive Deep Residual Network for Super Resolution (dADRSR). In the dADRSR 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 for 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 12 multi-scale adaptive residual blocks (M-ARBs), each including two sets of parallel dilated convolutional layers (with 120 and 10 filters, respectively) and a squeeze-and-excitation (SE) module that is 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 involve 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 a 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 to 20 here). 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 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 a reference in the loss function. Finally, the dADRSR training model 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 d.

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.

Firstly, we consider the classical root-mean-squared error (RMSE) given by

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

where x_i and y_i are the SST values of the ith pixel of images x and y, respectively, and N is the total number of pixels. This measure can be useful for evaluating 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, which is even worse than if it

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 that are 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 image I and the power of distorting noise that affects the quality of its representation, usually represented by the RMSE itself:

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

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), which is 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 RMSE comes out nearly identical for all the cases, the SSIM will capture the different perceptual quality, which is a weighted combination of luminance, contrast and structure measurements. The mathematical formulation of the SSIM between two images, x and y, is given by

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

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 by 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 reproduce the large-scale structure of the original SST field.

3 Results and discussion

Our aim is to verify whether the dADRSR 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 dADRSR network in comparison with other deep learning algorithms and a sensitivity analysis on the impact of different choices for the architecture are presented in Sect. 3.1. Finally, an additional test on an independent dataset built on a 1-year-long series of SST fields was carried out (Sect. 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 dADRSR network at $1/100°$ 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 Sentinel-3B (central panel) and the first-guess map used in the optimal interpolation algorithm (bottom panel) are shown below the network output. What is 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.

To highlight the ability of the CNN to more accurately capture small-scale features with respect to the statistical algorithm, we show in Fig. 5 three smaller panels corresponding to zoomed-in regions of the fields shown in Fig. 4, delimited by the coordinates 30–40° N, 19–36° E, with the correspondent SST spatial gradients (calculated using the Sobel operator) for the dADRSR reconstruction, the L3S Sentinel 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 L3S Sentinel 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 submesoscale. While the CNN is able to capture at least the most energetic structures (as shown in the top panels of Fig. 5), the first guess presents an extremely smooth field due to the OI algorithm applied where all the small-scale features have been filtered out even when high-resolution data are present. We recall that the Sobel gradient operates on a 3 pixel × 3 pixel kernel, a spatial scale not readily visible in the figures. This is why the first-guess gradient field is so much weaker than the CNN

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 Eq. (3) and the corresponding confidence interval calculated by the bootstrapping procedure, the peak signal-to-noise ratio obtained by Eq. (4), and the structural similarity index measure given by Eq. (5).

Model	RMSE [°C]	PSNR	SSIM
First guess	$0.33 \pm 7 \times 10^{-5}$	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

or SLSTR fields even though the large-scale structures are similar.

This visual analysis is quantitatively confirmed by the maps in Fig. 6, displaying the difference between the error made by the low-resolution approximation and the dADRSR model with respect to the original L3S Sentinel image, averaged to $1° \times 1°$ boxes. Here, red indicates an improvement of the CNN with respect to the first-guess image and blue a degradation. A clear predominance of red boxes is found in both the SST and the SST gradient error maps.

In Table 1, the comparison is summarized quantitatively, with the network reconstruction presenting a RMSE = 0.31 °C (versus the 0.33 °C obtained from the first-guess approximation) and a mean PSNR equal to 37.9 and a SSIM equal to 0.54, both larger than the first-guess result. To highlight the quality of the dADRSR reconstruction, we performed the same test using other deep learning super-resolution models: the enhanced deep residual network for super resolution (EDSR) proposed by Lim et al. (2017), the adaptive deep residual (ADR) network for super resolution developed by Liu et al. (2019) and the dADRSR proposed by Buongiorno Nardelli et al. (2022), setting the number of M-ARBs to be equal to 6 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 dADRSR output achieves the best value for all the evaluation methods presented.

