End-to-end neural data assimilation: application to ocean dynamics

R. Fablet, Prof. IMT Atlantique
Lab-STICC, INRIA team Odyssey, Brest, France
https://cia-oceanix.github.io/

with many colleagues

CEA-EDF-Inria Numerical Analysis school 2025, June 2025

General question

How to solve sampling gaps and extract high-level information for ocean modeling and forecasting ?

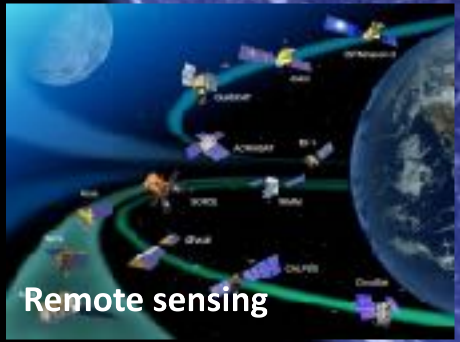

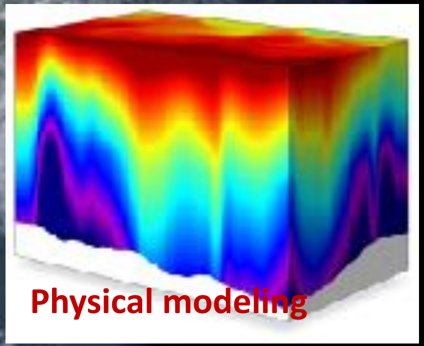

AI and Climate Science

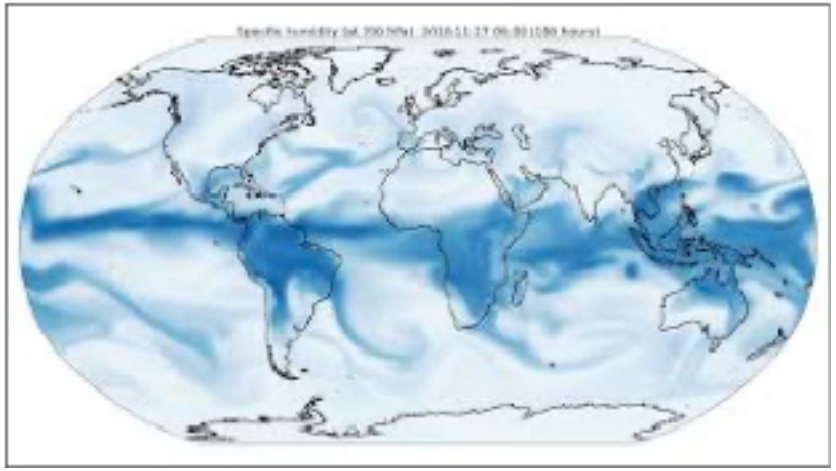

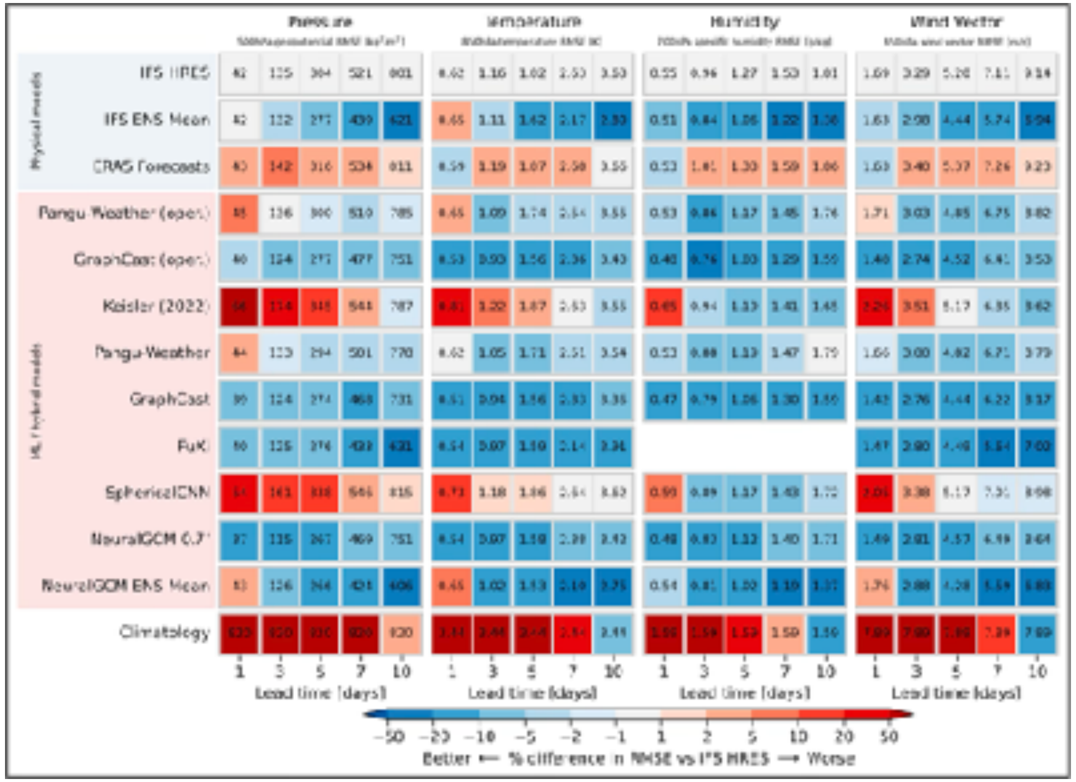

https://sites.research.google/weatherbench/

Ocean Modeling, Forecasting and Monitoring Systems

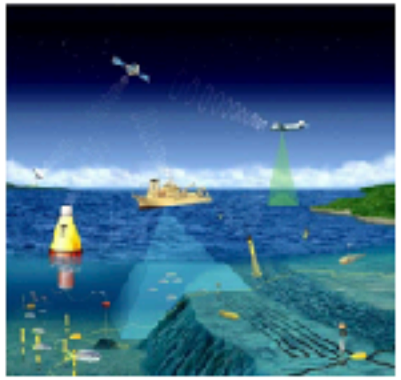

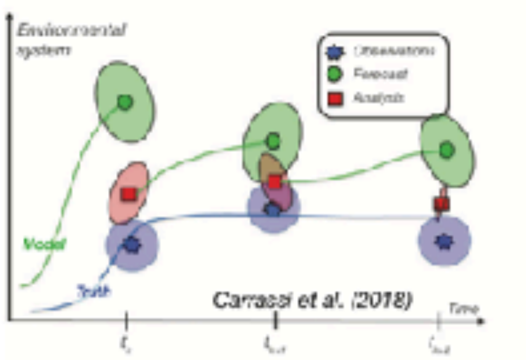

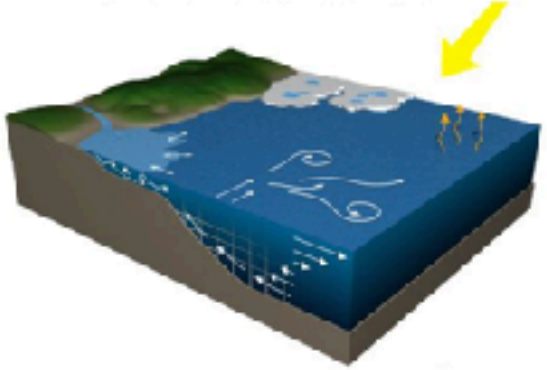

Observations → Data assimilation → Physics-based forecast → Post-processing, dissemination

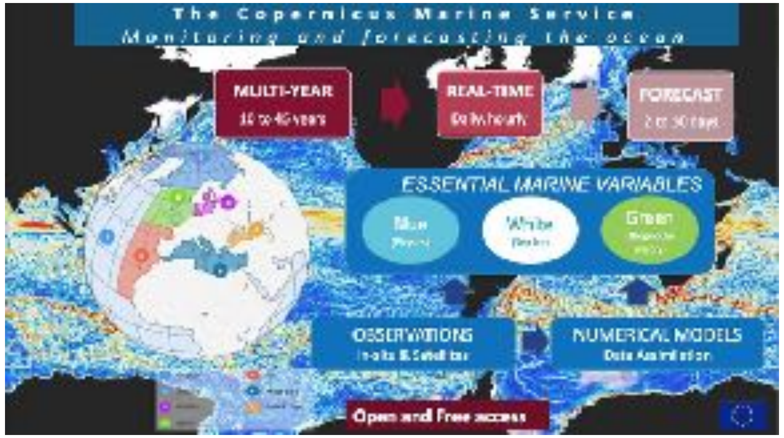

Operational products

Ocean Modeling, Forecasting and Monitoring Systems

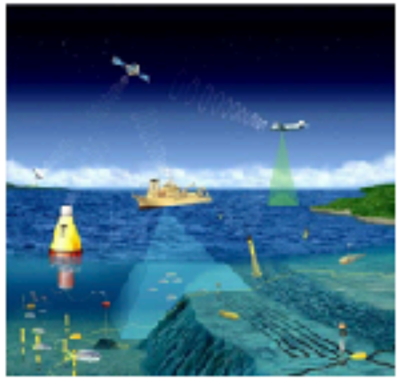

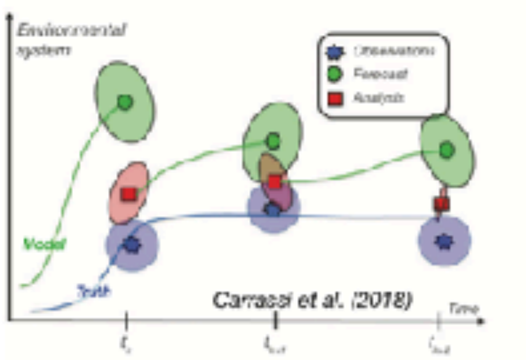

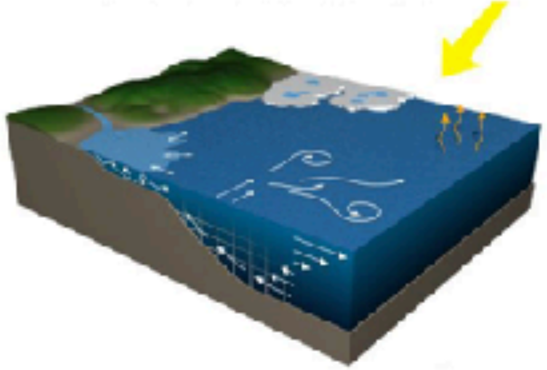

Observations → Data assimilation → Physics-based forecast → Post-processing, dissemination

