
Neural Fields for Fast and Scalable Interpolation of Geophysical Ocean Variables

J. Emmanuel Johnson

Univ. Grenoble Alpes, CNRS UMR IGE, Grenoble, France
johnsonj@univ-grenoble-alpes.fr

Redouane Lguensat

Institut Pierre Simon Laplace, IRD, Sorbonne Université, Paris, France

Ronan Fablet

IMT Atlantique, CNRS UMR Lab-STICC, Brest, France

Emmanuel Cosme

Univ. Grenoble Alpes, CNRS UMR IGE, Grenoble, France

Julien Le Sommer

Univ. Grenoble Alpes, CNRS UMR IGE, Grenoble, France

Abstract

Optimal Interpolation (OI) is a widely used, highly trusted algorithm for interpolation and reconstruction problems in geosciences. With the influx of more satellite missions, we have access to more and more observations and it is becoming more pertinent to take advantage of these observations in applications such as forecasting and reanalysis. With the increase of the volume of available data, scalability remains an issue for standard OI and it prevents many practitioners from effectively and efficiently taking advantage of these large sums of data to learn the model hyperparameters. In this work, we leverage recent advances in Neural Fields (NerFs) as an alternative to the OI framework where we show how they can be easily applied to standard reconstruction problems in physical oceanography. We illustrate the relevance of NerFs for gap-filling of sparse measurements of sea surface height (SSH) via satellite altimetry and demonstrate how NerFs are scalable with comparable results to the standard OI. We find that NerFs are a practical set of methods that can be readily applied to geoscience interpolation problems and we anticipate a wider adoption in the future.

1 Introduction

Maps of sea surface height (SSH) provide a crucial source of information for the oceanographic community. They are used in a number of applications including research involving mesoscale ocean dynamics [1, 2] and their role in global climate change [3, 4]. A good characterization could even lead to a better quantification of the transport of biogeochemical components [4]. Our understanding of SSH characteristics has gotten better not only through our improved understanding of the underlying physics, but also due to the fact that we have more available observations with the introduction of satellite altimetry [5, 6]. However, we only ever have partial observations of SSH through satellite

altimetry, so the community has put a lot of research effort into methods that can help interpolate the missing observations for use in subsequent SSH mapping products [7, 8].

There have been a number of different algorithmic developments that have had success in sea surface height reconstruction [9, 10]. In practice, the most notable of algorithms is the DUACS framework [11] which is a carefully handcrafted solution based on Optimal Interpolation (OI). Developed in 1997, this framework is capable of producing global daily maps of an effective resolution of 100-200km [12, 8]. This single method is solely responsible for producing many Copernicus products¹ such as global maps of sea surface height and sea surface temperature [7, 13, 14, 15]. However, a major downside of the OI framework is that it is computationally expensive with a very high memory consumption even at a GPU scale. This makes it very impractical for training and inference at very large scale especially with the onset of more and more available observation data [16].

In this work, we implement an alternative solution to address the computational limitations of OI: neural fields (NerF) [17]. NerFs are a family of coordinate-based, neural network algorithms which map a set of spatial-temporal coordinates, \mathbf{x}_ϕ , to a scalar or vector, \mathbf{y} , quantity of interest. These methods are parametric versions of the standard OI formulation where the parameters are *learned from data*. This is only possible due to the copious amounts of data we now have given the introduction of more altimetry satellites [16, 12]. We demonstrate that NerFs are a viable family of methods that can interpolate the SSH field given partial observations while also being orders of magnitudes faster and memory efficient than the standard OI solution. All of the preprocessing steps, machine learning modules and training regimes used in this submission is available on *Github*².

2 Methods

2.1 Data

In this work, we are looking at sea surface height as our observed variable of interest. The observation data is recovered from the NADIR track satellite altimetry. In particular, we parsed all data over the Gulf stream area during the time period of 01-01-2017 until 31-12-2017. We get an aggregated total of 7 altimetry tracks from the SARAL/Altika, Jason 2, Jason 3, Sentinel 3A, Cryosat-2, and Haiyang-2A satellite altimeters which is roughly ~ 1.8 million observations in total. The Cryosat-2 altimeter data was left out as an independent validation set. We extract the spatial coordinates (latitude, longitude) and the time coordinates (seconds) as the inputs, \mathbf{x}_ϕ , to our model and the corresponding SSH values as our outputs, y_{obs} . All data is available on the data challenge that has been setup on *Github*³ [18].

2.2 Baselines

As mentioned in the introduction, there are many methods in the literature that have been used for SSH reconstruction [19, 10, 9, 20, 21, 22]. However, the most prominent method that is foundational for the Copernicus products is the DUACS algorithm [7]. The DUACS algorithm is an OI-based scheme which tunes the kernel hyperparameters offline in a physically consistent manner with external altimetry validation datasets. While the standard OI algorithm is an excellent and trusted method for interpolation, this method suffers from scaling issues (training and inference) which can limit its use in practice. The training has a time complexity of $\mathcal{O}(N^3)$ and memory complexity of $\mathcal{O}(N^2)$ which makes this method infeasible without the appropriate hardware [23].

