
OCEANBENCH: The Sea Surface Height Edition

J. Emmanuel Johnson*

CNRS UMR IGE

johnsonj@univ-grenoble-alpes.fr

Quentin Febvre*

IMT Atlantique

quentin.febvre@imt-atlantique.fr

Anastasia Gorbunova

CNRS UMR IGE

Sammy Metref

DATLAS

Maxime Ballarotta

CLS

Julien Le Sommer

CNRS UMR IGE

Ronan Fablet

IMT Atlantique

Abstract

The ocean is a crucial component of the Earth’s system. It profoundly influences human activities and plays a critical role in climate regulation. Our understanding has significantly improved over the last decades with the advent of satellite remote sensing data, allowing us to capture essential sea surface quantities over the globe, e.g., sea surface height (SSH). Despite their ever-increasing abundance, ocean satellite data presents challenges for information extraction due to their sparsity and irregular sampling, signal complexity, and noise. Machine learning (ML) techniques have demonstrated their capabilities in dealing with large-scale, complex signals. Therefore we see an opportunity for these ML models to harness the full extent of the information contained in ocean satellite data. However, data representation and relevant evaluation metrics can be *the* defining factors when determining the success of applied ML. The processing steps from the raw observation data to a ML-ready state and from model outputs to interpretable quantities require domain expertise, which can be a significant barrier to entry for ML researchers. In addition, imposing fixed processing steps, like committing to specific variables, regions, and geometries, will narrow the scope of ML models and their potential impact on real-world applications. **OceanBench** is a unifying framework that provides standardized processing steps that comply with domain-expert standards. It is designed with a flexible and pedagogical abstraction: it a) provides plug-and-play data and pre-configured pipelines for ML researchers to benchmark their models w.r.t. ML and domain-related baselines and b) provides a transparent and configurable framework for researchers to customize and extend the pipeline for their tasks. In this work, we demonstrate the OceanBench framework through a first edition dedicated to SSH interpolation challenges. We provide datasets and ML-ready benchmarking pipelines for the long-standing problem of interpolating observations from simulated ocean satellite data, multi-modal and multi-sensor fusion issues, and transfer-learning to real ocean satellite observations. The OceanBench framework is available at github.com/jejohnson/oceanbench and the dataset registry is available at github.com/quentinf00/oceanbench-data-registry.

*These authors contributed equally to this work

1 Motivation

The ocean is vital to the Earth’s system [28]. It plays a significant role in climate regulation regarding carbon [40] and heat uptake [87]. It is also a primary driver of human activities (e.g., maritime traffic and world trade, marine resources and services) [106, 93]. However, monitoring the ocean is a critical challenge: the ocean state can only partially be determined because most of the ocean consists of subsurface quantities that we cannot directly observe. Thus, to quantify even a fraction of the physical or biochemical ocean state, we must often rely only on surface quantities that we can monitor from space, drifting buoys, or autonomous devices. Satellite remote sensing, in particular, is one of the most effective ways of measuring essential sea surface quantities [2] such as sea surface height (SSH) [95], sea surface temperature (SST) [77], and ocean color (OC) [53]. While these variables characterize only a tiny portion of the ocean ecosystem, they present a gateway to many other derived physical quantities [93].

Although we can access observable sea surface quantities, they are generally irregularly and extremely sparsely sampled. For instance, satellite-derived SSH data has less than 5% coverage of the globe daily [95]. These sampling gaps make the characterization of ocean processes highly challenging for operational products and downstream tasks that depend on relevant gap-free variables. This has motivated a rich literature in geoscience over the last decades, mainly using geostatistical kriging methods [95, 102] and model-driven data assimilation schemes [55, 60]. Despite significant progress, these schemes often need to improve their ability to leverage available observation datasets’ potential fully. This has naturally advocated for exploring data-driven approaches like shallow ML schemes [7, 6, 97, 71]. Very recently, deep learning schemes [116, 74, 9] have become appealing solutions to benefit from existing large-scale observation and simulation datasets and reach significant breakthroughs in the monitoring of upper ocean dynamics from scarcely and irregularly sampled observations. However, the heterogeneity and characteristics of the observation data present major challenges for effectively applying these methods beyond idealized case studies. A data source could have different variables, geometries, and noise levels, resulting in many domain-specific preprocessing procedures that can vastly change the solution outcome. Furthermore, the evaluation procedure of the methods and their effectiveness can be regionally-dependent as the physical phenomena vary in space and time, which adds another layer of complexity in convincing domain scientists of their trustworthiness. So the entire ML pipeline now requires a unified framework for dealing with heterogeneous data sources, different pre- and post-processing methodologies, and regionally-dependent evaluation procedures.

To address these challenges, we introduce **OceanBench**, a framework for co-designing machine-learning-driven high-level experiments from ocean observations. It consists of an end-to-end framework for piping data from its raw form to an ML-ready state and from model outputs to interpretable quantities. We regard OceanBench as a key facilitator for the uptake of MLOPs tools and research [66, 94] for ocean-related datasets and case studies. This first edition provides datasets and ML-ready benchmarking pipelines for SSH interpolation problems, an essential topic for the space oceanography community, related to ML communities dealing with issues like in-painting [111], denoising [99, 98], and super-resolution [107]. We expect OceanBench to facilitate new challenges to the applied machine learning community and contribute to meaningful ocean-relevant breakthroughs. The remainder of the paper is organized as follows: in §2, we outline some related work that was inspirational for this work; in §3, we formally outline OceanBench by highlighting the target audience, code structure, and problem scope; in §4, we outline the problem formulation of SSH interpolation and provide some insight into different tasks related to SSH interpolation where OceanBench could provide some helpful utility; and in §5 we give some concluding remarks while also informally inviting other researchers to help fill in the gaps.

2 Related Work

Machine learning applied to geosciences is becoming increasingly popular, but there are few examples of transparent pipelines involving observation data. After a thorough literature review, we have divided the field into three camps of ML applications that pertain to this work: 1) toy simulation datasets, 2) reanalysis datasets, and 3) observation datasets. We outline the literature for each of the three categories below.

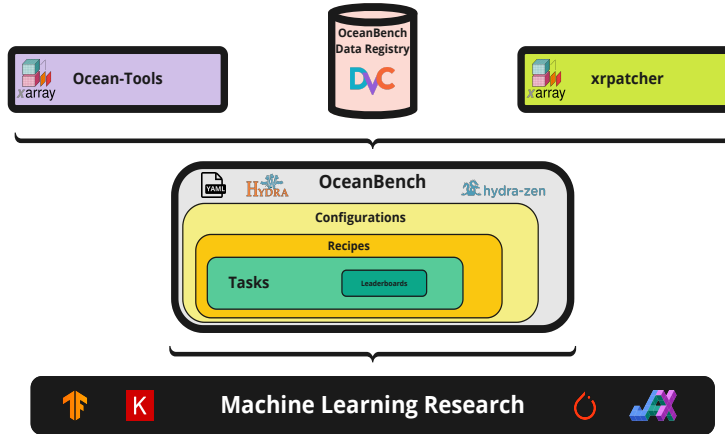


Figure 1: This figure showcases the OceanBench toolset. We have 1) OceanBench-Data-Registry which uses DVC to store and organize ML-ready datasets, 2) Ocean-Tools which features a suite of task-relevant geoprocessing functions with the xarray-backend, and 3) xrpatcher which can produce arbitrary subsets of xarray datastructures which nicely interface with dataloaders such as PyTorch. OceanBench provides an interface for ML researchers to parameterize arbitrary sequences of transformations to preprocess data from a domain-specific state to a ML-ready state.

Toy Simulation Data. One set of benchmarks focuses on learning surrogate models for well-defined but chaotic dynamical systems in the form of ordinary differential equations (ODEs) and partial differential equations (PDEs) and there are freely available code bases which implement different ODEs/PDEs [52, 96, 3, 64, 8, 103, 56, 85]. This is a great testing ground for simple toy problems that better mimic the structures we see in real-world observations. Working with simulated data is excellent because it is logistically simple and allows users to test their ideas on toy problems without increasing the complexity when dealing with real-world data. However, these are ultimately simple physical models that often do not reflect the authentic structures we see in real-world, observed data.

Reanalysis Data. This is assimilated data of real observations and model simulations. There are a few major platforms that host ocean reanalysis data like the Copernicus Marine Data Store [36, 33, 34, 37], the Climate Data Store [25], the BRAN2020 Model [26], and the NOAA platform [15]. However, to our knowledge, there is no standard ML-specific ocean-related tasks to accompany the data. On the atmospheric side, platforms like WeatherBench [86], ClimateBench [108], ENS10 [10] were designed to assess short-term and medium-term forecasting using ML techniques with recent success of ML [69, 84]. The clarity of the challenges set by the benchmark suites has inspired the idea of OceanBench, where we directly focus on problems dealing with ocean observation data.

Observation Data. These observation datasets (typically sparse) stem from satellite observations that measure surface variables or in-situ measurements that measure quantities within the water column. Some major platforms to host data include the Marine Data Store [32, 31], the Climate Data Store [23, 24, 22], ARGO [110], and the SOCAT platform [11]. However, it is more difficult to assess the efficacy of operational ML methods that have been trained only on observation data and, to our knowledge, there is no coherent ML benchmarking system for ocean state estimation. There has been significant effort by the *Ocean-Data-Challenge Group*¹ which provides an extensive suite of datasets and metrics for SSH interpolation. Their efforts heavily inspired our work, and we hope that OceanBench can build upon their work by adding cohesion and facilitating the ease of use for ML research and providing a high-level framework for providing ML-related data products.

3 OceanBench

3.1 Why OceanBench?

There is a high barrier to entry in working with ocean observations for researchers in applied machine learning as there are many processing steps for both the observation data and the domain-specific

¹Ocean Data Challenge group: Freely associated scientist for oceanographic algorithm and product improvements (ocean-data-challenges.github.io)

evaluation procedures. OceanBench aims to lower the barrier to entry cost for ML researchers to make meaningful progress in the field of state prediction. We distribute a standardized, transparent, and flexible procedure for defining data and evaluation pipelines for data-intensive geoscience applications. Proposed examples and case studies provide a plug-and-play framework to benchmark novel ML schemes w.r.t. state-of-the-art, domain-specific ML baselines (see figure 1). In addition, we adopt a pedagogical abstraction that allows users to customize and extend the pipelines for their specific tasks. To our knowledge, no framework embeds processing steps for earth observation data in a manner compatible with MLOps abstractions and standards regarding reproducibility and evaluation procedures. Ultimately, we aim to facilitate the uptake of ML schemes to address ocean observation challenges and to bring new challenges to the ML community to extend additional ML tools and methods for irregularly-sampled and partially-observed high-dimensional space-time dynamics. The abstractions proposed here apply beyond ocean sciences and SSH interpolation to other geosciences with similar tasks that intersect with machine learning.