To analyze the effectiveness of the high-resolution reconstruction, we compare the power spectral density (PSD) of the network output with the L3S Sentinel product and the first-guess map over three selected zones. We compute the PSD via a fast Fourier transform (FFT) with a Blackman–Harris window over the three areas corresponding to the boxes labeled a, b and c in Fig. 7. The zones are chosen in order to have the maximum number of valid pixels available
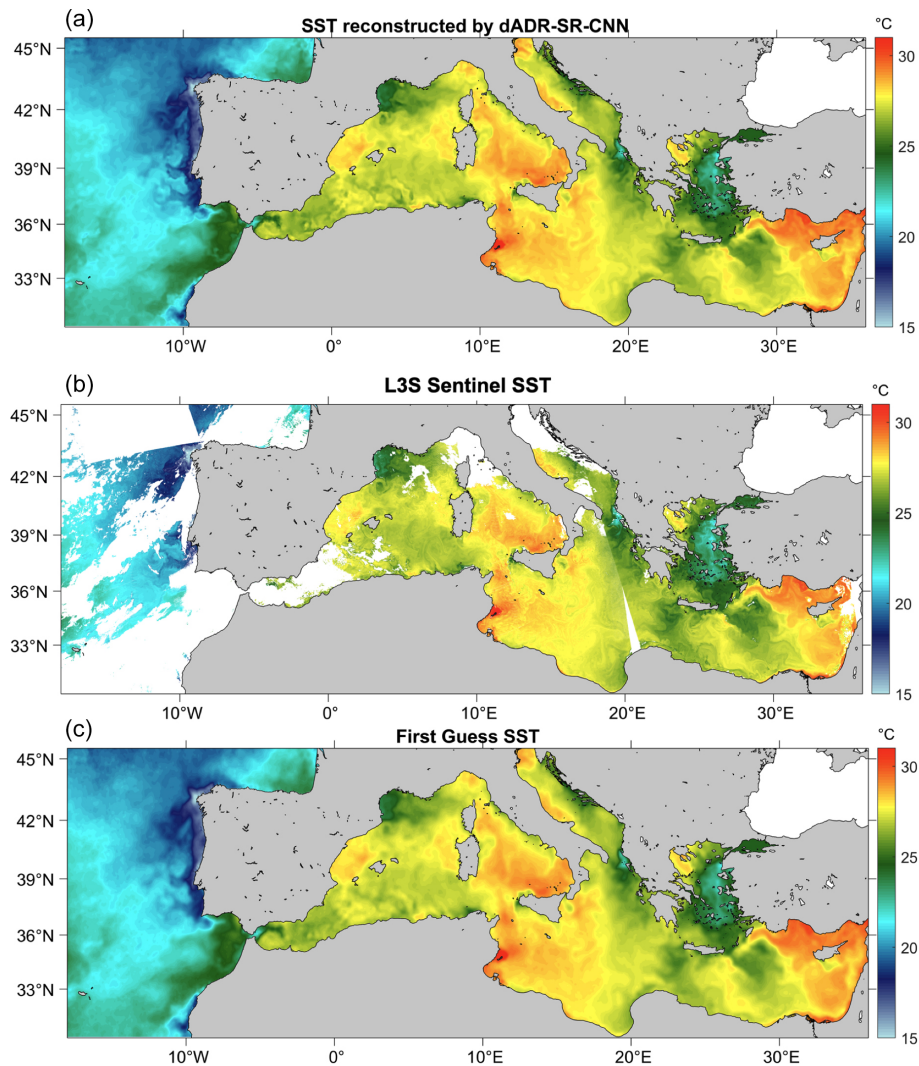
Figure 4. Comparison of the SST fields of 1 August 2020 provided by the reconstruction of the dADRSR **(a)**, the L3S data measured by Sentinel-3A and Sentinel-3B **(b)**, and the first-guess SST field **(c)**.

in the L3S Sentinel observations and represent the following different dynamical regimes:

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, and

c. a region over the Levantine Sea characterized by small-scale structures.

The three central panels of Fig. 7 show the PSD, presented as a function of the wavenumber, of the SST reconstructed by the dADRSR (in yellow), the first-guess map (in red) and the high-resolution L3S Sentinel observations (in blue) over the three zones delimited by the black rectangles over the L3S SST field on the top panel. In all these cases, note that