To alleviate this issue with scalability, users apply this algorithm locally by only conditioning on the observations available in the local region of interest and interpolating the remaining observations within that region. The algorithm is run along in a patch-like manner until covering the entire globe. This mosaic of many local patches results in a global map of interpolated SSH values. The local application of the OI method has several advantages: 1) It allows the OI framework to be scaled to larger datasets because it reduces the kernel matrix inversion to only the observations within the local region and 2) it removes spurious long-range correlations that are physically inconsistent [24].

¹DOI:10.48670/moi-00147

²<https://github.com/jejjohnson/ml4ssh>

³DOI:10.5281/zenodo.5511905

In this work, we apply the standard OI method with preconfigured hyperparameters via the local patch-based method described above to serve as the naive baseline which closely mimics the DUACS algorithm. In addition to our naive patch-based OI implementation, we also include the publicly available SSH maps, the solution to the DUACS method, as another baseline method.

2.3 Neural Fields

Neural Fields (NerFs) are implicit models that take coordinate-based inputs (e.g. spatial-temporal) and output a scalar value or vector field for a variable of interest (e.g. sea surface height). This family of neural networks have been very successful in many computer vision applications including image regression and 3D rendering [17] however, their adoption in Earth sciences is still at a very early stage, especially in ocean science applications [25].

Standard neural networks have problems modeling high frequency signals so there were many crucial modifications that were introduced to make them successful. The first attempt included a Fourier features transformation followed by any standard neural network, i.e. the Fourier Feature Network (FFN) [26]. This method was successfully able to capture the high frequency signals in 2D images however the gradients were noisy or nonsensical [27]. Another proposed approach is the SIREN method [27] which uses the sine function as an activation function within a fully connected neural network. This sinusoidal activation has been explored before [28] but the recent ingenuity is the special initialization for the weights and biases enabling the network converge much faster and avoid local minima. This method also results in calculable and relatively smooth first and second order gradients which are more representative of physical systems. This is a non-exhaustive list and there are many extensions to this including the modulated version [29] as well as Multiplicative Filter Networks (MFN) [30]. In this work, we chose the SIREN method for its simplicity and smooth gradients but a full comparison with other NerF variants for regression problems in physical systems would be an excellent follow-up study.

3 Experiments

3.1 Data Preparation

NerFs. We used a coordinate transformation of the spatial data from spherical coordinates (lat,lon) to Cartesian coordinates (x,y,z). After we aggregated the altimetry datasets, we had ~ 1.6 million data points for training and ~ 179 K data points for validation. A final evaluation is done over an entire spatial-temporal grid that is equivalent to DUACS product. The grid resolution is 0.2 degrees spatially and 1 day temporally which results in ~ 1 million data points for the final evaluation. The statistical metrics are computed comparing this grid product to the left out validation altimetry dataset and we plot maps for each algorithm for a visual comparison.

3.2 Hyperparameters

OI. The OI method was trained via the local patch-based version that was outlined in the previous section. A Radial Basis Function (RBF) kernel was used with a length scale per dimension: A length scale of 7 days for the temporal dimension and a length scale of 1 degree (~ 100 km) for the latitude and longitude length scale. The noise likelihood was set to 0.05.

NerF. The SIREN method was trained using the standard Adam optimizer with a base learning rate of $1e-4$. We used a cosine annealing scheduler with a linear warm-up. We used 100 epochs for the linear warm-up with a learning rate between 0 and $1e-4$ followed by 2,000 epochs for the subsequent cosine annealing with a learning rate between $1e-4$ and 0. We used *Weights and Biases* [31] to perform hyperparameters sweeps on the batch size, the number of layers and the width of layers. We used the standard mean squared error as the loss function and all training was done with one Nvidia Tesla V100 GPU.

4 Results

Statistics. A snapshot of predictions, the norm of the gradients and the Laplacian can be seen in figure 1. Each of the SSH maps are very similar as SIREN method can also capture some fine scale

Table 1: Metrics

Algorithm	Normalized RMSE	Effective Resolution	CPU	GPU
OI	0.85 ± 0.09	140 km	1 hr	40 mins
DUACs	0.88 ± 0.07	152 km	N/A	N/A
SIREN	0.88 ± 0.08	136 km	30 secs	5 secs

structures like the OI methods. This is even more evident in the Laplacian maps which result in more non-smooth structures for the SIREN method. These higher order gradients are important for other quantities of interest like the kinetic energy and vorticity and more evaluation is needed to check how each compares method compares in respective energy levels in the spectral domain. In table 1, we see the metrics computed with the independent validation dataset; i.e. the left out Cryosat-2 altimetry track that was unseen by any model. The normalized root mean squared error (NRMSE) and effective resolution (calculated from the signal-to-noise ratio) is comparable between all methods. For more details on the evaluation metrics, see section 3d in [32]. Note that this is simply a local region and more evaluation is needed in other regions and at a global scale.