3.2 Code Structure

OceanBench is lightweight in terms of the core functionality. We keep the code base simple and focus more on how the user can combine each piece. We adopt a strict functional style because it is easier to maintain and combine sequential transformations. There are five features we would like to highlight about OceanBench: 1) Data availability and version control, 2) an agnostic suite of geoprocessing tools for `xarray` datasets that were aggregated from different sources, 3) Hydra integration to pipe sequential transformations, 4) a flexible multi-dimensional array generator from `xarray` datasets that are compatible with common deep learning (DL) frameworks, and 5) a JupyterBook [38] that offers library tutorials and demonstrates use-cases. In the following section, we highlight these components in more detail.

OceanBench-Data-Registry. The most important aspect is the public availability of the datasets. We aggregate all pre-curated datasets from other sources, e.g. the *Ocean-Data-Challenge* [13, 12], and organize them to be publicly available from a single source [2]. We also offer a few derived datasets which can be used for demonstrations and evaluation. Data is never static in a pipeline setting, as one can have many derived datasets which stem from numerous preprocessing choices. In fact, in research, we often work with derived datasets that have already been through some preliminary preprocessing methods. To facilitate the ever-changing nature of data, we use the Data Version Control (DVC) tool [67], which offers a git-like version control of the datasets.

Ocean-Tools [3]. The Ocean-Tools library uses a core suite of functions specific to processing geo-centric data. While a few particular functionalities vary from domain to domain, many operations are standard, e.g., data variable selections, filtering/smoothing, regridding, coordinate transformations, and standardization. We almost work exclusively with the `xarray` [58] framework because it is a coordinate-aware, flexible data structure. In addition, the geoscience community has an extensive suite of specialized packages that operate in the `xarray` framework to accomplish many different tasks. Almost all Ocean-Tools toolsets are exclusively within the `xarray` framework to maintain compatibility with a large suite of tools already available from the community.

Hydra Integration. As discussed above, many specific packages accomplish many different tasks. However, what needs to be added is the flexibility to mix and match these operations as the users see fit. Hydra [112] and Hydra-Zen [92] provide a configurable way to aggregate and *pipe* many sequential operations together. It also maintains readability, robustness, and flexibility through the use of `.yaml` files which explicitly highlights the function used, the function parameters chosen, and the sequence of operations performed. In the ML software stack, Hydra is often used to manage the model, optimizer, and loss configurations which helps the user experiment with different options. We apply this same concept in preprocessing, geoprocessing, and evaluation steps, often more important than the model configuration in geoscience-related tasks.

XR Patcher [4]. Every machine learning pipeline will inevitably require moving data from the geo-specific data structure to a multi-dimensional array easily digestible for ML models. A rather underrated, yet critical, feature of ML frameworks such as PyTorch [83] (Lightning [45]) and

² Available at: [quentinf00/oceanbench-data-registry](https://github.com/quentinf00/oceanbench-data-registry)

³ Available at: [jejjohnson/ocn-tools](https://github.com/jejjohnson/ocn-tools)

⁴ Available at: [jejjohnson/xrpatcher](https://github.com/jejjohnson/xrpatcher)

TensorFlow [1] (Keras [30]) is the abstraction of the dataset, dataloader, datamodules, and data pipelines. In applied ML in geosciences, the data pipelines are often more important than the actual model [89]. The user can control the *patch*-size and the *stride*-step, which can generate arbitrary coordinate-aware items directly from the xarray data structure. In addition, XRPatcher provides a way to reconstruct the fields from an arbitrary patch configuration. This robust reconstruction step is convenient to extend the ML inference step where one can reconstruct entire fields of arbitrary dimensions beyond the training configuration, e.g., to account for the border effects within the field (see appendix E) or to reconstruct quantities in specific regions or globally.

JupyterBook. Building a set of tools is relatively straightforward; however, ensuring that it sees a broader adoption across a multi-disciplinary community is much more challenging. We invested heavily in showing use cases that appeal to different users with the JupyterBook platform [38]. Code with context is imperative for domain and ML experts as we need to explain and justify each component and give many examples of how they can be used in other situations. Thus, we have paid special attention to providing an extensive suite of tutorials, and we also highlight use cases for how one can effectively use the tools.

3.3 Problem Scope

There are many problems that are of great interest the ocean community [29] but we limit the scope to state estimation problems [21]. Under this scope, there are research questions that are relevant to operational centers which are responsible for generating the vast majority of global ocean state maps [36, 34, 33, 37] that are subsequently used for many downstream tasks [93]. For example: how can we effectively use heterogeneous observations to predict the ocean state on the sea surface [55, 62, 102, 44, 14, 77]; how can we incorporate prior physics knowledge into our predictions of ocean state trajectories [55, 29, 93]; and how can we use the current ocean state at time T to predict the future ocean state at time $T + \tau$ [42, 86, 16]. In the same vein, there are more research questions that are of interest to the academic modeling community. For example: is simulated or reanalysis data more effective for learning ML emulators that replace expensive ocean models [49, 114]; what metrics are more effective for assessing our ability to mimic ocean dynamics [75, 48]; and how much model error can we characterize when learning from observations [18, 68].

We have cited many potential applications of how ML can be applied to tackle the state estimation problem. However, to our knowledge there is no publicly available, standardized benchmark system that caters to ML-research standards. We believe that, irrespective of the questions posed above and the data we access, there are many logistical similarities for each of the problem formulations where we can start to set standards for a subset of tasks like interpolation or forecasting. On the front-end, we need a way to select regions, periods, variables, and a valid train-test split (see sec. D.1). On the back-end, we need a way to transform the predictions into more meaningful variables with appropriate metrics for validation (see sec. D.2 and D.3). OceanBench was designed to be an agnostic tool that is extensible to the types of datasets, processing techniques and metrics needed for working with a specific class of Ocean-related datasets. We strongly feel that a suite like this is the first step in designing task-specific benchmarks within the ocean community that is compatible with ML standards. In the remainder of the paper, we will demonstrate how OceanBench can be configured to facilitate a ML-ready data challenge involving our first edition to demonstrate OceanBench’s applicability: sea surface height interpolation.

4 Sea Surface Height Edition

Sea surface height (SSH) is one of the most critical, observable quantities when determining the ocean state. It is widely used to study ocean dynamics and the adverse impact on global climate and human activities [78]. SSH enables us to track phenomena such as currents and eddies [78, 27, 82], which leads to a better quantification of the transport of energy, heat, and salt. In addition, SSH helps us quantify sea level rise at regional and global scales [4, 39], which is used for operational monitoring of the marine environment [106]. Furthermore, SSH characterization provides a plethora of data products that downstream tasks can use for many other applications [79, 20]. Due to the irregular sampling delivered by satellite altimeter, state-of-the-art operational methods using optimal interpolation schemes [95, 102] or model-driven data assimilation [7, 6, 71, 97] fail to fully retrieve SSH dynamics at fine scales below 100-200km on a global or regional scale, so improving the

space-time resolution of SSH fields has been a critical challenge in ocean science. Beyond some technological developments [51], recent studies support the critical role of ML-based schemes in overcoming the current limitations of the operational systems [14, 55, 116]. The rest of this section gives an overview of the general problem definition for SSH interpolation, followed by a brief ontology for ML approaches to address the problem. We also give an overview of some experimental designs and datasets with a demonstration of metrics and plots generated by the OceanBench platform.

4.1 Problem Definition

We are dealing with satellite observations, so we are interested in the domain across the Earth’s surface. Let us define the Earth’s domain by some spatial coordinates, $\mathbf{x} = [\text{Longitude}, \text{Latitude}]^\top \in \mathbb{R}^{D_s}$, and temporal coordinates, $t = [\text{Time}] \in \mathbb{R}^+$, where D_s is the dimensionality of the coordinate vector. We can define some spatial (sub-)domain, $\Omega \subseteq \mathbb{R}^{D_s}$, and a temporal (sub-)domain, $\mathcal{T} \subseteq \mathbb{R}^+$. This domain could be the entire globe for 10 years or a small region within the North Atlantic for 1 year.

$$\begin{aligned} \text{Spatial Coordinates :} & \quad \mathbf{x} \in \Omega \subseteq \mathbb{R}^{D_s} & (1) \\ \text{Temporal Coordinates :} & \quad t \in \mathcal{T} \subseteq \mathbb{R}^+. & (2) \end{aligned}$$

In this case $D_s = 2$ because we only have a two coordinates, however we can do some coordinate transformations like spherical to Cartesian. Likewise, we can do some coordinate transformation for the temporal coordinates like cyclic transformations or sinusoidal embeddings [105]. We have two fields of interest from these spatiotemporal coordinates: the state and the observations.

$$\begin{aligned} \text{State :} & \quad \mathbf{u}(\mathbf{x}, t) : \Omega \times \mathcal{T} \rightarrow \mathbb{R}^{D_u} & (3) \\ \text{Observations :} & \quad \mathbf{y}_{obs}(\mathbf{x}, t) : \Omega \times \mathcal{T} \rightarrow \mathbb{R}^{D_{obs}} & (4) \end{aligned}$$

The state domain, $u \in \mathcal{U}$, is a scalar or vector-valued field of size D_u which is typically the quantity of interest and the observation domain, $y_{obs} \in \mathcal{Y}_{obs}$, is the observable quantity which is also a scalar or vector-valued field of size D_{obs} . Now, we make the assumption that we have an operator \mathcal{H} that transforms the field from the state space, \mathbf{u} , to the observation space, \mathbf{y}_{obs} .

$$\mathbf{y}_{obs}(\mathbf{x}, t) = \mathcal{H}(\mathbf{u}(\mathbf{x}, t), t, \boldsymbol{\varepsilon}, \boldsymbol{\mu}) \quad (5)$$