---

Hardware and Software Ecosystem

+

Ocean Modeling, Forecasting and Monitoring Systems

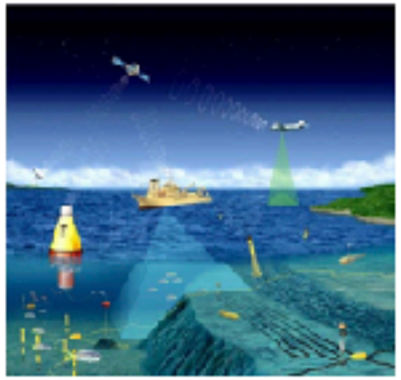

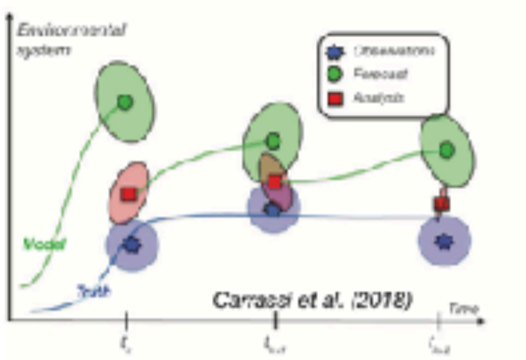

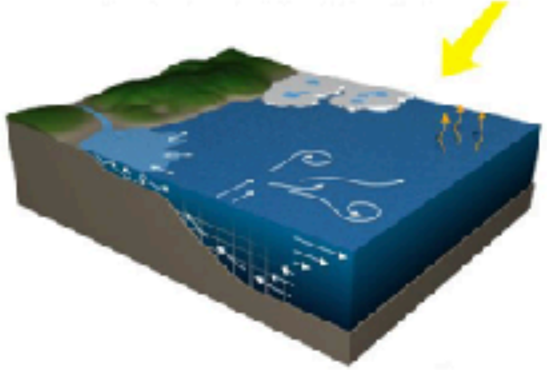

Observations → Data assimilation → Physics-based forecast → Post-processing, dissemination

Hardware and Software Ecosystem

F + 

Challenges

Computational Complexity

Under exploitation of observation datasets

Uncertainty Quantification

Challenges? (1) Computational complexity

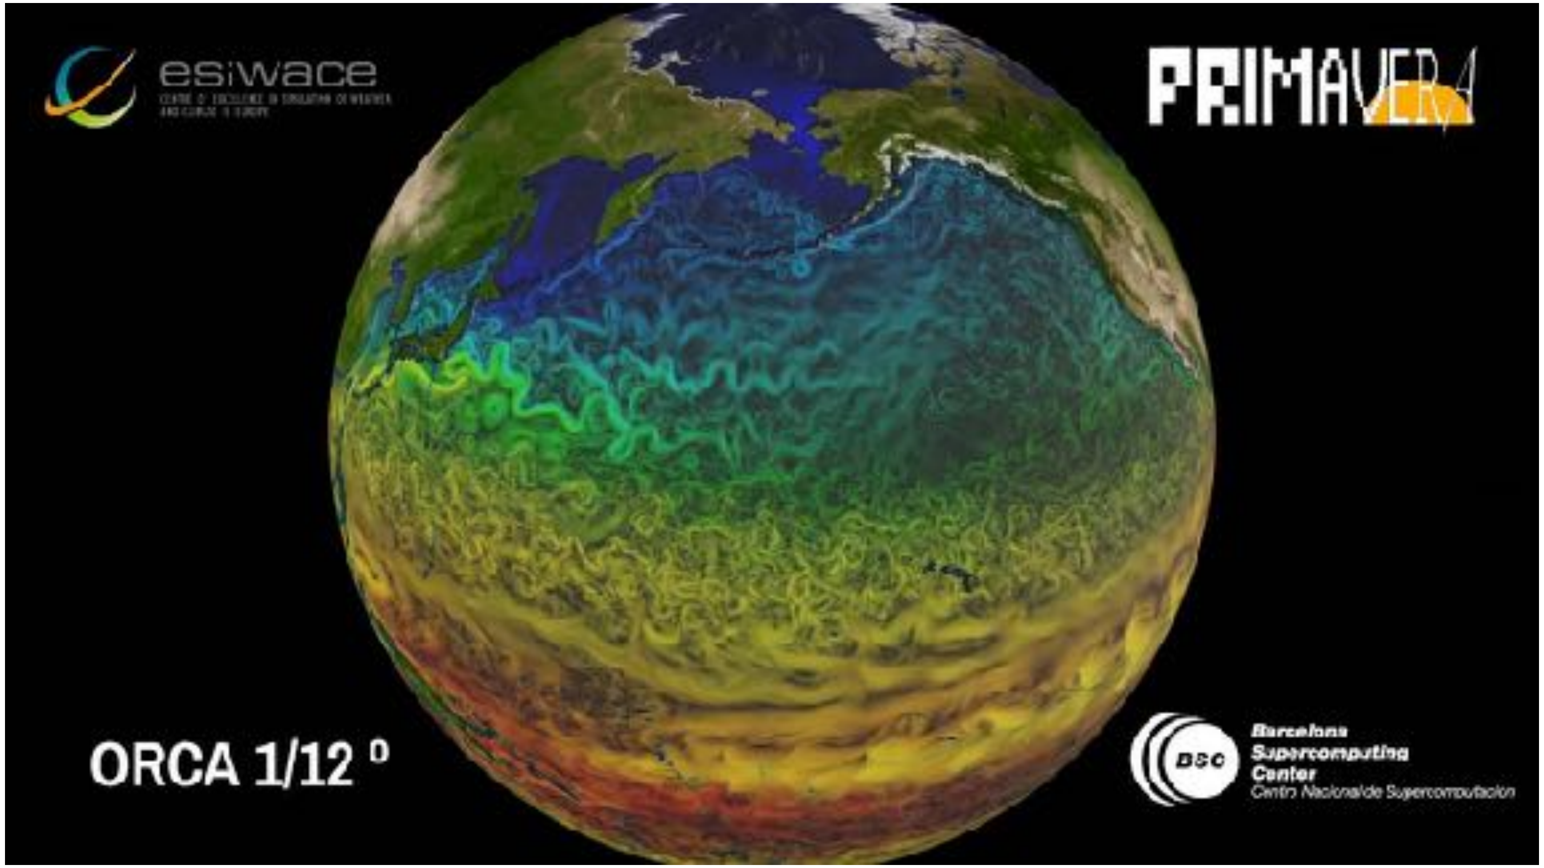

1/4° ~ 10kh CPU
1/60° ~ 40Mh CPU

Consequences:

- Strong constraints on representativeness vs. complexity
- Hardware and software solutions based HPC +

Challenges? (2) How to make the most of observations?

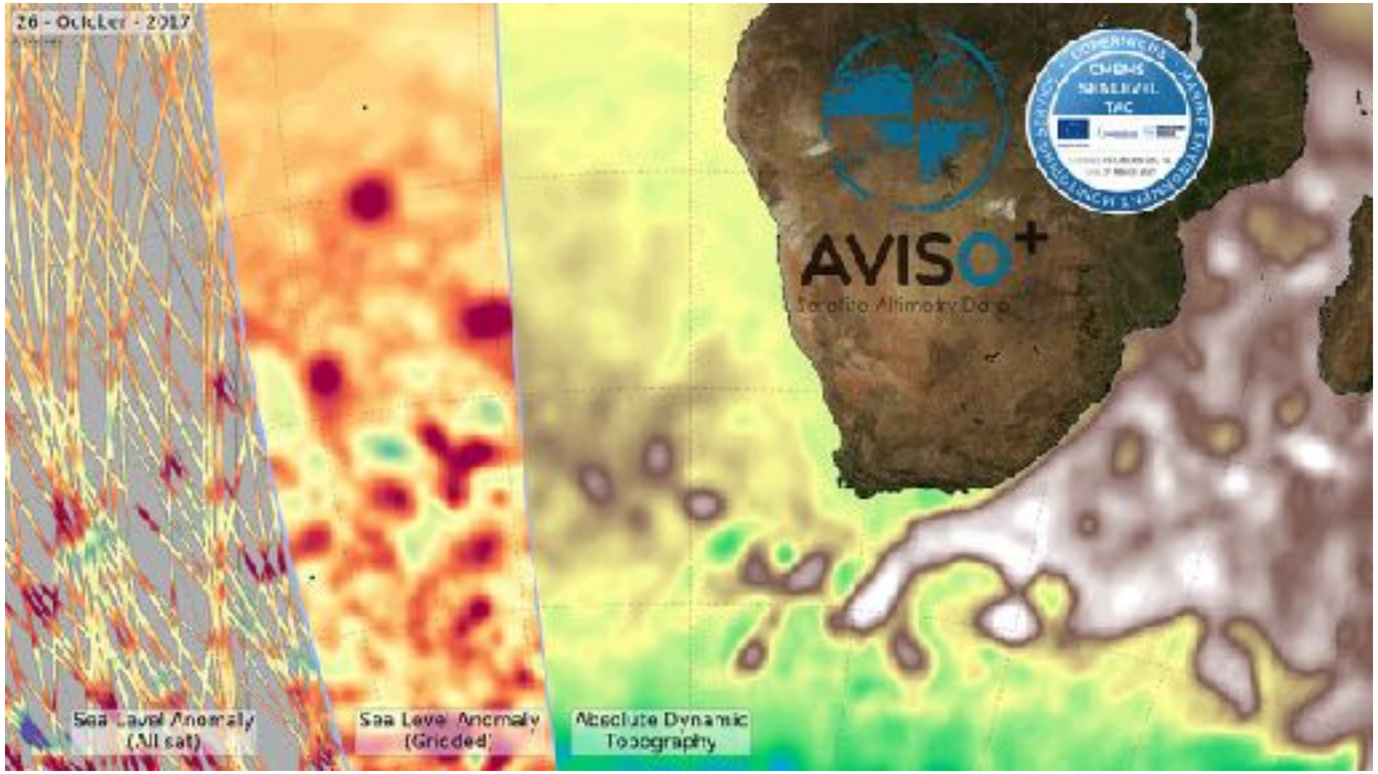

Altimetry data (Sea Surface Height/Sea Surface Anomaly)