Timings. In table 1, we see that the SIREN method is the most scalable in terms of inference⁴. The naive OI framework on ~ 1 million data points results in 1 hr on a CPU and 40 minutes on a GPU. However, the NerF algorithm predicts the same amount of data in 30 seconds on a CPU and 5 seconds on a GPU. This results in $\times 120$ speedup and $\times 480$ speedup respectively. This is expected due to the time complexity of matrix inversions, $\mathcal{O}(N^3)$, versus matrix multiplications, $\mathcal{O}(N)$. In addition, the OI method requires one to condition on the observations which results in an additional time and memory footprint that inhibits learning with more observational data.

5 Conclusions/Future Work

Discussion. In this work, we introduced Neural Fields as a viable interpolation method for sea surface height. We demonstrated that we were able to match the results of the standard Optimal Interpolation algorithm by training on observations. Not only did we get results that were statistically comparable to the standard DUACS, we also achieve inference speeds and memory constraints that are viable even with a modest computational budget. Neural networks are often only good when we have a large amount of training data available. However, this work has made a case for NerFs coupled with altimetry data as a candidate algorithm for interpolation at scale.

Future Work. Aside from the more detailed ablation studies related to noise and sampling density, there is much to be done to make these models more appealing and more trustworthy for the broader community. They lack immediate interpretability because they are fully parameterized machines from data. Techniques such as Explainable AI [33] or Physics Informed Neural Nets [34] could be of use to understand and constrain the NN predictions. Quantifying the predictive uncertainty is another main future direction, in fact, there are extensions that have been explored in the literature [35, 36, 37]. Lastly, we observed that the bulk of the computational limit is training these NNs. Under a computational budget, faster training strategies [38] and better hardware would be necessary to make these methods scale to global datasets with billions of data points.

Ethical Considerations

The spatial and temporal resolution of the altimetry datasets used in this study is too coarse to affect the privacy of any individual and we only consider resolutions of $\sim 100 - 200$ km. In addition, all datasets used were publically available via the Copernicus Marine Service. However, we believe our method would be useful in higher resolution scenarios ($\sim 1 - 50$ km) which would be touching upon the territory of privacy in certain regions of the globe. The dataset itself is unlikely to enable activities that are harmful for the environment or individuals. However, if we consider that one could

⁴Note that we did not showcase the training times for each method because our OI is unoptimized and the DUACS methods is closed-source. However, the SIREN method is a coordinate-based method which can take a long time to train. Our model took an average of 6 hours to train not including the hyper-parameter sweeps. We recommend multi-scale training strategies and multi-node GPU architectures to get drastic speedups [25].