This equation is the continuous function defined over the entire spatiotemporal domain. The operator, $\mathcal{H}(\cdot)$, is flexible and problem dependent. For example, in a some discretized setting there are 0’s wherever there are no observations, and 1’s wherever there are observations, and in other discretized settings it takes a weighted average of the neighboring pixels. We also include a generic noise function, $\boldsymbol{\varepsilon}(\mathbf{x}, t)$. This could stem from a distribution, it could stationary noise operator, $\boldsymbol{\varepsilon}(\mathbf{x})$, or it could be constant in space but vary with Time, $\boldsymbol{\varepsilon}(t)$. We also include a control parameter, $\boldsymbol{\mu}$, representing any external factors or latent variables that could connect the state vector to the observation vector, e.g., sea surface temperature. Our quantity of interest is SSH, η , a scalar-valued field defined everywhere on the domain. In our application, we assume that the SSH we observe from satellite altimeters, η_{obs} , is the same as the SSH state, except it could be missing for some coordinates due to incomplete coverage from the satellite. So our transformation is defined as follows:

$$\boldsymbol{\eta}_{obs}(\mathbf{x}, t) = \mathcal{H}(\boldsymbol{\eta}(\mathbf{x}, t), t, \boldsymbol{\varepsilon}, \boldsymbol{\mu}) \quad (6)$$

In practice, the satellite providers have a reasonable estimation of the amount of structured noise level we can expect from the satellite altimetry data; however, unresolved noise could still be present. Finally, we are interested in finding some model, \mathcal{M} , that maps the SSH we observe to the true SSH given by

$$\mathcal{M} : \boldsymbol{\eta}_{obs}(\mathbf{x}, t, \boldsymbol{\mu}) \rightarrow \boldsymbol{\eta}(\mathbf{x}, t), \quad (7)$$

which is essentially an inverse problem that maps the observations to the state. One could think of it as trying to find the inverse operator, $\mathcal{M} = \mathcal{H}^{-1}$, but this could be some other arbitrary operator.

4.2 Machine Learning Model Ontology

In general, we are interested in finding some parameterized operator, \mathcal{M}_θ , that maps the incomplete SSH field to the complete SSH field

$$\mathcal{M}_\theta : \boldsymbol{\eta}_{obs}(\mathbf{x}, t, \boldsymbol{\mu}) \rightarrow \boldsymbol{\eta}(\mathbf{x}, t), \quad (8)$$

whereby we learn the parameters from data. The two main tasks we can define from this problem setup are 1) interpolation and 2) extrapolation. We define *interpolation* as the case when the boundaries of the inferred state domain lie within a predefined shape for the boundaries of the spatiotemporal observation domain. For example, the shape of the spatial domain could be a line, box, or sphere, and the shape of the temporal domain could be a positive real number line. We define *extrapolation* as the case where the boundaries of the inferred state domain are outside the boundaries of the spatiotemporal observation domain. In this case, the inferred state domain could be outside of either domain or both. A prevalent specific case of extrapolation is *hindcasting* or *forecasting*, where the inferred state domain lies within the spatial observation domain’s boundaries but outside of the temporal observation domain’s. In the rest of this paper, we will look exclusively at the interpolation problem. However, we refer the reader to appendix F for a more detailed look at other subtasks that can arise.

From a ML point of view, we can explore various ways to define the operator in equation (7). We may distinguish three main categories: (i) coordinate-based methods that learn a parameterized continuous function to map the domain coordinates to the scalar values, (ii) the explicit mapping of the state from the observation, (iii) implicit methods defined as the solution of an optimization problem. The first category comprises of kriging approaches, which have been used operationally with historical success [113, 95]. Beyond such covariance-based approaches, recent contributions explore more complex trainable functional models [72], basis functions [102], and neural networks [62]. The second category of schemes bypasses the physical modeling aspect and amortizes the prediction directly using state-of-the-art neural architectures such as UNets and ConvLSTMs [116, 74, 9]. This category may straightforwardly benefit from available auxiliary observations [23, 24, 22] to state the interpolation problem as a super-resolution [107] or image-to-image translation problem [81, 59]. The third category relates to inverse problem formulations and associated deep learning schemes, for example deep unfolding methods and plug-and-play priors [115]. Interestingly, recent contributions explore novel neural schemes which combine data assimilation formulations [21] and learned optimizer strategies [14, 44]. We provide a more detailed ontology of methods used for interpolation problems in appendix G. We consider at least one baseline approach from each category for each data challenge described in section 4.4. While all these methods have pros and cons, we expect the OceanBench platform to showcase to new experimental evidence and understanding regarding their applicability to SSH interpolation problems.

4.3 Experimental Design

Table 1: This table gives a brief overview of the datasets provided to complete the data challenges listed in 4.4 and A. Note that the OSSE datasets are all gridded products whereas the OSE NADIR is an alongtrack product. See figure 2 for an example of the OSSE NEMO Simulations for SSH and SST and pseudo-observations for NADIR & SWOT.

	OSSE	OSSE NADIR + SWOT	OSSE SST	OSE NADIR
Data Type	Simulations	Pseudo-Observations	Simulations	Observations
Source	NEMO [5]	NEMO [5]	NEMO [5]	Altimetry [32]
Region	GulfStream	GulfStream	GulfStream	GulfStream
Domain Size	$10 \times 10^\circ$	$10 \times 10^\circ$	$10 \times 10^\circ$	$10 \times 10^\circ$
Longitude Extent	$[-65^\circ, -55^\circ]$	$[-65^\circ, -55^\circ]$	$[-65^\circ, -55^\circ]$	$[-65^\circ, -55^\circ]$
Latitude Extent	$[33^\circ, 43^\circ]$	$[33^\circ, 43^\circ]$	$[33^\circ, 43^\circ]$	$[33^\circ, 43^\circ]$
Resolution	$0.05^\circ \times 0.05^\circ$	$0.05^\circ \times 0.05^\circ$	$0.05^\circ \times 0.05^\circ$	7 km
Grid Size	200×200	200×200	200×200	N/A
Num Datapoints	$\sim 14.6M$	$\sim 14.6M$	$\sim 14.6M$	$\sim 1.6M$
Period Start	2012-10-01	2012-10-01	2012-10-01	2016-12-01
Period End	2013-09-30	2013-09-30	2013-09-30	2018-01-31
Frequency	Daily	Daily	Daily	1 Hz

The availability of multi-year simulation and observation datasets naturally advocates for the design of synthetic (or twin) experiments, referred to as observing system simulation experiments (OSSE), and of real-world experiments, referred to as observing system experiments (OSE). We outline these two experimental setups below.

Observing System Simulation Experiments (OSSE). A staple and groundtruthed experimental setup uses a reference simulation dataset to simulate the conditions we can expect from actual satellite observations. This setup allows researchers and operational centers to create a fully-fledged pipeline that mirrors the real-world experimental setting. An ocean model simulation is deployed over a specified spatial domain and period, and a satellite observation simulator is deployed to simulate satellite observations over the same domain and period. This OSSE setup has primarily been considered for performance evaluation, as one can assess a reconstruction performance over the entire space-time domain. It also provides the basis for the implementation of classic supervised learning strategies [9, 74, 116]. The domain expert can vary the experimental conditions depending on the research question. For example, one could specify a region based on the expected dynamical regime [12] or add a certain noise level to the observation tracks based on the satellite specifications. The biggest downside to OSSE experiments is that we train models exclusively with ocean simulations which could produce models that fail to generalize to the actual ocean state. Furthermore, the simulations are often quite expensive, which prevents the community from having high spatial resolution over very long periods, which would be essential to capture as many dynamical regimes as possible.

Observing System Experiments (OSE). As more observations have become available over the past few decades, we can also design experiments using real data. This involves aggregating as many observations from real ocean altimetry satellites as possible with some specific independent subset left out for evaluation purposes. A major downside to OSE experiments is that the sparsity and spatial coverage of the observations narrow the possible scope of performance metrics and make it very challenging to learn directly from observation datasets. The current standard altimetry data are high resolution but cover a tiny area. As such, it can only inform fine-scale SSH patterns in the along-track satellite direction and cannot explicitly reveal two-dimensional patterns. Despite these drawbacks, it provides a quantitative evaluation of the generalizability of the ML methods concerning the true ocean state.

4.4 Data Challenges

We rely on existing OSSE and OSE experiments for SSH interpolation designed by domain experts [13, 12] and recast them into OceanBench framework to deliver a ML-ready benchmarking suites. The selected data challenges for this first edition address SSH interpolation for a 1000km×1000km Gulfstream region. We briefly outline them below.

Experiment I (OSSE NADIR) addresses SSH interpolation using NADIR altimetry tracks which are very fine, thin ocean satellite observations (see Figure 2). It relies on an OSSE using high-resolution (1/60° resolution) ocean simulations generated by the NEMO model over one year with a whole field every day.

Experiment II (OSSE SWOT) addresses SSH interpolation using jointly NADIR and SWOT altimetry data where we complement the **OSSE NADIR** configuration with simulated SWOT observations. SWOT is a new satellite altimetry mission with a much higher spatial coverage but a much lower temporal resolution as illustrated in Figure 2. The higher spatial resolution allows us to see structures at a smaller resolution but at the cost of a massive influx of observations (over ×100).

Experiment III (OSSE SST) addresses SSH interpolation using altimetry and SST satellite data jointly. We complement the **OSSE SWOT** challenge with simulated SST observations. Satellite-derived SST observations are more abundantly available in natural operational settings than SSH at a finer resolution, and structures have visible similarities [51, 55]. So this challenge allows for methods to take advantage of multi-modal learning [44, 116].

Experiment IV (OSE NADIR) addresses SSH interpolation for real NADIR altimetry data. In contrast to the three OSSE data challenges, it only looks at actual observations aggregated from the currently available ocean altimetry data from actual satellites. It involves a similar space-time sampling as Experiment (**OSSE NADIR**) to evaluate the generalization of ML methods trained in Experiment I to real altimetry data. The training problem’s complexity increases significantly due to the reference dataset’s sparsity compared with the **OSSE NADIR** dataset. One may also explore transfer learning or fine-tuning strategies from the available OSSE dataset.