the PSD follows the same behavior for wavenumbers smaller than 1 per degree, i.e., for scales larger than 100 km. This means that for such scales, the SST fields obtained by both the CNN and the first-guess characterize large mesoscale features well. Conversely, for regions a and b, both the first-guess and the network reconstruction exhibit a significant PSD decrease for wavenumbers higher than 1 per degree, but starting from approximately 10 per degree, the first-guess spectrum separates from the dADRSR spectrum (and consequently increases the distance from the L3S Sentinel one), indicating a poorer reconstruction of spatial features below 10 km. The network spectrum, on the other hand, shows that the machine learning algorithm captures those small-scale features in all the cases better. Analogous behavior is found for the spectra of the SST gradients over the same three regions (three bottom panels). The abrupt decreases seen in the L3S Sentinel and first-guess spectra starting at 3×10 per
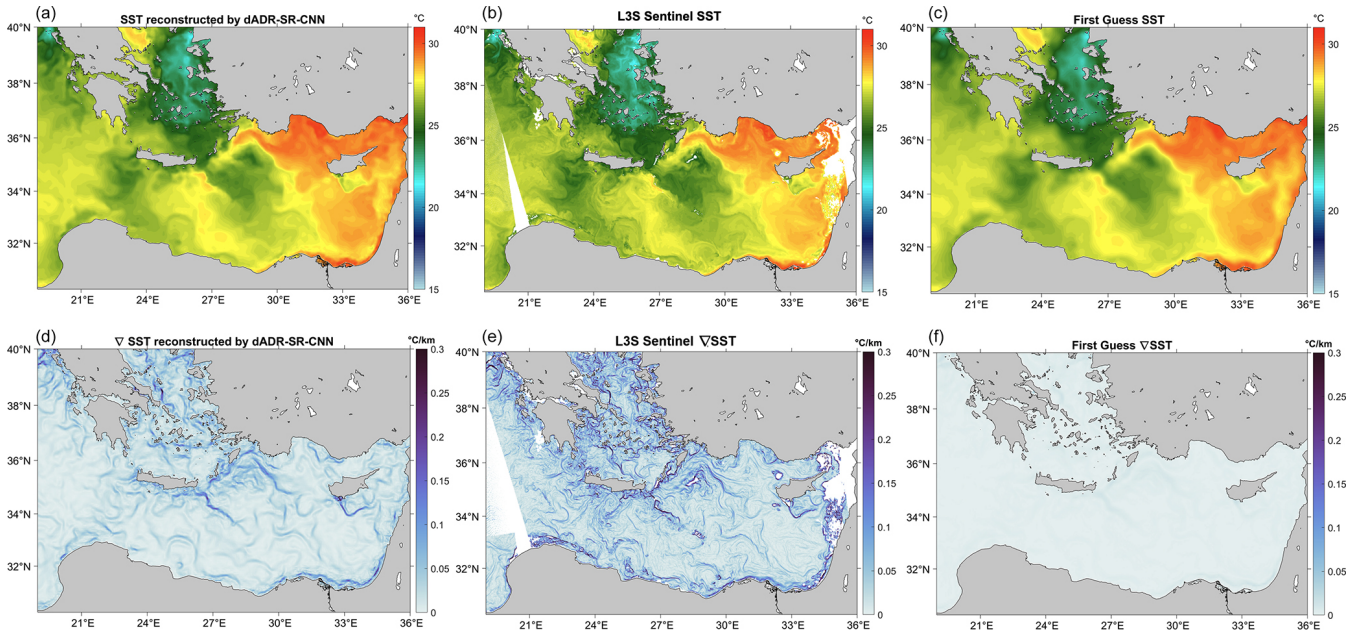
Figure 5. Comparison of the SST **(a, d)** and SST gradients **(c, f)** provided by the reconstruction of the dADRSR **(a–c)**, the L3S data measured by Sentinel-3A and Sentinel-3B **(b, e)**, and the optimal interpolated first guess **(d–f)** in the selected region delimited by the coordinates 30–40° N, 19–36° E.
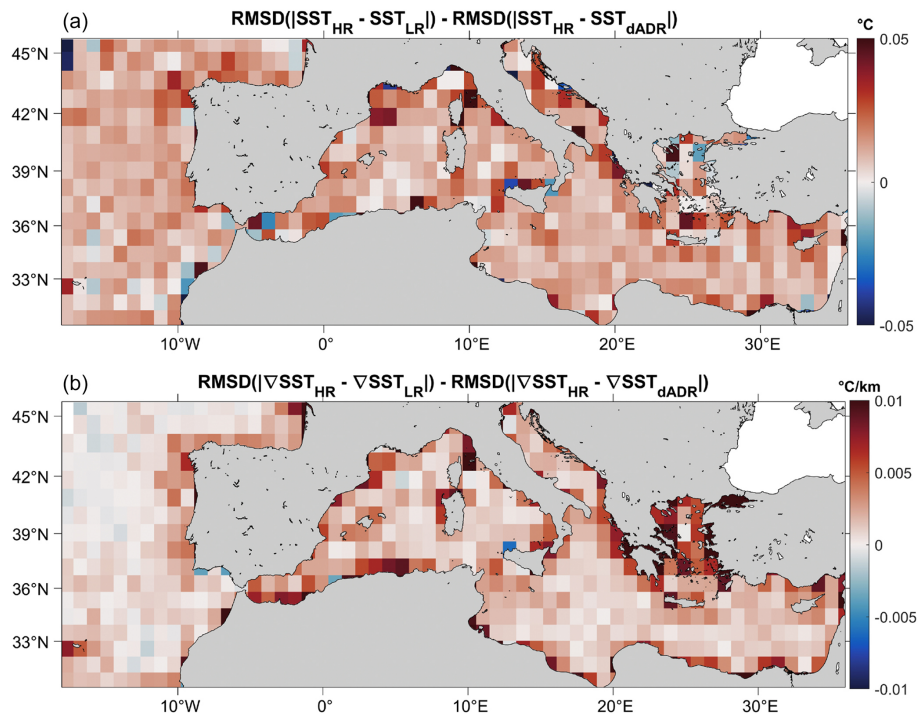