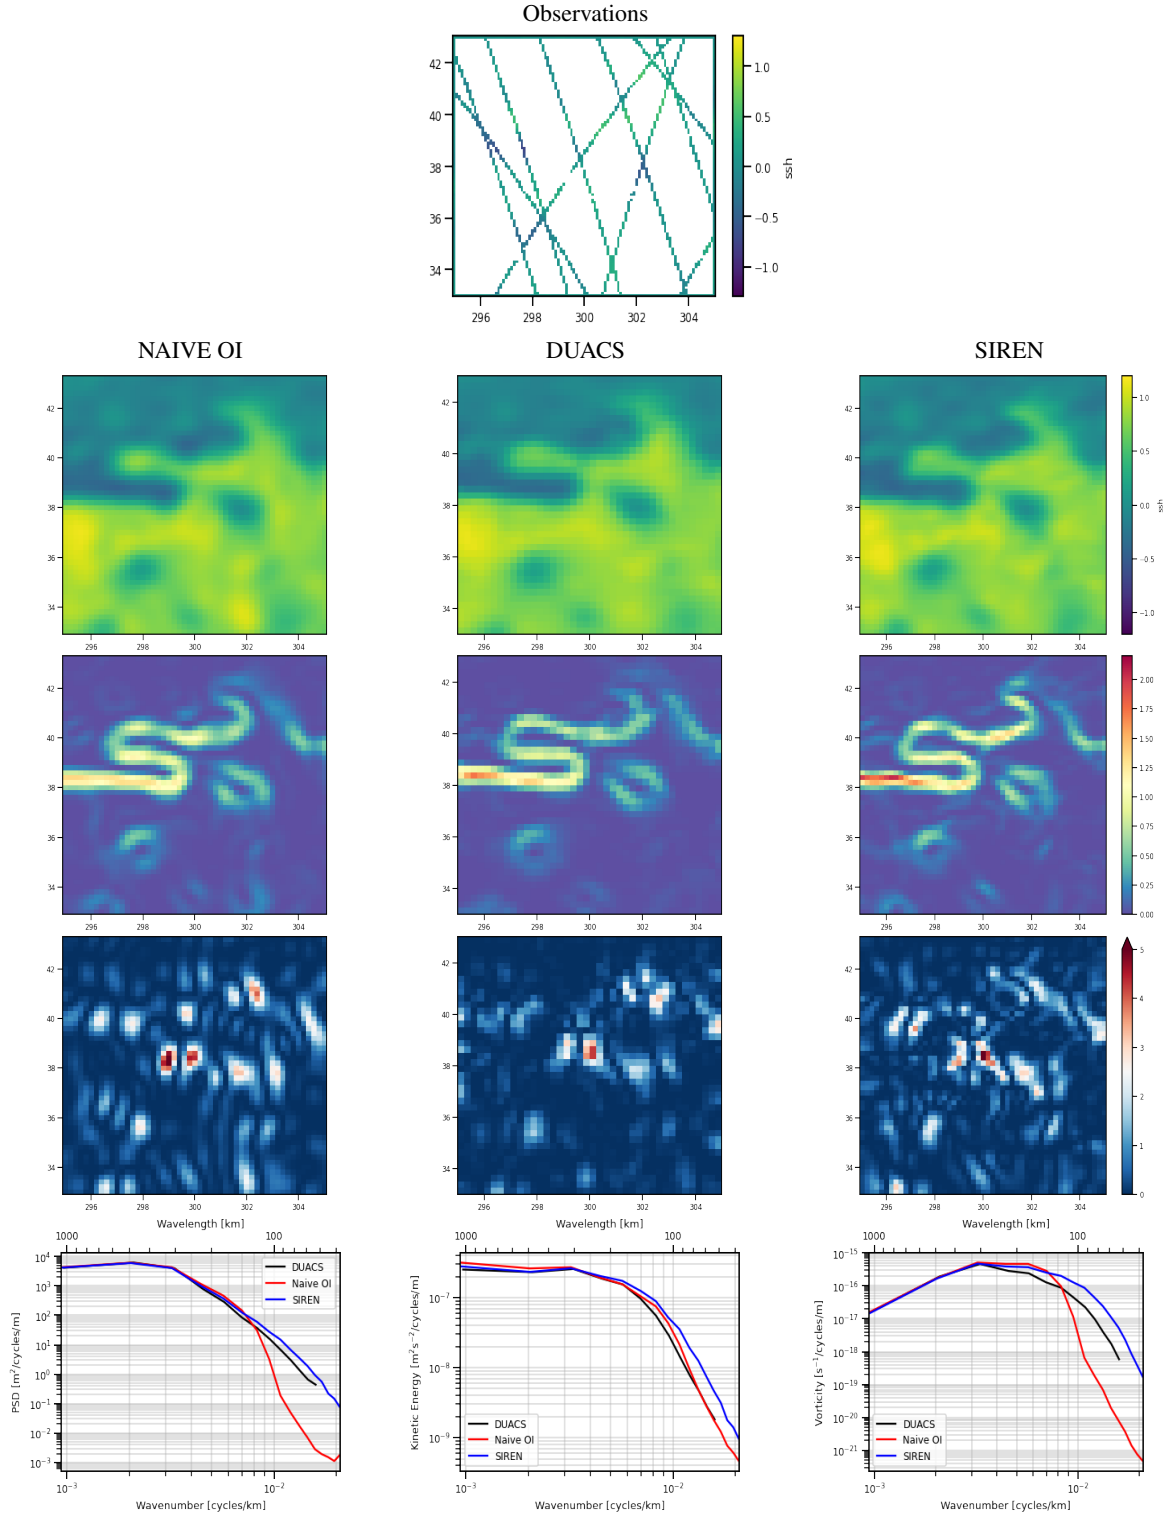


Figure 1: Snapshots of the daily aggregated altimetry observations (top row), the SSH field (2nd row), the norm of the gradient of SSH field (3rd row) and the Laplacian of the SSH field (4th row) over the Gulf-stream at time step 20-01-2017. Each column showcases each of the methods presented in this work: 1) Naive implementation of OI, 2) the OI in production (DUACS), and 3) the SIREN NerF. The last row showcases the spectrum for each of the variables: bottom row (left) is SSH, bottom row (middle) is the Kinetic energy and bottom row (left) is the Vorticity.

further develop these models for interpolating finer resolutions, one could imagine some downstream potential adverse effects on the environmental applications like overfishing or security concerns. In addition, these models are not limited to geophysical variables so one could easily apply these to land variables or human-made objects. So one should always take care when working with high spatial and temporal resolution maps.

Broader Impact

The theme of interpolation is present in many applied communities with different names, e.g. kriging in ecology/hydrology, Optimal interpolation in oceanography, and Gaussian processes in statistics. We hope that this work bridges this gap between the communities and we invite other works to try to highlight concrete ways that machine learning and classic physics have commonalities.

In the oceanography community in particular, we especially hope to see more adoption of machine learning methods for interpolation. DUACS is ultimately a closed-system so the wider scientific community does not have access to the algorithm. Our OI baseline hopefully unveils some of the finer details of the method. However, in general, the standard OI methods used in the applied community cannot keep up with the massive influxes of observations we receive. So this work is a first step in demonstrating that neural networks (in particular NerFs) are a viable, simpler, and scalable alternative.

Acknowledgments and Disclosure of Funding

This work was supported by ANR (French National Research Agency, project number ANR-17-CE01-0009-01) and CNES through the SWOT Science Team program. This work was also supported by LEFE program (LEFE MANU project IA-OAC), CNES (OSTST DUACS-HR and SWOT ST DIEGO) and ANR Projects Melody (ANR-19-CE46-0011) and OceaniX (ANR-19-CHIA-0016). It benefited from HPC and GPU resources from GENCI-IDRIS (Grant 2021-101030). We would like to thank Weights & Biases for their support and for providing us with an academic license to use their experiment tracking tools for our work.