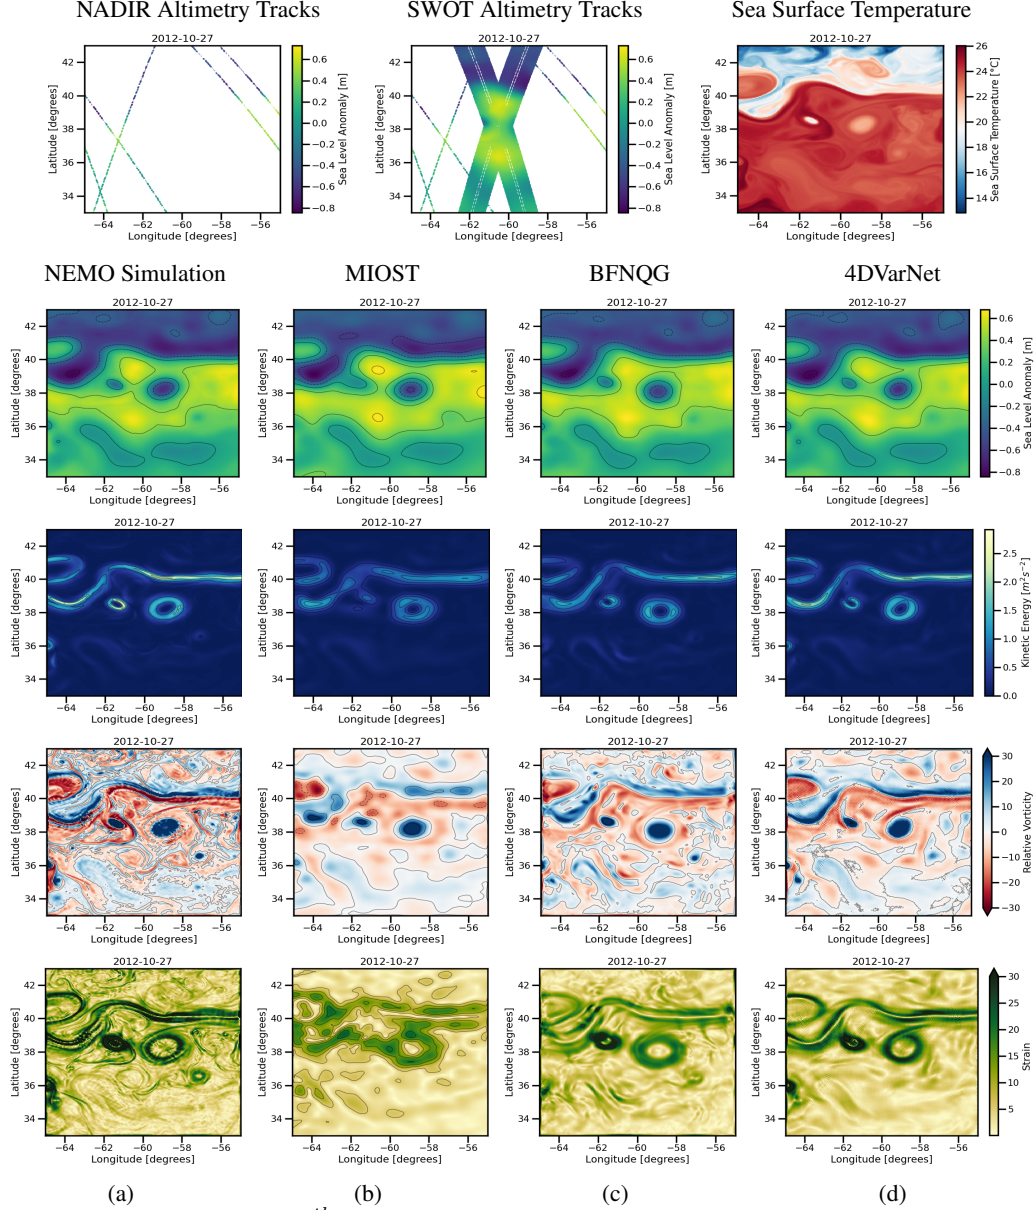


Figure 2: A snapshot at 27th October, 2012 of the sea level anomaly (SLA) from the NEMO simulation for the OSSE experiment outlined in section 4.3. The top row showcases the aggregated NADIR altimetry tracks and the aggregated SWOT altimetry tracks (12 hours before and 12 hours after) as well as the SST from the NEMO simulation. Each subsequent row showcases the following physical variables found in appendix B: (a) Sea Level Anomaly, (b) Kinetic Energy, (c) Relative Vorticity, and (d) Strain. Each column in the subsequent rows showcase the following reconstructed field from the NEMO simulation found in column (a): (b) MIOST [102], (c) BFN-QG [55], and (d) 4DVarNet [14].

4.5 OceanBench Pipelines

For the four data challenges presented in the previous section, we used OceanBench pipelines to deliver a ML-ready benchmarking framework. We used the `hydra` and the `geoprocessing` tools outlined in section 3.2 with specialized routines for regriding the ocean satellite data to a uniformly gridded product and vice versa when necessary. Appendix D showcases an example of the `hydra` integration for the preprocessing pipeline. A key feature is the creation of a custom patcher for the appropriate geophysical variables using our `XRPatcher` tool, which is later integrated into custom datasets and dataloaders for the appropriate model architecture, e.g., coordinate-based or grid-based.




Figure 3: This figure showcases some statistics for evaluation of the SSH field reconstructions for the OSSE NADIR experiment outlined in section 4. Subfigure (a) showcases the normalized root mean squared error (nRMSE), (b) showcases the isotropic power spectrum decomposition (PSD), (c) showcases isotropic PSD scores. The bottom row showcases the space-time PSD for the NEMO simulation (subfigure (d)) and the PSD scores for three reconstruction models: (e) the MIOST model [102], (f) the BFN-QG model [55], and (g) the 4DVarNet model [14].

Table 2: This table highlights some of the results for the OSSE NADIR experiment outlined in section 4.4 and appendix A. This table highlights the performance statistically in the real and spectral space; the normalized RMSE score for the real space and the minimum spatial and temporal scales resolved in the spectral domain. For more information about the class of models displayed and class of metrics, see appendix G and appendix C respectively. We only showcase the model performance on the alongtrack NADIR data available. For the extended table for each of the challenges, see Table 3.

Experiment	Algorithm	Algorithm Class	nRMSE Score	λ_x [km]	λ_t [days]
OSSE NADIR	OI [95]	Coordinate-Based	0.92 ± 0.01	175	10.8
OSSE NADIR	MIOST [102]	Coordinate-Based	0.93 ± 0.01	157	10.1
OSSE NADIR	BFNQG [55]	Hybrid Model	0.93 ± 0.01	139	10.6
OSSE NADIR	4DVarNet [14]	Bi-Level Opt.	0.95 ± 0.01	117	7.7

We provide an example snippet of how this can be done easily in section E. OceanBench also features some tools specific to the analysis of SSH. For example, physically-interpretable variables like geostrophic currents and relative vorticity, which can be derived from first-order and second-order derivatives of the SSH, are essential for assessing the quality of the reconstructions generated by the models. Figure 2 showcases some fields of the most common physical variables used in the oceanography literature for the SSH-based analysis of sea surface dynamics. For more details regarding the nature of the physical variables, see appendix B.

Regarding the evaluation framework, we include domain-relevant performance metrics beyond the standard ML loss and accuracy functions. They account for the sampling patterns of the evaluation data. Spectral analytics are widely used in geoscience [55], and here, we consider spectral scores computed as the minimum spatial and temporal scales resolved by the reconstruction methods proposed in [55]. For example, figure 3 showcases how OceanBench generated the isotropic power spectrum and score and the space-time power spectrum decomposition and score. Table 2 outlines some standard and domain-specific scores for the experiments outlined in section 4.3. We give a more detailed description of the rationale and construction of the power-spectrum-specific metrics in

appendix [C](#). In terms of baselines, we report for each data challenge the performance of at least one approach for each of the category outlined in Section [4.2](#).

5 Conclusions

The ocean community faces technological and algorithmic challenges to make the most of available observation and simulation datasets. In this context, recent studies evidence the critical role of ML schemes in reaching breakthroughs in our ability to monitor ocean dynamics for various space-time scales and processes. Nevertheless, domain-specific preprocessing steps and evaluation procedures slow down the uptake of ML toward real-world applications. Our application of choice was SSH mapping which facilitates the production of many crucial derived products that are used in many downstream tasks like subsequent modeling [\[93\]](#), ocean health monitoring [\[101, 73, 47\]](#) and maritime risk assessment [\[106\]](#).

Through OceanBench framework, we embed domain-level requirements into the MLOPs considerations by building a flexible framework that adds this into the hyperparameter considerations for ML models. We proposed four challenges towards a ML-ready benchmarking suite for ocean observation challenges. We outlined the inner workings OceanBench and demonstrated its usefulness by recreating some preprocessing and analysis pipelines from a few data challenges involving SSH interpolation. We firmly believe that the OceanBench platform is a crucial step to lowering the barrier of entry for new ML researchers interested in applying and developing their methods to relevant problems in the ocean sciences.

Acknowledgments and Disclosure of Funding

This work was supported by the French National Research Agency (ANR), through projects number ANR-17-CE01-0009-01, ANR-19-CE46-0011 and ANR-19-CHIA-0016); by the French National Space Agency (CNES) through the SWOT Science Team program (projects MIDAS and DIEGO) and the OSTST program (project DUACS-HR); by the French National Centre for Scientific Research (CNRS) through the LEFE-MANU program (project IA-OAC). This project also received funding from the European Union’s Horizon Europe research and innovation programme under the grant No 101093293 (EDITO-Model Lab project). This project benefited from HPC and GPU computing resources from GENCI-IDRIS (Grant 2021-101030).