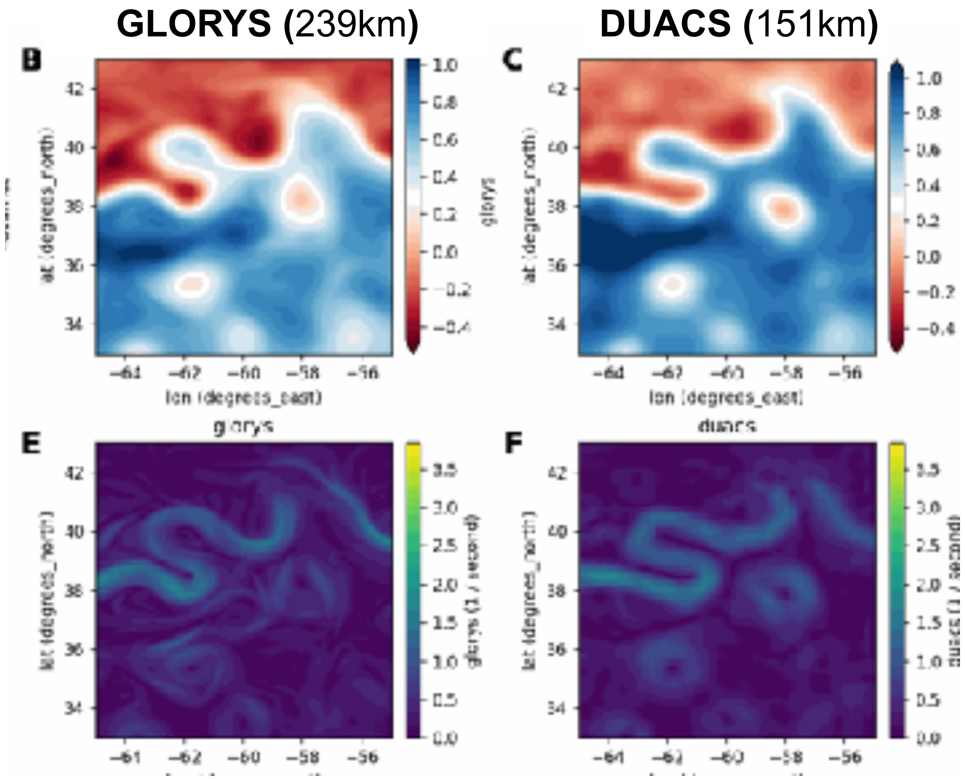

Operational altimeter-derived products

Challenges? (2) How to make the most of observations?