References

- [1] Josué Martínez-Moreno, Andrew McC. Hogg, Matthew H. England, Navid C. Constantinou, Andrew E. Kiss, and Adele K. Morrison. Global changes in oceanic mesoscale currents over the satellite altimetry record. *Nature Climate Change*, 11(5):397–403, May 2021.
- [2] Peter B. Rhines. Mesoscale eddies. In J. Kirk Cochran, Henry J. Bokuniewicz, and Patricia L. Yager, editors, *Encyclopedia of Ocean Sciences (Third Edition)*, pages 115–127. Academic Press, Oxford, third edition edition, 2019.
- [3] Josué Martínez Moreno, Andrew Hogg, Matthew England, Navid Constantinou, Andrew Kiss, and Adele Morrison. Global changes in oceanic mesoscale currents over the satellite altimetry record. 10 2020.
- [4] Julius J. M. Busecke and Ryan P. Abernathey. Ocean mesoscale mixing linked to climate variability. *Science Advances*, 5(1):eaav5014, 2019.
- [5] Saleh Abdalla, Abdolnabi Abdeh Kolahchi, Michaël Ablain, Susheel Adusumilli, Suchandra Aich Bhowmick, Eva Alou-Font, Laiba Amarouche, Ole Baltazar Andersen, Helena Antich, Lotfi Aouf, Brian Arbic, Thomas Armitage, Sabine Arnault, Camila Artana, Giuseppe Aulicino, Nadia Ayoub, Sergei Badulin, Steven Baker, Chris Banks, Lifeng Bao, Silvia Barbeta, Bàrbara Barceló-Llull, François Barlier, Sujit Basu, Peter Bauer-Gottwein, Matthias Becker, Brian Beckley, Nicole Bellefond, Tatyana Belonenko, Mounir Benkiran, Touati Benkouider, Ralf Bennartz, Jérôme Benveniste, Nicolas Bercher, Muriel Berge-Nguyen, Joao Bettencourt, Fabien Blarel, Alejandro Blazquez, Denis Blumstein, Pascal Bonnefond, Franck Borde, Jérôme Bouffard, François Boy, Jean-Paul Boy, Cédric Brachet, Pierre Brasseur, Alexander Braun, Luca Brocca, David Brockley, Laurent Brodeau, Shannon Brown, Sean Bruinsma, Anna Bulczak, Sammie Buzzard, Madeleine Cahill, Stéphane Calmant, Michel Calzas, Stefania Camici, Mathilde Cancet, Hugues Capdeville, Claudia Cristina Carabajal, Loren Carrere, Anny