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] All the contributions listed in the abstract are elaborated in sections 3.2, 4.4 and 5
 - (b) Did you describe the limitations of your work? [Yes] See the last paragraph of section 5 and the appendix as well.
 - (c) Did you discuss any potential negative societal impacts of your work? [Yes] We do not believe that our work has any potential negative societal impacts directly as we do not deal with any confidential or private data. However, we do outline in the appendix how there may be some adverse effects related to downstream uses which could have some negative societal impacts.
 - (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes] We do not include any confidential or private data. We only include numerical values which stem from general physical systems or machine learning models. We do not believe they hold any ethical issues. However, we do acknowledge that there would be environmental damage should users go forward and explore methods which obscenely high computing hours. This discussion outlined in the appendix.
2. If you are including theoretical results...
 - (a) Did you state the full set of assumptions of all theoretical results? [N/A] We do not include any theoretical results.
 - (b) Did you include complete proofs of all theoretical results? [N/A] We do not include any theoretical results.
3. If you ran experiments (e.g. for benchmarks)...
 - (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] We include the parameters used to reproduce the dataset preprocessing and evaluation procedure in Appendix A and instructions are given to download the data via <https://github.com/quentinf00/oceanbench-data-registry> and rerun the evaluation procedure in our code repository which is available at <https://github.com/jejjohnson/oceanbench>
 - (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes] We showcase all preprocessing steps necessary to reproduce the experimental configurations in Appendix A and the configuration files are available in our code repository at <https://github.com/jejjohnson/oceanbench>
 - (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [N/A] This is not applicable for this instantiation because we do not include any randomness within the experiment procedure nor the results.
 - (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] We do not do any model training and leave it up the user for their local or cloud machine. However, we do provide the cloud provider for the data found the the data registry which can be found at <https://github.com/quentinf00/oceanbench-data-registry>
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] We adopted the implementation of the preprocessing procedures and evaluation steps with some modifications. We give proper citation and credit to the authors as well as all other existing software packages included in this work.
 - (b) Did you mention the license of the assets? [Yes] The appropriate license notices are included in the source code files.
 - (c) Did you include any new assets either in the supplemental material or as a URL? [Yes] All the processing and evaluation scripts are included in the GitHub repository.
 - (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [Yes] We only include data that is already publicly available. We also discussed with the original generators of the datasets and keep the appropriate licenses.
 - (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A] We do not include any personal information or offensive content in our datasets.
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] We do not use crowdsourcing and we do not conduct research with human subjects.
 - (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A] See the previous point.
 - (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A] See the previous point.

References

- [1] Martín Abadi, Paul Barham, Jianmin Chen, Zhifeng Chen, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Geoffrey Irving, Michael Isard, et al. Tensorflow: A system for large-scale machine learning. In *OSDI*, volume 16, pages 265–283, 2016.
- [2] Saleh Abdalla, Abdolnabi Abdeh Kolahchi, (...), 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.
- [3] Ryan Abernathey, rochanotes, Andrew Ross, Malte Jansen, Ziwei Li, Francis J. Poulin, Navid C. Constantinou, Anirban Sinha, Dhruv Balwada, SalahKouhen, Spencer Jones, Cesar B Rocha, Christopher L. Pitt Wolfe, Chuizheng Meng, Hugo van Kemenade, James Bourbeau, James Penn, Julius Busecke, Mike Bueti, and Tobias. pyqg/pyqg: v0.7.2, May 2022.
- [4] M. Ablain, J. F. Legeais, P. Prandi, M. Marcos, L. Fenoglio-Marc, H. B. Dieng, J. Benveniste, and A. Cazenave. Satellite Altimetry-Based Sea Level at Global and Regional Scales. *Surveys in Geophysics*, 38(1):7–31, January 2017.
- [5] Adekunle Ajayi, Julien Le Sommer, Eric Chassignet, Jean-Marc Molines, Xiaobiao Xu, Aurelie Albert, and Emmanuel Cosme. Spatial and temporal variability of the north atlantic eddy field from two kilometeric-resolution ocean models. *Journal of Geophysical Research: Oceans*, 125(5):e2019JC015827, 2020. e2019JC015827 10.1029/2019JC015827.
- [6] Aida Alvera Azcarate, Alexander Barth, Jean-Marie Beckers, and Robert H Weisberg. Multivariate reconstruction of missing data in sea surface temperature, chlorophyll, and wind satellite fields. *Journal of Geophysical Research. Oceans*, 112(C3), 2007.
- [7] 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.
- [8] Jeffrey Anderson, Tim Hoar, Kevin Raeder, Hui Liu, Nancy Collins, Ryan Torn, and Avelino Avellano. The data assimilation research testbed: A community facility. *Bulletin of the American Meteorological Society*, 90(9):1283 – 1296, 2009.

- [9] Théo Archambault, Arthur Filoche, Anastase Charantonnis, and Dominique Béréziat. Multimodal Unsupervised Spatio-Temporal Interpolation of satellite ocean altimetry maps. In *VISAPP*, Lisboa, Portugal, February 2023.
- [10] Saleh Ashkboos, Langwen Huang, Nikoli Dryden, Tal Ben-Nun, Peter Dueben, Lukas Gianinazzi, Luca Kummer, and Torsten Hoefler. Ens-10: A dataset for post-processing ensemble weather forecasts. In S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh, editors, *Advances in Neural Information Processing Systems*, volume 35, pages 21974–21987. Curran Associates, Inc., 2022.
- [11] D. C. E. Bakker, B. Pfeil, C. S. Landa, N. Metzl, K. M. O’Brien, A. Olsen, K. Smith, C. Cosca, S. Harasawa, S. D. Jones, S. Nakaoka, Y. Nojiri, U. Schuster, T. Steinhoff, C. Sweeney, T. Takahashi, B. Tilbrook, C. Wada, R. Wanninkhof, S. R. Alin, C. F. Balestrini, L. Barbero, N. R. Bates, A. A. Bianchi, F. Bonou, J. Boutin, Y. Bozec, E. F. Burger, W.-J. Cai, R. D. Castle, L. Chen, M. Chierici, K. Currie, W. Evans, C. Featherstone, R. A. Feely, A. Fransson, C. Goyet, N. Greenwood, L. Gregor, S. Hankin, N. J. Hardman-Mountford, J. Harlay, J. Hauck, M. Hoppema, M. P. Humphreys, C. W. Hunt, B. Huss, J. S. P. Ibáñez, T. Johannessen, R. Keeling, V. Kitidis, A. Körtzinger, A. Kozyr, E. Krasakopoulou, A. Kuwata, P. Landschützer, S. K. Lauvset, N. Lefèvre, C. Lo Monaco, A. Manke, J. T. Mathis, L. Merlivat, F. J. Millero, P. M. S. Monteiro, D. R. Munro, A. Murata, T. Newberger, A. M. Omar, T. Ono, K. Paterson, D. Pearce, D. Pierrot, L. L. Robbins, S. Saito, J. Salisbury, R. Schlitzer, B. Schneider, R. Schweitzer, R. Sieger, I. Skjelvan, K. F. Sullivan, S. C. Sutherland, A. J. Sutton, K. Tadokoro, M. Telszewski, M. Tuma, S. M. A. C. van Heuven, D. Vandemark, B. Ward, A. J. Watson, and S. Xu. A multi-decade record of high-quality f_{CO_2} data in version 3 of the surface ocean CO_2 atlas (socat). *Earth System Science Data*, 8(2):383–413, 2016.
- [12] Maxime Ballarotta, Emmanuel Cosme, and Aurélie Albert. ocean-data-challenges/2020a_SSH_mapping_NATL60: Material for SSH mapping data challenge, September 2020. This challenge is part of the BOOST-SWOT project funded by ANR (project number ANR-17-CE01-0009-01) and a contribution to the MIDAS project funded by CNES for the NASA/CNES SWOT Science Team.
- [13] Maxime Ballarotta and Florian Le Guillou. ocean-data-challenges/2021a_SSH_mapping_OSE: Material for SSH mapping OSE data challenge, September 2021.
- [14] Maxime Beauchamp, Quentin Febvre, Hugo Geogrentum, and Ronan Fablet. 4dvarnet-ssh: end-to-end learning of variational interpolation schemes for nadir and wide-swath satellite altimetry. *Geoscientific Model Development*, 2022.
- [15] David W. Behringer, Ming Ji, and Ants Leetmaa. An improved coupled model for ENSO prediction and implications for ocean initialization. part i: The ocean data assimilation system. *Monthly Weather Review*, 126(4):1013 – 1021, 1998.
- [16] Renato Berlinghieri, Brian L. Trippe, David R. Burt, Ryan James Giordano, Kaushik Srinivasan, Tamay Özgökmen, Junfei Xia, and Tamara Broderick. Gaussian processes at the helm(holtz): A more fluid model for ocean currents. In Andreas Krause, Emma Brunskill, Kyunghyun Cho, Barbara Engelhardt, Sivan Sabato, and Jonathan Scarlett, editors, *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pages 2113–2163. PMLR, 23–29 Jul 2023.
- [17] Thomas Bolton and Laure Zanna. Applications of deep learning to ocean data inference and sub-grid parameterisation. *Journal of Advances in Modeling Earth Systems*, 2019.
- [18] Massimo Bonavita and Patrick Laloyaux. Estimating model error covariances with artificial neural networks, 2022.
- [19] Boris Bonev, Thorsten Kurth, Christian Hundt, Jaideep Pathak, Maximilian Baust, Karthik Kashinath, and Anima Anandkumar. Spherical fourier neural operators: Learning stable dynamics on the sphere, 2023.
- [20] B. Buongiorno Nardelli. A multi-year time series of observation-based 3d horizontal and vertical quasi-geostrophic global ocean currents. *Earth System Science Data*, 12(3):1711–1723, 2020.