Figure 6. Comparison of the performance of the SST **(a)** and ∇SST **(b)** dADRSR reconstruction and the first guess with respect to the L3S data measured by Sentinel-3A and Sentinel-3B satellites. Positive red values show an improvement of the network reconstruction with respect to the optimal interpolated first guess.
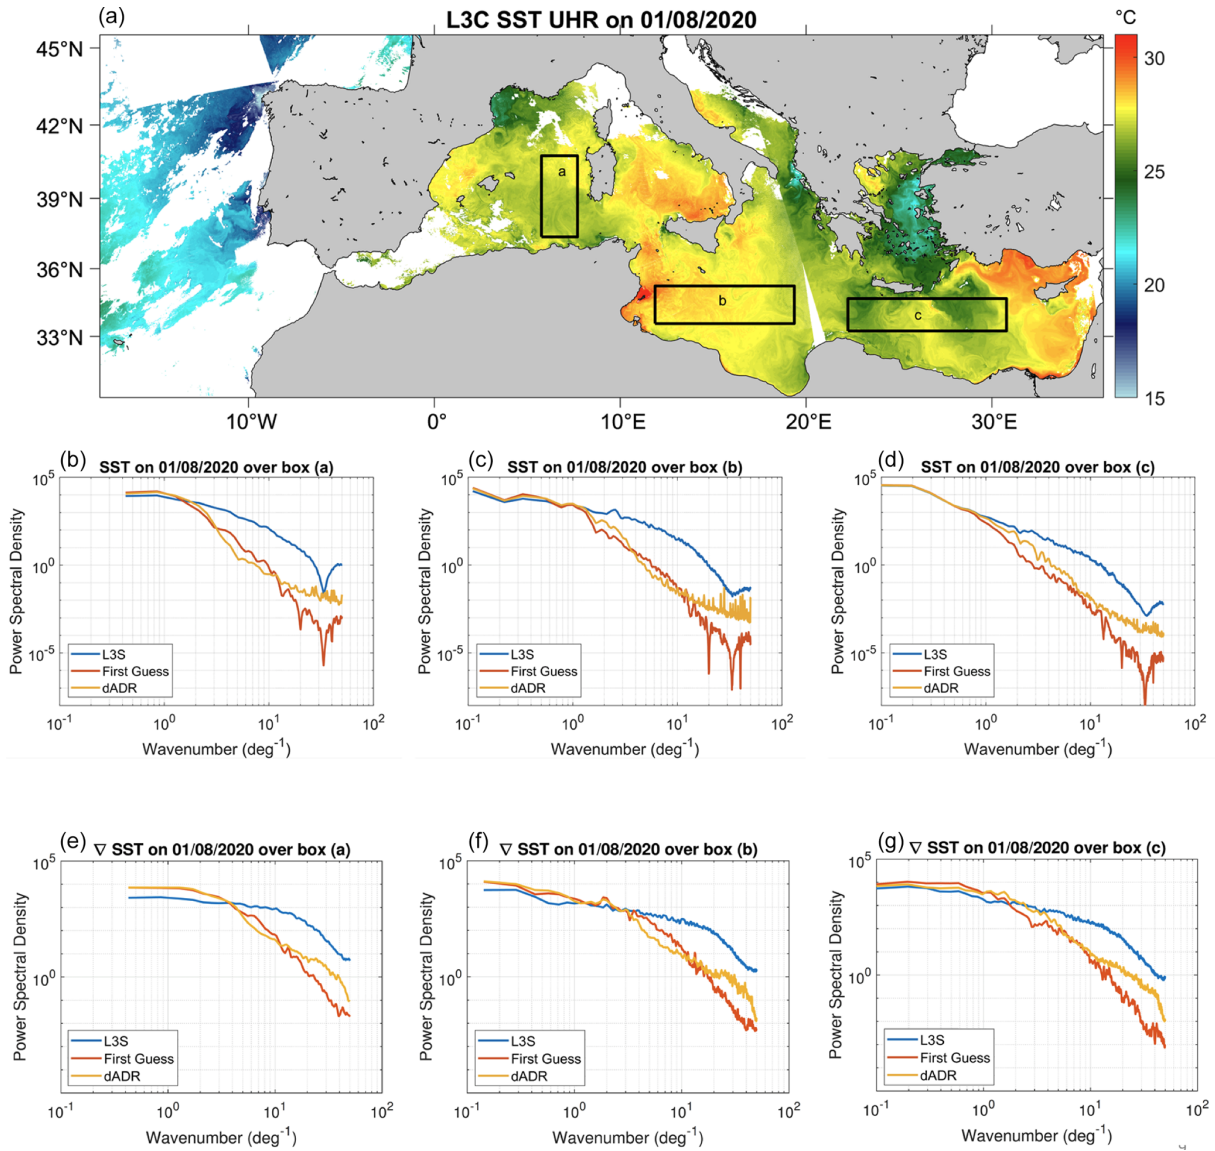
Figure 7. The PSD of the SST **(b–d)** and SST gradients **(e–g)** reconstructed by the dADRSR (in yellow), the first-guess map (in red) and the L3S Sentinel high-resolution observations (in blue) over the three zones delimited by black rectangles on the L3S SST field on 1 August 2020 in panel **(a)**.

degree are probably due to artifacts introduced by the regridding, as already discussed by Liberti et al. (2023), leading to the lack of physical meaning for the spectra from this point onwards (i.e., for the highest wavenumbers). Overall, we can conclude that the CNN spectra tracks the observed spectra from 1 to 10 per degree quite well, while the first-guess spectra are more energetic in this region. On the other hand, although the CNN spectra are less energetic than the L3S Sentinel spectra for wave numbers greater than 10 per degree, they are more representative of the true spectra in this region than the first-guess spectra.

To quantify the differences between the effective resolution of the products, we calculated the PSD ratio