Cazenave, Eric P. Chassignet, Prakash Chauhan, Selma Cherchali, Teresa Chereskin, Cecile Cheymol, Daniele Ciani, Paolo Cipollini, Francesca Cirillo, Emmanuel Cosme, Steve Coss, Yuri Cotroneo, David Cotton, Alexandre Couhert, Sophie Coutin-Faye, Jean-François Crétaux, Frederic Cyr, Francesco d'Ovidio, José Darrozes, Cedric David, Nadim Dayoub, Danielle De Staerke, Xiaoli Deng, Shailen Desai, Jean-Damien Desjonqueres, Denise Dettmering, Alessandro Di Bella, Lara Díaz-Barroso, Gerald Dibarboure, Habib Boubacar Dieng, Salvatore Dinardo, Henryk Dobslaw, Guillaume Dodet, Andrea Doglioli, Alessio Domeneghetti, David Donahue, Shenfu Dong, Craig Donlon, Joël Dorandeu, Christine Drezen, Mark Drinkwater, Yves Du Penhoat, Brian Dushaw, Alejandro Egido, Svetlana Erofeeva, Philippe Escudier, Saskia Esselborn, Pierre Exertier, Ronan Fablet, Cédric Falco, Sinead Louise Farrell, Yannice Faugere, Pierre Femenias, Luciana Fenoglio, Joana Fernandes, Juan Gabriel Fernández, Pascale Ferrage, Ramiro Ferrari, Lionel Fichen, Paolo Filippucci, Stylianos Flampouris, Sara Fleury, Marco Fornari, Rene Forsberg, Frédéric Frappart, Marie laure Frery, Pablo Garcia, Albert Garcia-Mondejar, Julia Gaudelli, Lucile Gaultier, Augusto Getirana, Ferran Gibert, Artur Gil, Lin Gilbert, Sarah Gille, Luisella Giulicchi, Jesús Gómez-Enri, Laura Gómez-Navarro, Christine Gommenginger, Lionel Gourdeau, David Griffin, Andreas Groh, Alexandre Guerin, Raul Guerrero, Thierry Guinle, Praveen Gupta, Benjamin D. Gutknecht, Mathieu Hamon, Guoqi Han, Danièle Hauser, Veit Helm, Stefan Hendricks, Fabrice Hernandez, Anna Hogg, Martin Horwath, Martina Idžanović, Peter Janssen, Eric Jeansou, Yongjun Jia, Yuanyuan Jia, Liguang Jiang, Johnny A. Johannessen, Masafumi Kamachi, Svetlana Karimova, Kathryn Kelly, Sung Yong Kim, Robert King, Cecile M.M. Kittel, Patrice Klein, Anna Klos, Per Knudsen, Rolf Koenig, Andrey Kostianoy, Alexei Kouraev, Raj Kumar, Sylvie Labroue, Loreley Selene Lago, Juliette Lambin, Léa Lasson, Olivier Laurain, Rémi Laxenaire, Clara Lázaro, Sophie Le Gac, Julien Le Sommer, Pierre-Yves Le Traon, Sergey Lebedev, Fabien Léger, Benoît Legresy, Frank Lemoine, Luc Lenain, Eric Leuliette, Marina Levy, John Lillibridge, Jianqiang Liu, William Llovel, Florent Lyard, Claire Macintosh, Eduard Makhoul Varona, Cécile Manfredi, Frédéric Marin, Evan Mason, Christian Massari, Constantin Mavrocordatos, Nikolai Maximenko, Malcolm McMillan, Thierry Medina, Angelique Melet, Marco Meloni, Stelios Mertikas, Sammy Metref, Benoit Meyssignac, Jean-François Minster, Thomas Moreau, Daniel Moreira, Yves Morel, Rosemary Morrow, John Moyard, Sandrine Mulet, Marc Naeije, Robert Steven Nerem, Hans Ngodock, Karina Nielsen, Jan Even Øie Nilsen, Fernando Niño, Carolina Nogueira Lodo, Camille Noûs, Estelle Obligis, Inès Ootosaka, Michiel Otten, Berguzar Oztunali Ozbahceci, Roshin P. Raj, Rodrigo Paiva, Guillermina Paniagua, Fernando Paolo, Adrien Paris, Ananda Pascual, Marcello Passaro, Stephan Paul, Tamlin Pavelsky, Christopher Pearson, Thierry Penduff, Fukai Peng, Felix Perosanz, Nicolas Picot, Fanny Piras, Valerio Poggiali, Étienne Poirier, Sonia Ponce de León, Sergey Prants, Catherine Prigent, Christine Provost, M-Isabelle Pujol, Bo Qiu, Yves Quilfen, Ali Rami, R. Keith Raney, Matthias Raynal, Elisabeth Remy, Frédérique Rémy, Marco Restano, Annie Richardson, Donald Richardson, Robert Ricker, Martina Ricko, Eero Rinne, Stine Kildegaard Rose, Vinca Rosmorduc, Sergei Rudenko, Simón Ruiz, Barbara J. Ryan, Corinne Salaün, Antonio Sanchez-Roman, Louise Sandberg Sørensen, David Sandwell, Martin Saraceno, Michele Scagliola, Philippe Schaeffer, Martin G. Scharffenberg, Remko Scharroo, Andreas Schiller, Raphael Schneider, Christian Schwatke, Andrea Scozzari, Enrico Ser-giacomi, Frederique Seyler, Rashmi Shah, Rashmi Sharma, Andrew Shaw, Andrew Shepherd, Jay Shriver, C.K. Shum, Wim Simons, Sebatian B. Simonsen, Thomas Slater, Walter Smith, Saulo Soares, Mikhail Sokolovskiy, Laurent Soudarin, Ciprian Spatar, Sabrina Speich, Margaret Srinivasan, Meric Srokosz, Emil Stanev, Joanna Staneva, Nathalie Steunou, Julienne Stroeve, Bob Su, Yohanes Budi Sulistioadi, Debadatta Swain, Annick Sylvestre-baron, Nicolas Taburet, Rémi Tailleux, Katsumi Takayama, Byron Tapley, Angelica Tarpanelli, Gilles Tavernier, Laurent Testut, Praveen K. Thakur, Pierre Thibaut, LuAnne Thompson, Joaquín Tintoré, Céline Tison, Cédric Tourain, Jean Tournadre, Bill Townsend, Ngan Tran, Sébastien Trilles, Michel Tsamados, Kuo-Hsin Tseng, Clément Ubelmann, Bernd Uebbing, Oscar Vergara, Jacques Verron, Telmo Vieira, Stefano Vignudelli, Nadya Vinogradova Shiffer, Pieter Visser, Frederic Vivier, Denis Volkov, Karina von Schuckmann, Valerii Vuglinskii, Pierrick Vuilleumier, Blake Walter, Jida Wang, Chao Wang, Christopher Watson, John Wilkin, Josh Willis, Hilary Wilson, Philip Woodworth, Kehan Yang, Fangfang Yao, Raymond Zaharia, Elena Zakharova, Edward D. Zaron, Yongsheng Zhang, Zhongxiang Zhao, Vadim Zinchenko, and Victor Zlotnicki. *Altimetry for the future: Building on 25 years of progress. Advances in Space Research*, 68(2):319–363, 2021. 25 Years of Progress in Radar Altimetry.