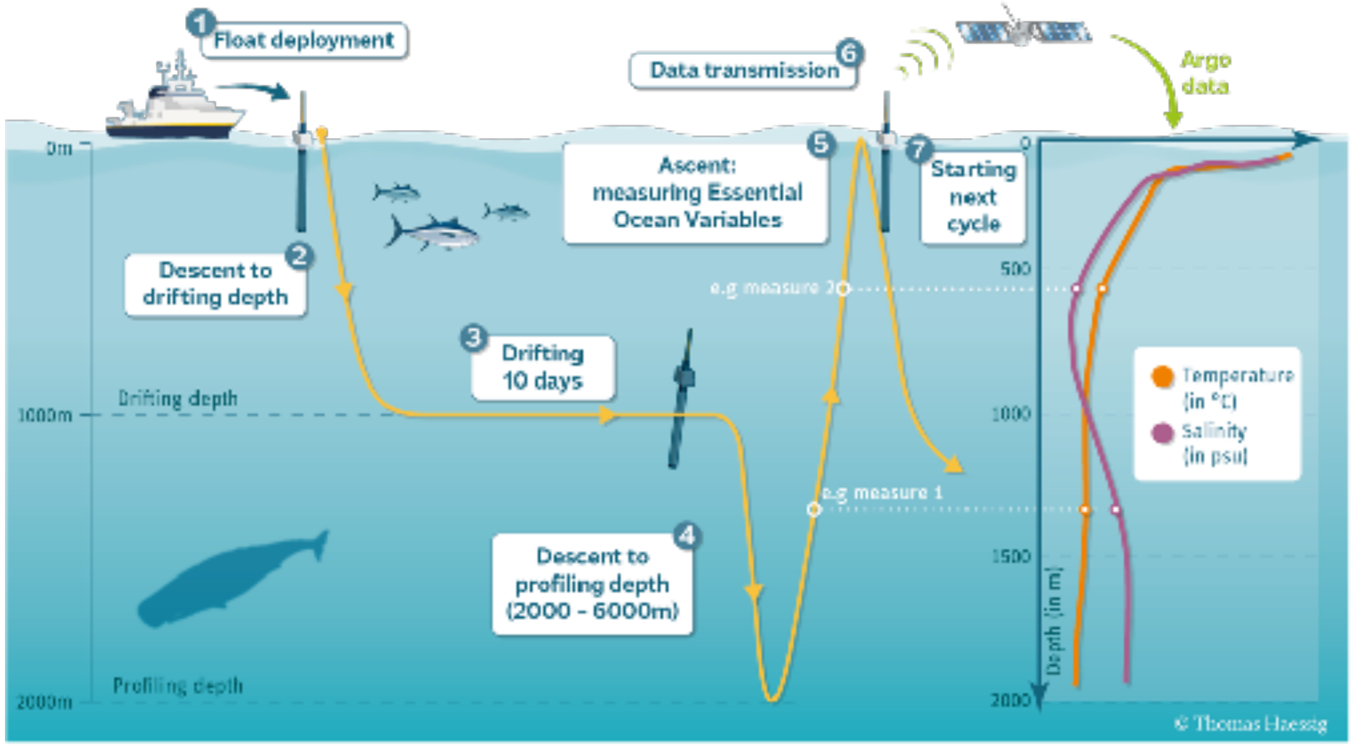

Deployment of ARGO floats

ARGO data for one day (April, 16, 2024**)**

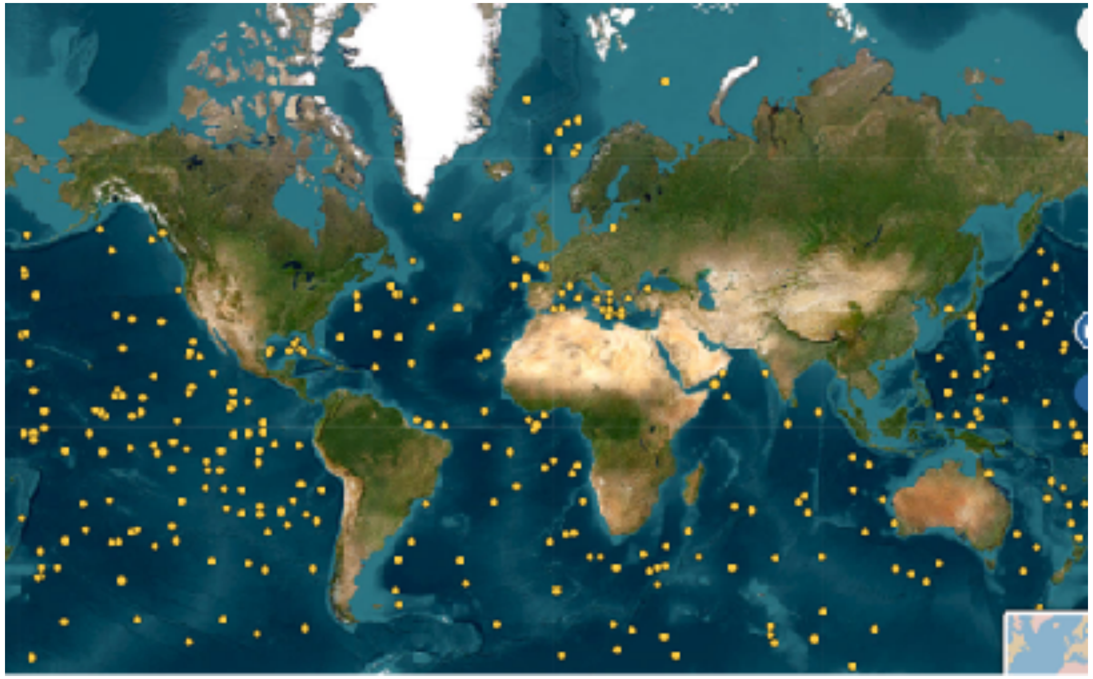

Challenges? (3) Uncertainty Quantification

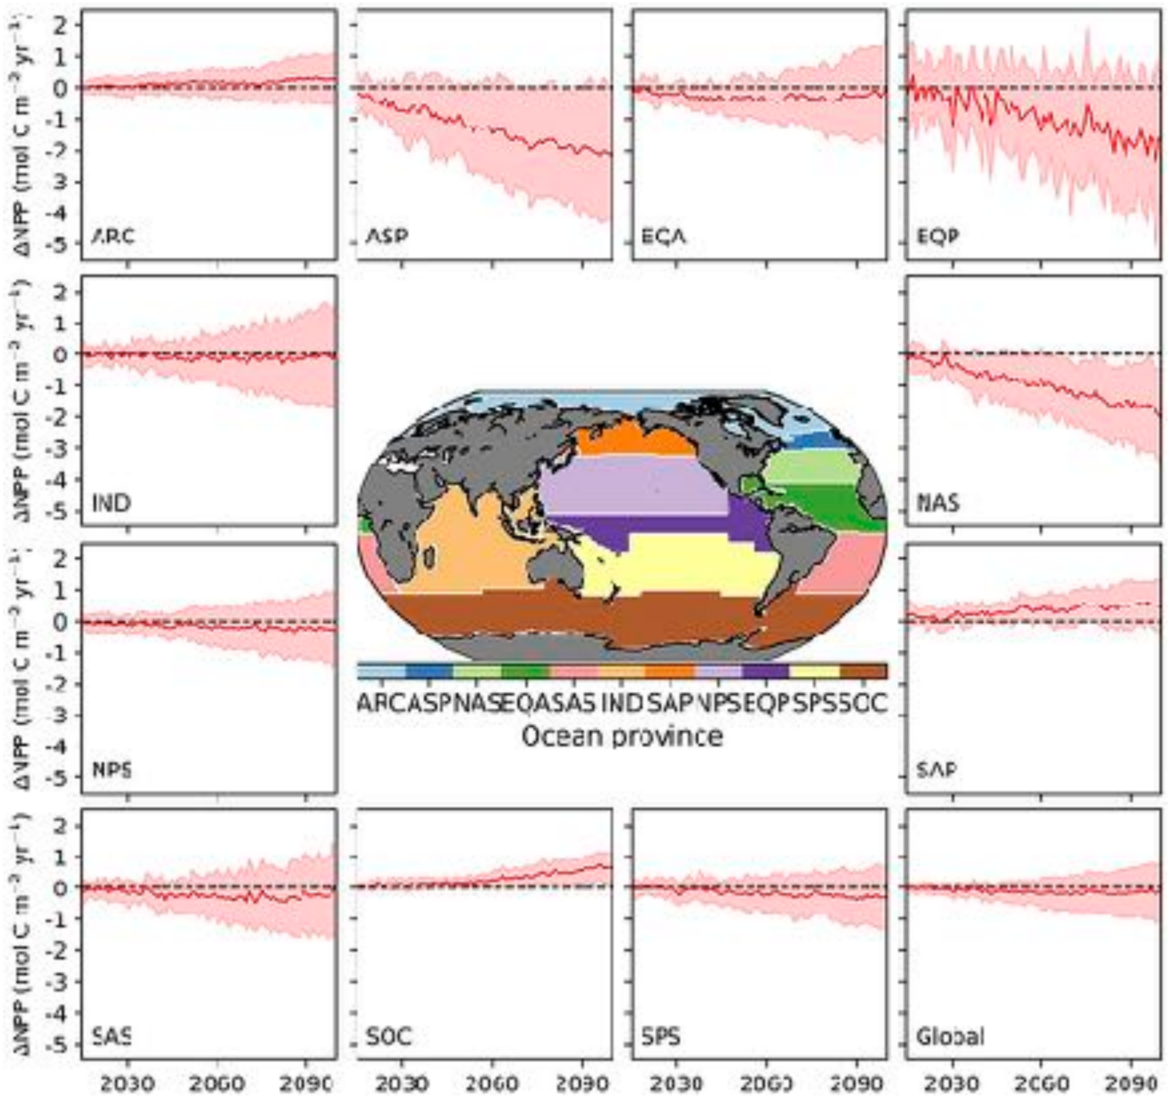

From Tagliabue et al., 2021, https://doi.org/10.3389/fclim.2021.738224

Illustration of uncertainties in the simulation of the primary production of the ocean (CMIP6 simulation ensembles)

Challenges and shortcomings : model calibration, parametrisation of unresolved uncertainties on forcings, ensembles' size ….

How can AI advance operational oceanography?

AI to complement operational systems (new tasks)

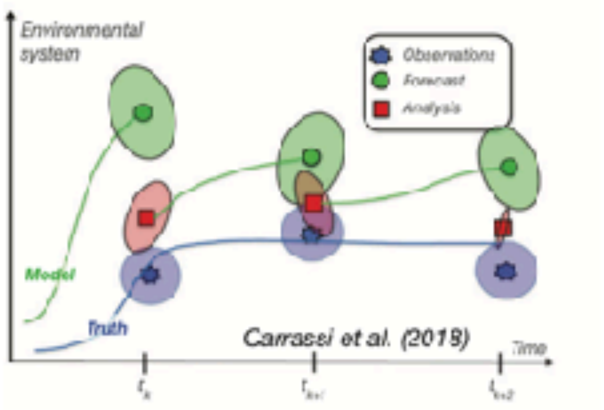

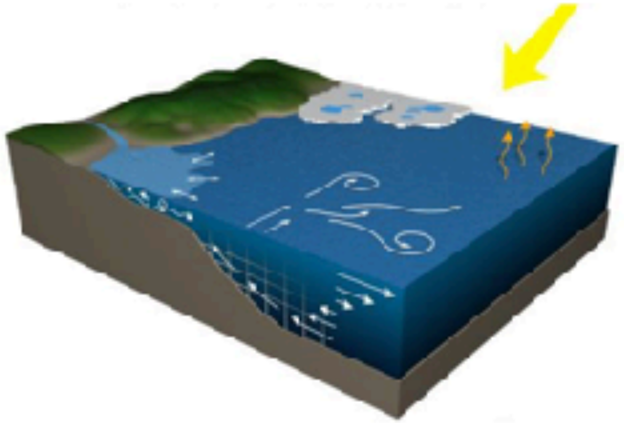

Observations → Data assimilation → Physics-based forecast → Post-processing, dissemination

AI

AI in Operational Oceanography Systems

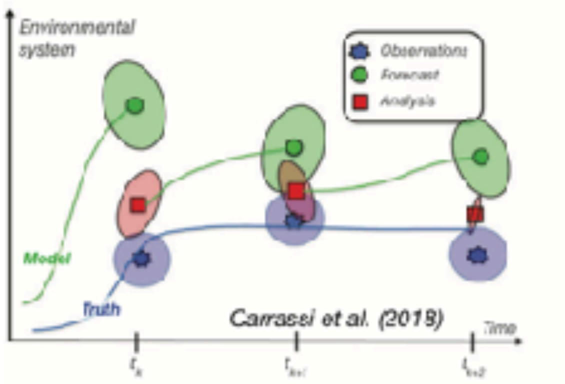

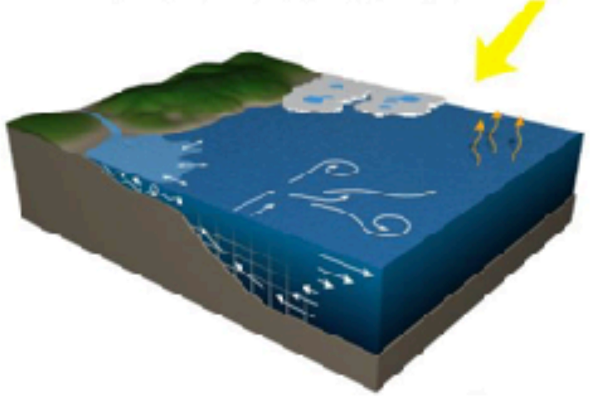

Observations → **Data assimilation** (AI) → **Physics-based forecast** (AI) → **Post-processing, dissemination**

Towards hybrid ocean models

(End-to-end) learning of reduced-order models

High-resolution
vs.
Reduce-order
(Low-resolution)

$$\begin{cases} \dfrac{\partial X_t}{\partial t} = f\left(X_t\right) + g\left(Z_t\right) \\[2ex] \dfrac{\partial \overline{X}_t}{\partial t} = f\left(\overline{X}_t\right) + g\left(Z_t\right) + \boxed{h\left(\overline{X}_t\right)} \end{cases}$$

Known terms

Unknown term to be represented by a neural model (Closure)