(RPSD) between the spectral content of the mapping error (e.g., $SST_{L3S} - SST_{dADRSR}$) and the spectral content of the L3S Sentinel observation of the SST for the dADRSR-reconstructed fields and the first-guess map (see Fig. 8) as follows:

$$RPSD(\lambda) = \frac{PSD_{diff}(\lambda)}{PSD_{L3S}(\lambda)}, \tag{6}$$

where λ is the wavelength, PSD_{L3S} is the spectra of the L3S Sentinel observations, and PSD_{diff} is the spectra of the differences between the L3S data and either the output of the CNN or the first-guess field. The RPSD defined in Eq. (6) is calculated over the three zones delimited by black rectangles
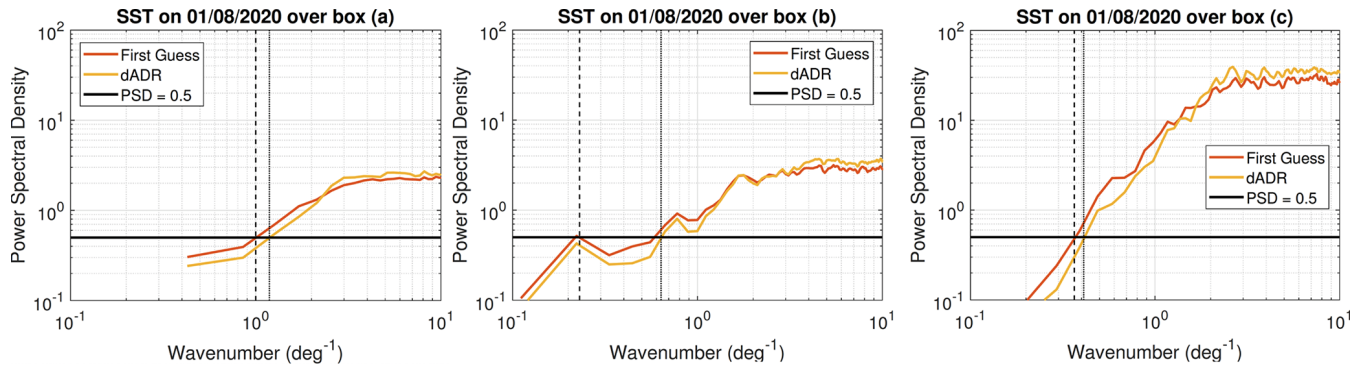
Figure 8. The ratio between the spectral content of the mapping error and the spectral content of the L3S observation of the SST for the dADRSR reconstructed fields (in yellow) and the first-guess map (in red) over the three zones delimited by black rectangles on the L3S SST field on 1 August 2020 presented in Fig. 7a. The black line represents the ratio equal to 0.5 and the dashed black line and the dotted black line represent the intersection between PSD ratio of the first guess and the dADRSR, respectively, showing the effective resolution of the products.
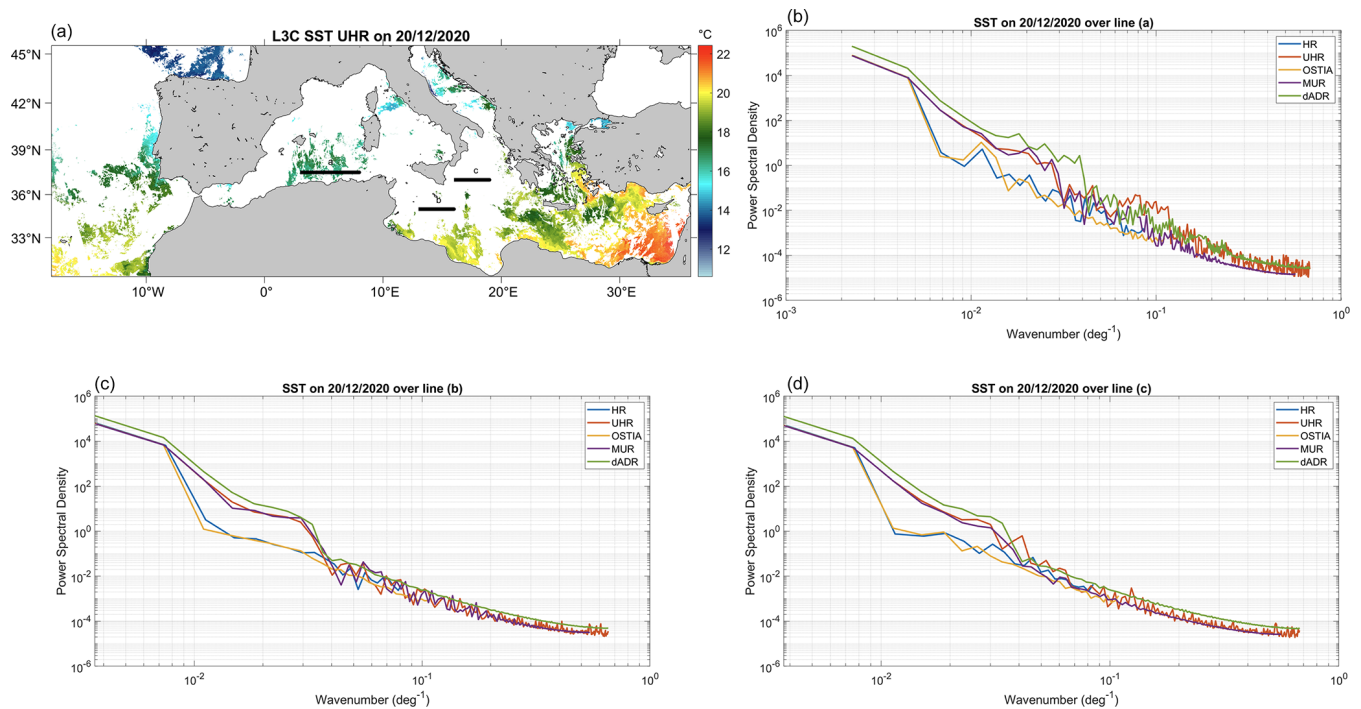
Figure 9. The power spectral density profiles **(c, d)** of different SST products under cloudy conditions calculated along the three transects (black lines) on the L3S SST field on 20 December 2020 **(a, b)**. The MED L4 NRT HR product at 1/16° (Copernicus) is in blue, the MED L4 NRT UHR at 1/100° (Copernicus) is in red, the GLOBAL OCEAN OSTIA product at 0.05° (Copernicus) is in yellow, the Multi-scale Ultra-high Resolution (MUR) product at 0.01° (NASA-JPL) is in purple and the dADRSR reconstruction developed in this work is in green.

on the L3S SST field on 1 August 2020 presented in the bottom panel of Fig. 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 dADRSR network to resolve smaller scales.