- [21] Alberto Carrassi, Marc Bocquet, Laurent Bertino, and Geir Evensen. Data assimilation in the geosciences: An overview of methods, issues, and perspectives. *WIREs Climate Change*, 9(5):e535, 2018.
- [22] E.U. Copernicus Climate Change Service (CCCS). Ocean colour daily data from 1997 to present derived from satellite observations.
- [23] E.U. Copernicus Climate Change Service (CCCS). Sea surface temperature daily data from 1981 to present derived from satellite observations.
- [24] E.U. Copernicus Climate Change Service (CCCS). Sea surface temperature daily gridded data from 1981 to 2016 derived from a multi-product satellite-based ensemble.
- [25] E.U. Copernicus Climate Change Service (CCCS). Oras5 global ocean reanalysis monthly data from 1958 to present, 2021.
- [26] M. A. Chamberlain, P. R. Oke, R. A. S. Fiedler, H. M. Beggs, G. B. Brassington, and P. Divakaran. Next generation of bluelink ocean reanalysis with multiscale data assimilation: Bran2020. *Earth System Science Data*, 13(12):5663–5688, 2021.
- [27] Dudley B. Chelton, Michael G. Schlax, and Roger M. Samelson. Global observations of nonlinear mesoscale eddies. *Progress in Oceanography*, 91(2):167–216, 2011.
- [28] Lijing Cheng, Karina Schuckmann, John Abraham, Kevin Trenberth, Michael Mann, Laure Zanna, Matthew England, Jan Zika, John Fasullo, Yongqiang Yu, Yuying Pan, Jiang Zhu, Emily Newsom, Ben Bronselaer, and Xiaopei Lin. Past and future ocean warming. *Nature Reviews Earth and Environment*, pages 1–19, 10 2022.
- [29] Sibó Cheng, César Quilodrán-Casas, Said Ouala, Alban Farchi, Che Liu, Pierre Tandeo, Ronan Fablet, Didier Lucor, Bertrand Iooss, Julien Brajard, et al. Machine learning with data assimilation and uncertainty quantification for dynamical systems: a review. *IEEE/CAA Journal of Automatica Sinica*, 10(6):1361–1387, 2023.
- [30] François Chollet et al. Keras. <https://keras.io>, 2015.
- [31] E.U. Copernicus Marine Service Information (CMEMS). Global ocean- in-situ near real time observations of ocean currents.
- [32] E.U. Copernicus Marine Service Information (CMEMS). Global ocean along-track 13 sea surface heights reprocessed (1993-ongoing) tailored for data assimilation.
- [33] E.U. Copernicus Marine Service Information (CMEMS). Global ocean biogeochemistry analysis and forecast.
- [34] E.U. Copernicus Marine Service Information (CMEMS). Global ocean ensemble physics reanalysis.
- [35] E.U. Copernicus Marine Service Information (CMEMS). Global ocean gridded 14 sea surface heights and derived variables reprocessed 1993 ongoing.
- [36] E.U. Copernicus Marine Service Information (CMEMS). Global ocean physics analysis and forecast.
- [37] E.U. Copernicus Marine Service Information (CMEMS). Global ocean waves reanalysis.
- [38] Executable Books Community. Jupyter book, February 2020.
- [39] Matthew P Couldrey, Jonathan M Gregory, Fabio Boeira Dias, Peter Dobrohotoff, Catia M Domingues, Oluwayemi Garuba, Stephen M Griffies, Helmuth Haak, Aixue Hu, Masayoshi Ishii, et al. What causes the spread of model projections of ocean dynamic sea-level change in response to greenhouse gas forcing? *Climate Dynamics*, pages 1–33, 2020.
- [40] Tim DeVries. The ocean carbon cycle. *Annual Review of Environment and Resources*, 47(1):317–341, 2022.

- [41] Emilien Dupont, Hyunjik Kim, S. M. Ali Eslami, Danilo Jimenez Rezende, and Dan Rosenbaum. From data to functa: Your data point is a function and you can treat it like one. In *International Conference on Machine Learning*, 2022.
- [42] Lasse Espeholt, Shreya Agrawal, Casper Sønderby, Manoj Kumar, Jonathan Heek, Carla Bromberg, Cenk Gizen, Rob Carver, Marcin Andrychowicz, Jason Hickey, Aaron Bell, and Nal Kalchbrenner. Deep learning for twelve hour precipitation forecasts. *Nature Communications*, 13(1):5145, 2022.
- [43] Ronan Fablet, Bertrand Chapron, Julien Le Sommer, and Florian S’avellec. Inversion of sea surface currents from satellite-derived sst-ssh synergies with 4dvarnets. *ArXiv*, abs/2211.13059, 2022.
- [44] Ronan Fablet, Quentin Febvre, and Bertrand Chapron. Multimodal 4dvarnets for the reconstruction of sea surface dynamics from sst-ssh synergies. *IEEE Transactions on Geoscience and Remote Sensing*, 61:1–14, 2022.
- [45] William Falcon and The PyTorch Lightning team. PyTorch Lightning, March 2019.
- [46] Quentin Febvre, Clément Ubelmann, Julien Le Sommer, and Ronan Fablet. Scale-aware neural calibration for wide swath altimetry observations. *ArXiv*, abs/2302.04497, 2023.
- [47] Andrea Franke, Thorsten Blenckner, Carlos M. Duarte, Konrad Ott, Lora E. Fleming, Avan Antia, Thorsten B.H. Reusch, Christine Bertram, Jonas Hein, Ulrike Kronfeld-Goharani, Jan Dierking, Annegret Kuhn, Chie Sato, Erik van Doorn, Marlene Wall, Markus Schartau, Rolf Karez, Larry Crowder, David Keller, Anja Engel, Ute Hentschel, and Enno Prigge. Operationalizing ocean health: Toward integrated research on ocean health and recovery to achieve ocean sustainability. *One Earth*, 2(6):557–565, 2020.
- [48] Hugo Frezat, Guillaume Balarac, Julien Le Sommer, Ronan Fablet, and Redouane Lguensat. Physical invariance in neural networks for subgrid-scale scalar flux modeling. *Phys. Rev. Fluids*, 6:024607, Feb 2021.
- [49] Hugo Frezat, Julien Le Sommer, Ronan Fablet, Guillaume Balarac, and Redouane Lguensat. A posteriori learning for quasi-geostrophic turbulence parametrization. *Journal of Advances in Modeling Earth Systems*, 14(11):e2022MS003124, 2022. e2022MS003124 2022MS003124.
- [50] Martin Gauch, Frederik Kratzert, Oren Gilon, Hoshin Gupta, Juliane Mai, Grey Nearing, Bryan Tolson, Sepp Hochreiter, and Daniel Klotz. In defense of metrics: Metrics sufficiently encode typical human preferences regarding hydrological model performance. *Water Resources Research*, 59(6):e2022WR033918, 2023. e2022WR033918 2022WR033918.
- [51] Lucile Gaultier, Clément Ubelmann, and Lee-Lueng Fu. The challenge of using future swot data for oceanic field reconstruction. *Journal of Atmospheric and Oceanic Technology*, 33(1):119 – 126, 2016.
- [52] William Gilpin. Chaos as an interpretable benchmark for forecasting and data-driven modelling. In J. Vanschoren and S. Yeung, editors, *Proceedings of the Neural Information Processing Systems Track on Datasets and Benchmarks*, volume 1. Curran, 2021.
- [53] Steve Groom, Shubha Sathyendranath, Yai Ban, Stewart Bernard, Robert Brewin, Vanda Brotas, Carsten Brockmann, Prakash Chauhan, Jong-kuk Choi, Andrei Chuprin, Stefano Ciavatta, Paolo Cipollini, Craig Donlon, Bryan Franz, Xianqiang He, Takafumi Hirata, Tom Jackson, Milton Kampel, Hajo Krasemann, Samantha Lavender, Silvia Pardo-Martinez, Frédéric Mélin, Trevor Platt, Rosalia Santoleri, Jozef Skakala, Blake Schaeffer, Marie Smith, Francois Steinmetz, Andre Valente, and Menghua Wang. Satellite ocean colour: Current status and future perspective. *Frontiers in Marine Science*, 6, 2019.
- [54] Arthur P. Guillaumin and Laure Zanna. Stochastic-deep learning parameterization of ocean momentum forcing. *Journal of Advances in Modeling Earth Systems*, n/a(n/a):e2021MS002534, 2021. e2021MS002534 2021MS002534.

- [55] 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.
- [56] D. Häfner, R. L. Jacobsen, C. Eden, M. R. B. Kristensen, M. Jochum, R. Nuterman, and B. Vinter. Veros v0.1 – a fast and versatile ocean simulator in pure python. *Geoscientific Model Development*, 11(8):3299–3312, 2018.
- [57] Helene T Hewitt, Malcolm Roberts, Pierre Mathiot, Arne Biastoch, Ed Blockley, Eric P Chassignet, Baylor Fox-Kemper, Pat Hyder, David P Marshall, Ekaterina Popova, et al. Resolving and parameterising the ocean mesoscale in earth system models. *Current Climate Change Reports*, pages 1–16, 2020.
- [58] S. Hoyer and J. Hamman. xarray: N-D labeled arrays and datasets in Python. *Journal of Open Research Software*, 5(1), 2017.
- [59] Henri Hoyez, Cédric Schockaert, Jason Rambach, Bruno Mirbach, and Didier Stricker. Unsupervised image-to-image translation: A review. *Sensors*, 22(21), 2022.
- [60] Lellouche Jean-Michel, Greiner Eric, Bourdallé-Badie Romain, Garric Gilles, Melet Angélique, Drévillon Marie, Bricaud Clément, Hamon Mathieu, Le Galloudec Olivier, Regnier Charly, Candela Tony, Testut Charles-Emmanuel, Gasparin Florent, Ruggiero Giovanni, Benkiran Mounir, Drillet Yann, and Le Traon Pierre-Yves. The copernicus global 1/12 ° oceanic and sea ice glorys12 reanalysis. *Frontiers in Earth Science*, 9, 2021.
- [61] J. Emmanuel Johnson, Valero Laparra, María Piles, and Gustau Camps-Valls. Gaussianizing the earth: Multidimensional information measures for earth data analysis. *IEEE Geoscience and Remote Sensing Magazine*, 9:191–208, 2020.
- [62] J. Emmanuel Johnson, Redouane Lguensat, Ronan Fablet, Emmanuel Cosme, and Julien Le Sommer. Neural fields for fast and scalable interpolation of geophysical ocean variables. *ArXiv*, abs/2211.10444, 2022.
- [63] George Em Karniadakis, Ioannis G. Kevrekidis, Lu Lu, Paris Perdikaris, Sifan Wang, and Liu Yang. Physics-informed machine learning. *Nature Reviews Physics*, 3(6), 5 2021.
- [64] Dmitrii Kochkov, Jamie A. Smith, Ayya Alieva, Qing Wang, Michael P. Brenner, and Stephan Hoyer. Machine learning–accelerated computational fluid dynamics. *Proceedings of the National Academy of Sciences*, 118(21), 2021.
- [65] Simon Kornblith, Mohammad Norouzi, Honglak Lee, and Geoffrey Hinton. Similarity of neural network representations revisited. In Kamalika Chaudhuri and Ruslan Salakhutdinov, editors, *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pages 3519–3529. PMLR, 09–15 Jun 2019.
- [66] Dominik Kreuzberger, Niklas Köhl, and Sebastian Hirschl. Machine learning operations (mlops): Overview, definition, and architecture, 2022.
- [67] Ruslan Kuprieiev, skshetry, Peter Rowlands, Dmitry Petrov, Paweł Redzyński, Casper da Costa-Luis, David de la Iglesia Castro, Alexander Schepanovski, Gao, Ivan Shcheklein, Batuhan Taskaya, Jorge Orpinel, Dave Berenbaum, Fábio Santos, daniele, Ronan Lamy, Aman Sharma, Zhanibek Kaimuldenov, Dani Hodovic, Nikita Kodenko, Andrew Grigorev, Earl, Nabanita Dash, George Vyshnya, maykulkarni, Max Hora, Vera, and Sanidhya Mangal. Dvc: Data version control - git for data & models, May 2023.
- [68] Patrick Laloyaux, Thorsten Kurth, Peter Dominik Dueben, and David Hall. Deep learning to estimate model biases in an operational nwp assimilation system. *Journal of Advances in Modeling Earth Systems*, 14(6):e2022MS003016, 2022. e2022MS003016 2022MS003016.
- [69] Rémi R. Lam, Alvaro Sanchez-Gonzalez, Matthew Willson, Peter Wirnsberger, Meire Fortunato, Alexander Pritzel, Suman V. Ravuri, Timo Ewalds, Ferran Alet, Zach Eaton-Rosen, Weihua Hu, Alexander Merose, Stephan Hoyer, George Holland, Jacklynn Stott, Oriol Vinyals, Shakir Mohamed, and Peter W. Battaglia. Graphcast: Learning skillful medium-range global weather forecasting. *ArXiv*, abs/2212.12794, 2022.