SQG case-study: example of simulated vorticity state using different closures

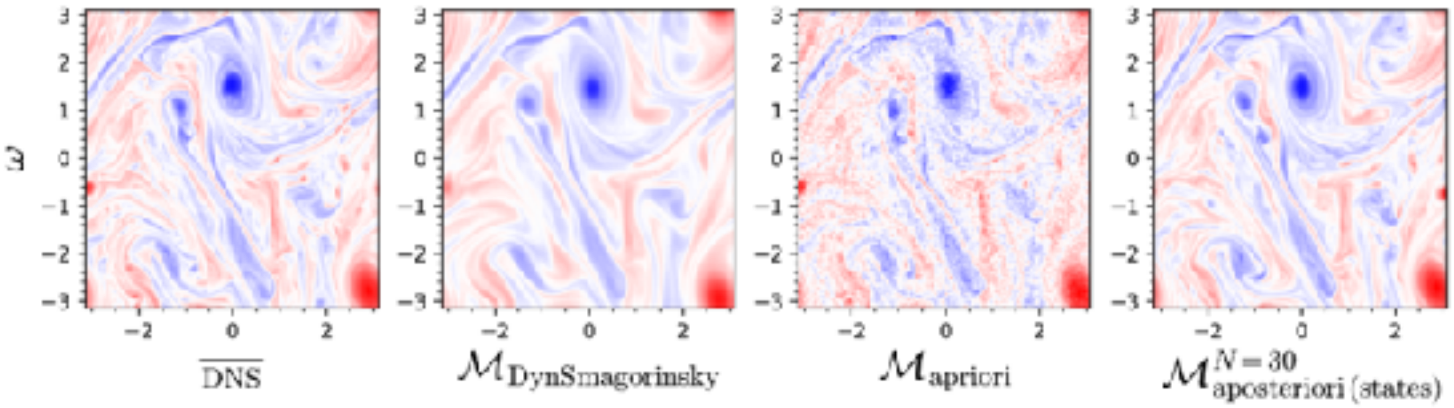

Frezat et al., 2022

Towards hybrid ocean models

Neural Networks

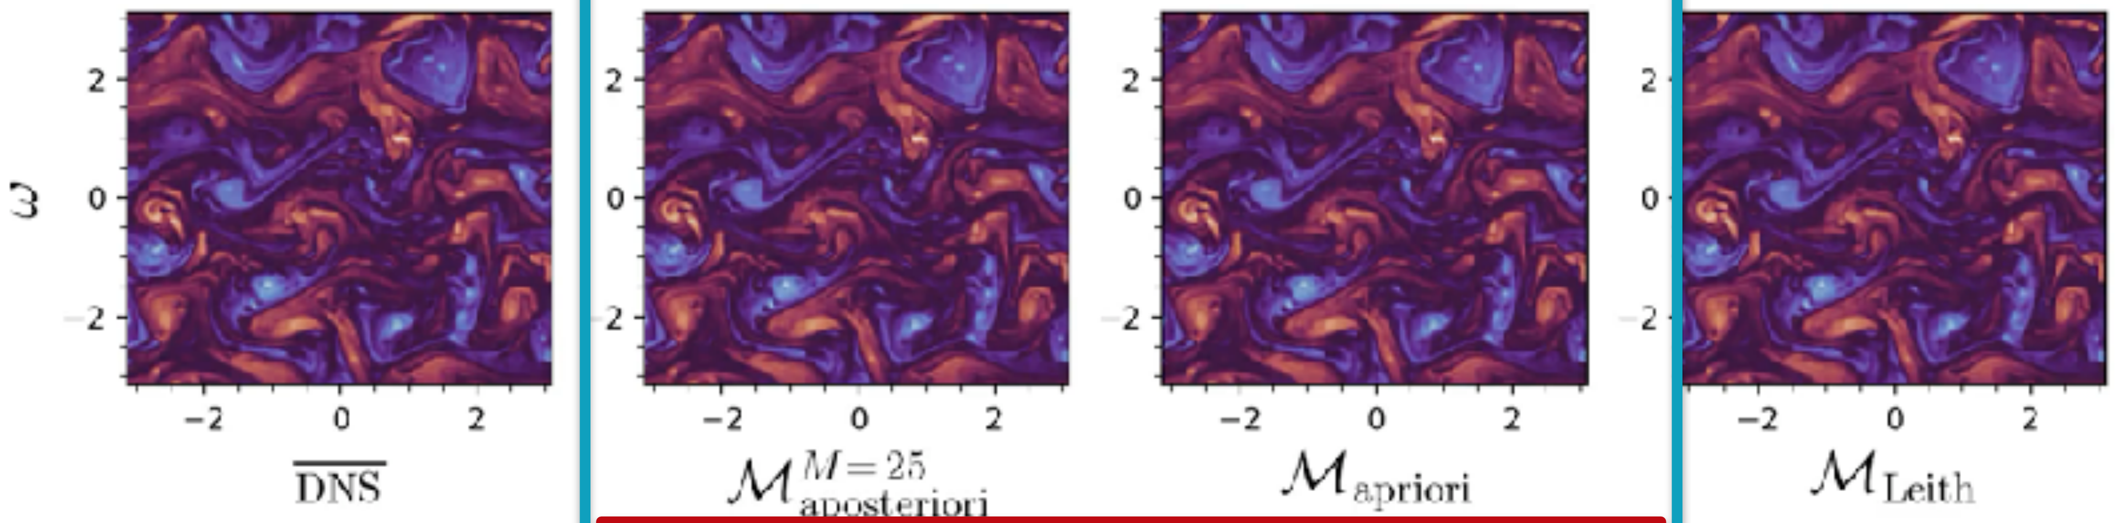

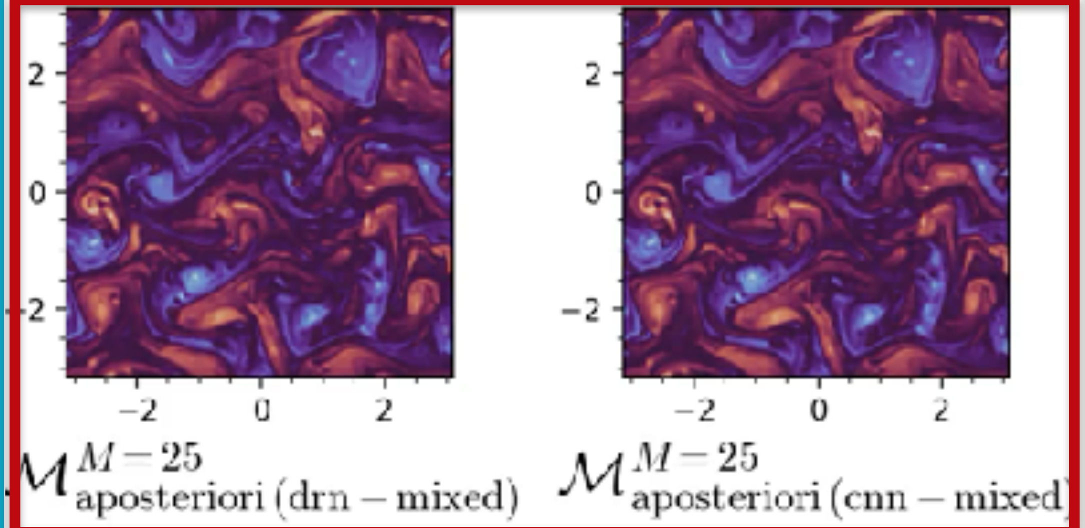

A posteriori learning can improve long-term stability

Emulators to deal with non-differentiable LES models.