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 (Fig. 9). The products used are the L4 NRT HR at a nominal 1/16° resolution and the L4 NRT UHR at a nominal 1/100° 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 a 0.05° resolution (https://doi.org/10.48670/moi-00165), the Multi-scale Ultra-high Resolution (MUR) prod-
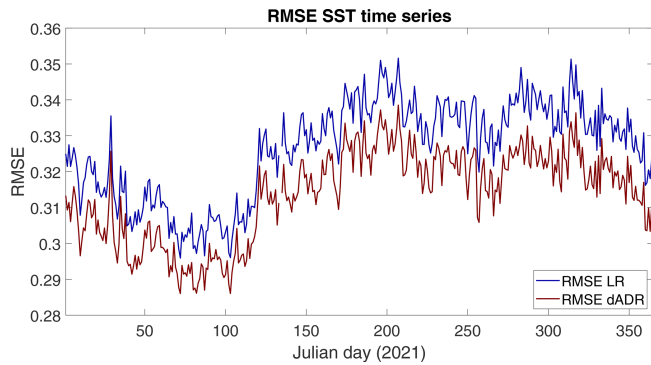
Figure 10. Comparison of the daily RMSE time series of the SST obtained via dADRSR reconstruction (red) and the L4 first guess (blue) with respect to the L3S data measured by Sentinel-3A and Sentinel-3B satellites during the year 2021. The jump at Julian day 134 is due to the absence of L3 observation for that day, which made the comparison with the input and the output of the network impossible.

uct provided by the NASA-JPL at a 0.01° resolution (Chin et al., 2017), and the super-resolved SST field obtained by the application of the dADRSR network developed in this work at a 0.01° resolution. The PSD are computed along the three transects (a, b and c in Fig. 9a), chosen in three areas affected (at different levels) by cloud coverage. In all cases, the slope of the spectra is very similar for the three products at higher resolution (namely, the UHR, the MUR and the dADRSR approximations), while lower values are observed for the HR and OSTIA products (which do not reach wavenumbers larger than 0.1 per degree by construction). The PSD analysis for section a (Fig. 9b) shows that the dADRSR approximation is not affected by excessive noise as are the other products, presenting PSD values that are higher than the corresponding values of the MUR and UHR spectra (especially at small scales). For sections b (Fig. 9c) and c (Fig. 9d), the green line representing the dADRSR spectrum lies above all the other lines for almost the entire wavenumber range.

3.2 An extensive test dataset

Having established the potential ability of the network, we performed an additional validation test with the aim to strengthen the robustness of the statistics. We built a new totally independent dataset exploiting the L4 SST fields used as an initial guess for the OI algorithm, which produces gap-free SST maps over the Mediterranean Sea for the year 2021 at a 0.01° × 0.01° spatial resolution. Before applying the dADRSR network to this dataset, the fields were preprocessed following the same steps described in Sect. 2.1. The comparison between the evolution of the daily basin-scale RMSE for the input first guesses and the dADRSR reconstructions for the entire year confirms the improvement of the network output with respect to the first-guess fields

(Fig. 10). The mean values of the RMSE for 2021 are in line with the ones found for the previous test (i.e., 0.31 °C for the dADRSR output versus 0.33 °C for the first-guess approximation).

Analogous to Fig. 6, the differences (averaged on 1° × 1° boxes) between the errors in the two methods approximating the L3S Sentinel-observed data (Fig. 11) present an overall improvement of the output of the dADRSR network with respect to the interpolated maps. The positive red values are found almost everywhere for both SST and SST gradient 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 to 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 machine learning approach used here exploits the progress made in the field of computer vision for extrapolating high-resolution features even when a direct measurement is missing. Learning directly from L3S SST fields and taking advantage of both dilated convolution and attention mechanisms, the deep neural network employed here proved to be able to reproduce small-scale signals that are generally smoothed out by optimal interpolation algorithms. The strong variability in the SST in the Mediterranean Sea allowed us to obtain excellent results even by only considering 1 year of data during the training phase. However, it is 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 in the SST gradients in the loss function. Another aspect that deserves further investigation concerns the applicability of the dADRSR network to differentiating between sea and ocean areas even though fine-tuning the model would probably be needed.