- [6] Lee-Lueng Fu, Tong Lee, W Timothy Liu Liu, and Ronald Kwok. Fifty years of satellite remote sensing of the ocean. *Meteorological Monographs*, 59, 02 2019.
- [7] M.-I. Pujol, Y. Faugère, G. Taburet, S. Dupuy, C. Pelloquin, M. Ablain, and N. Picot. Duacs dt2014: the new multi-mission altimeter data set reprocessed over 20 years. *Ocean Science*, 12(5):1067–1090, 2016.
- [8] P. Y. Le Traon. From satellite altimetry to argo and operational oceanography: three revolutions in oceanography. *Ocean Science*, 9(5):901–915, 2013.
- [9] Redouane Lguensat, Phi Huynh Viet, Miao Sun, Ge Chen, Tian Fenglin, Bertrand Chapron, and Ronan Fablet. Data-driven interpolation of sea level anomalies using analog data assimilation. *Remote Sensing*, 11(7), 2019.
- [10] Maxime Beauchamp, Mohamed Mahmoud Amar, Quentin Febvre, and Ronan Fablet. End-to-end learning of variational interpolation schemes for satellite-derived ssh data. In *2021 IEEE International Geoscience and Remote Sensing Symposium IGARSS*, pages 7418–7421, 2021.
- [11] G. Taburet, A. Sanchez-Roman, M. Ballarotta, M.-I. Pujol, J.-F. Legeais, F. Fournier, Y. Faugere, and G. Dibarboure. Duacs dt2018: 25 years of reprocessed sea level altimetry products. *Ocean Science*, 15(5):1207–1224, 2019.
- [12] M. Ballarotta, C. Ubelmann, M.-I. Pujol, G. Taburet, F. Fournier, J.-F. Legeais, Y. Faugère, A. Delepouille, D. Chelton, G. Dibarboure, and N. Picot. On the resolutions of ocean altimetry maps. *Ocean Science*, 15(4):1091–1109, 2019.
- [13] A. Pisano, D. Ciani, S. Marullo, R. Santoleri, and B. Buongiorno Nardelli. A new operational mediterranean diurnal optimally interpolated sea surface temperature product within the copernicus marine service. *Earth System Science Data*, 14(9):4111–4128, 2022.
- [14] Chunxue Yang, Francesca Elisa Leonelli, Salvatore Marullo, Vincenzo Artale, Helen Beggs, Bruno Buongiorno Nardelli, Toshio M. Chin, Vincenzo De Toma, Simon Good, Boyin Huang, Christopher J. Merchant, Toshiyuki Sakurai, Rosalia Santoleri, Jorge Vazquez-Cuervo, Huai-Min Zhang, and Andrea Pisano. Sea surface temperature intercomparison in the framework of the copernicus climate change service (c3s). *Journal of Climate*, 34(13):5257 – 5283, 2021.
- [15] V. Banzon, T. M. Smith, T. M. Chin, C. Liu, and W. Hankins. A long-term record of blended satellite and in situ sea-surface temperature for climate monitoring, modeling and environmental studies. *Earth System Science Data*, 8(1):165–176, 2016.
- [16] Rosemary Morrow, Lee-Lueng Fu, Fabrice Ardhuin, Mounir Benkiran, Bertrand Chapron, Emmanuel Cosme, Francesco d’Ovidio, J. Thomas Farrar, Sarah T. Gille, Guillaume Lapeyre, Pierre-Yves Le Traon, Ananda Pascual, Aurélien Ponte, Bo Qiu, Nicolas Rascle, Clement Ubelmann, Jinbo Wang, and Edward D. Zaron. Global observations of fine-scale ocean surface topography with the surface water and ocean topography (swot) mission. *Frontiers in Marine Science*, 6, 2019.
- [17] Yiheng Xie, Towaki Takikawa, Shunsuke Saito, Or Litany, Shiqin Yan, Numair Khan, Federico Tombari, James Tompkin, Vincent Sitzmann, and Srinath Sridhar. Neural fields in visual computing and beyond, 2021.
- [18] Maxime Ballarotta and Florian Le Guillou. ocean-data-challenges/2021a_SSH_mapping_OSE: Material for SSH mapping OSE data challenge, September 2021.
- [19] Rossi Hamzah and Teguh Prayogo. Interpolation methods for sea surface height mapping from altimetry satellites in indonesian seas. *International Journal of Remote Sensing and Earth Sciences (IJReSES)*, 11:33, 04 2017.
- [20] Aida Alvera Azcarate, Alexander Barth, Michel Rixen, and Jean-Marie Beckers. Reconstruction of incomplete oceanographic data sets using empirical orthogonal functions: application to the adriatic sea surface temperature. *Ocean Modelling*, 9(4), 2005.