- [70] Valero Laparra, J. Emmanuel Johnson, Gustau Camps-Valls, Raúl Santos-Rodríguez, and Jesús Malo. Information theory measures via multidimensional gaussianization. *ArXiv*, abs/2010.03807, 2020.
- [71] Redouane Lguensat, Pierre Tandeo, Pierre Ailliot, Manuel Pulido, and Ronan Fablet. The analog data assimilation. *Monthly Weather Review*, 145(10):4093 – 4107, 2017.
- [72] Haitao Liu, Y. Ong, Xiaobo Shen, and Jianfei Cai. When gaussian process meets big data: A review of scalable gps. *IEEE Transactions on Neural Networks and Learning Systems*, 31:4405–4423, 2018.
- [73] Catherine S. Longo, Melanie Frazier, Scott C. Doney, Jennie E. Rheuban, Jennifer Macy Humberstone, and Benjamin S. Halpern. Using the ocean health index to identify opportunities and challenges to improving southern ocean ecosystem health. *Frontiers in Marine Science*, 4, 2017.
- [74] Scott A. Martin, Georgy E. Manucharyan, and Patrice Klein. Synthesizing sea surface temperature and satellite altimetry observations using deep learning improves the accuracy and resolution of gridded sea surface height anomalies. *Journal of Advances in Modeling Earth Systems*, 15(5):e2022MS003589, 2023. e2022MS003589 2022MS003589.
- [75] Scott A. Martin, Georgy E. Manucharyan, and Patrice Klein. Synthesizing sea surface temperature and satellite altimetry observations using deep learning improves the accuracy and resolution of gridded sea surface height anomalies. *Journal of Advances in Modeling Earth Systems*, 15(5):e2022MS003589, 2023. e2022MS003589 2022MS003589.
- [76] Daisuke Matsuoka, Fumiaki Araki, Yumi Inoue, and Hideharu Sasaki. A new approach to ocean eddy detection, tracking, and event visualization –application to the northwest pacific ocean. *Procedia Computer Science*, 80:1601–1611, 2016. International Conference on Computational Science 2016, ICCS 2016, 6-8 June 2016, San Diego, California, USA.
- [77] P.J. Minnett, A. Alvera-Azcárate, T.M. Chin, G.K. Corlett, C.L. Gentemann, I. Karagali, X. Li, A. Marsouin, S. Marullo, E. Maturi, R. Santoleri, S. Saux Picart, M. Steele, and J. Vazquez-Cuervo. Half a century of satellite remote sensing of sea-surface temperature. *Remote Sensing of Environment*, 233:111366, 2019.
- [78] Rosemary Morrow and Pierre-Yves Le Traon. Recent advances in observing mesoscale ocean dynamics with satellite altimetry. *Advances in Space Research*, 50(8):1062–1076, 2012. Oceanography, Cryosphere and Freshwater Flux to the Ocean.
- [79] S. Mulet, M.-H. Rio, A. Mignot, S. Guinehut, and R. Morrow. A new estimate of the global 3d geostrophic ocean circulation based on satellite data and in-situ measurements. *Deep Sea Research Part II: Topical Studies in Oceanography*, 77-80:70–81, 2012. Satellite Oceanography and Climate Change.
- [80] Akira Okubo. Horizontal dispersion of floatable particles in the vicinity of velocity singularities such as convergences. *Deep Sea Research and Oceanographic Abstracts*, 17(3):445–454, 1970.
- [81] Yingxue Pang, Jianxin Lin, Tao Qin, and Zhibo Chen. Image-to-image translation: Methods and applications. *IEEE Transactions on Multimedia*, 24:3859–3881, 2021.
- [82] Ananda Pascual, Yannice Faugère, Gilles Larnicol, and Pierre-Yves Le Traon. Improved description of the ocean mesoscale variability by combining four satellite altimeters. *Geophysical Research Letters*, 33(2), 2006.
- [83] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library, 2019.

- [84] Jaideep Pathak, Shashank Subramanian, Peter Harrington, Sanjeev Raja, Ashesh Chattopadhyay, Morteza Mardani, Thorsten Kurth, David Hall, Zongyi Li, Kamyar Azizzadenesheli, Pedram Hassanzadeh, Karthik Kashinath, and Animashree Anandkumar. Fourcastnet: A global data-driven high-resolution weather model using adaptive fourier neural operators, 2022.
- [85] Ali Ramadhan, Gregory LeClaire Wagner, Chris Hill, Jean-Michel Campin, Valentin Churavy, Tim Besard, Andre Souza, Alan Edelman, Raffaele Ferrari, and John Marshall. Oceananigans.jl: Fast and friendly geophysical fluid dynamics on gpus. *Journal of Open Source Software*, 5(53):2018, 2020.
- [86] Stephan Rasp, Peter D. Dueben, Sebastian Scher, Jonathan A. Weyn, Soukayna Mouatadid, and Nils Thuerey. Weatherbench: A benchmark data set for data-driven weather forecasting. *Journal of Advances in Modeling Earth Systems*, 12(11):e2020MS002203, 2020. e2020MS002203 10.1029/2020MS002203.
- [87] L. Resplandy, R. F. Keeling, Y. Eddebbar, M. K. Brooks, R. Wang, L. Bopp, M. C. Long, J. P. Dunne, W. Koeve, and A. Oschlies. Quantification of ocean heat uptake from changes in atmospheric o₂ and co₂ composition. *Nature*, 563(7729):105–108, 2018.
- [88] Andrew Ross, Ziwei Li, Pavel Perezhugin, Carlos Fernandez-Granda, and Laure Zanna. Benchmarking of machine learning ocean subgrid parameterizations in an idealized model. *Journal of Advances in Modeling Earth Systems*, 15(1):e2022MS003258, 2023.
- [89] Nithya Sambasivan, Shivani Kapania, Hannah Highfill, Diana Akrong, Praveen Paritosh, and Lora M Aroyo. “everyone wants to do the model work, not the data work”: Data cascades in high-stakes ai. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, CHI ’21, New York, NY, USA, 2021. Association for Computing Machinery.
- [90] B K Shivamoggi, G J F van Heijst, and L P J Kamp. The okubo–weiss criterion in hydrodynamic flows: geometric aspects and further extension. *Fluid Dynamics Research*, 54(1):015505, jan 2022.
- [91] Anirban Sinha and Ryan Abernathey. Estimating ocean surface currents with machine learning. *Frontiers in Marine Science*, 8, 2021.
- [92] Ryan Soklaski, Justin Goodwin, Olivia Brown, Michael Yee, and Jason Matterer. Tools and practices for responsible ai engineering. *arXiv preprint arXiv:2201.05647*, 2022.
- [93] Maike Sonnewald, Redouane Lguensat, Daniel C Jones, Peter D Dueben, Julien Brajard, and V Balaji. Bridging observations, theory and numerical simulation of the ocean using machine learning. *Environmental Research Letters*, 16(7):073008, July 2021.
- [94] Georgios Symeonidis, Evangelos Nerantzis, Apostolos Kazakis, and George A. Papakostas. Mlops - definitions, tools and challenges. *CoRR*, abs/2201.00162, 2022.
- [95] 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.
- [96] Makoto Takamoto, Timothy Praditia, Raphael Leiteritz, Daniel MacKinlay, Francesco Alesiani, Dirk Pflüger, and Mathias Niepert. Pdebench: An extensive benchmark for scientific machine learning. In S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh, editors, *Advances in Neural Information Processing Systems*, volume 35, pages 1596–1611. Curran Associates, Inc., 2022.
- [97] Pierre Tandeo, Pierre Ailliot, Juan Ruiz, Alexis Hannart, Bertrand Chapron, Anne Cuzol, Valérie Monbet, Robert Easton, and Ronan Fablet. Combining analog method and ensemble data assimilation: Application to the lorenz-63 chaotic system. In Valliappa Lakshmanan, Eric Gilleland, Amy McGovern, and Martin Tingley, editors, *Machine Learning and Data Mining Approaches to Climate Science*, pages 3–12, Cham, 2015. Springer International Publishing.
- [98] Rini Smita Thakur, Shubhojeet Chatterjee, Ram Narayan Yadav, and Lalita Gupta. Image de-noising with machine learning: A review. *IEEE Access*, 9:93338–93363, 2021.