In the future, we also plan to study other super-resolution techniques which have recently become popular in the field of computer vision. Firstly, we are currently investigating the possibility of improving the reconstruction of small-scale features in SST fields via other successful generative AI, such as GANs or diffusion models. The former exploits the
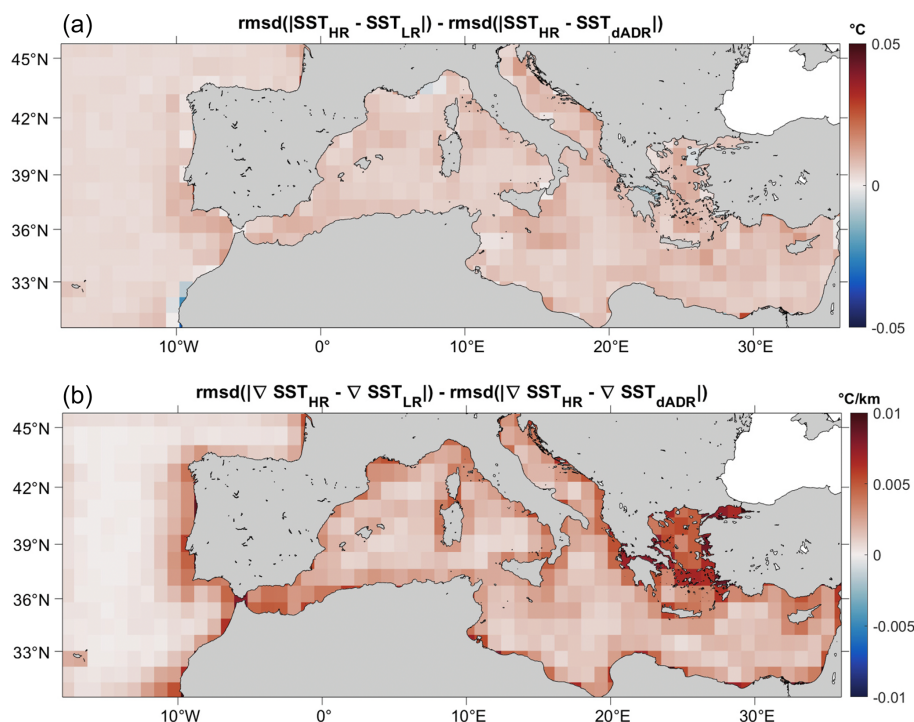
Figure 11. Comparison of the performance of the SST (**a**) and ∇SST (**b**) dADRSR reconstruction and the L4 first guess with respect to the L3S data measured by Sentinel-3A and Sentinel-3B satellites during the year 2021. Positive red values show an improvement of the network reconstruction with respect to the optimal interpolated first guess.

outcomes of a min–max 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, only to then reverse this process until it converges to the desired distribution. However, although such networks applied to remote sensing data have proven to be able to reconstruct very realistic small-scale structures, they seem to fail to optimize a point-match evaluation of reconstructed remotely sensed SST fields. 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 two-dimensional/three-dimensional 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). Enhancement of the effective resolution of SST data and especially SST gradients may also benefit data assimilation in forecast modeling given their proven sensibility to small structures of sea surface temperature (Maloney and Chelton, 2006; Woollings et al., 2010). 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. CF, AP and BBN: conceptualization; CF, DC, AP and BBN: methodology; CF, DC, AP and BBN: investigation; CF: visualization and writing (original draft); and CF, DC, AP and BBN: writing (review and editing). All authors have read and agreed to the published version of the paper.

References

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, Int. J. Remote Sens., 42, 1785–1800, 2021.

Ballarotta, M., Ubelmann, C., Pujol, M.-I., Taburet, G., Fournier, F., Legeais, J.-F., Faugère, Y., Delepoulle, 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.

Bolton, T. and Zanna, L.: Applications of deep learning to ocean data inference and subgrid parameterization, J. Adv. Model. Earth Sy., 11, 376–399, 2019.

Bowen, M. M., Emery, W. J., Wilkin, J. L., Tildesley, P. C., Barton, I. J., and Knewtson, R.: Extracting multiyear surface currents from sequential thermal imagery using the maximum cross-correlation technique, J. Atmos. Ocean. Tech., 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 MODE-73, Deep Sea Res., 23, 559–582, 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 Sens. Environ., 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 Sens.-Basel, 14, 1159, https://doi.org/10.3390/rs14051159, 2022.

Castro, S. L., Emery, W. J., Wick, G. A., and Tandy, W.: Submesoscale sea surface temperature variability from UAV and satellite measurements, Remote Sens.-Basel, 9, 1089, https://doi.org/10.3390/rs9111089, 2017.

Chang, Y. and Cornillon, P.: A comparison of satellite-derived sea surface temperature fronts using two edge detection algorithms, Deep-Sea Res. Pt. II, 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 Sens. Environ., 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 Sens.-Basel, 12, 1601, https://doi.org/10.3390/rs12101601, 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 satellite, Appl. 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 Sens.-Basel, 14, 710, https://doi.org/10.3390/rs14030710, 2022.

Deo, M. and Naidu, C. S.: Real time wave forecasting using neural networks, Ocean Eng., 26, 191–203, 1998.

Deser, C., Alexander, M. A., Xie, S.-P., and Phillips, A. S.: Sea surface temperature variability: Patterns and mechanisms, Annu. Rev. Mar. Sci., 2, 115–143, 2010.