- [21] Aida Alvera Azcarate, Alexander Barth, Jean-Marie Beckers, and Robert H Weisberg. Multi-variate reconstruction of missing data in sea surface temperature, chlorophyll, and wind satellite fields. *Journal of Geophysical Research. Oceans*, 112(C3), 2007.
- [22] A. Barth, A. Alvera-Azcárate, M. Licer, and J.-M. Beckers. Dincae 1.0: a convolutional neural network with error estimates to reconstruct sea surface temperature satellite observations. *Geoscientific Model Development*, 13(3):1609–1622, 2020.
- [23] Geoff Pleiss, Jacob R. Gardner, Kilian Q. Weinberger, and Andrew Gordon Wilson. Constant-time predictive distributions for gaussian processes, 2018.
- [24] Cari G. Kaufman, Mark J. Schervish, and Douglas W. Nychka. Covariance tapering for likelihood-based estimation in large spatial data sets. *Journal of the American Statistical Association*, 103(484):1545–1555, 2008.
- [25] Shaowu Pan, Steven L. Brunton, and J. Nathan Kutz. Neural implicit flow: a mesh-agnostic dimensionality reduction paradigm of spatio-temporal data, 2022.
- [26] Matthew Tancik, Pratul P. Srinivasan, Ben Mildenhall, Sara Fridovich-Keil, Nithin Raghavan, Utkarsh Singhal, Ravi Ramamoorthi, Jonathan T. Barron, and Ren Ng. Fourier features let networks learn high frequency functions in low dimensional domains. *NeurIPS*, 2020.
- [27] Vincent Sitzmann, Julien N.P. Martel, Alexander W. Bergman, David B. Lindell, and Gordon Wetzstein. Implicit neural representations with periodic activation functions. In *Proc. NeurIPS*, 2020.
- [28] Michael S. Gashler and Stephen C. Ashmore. Training deep fourier neural networks to fit time-series data. In *ICIC*, 2014.
- [29] Ishit Mehta, Michaël Gharbi, Connelly Barnes, Eli Shechtman, Ravi Ramamoorthi, and Manmohan Chandraker. Modulated periodic activations for generalizable local functional representations. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 14214–14223, October 2021.
- [30] Rizal Fathony, Anit Kumar Sahu, Devin Willmott, and J. Zico Kolter. Multiplicative filter networks. In *ICLR*, 2021.
- [31] Lukas Biewald. Experiment tracking with weights and biases, 2020. Software available from wandb.com.
- [32] Florian Le Guillou, Sammy Metref, Emmanuel Cosme, Clément Ubelmann, Maxime Ballarotta, Julien Le Sommer, and Jacques Verron. Mapping altimetry in the forthcoming swot era by back-and-forth nudging a one-layer quasigeostrophic model. *Journal of Atmospheric and Oceanic Technology*, 38(4):697 – 710, 2021.
- [33] Simon Letzgus, Patrick Wagner, Jonas Lederer, Wojciech Samek, Klaus-Robert Müller, and Grégoire Montavon. Toward explainable artificial intelligence for regression models: A methodological perspective. *IEEE Signal Processing Magazine*, 39(4):40–58, 2022.
- [34] Mengyu Chu, Lingjie Liu, Quan Zheng, Erik Franz, Hans-Peter Seidel, Christian Theobalt, and Rhaleb Zayer. Physics informed neural fields for smoke reconstruction with sparse data. *ACM Trans. Graph.*, 41(4), jul 2022.
- [35] J. Shen, A. Ruiz, A. Agudo, and F. Moreno-Noguer. Stochastic neural radiance fields: Quantifying uncertainty in implicit 3d representations. In *2021 International Conference on 3D Vision (3DV)*, pages 972–981, Los Alamitos, CA, USA, dec 2021. IEEE Computer Society.
- [36] Michael Niemeyer, Jonathan T. Barron, Ben Mildenhall, Mehdi S. M. Sajjadi, Andreas Geiger, and Noha Radwan. Regnerf: Regularizing neural radiance fields for view synthesis from sparse inputs. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2022.
- [37] Jianxiong Shen, Antonio Agudo, Francesc Moreno-Noguer, and Adria Ruiz. Conditional-flow nerf: Accurate 3d modelling with reliable uncertainty quantification. In *Arxiv*, 2022.

- [38] Haithem Turki, Deva Ramanan, and Mahadev Satyanarayanan. Mega-nerf: Scalable construction of large-scale nerfs for virtual fly-throughs. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 12922–12931, June 2022.

Checklist

1. For all authors...
 - (a) Do the main claims made in the abstract and introduction accurately reflect the paper’s contributions and scope? [Yes] See sections 4 and 5
 - (b) Did you describe the limitations of your work? [Yes] See section 5
 - (c) Did you discuss any potential negative societal impacts of your work? [Yes] See section 5
 - (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes] See section 5
2. If you are including theoretical results...
 - (a) Did you state the full set of assumptions of all theoretical results? [N/A]
 - (b) Did you include complete proofs of all theoretical results? [N/A]
3. If you ran experiments...
 - (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes]
 - (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes] See section 3.2
 - (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [Yes] See table 1 for the NRMSE statistics.
 - (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes] See section 3.2
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
 - (a) If your work uses existing assets, did you cite the creators? [Yes] See section 3.1
 - (b) Did you mention the license of the assets? [Yes] See section 3.1
 - (c) Did you include any new assets either in the supplemental material or as a URL? [N/A]
 - (d) Did you discuss whether and how consent was obtained from people whose data you’re using/curating? [N/A]
 - (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]
5. If you used crowdsourcing or conducted research with human subjects...
 - (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
 - (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
 - (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]