Frezat et al., 2022. Ouala et al., 2024

AI in Operational Oceanography Systems

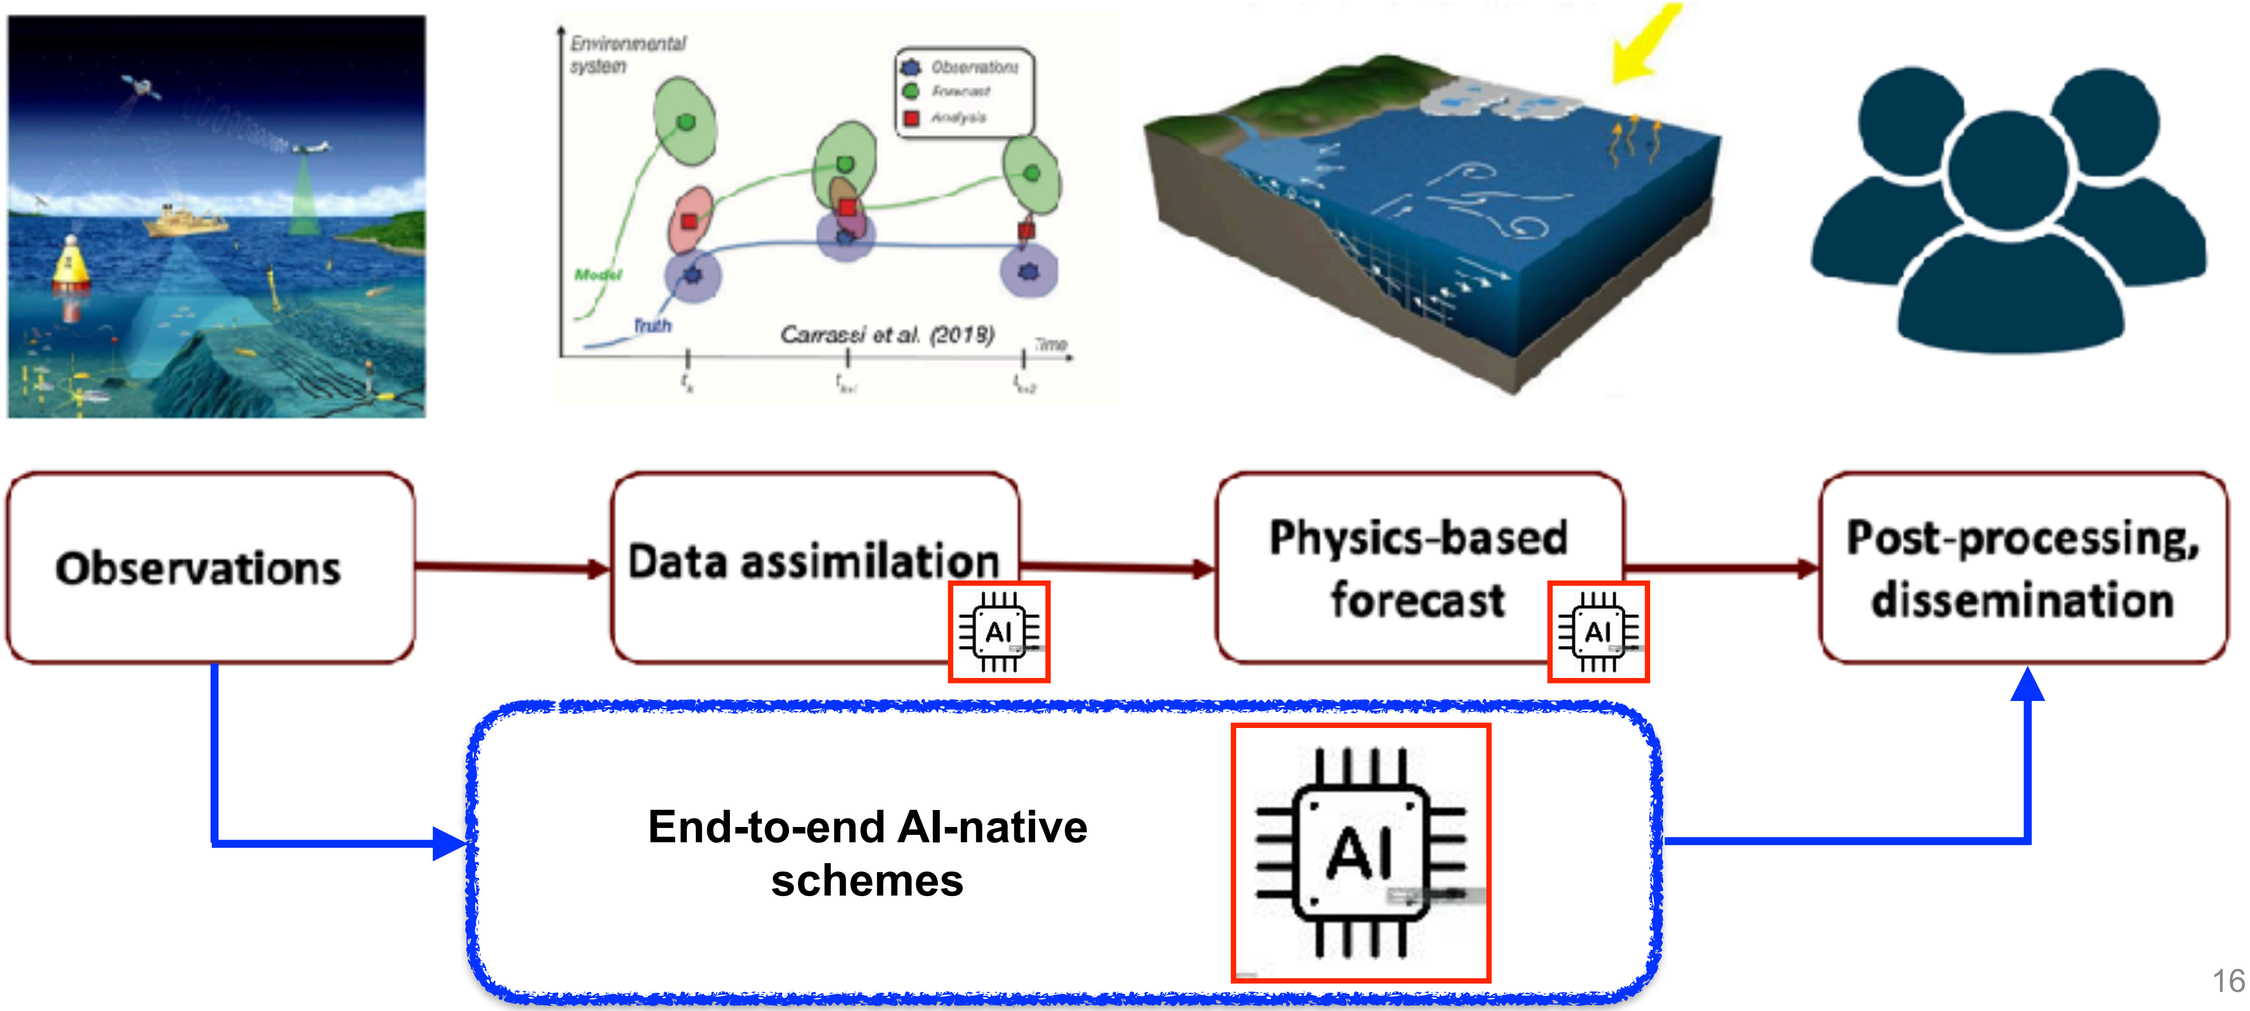

AI to replace existing ocean models (here, MEMO)

Short-term neural ocean forecast (eg, GLONET)

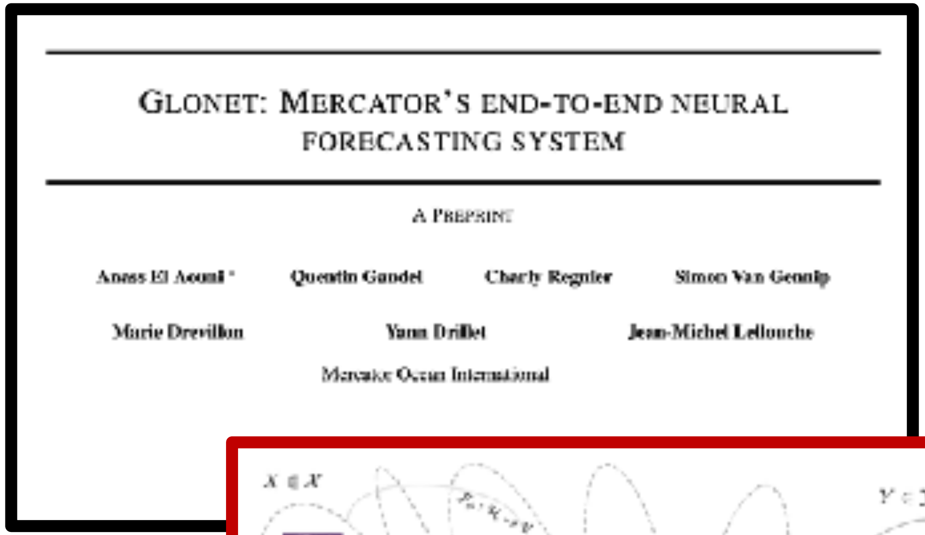
GLONET: MERCATOR'S END-TO-END NEURAL FORECASTING SYSTEM

A PREPRINT

Anass El Aouni * Quentin Gaudel Charly Regnier Simon Van Gennip

Marie Drevillon Yann Drillet Jean-Michel Lellouche

Mercator Ocean International

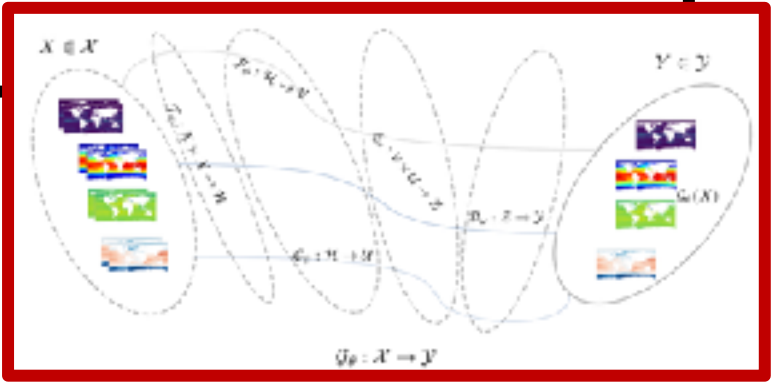

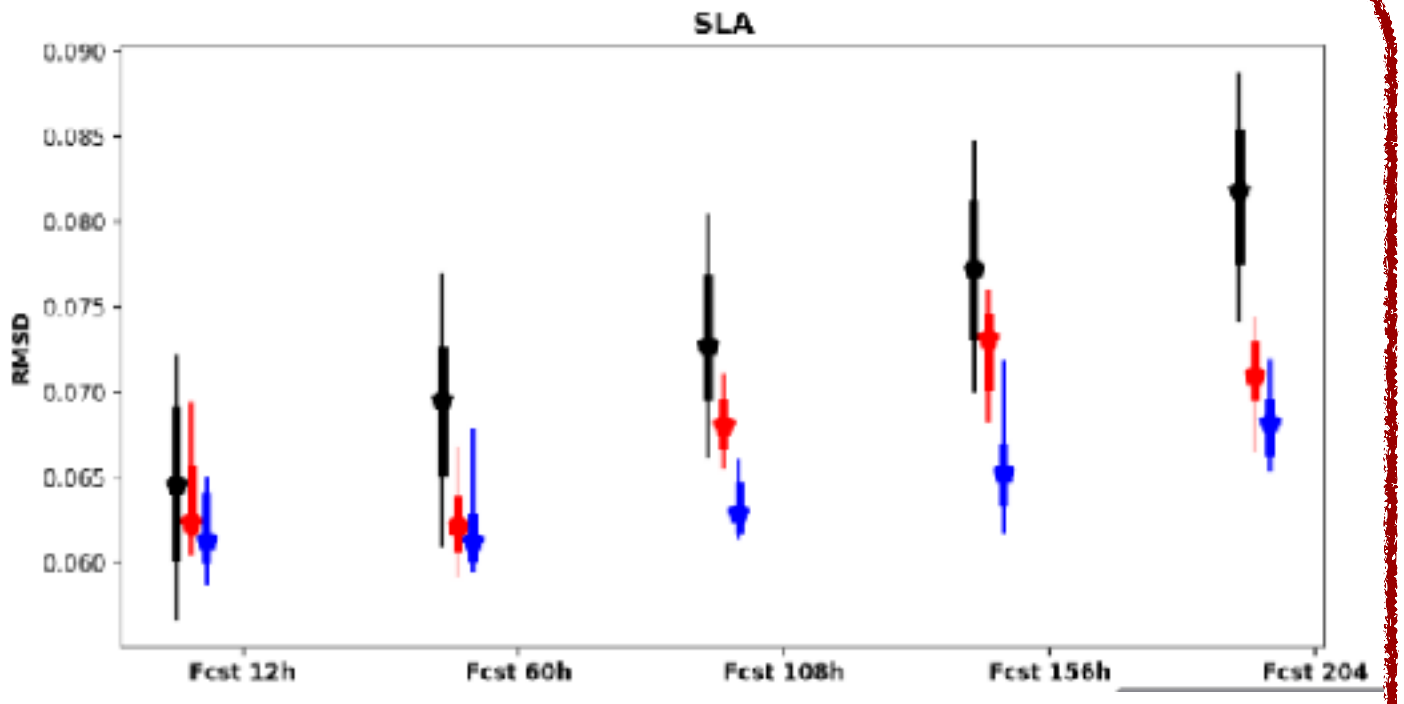

https://glonet.lab.dive.edito.eu/

AI to replace ocean Data Assimilation Ssytem

Application of end-to-end neural DA to SSH mapping (altimetry data)