- [99] Chunwei Tian, Lunke Fei, Wenxian Zheng, Yong Xu, Wangmeng Zuo, and Chia-Wen Lin. Deep learning on image denoising: An overview. *Neural networks : the official journal of the International Neural Network Society*, 131:251–275, 2019.
- [100] Anaëlle Tréboutte, Elisa Carli, Maxime Ballarotta, Benjamin Carpentier, Yannice Faugère, and Gérald Dibarboure. Karin noise reduction using a convolutional neural network for the swot ocean products. *Remote Sensing*, 15(8), 2023.
- [101] Devis Tuia, Benjamin Kellenberger, Sara Beery, Blair R. Costelloe, Silvia Zuffi, Benjamin Risse, Alexander Mathis, Mackenzie W. Mathis, Frank van Langevelde, Tilo Burghardt, Roland Kays, Holger Klinck, Martin Wikelski, Iain D. Couzin, Grant van Horn, Margaret C. Crofoot, Charles V. Stewart, and Tanya Berger-Wolf. Perspectives in machine learning for wildlife conservation. *Nature Communications*, 13(1):792, 2022.
- [102] Clement Ubelmann, Gérald Dibarboure, Lucile Gaultier, Aurélien Ponte, Fabrice Arduin, Maxime Ballarotta, and Yannice Faugère. Reconstructing ocean surface current combining altimetry and future spaceborne doppler data. *Journal of Geophysical Research: Oceans*, 126(3):e2020JC016560, 2021. e2020JC016560 2020JC016560.
- [103] UCAR/NCAR/CISL/DAReS. The data assimilation research testbed. <https://keras.io>, 2023.
- [104] Takaya Uchida, Quentin Jamet, Andrew C. Poje, Nico Wienders, William K. Dewar, and Bruno Deremble. Wavelet-based wavenumber spectral estimate of eddy kinetic energy: Idealized quasi-geostrophic flow. *Journal of Advances in Modeling Earth Systems*, 15(3):e2022MS003399, 2023. e2022MS003399 2022MS003399.
- [105] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017.
- [106] Karina von Schuckmann, Pierre-Yves Le Traon, Neville Smith, Ananda Pascual, Pierre Brasseur, Katja Fennel, Samy Djavidnia, Signe Aaboe, Enrique Alvarez Fanjul, Emmanuelle Autret, Lars Axell, Roland Aznar, Mario Benincasa, Abderahim Bentamy, Fredrik Boberg, Romain Bourdallé-Badie, Bruno Buongiorno Nardelli, Vittorio E. Brando, Clément Bricaud, Lars-Anders Breivik, Robert J.W. Brewin, Arthur Capet, Adrien Ceschin, Stefania Ciliberti, Gianpiero Cossarini, Marta de Alfonso, Alvaro de Pascual Collar, Jos de Kloe, Julie Deshayes, Charles Desportes, Marie Drévillon, Yann Drillet, Riccardo Droghei, Clotilde Dubois, Owen Embury, Héléne Etienne, Claudia Fratianni, Jesús García Lafuente, Marcos Garcia Sotillo, Gilles Garric, Florent Gasparin, Riccardo Gerin, Simon Good, Jérôme Goussier, Marilaure Grégoire, Eric Greiner, Stéphanie Guinehut, Elodie Gutknecht, Fabrice Hernandez, Olga Hernandez, Jacob Høyer, Laura Jackson, Simon Jandt, Simon Josey, Mélanie Juza, John Kennedy, Zoi Kokkini, Gerasimos Korres, Mariliis Kõuts, Priidik Lagemaa, Thomas Lavergne, Bernard le Cann, Jean-François Legeais, Benedicte Lemieux-Dudon, Bruno Levier, Vidar Lien, Ilja Maljutenko, Fernando Manzano, Marta Marcos, Veselka Marinova, Simona Masina, Elena Mauri, Michael Mayer, Angélique Melet, Frédéric Mélin, Benoit Meyssignac, Maeva Monier, Malte Müller, Sandrine Mulet, Cristina Naranjo, Giulio Notarstefano, Aurélien Paulmier, Begoña Pérez Gomez, Irene Pérez Gonzalez, Elisaveta Peneva, Coralie Perruche, K. Andrew Peterson, Nadia Pinardi, Andrea Pisano, Silvia Pardo, Pierre-Marie Poulain, Roshin P. Raj, Urmas Raudsepp, Michaelis Ravdas, Rebecca Reid, Marie-Hélène Rio, Stefano Salon, Annette Samuelsen, Michela Sammartino, Simone Sammartino, Anne Britt Sandø, Rosalia Santoleri, Shubha Sathyendranath, Jun She, Simona Simoncelli, Cosimo Solidoro, Ad Stoffelen, Andrea Storto, Tanguy Szerkely, Susanne Tamm, Steffen Tietsche, Jonathan Tinker, Joaquín Tintore, Ana Trindade, Daphne van Zanten, Luc Vandenbulcke, Anton Verhoef, Nathalie Verbrugge, Lena Viktorsson, Karina von Schuckmann, Sarah L. Wakelin, Anna Zacharioudaki, and Hao Zuo. Copernicus marine service ocean state report. *Journal of Operational Oceanography*, 11(sup1):S1–S142, 2018.
- [107] Zhihao Wang, Jian Chen, and Steven C. H. Hoi. Deep learning for image super-resolution: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 43:3365–3387, 2019.

- [108] D. Watson-Parris, Y. Rao, D. Olivié, Ø. Seland, P. Nowack, G. Camps-Valls, P. Stier, S. Bouabid, M. Dewey, E. Fons, J. Gonzalez, P. Harder, K. Jeggle, J. Lenhardt, P. Manshausen, M. Novitasari, L. Ricard, and C. Roesch. Climatebench v1.0: A benchmark for data-driven climate projections. *Journal of Advances in Modeling Earth Systems*, 14(10):e2021MS002954, 2022. e2021MS002954 2021MS002954.
- [109] John Weiss. The dynamics of enstrophy transfer in two-dimensional hydrodynamics. *Physica D: Nonlinear Phenomena*, 48(2):273–294, 1991.
- [110] Annie P. S. Wong, Susan E. Wijffels, Stephen C. Riser, Sylvie Pouliquen, Shigeki Hosoda, Dean Roemmich, John Gilson, Gregory C. Johnson, Kim Martini, David J. Murphy, Megan Scanderbeg, T. V. S. Udaya Bhaskar, Justin J. H. Buck, Frederic Merceur, Thierry Carval, Guillaume Maze, Cécile Cabanes, Xavier André, Noé Poffa, Igor Yashayaev, Paul M. Barker, Stéphanie Guinehut, Mathieu Belbéoch, Mark Ignaszewski, Molly O’Neil Baringer, Claudia Schmid, John M. Lyman, Kristene E. McTaggart, Sarah G. Purkey, Nathalie Zilberman, Matthew B. Alkire, Dana Swift, W. Brechner Owens, Steven R. Jayne, Cora Hersh, Pelle Robbins, Deb West-Mack, Frank Bahr, Sachiko Yoshida, Philip J. H. Sutton, Romain Cancouët, Christine Coatanoan, Delphine Dobbler, Andrea Garcia Juan, Jérôme Gourrion, Nicolas Kolodziejczyk, Vincent Bernard, Bernard Boulès, Hervé Claustre, Fabrizio D’Ortenzio, Serge Le Reste, Pierre-Yve Le Traon, Jean-Philippe Rannou, Carole Saout-Grit, Sabrina Speich, Virginie Thierry, Nathalie Verbrugge, Ingrid M. Angel-Benavides, Birgit Klein, Giulio Notarstefano, Pierre-Marie Poulain, Pedro Vélez-Belchí, Toshio Suga, Kentaro Ando, Naoto Iwasaka, Taiyo Kobayashi, Shuhei Masuda, Eitarou Oka, Kanako Sato, Tomoaki Nakamura, Katsunari Sato, Yasushi Takatsuki, Takashi Yoshida, Rebecca Cowley, Jenny L. Lovell, Peter R. Oke, Esmee M. van Wijk, Fiona Carse, Matthew Donnelly, W. John Gould, Katie Gowers, Brian A. King, Stephen G. Loch, Mary Mowat, Jon Turton, E. Pattabhi Rama Rao, M. Ravichandran, Howard J. Freeland, Isabelle Gaboury, Denis Gilbert, Blair J. W. Greenan, Mathieu Ouellet, Tetjana Ross, Anh Tran, Mingmei Dong, Zenghong Liu, Jianping Xu, KiRyong Kang, HyeongJun Jo, Sung-Dae Kim, and Hyuk-Min Park. Argo data 1999–2019: Two million temperature-salinity profiles and subsurface velocity observations from a global array of profiling floats. *Frontiers in Marine Science*, 7, 2020.
- [111] Hanyu Xiang, Qin Zou, Muhammad Ali Nawaz, Xianfeng Huang, Fan Zhang, and Hongkai Yu. Deep learning for image inpainting: A survey. *Pattern Recognition*, 134:109046, 2023.
- [112] Omry Yadan. Hydra - a framework for elegantly configuring complex applications. Github, 2019.
- [113] Fatemeh Zakeri and Gregoire Mariethoz. A review of geostatistical simulation models applied to satellite remote sensing: Methods and applications. *Remote Sensing of Environment*, 259:112381, 2021.
- [114] Laure Zanna and Thomas Bolton. Data-driven equation discovery of ocean mesoscale closures. *Geophysical Research Letters*, 47(17):e2020GL088376, 2020.
- [115] Kai Zhang and Radu Timofte. Chapter 14 - deep plug-and-play and deep unfolding methods for image restoration. In E.R. Davies and Matthew A. Turk, editors, *Advanced Methods and Deep Learning in Computer Vision*, Computer Vision and Pattern Recognition, pages 481–509. Academic Press, 2022.
- [116] Ke Zhang, Lei Huang, Zhiqiang Wei, Chen An, and Xianqing Lv. Sea surface height data reconstruction via inter and intra layer features based on dual attention. *Neurocomputing*, 545:126313, 2023.
- [117] Y. Zhang, X. Wang, Y. Sun, C. Ning, S. Xu, H. An, D. Tang, H. Guo, H. Yang, Y. Pu, B. Jiang, and B. Wang. Ocean modeling with adaptive resolution (omare, version 1.0) – refactoring nemo model (version 4.0.1) with the parallel computing framework of jasmin. part 1: adaptive grid refinement in an idealized double-gyre case. *EGUsphere*, 2022:1–37, 2022.