Dong, C., Loy, C. C., He, K., and Tang, X.: Image super-resolution using deep convolutional networks, IEEE T. Pattern Anal., 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., Lin, X., and Zou, B.: The near-global ocean mesoscale eddy atmospheric-oceanic-biological interaction observational dataset, Scientific Data, 9, 436, https://doi.org/10.1038/s41597-022-01550-9, 2022a.

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, 9870950, https://doi.org/10.34133/2022/9870950, 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 Recogniton in Remote Sensing (PRRS), 4 December 2016, Cancun, Mexico, IEEE, 1–6, https://doi.org/10.1109/PRRS.2016.7867019, 2016.

Duo, Z., Wang, W., and Wang, H.: Oceanic mesoscale eddy detection method based on deep learning, Remote Sens.-Basel, 11, 1921, https://doi.org/10.3390/rs11161921, 2019.

European Union-Copernicus Marine Service: Mediterranean Sea High Resolution and Ultra High Resolution Sea Surface Temperature Analysis, Mercator Ocean International [data set], https://doi.org/10.48670/moi-00172, 2008.

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 T. Geosci. Remote, IEEE Transactions on Geoscience and Remote Sensing, 61, 1–14, 2023.

Ghiasi, G., Lin, T.-Y., and Le, Q. V.: Dropblock: A regularization method for convolutional networks, Adv. Neur. In., 31, https://doi.org/10.48550/arXiv.1810.12890, 2018.

González-Haro, C. and Isern-Fontanet, J.: Global ocean current reconstruction from altimetric and microwave SST measurements, J. Geophys. Res.-Oceans, 119, 3378–3391, 2014.

Goodfellow, I., Bengio, Y., and Courville, A.: Deep learning, MIT Press, ISBN: 9780262035613, 2016.

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, 27–30 June 2016, Las Vegas, NV, USA, IEEE, 770–778, https://doi.org/10.1109/CVPR.2016.90, 2016.

Hu, J., Shen, L., and Sun, G.: Squeeze-and-excitation networks, in: Proceedings of the IEEE conference on computer vision and pattern recognition, 18–23 June 2018, Salt Lake City, UT, USA, IEEE, 7132–7141, https://doi.org/10.1109/CVPR.2018.00745, 2018.

Isern-Fontanet, J., Chapron, B., Lapeyre, G., and Klein, P.: Potential use of microwave sea surface temperatures for the estimation of ocean currents, Geophys. Res. Lett., 33, L24608, 2006.

Jha, B., Hu, Z.-Z., and Kumar, A.: SST and ENSO variability and change simulated in historical experiments of CMIP5 models, Clim. Dynam., 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.

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, 27–30 June 2016, Las Vegas, NV, USA, IEEE, 1637–1645, https://doi.org/10.1109/CVPR.2016.181, 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.

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, Russ. J. Earth Sci., 20, ES4008-4008, https://doi.org/10.2205/2020ES000737, 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, 22–27 July 2018, Valencia, Spain, IEEE, 1764–1767, https://doi.org/10.1109/IGARSS.2018.8518411, 2018.

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 Sens.-Basel, 15, 2453, https://doi.org/10.3390/rs15092453, 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, 21–26 July 2017, Honolulu, HI, USA, IEEE, 136–144, https://doi.org/10.1109/CVPRW.2017.151, 2017.

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 T. Geosci. Remote, 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, J. Marine Syst., 68, 405–420, 2007.

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, J. 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 T. Neur. Net. Lear., 34, 3357–3370, https://doi.org/10.1109/TNNLS.2021.3123968, 2021.

Minnett, P., Alvera-Azcárate, A., Chin, T., Corlett, G., Gentemann, C., Karagali, I., Li, X., Marsouin, A., Marullo, S., Maturi, E., Santoleri, R., Saux Picart, S., Steele, M., and Vazquez-Cuervo, J.: Half a century of satellite remote sensing of sea-surface temperature, Remote Sens. Environ., 233, 111366, 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 Sens.-Basel, 11, 2362, https://doi.org/10.3390/rs11202362, 2019.

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 Sens.-Basel, 12, 132, https://doi.org/10.3390/rs12010132, 2020.

Rio, M.-H., Santoleri, R., Bourdalle-Badie, R., Griffa, A., Piterbarg, L., and Taburet, G.: Improving the altimeter-derived sur-

face currents using high-resolution sea surface temperature data: a feasability study based on model outputs, J. Atmos. Ocean. Tech., 33, 2769–2784, 2016.

Singha, S., Bellerby, T. J., and Trieschmann, O.: Satellite oil spill detection using artificial neural networks, IEEE J. Sel. Top. Appl., 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 T. Image Process., 13, 600–612, 2004.

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, Mon. Weather Rev., 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, Clim. Dynam., 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., Merchant, C. J., Sakurai, T., Santoleri, R., Vazquez-Cuervo, J., Zhang, H.-M., and Pisano, A.: Sea surface temperature intercomparison in the framework of the Copernicus Climate Change Service (C3S), J. Climate, 34, 5257–5283, 2021.

Zanna, L. and Bolton, T.: Data-driven equation discovery of ocean mesoscale closures, Geophys. Res. Lett., 47, e2020GL088376, https://doi.org/10.1029/2020GL088376, 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, Q. J. Roy. Meteor. Soc., 145, 160–175, 2019.