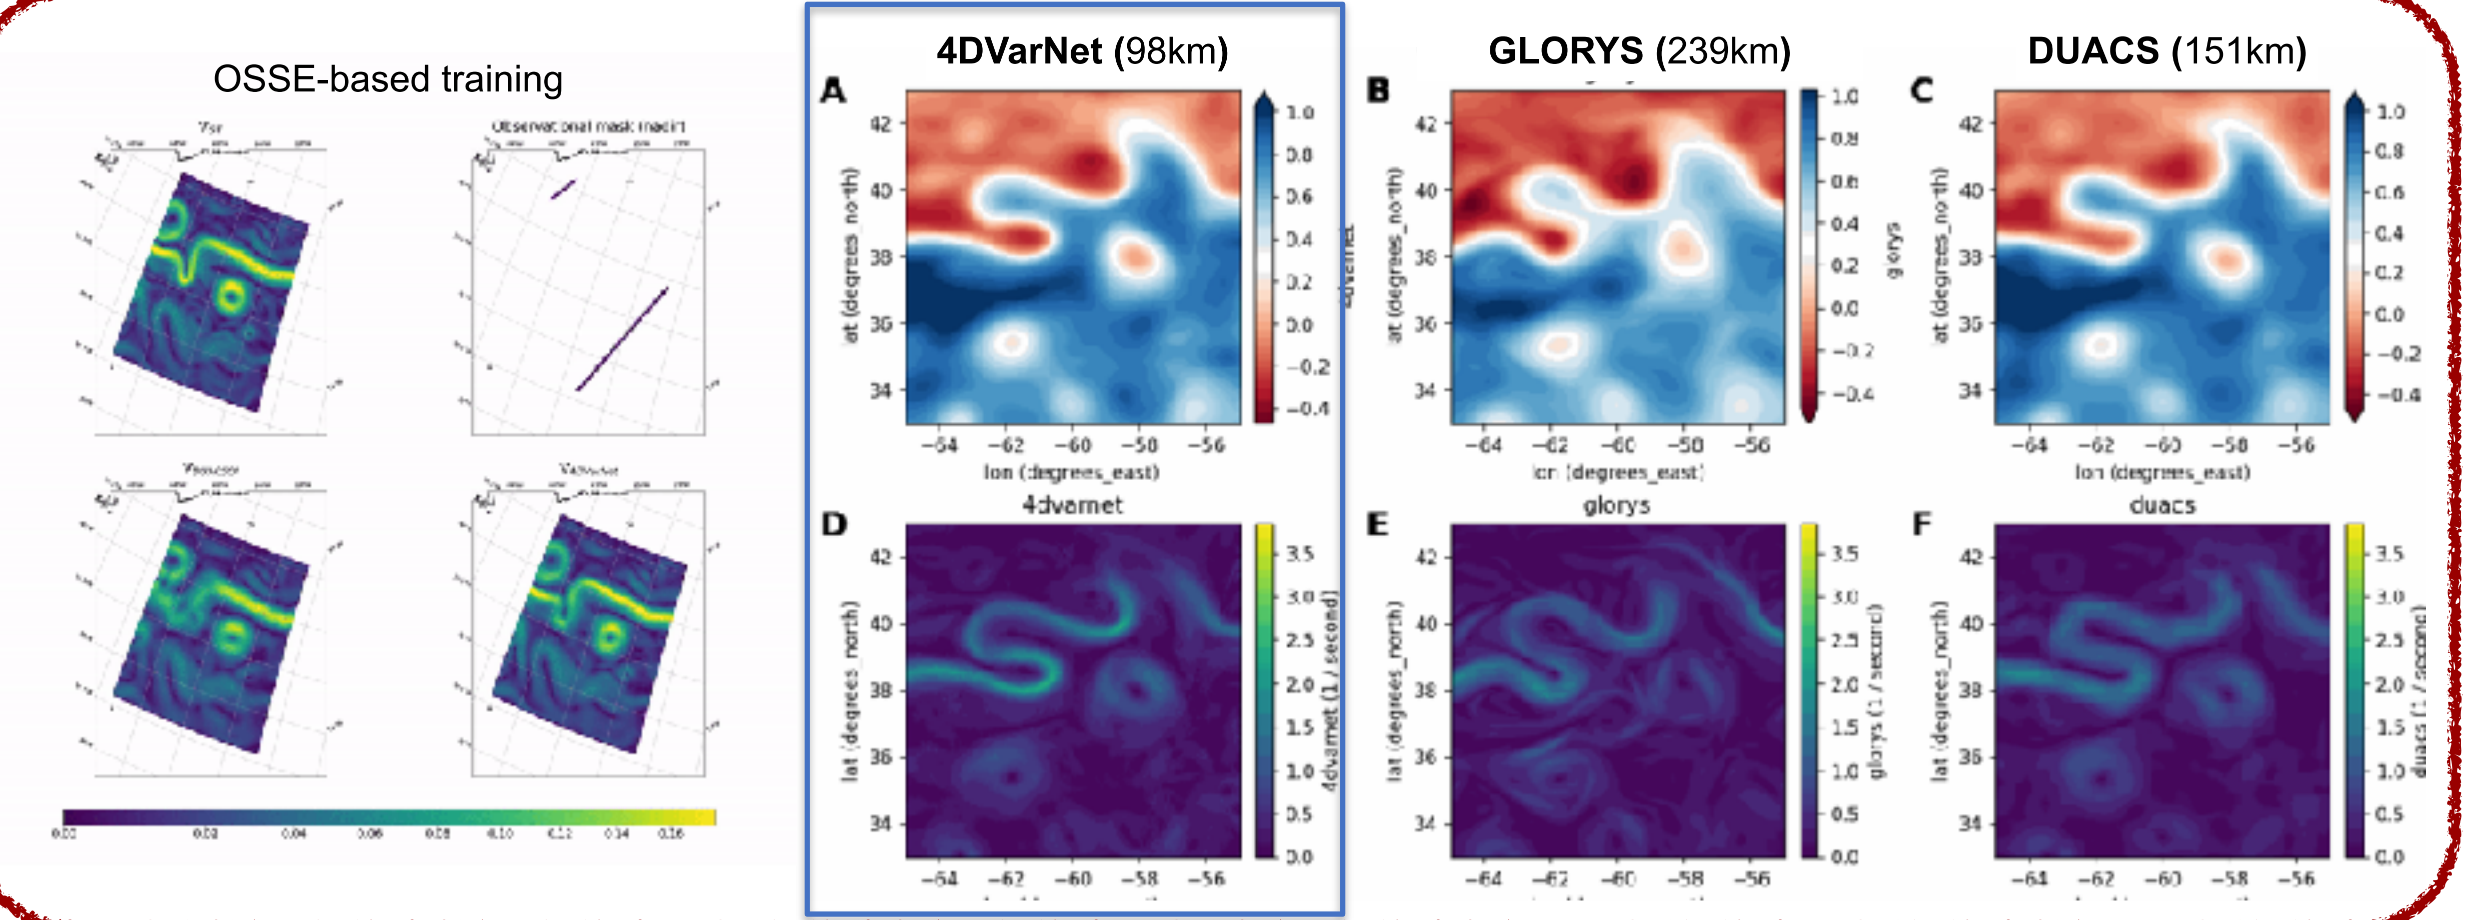

Febvre et al., 2024

Lessons learnt so far

- ***Growing number of (operational) demonstrations of neural schemes in operational oceanography***
- ***State-of-the-art performance of neural schemes for "at-scale" demonstrations (**eg, global ocean**)***
- ***Training schemes from simulation, reanalysis and/or observation datasets***
- ***Key to develop collaborative domain-relevant datasets and benchmarks [**OceanBench, arxiv:2309.15599**]***
- ***Towards ocean foundation models***

Data Assimilation

How to reconstruct processes ?

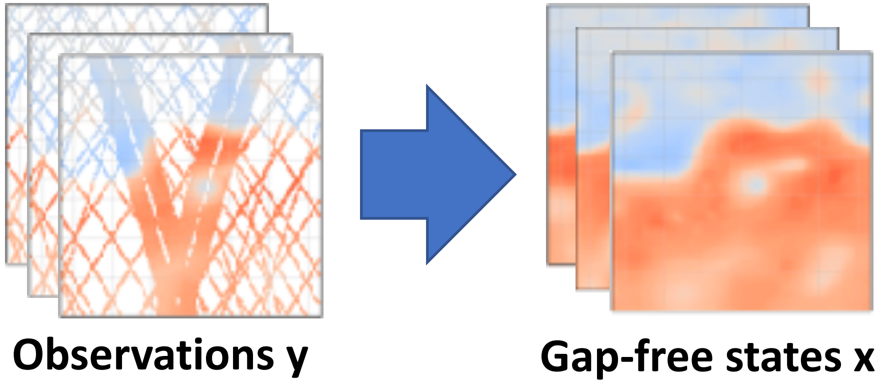

How to simulate/forecast?

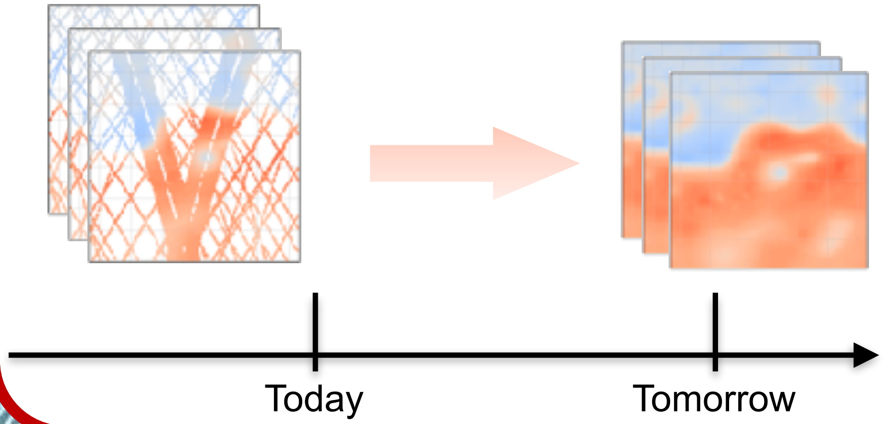

Data assimilation (for a ML scientist)

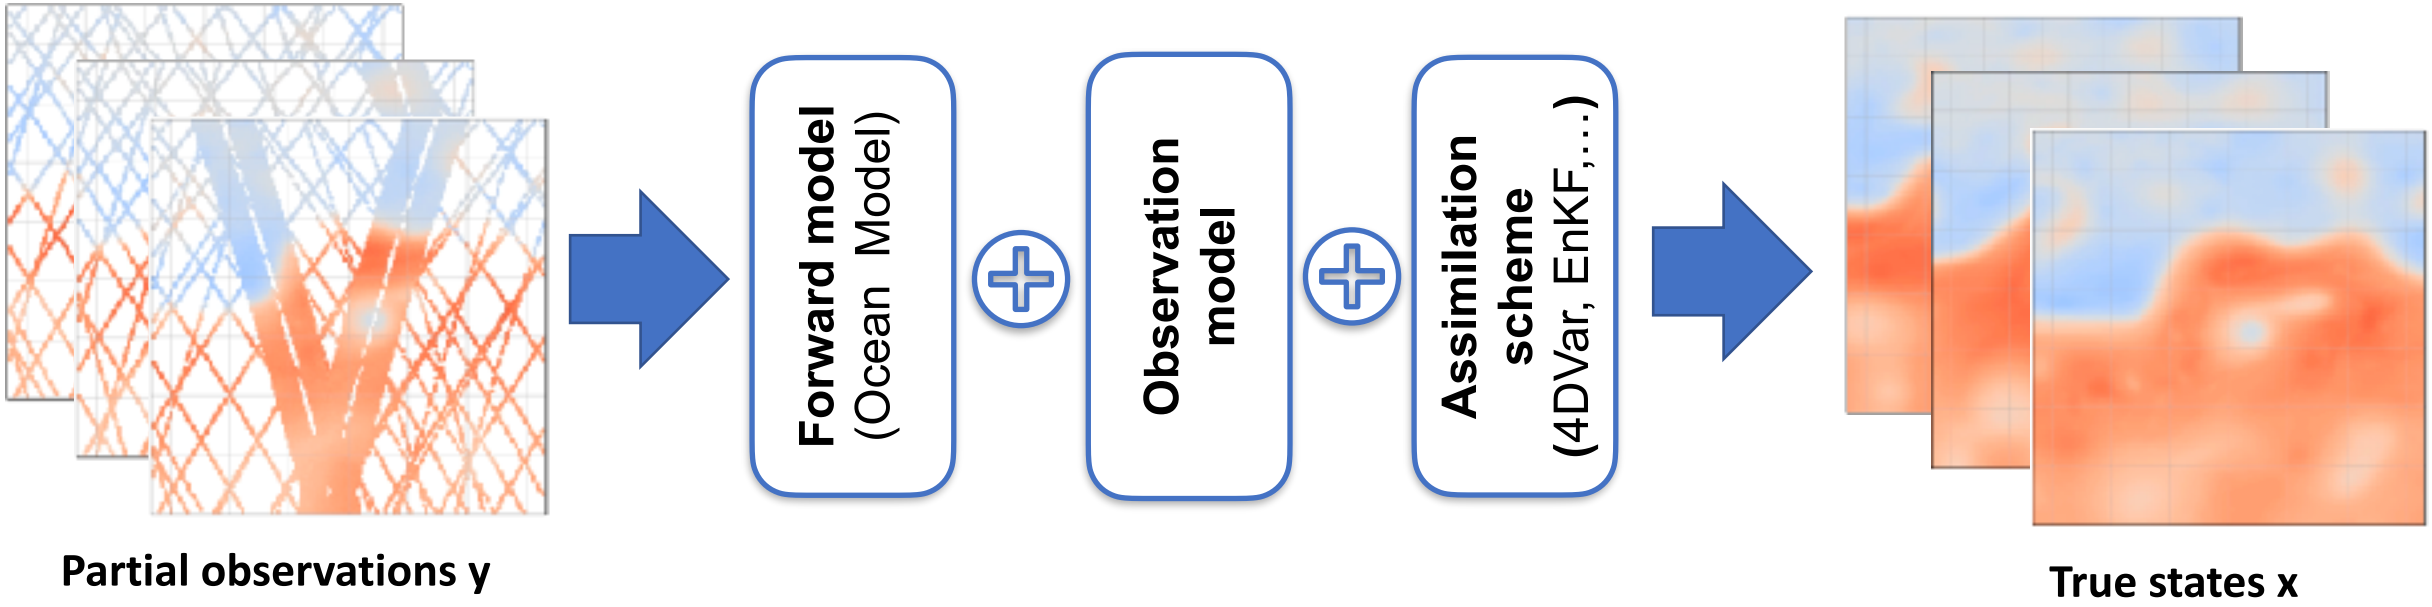

DA scheme regarded as the composition of elementary components based on model-driven principles

(Weak constraint) **4DVar Data Assimilation (DA) formulation**

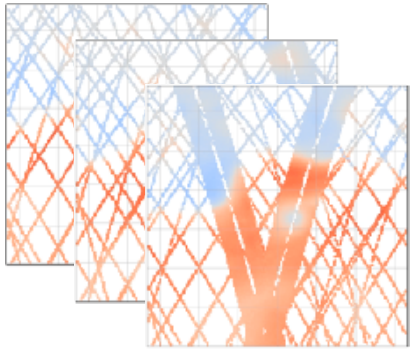

Partial observations y

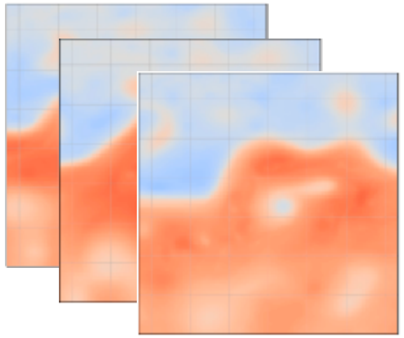

True states x

State-space formulation:

$$\left\{\begin{array}{lcl} \dfrac{\partial x(t)}{\partial t} & = & \mathcal{M}\left(x(t)\right) \\ \\ y(t) & = & x(t) + \epsilon(t), \forall t \in \{t_0, t_0 + \Delta t ...., , t_0 + N\Delta t\} \end{array}\right.$$

Associated variational formulation:

$$\arg\min_x \lambda_1 \sum_i \|x(t_i) - y(t_i)\|^2_{\Omega_{t_i}} + \lambda_2 \sum_n \|x(t_i) - \Phi(x)(t_i)\|^2$$

with $\Phi(x)(t) = x(t-\Delta) + \int_{t-\Delta}^{t} \mathcal{M}\left(x(u)\right) du$

$$\boxed{\arg\min_x \lambda_1 \|x - y\|^2_\Omega + \lambda_2 \|x - \Phi(x)\|^2}$$

(Weak constraint) **4DVar Data Assimilation (DA) formulation**

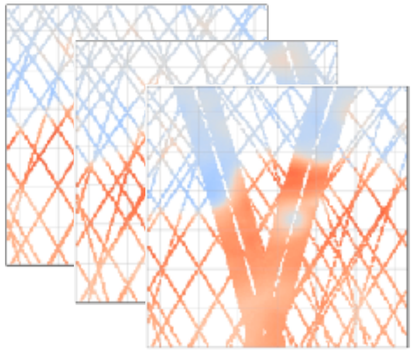

Partial observations y

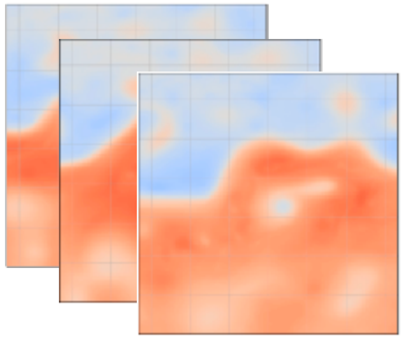

True states x

$$\arg\min_{x} \|y - \mathbf{H}x\|^2 + \lambda\|x - \Phi(x)\|^2$$

$$x^{(k+1)} = x^{(k)} - \alpha \nabla_x \mathcal{J}_{\mathbf{H},\Phi}\left(x^{(k)}, y\right)$$

Challenges to solve the 4DVar minimisation

- High-dimensional state ($>10^{11}$)
- Availability of the gradient of the 4DVar cost
- Non-convex optimisation / local minima
- Computational complexity

Solutions in the DA literature:

- Sequential optimisation
- Gradient-free approaches (e.g., EnKF)
- Reduced-rank approaches

The example of Optimal Interpolation/Kalman Filter

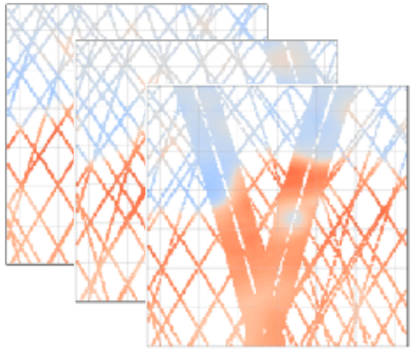

Partial observations y

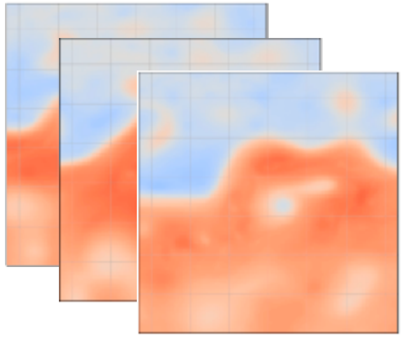

True states x

State-space formulation:
$$\begin{cases} x \sim \mathcal{N}(\mu, \mathbf{\Sigma}) \\ y = \mathbf{H}x + \epsilon \end{cases}$$

Resulting variational formulation:

$$\widehat{x} = \arg\min_x 1/\sigma^2 \|y - \mathbf{H}x\|^2 + (x - \mu)^t \mathbf{\Sigma}^{-1} (x - \mu)$$

$$\Rightarrow \quad \widehat{x} = \mu + \mathbf{K}\left(y - \mathbf{H} \cdot \mu\right) \text{ with } \mathbf{K} = \mathbf{\Sigma}\mathbf{H}^t \left[\mathbf{H}\mathbf{\Sigma}\mathbf{H}^t + \sigma^2 \mathbf{I}\right]^{-1}$$

Equivalent variational formulation:

$$\widehat{x} = \arg\min_x \|y - \mathbf{H}x\|^2 + \lambda \|x - \Phi(x)\|^2$$

$$\text{subject to } \Phi(x) = [\mathbf{I} - \mathbf{\Sigma}^{-1/2}] \cdot x + \mathbf{\Sigma}^{-1/2} \cdot \mu$$

Application in oceanography (satellite altimetry)

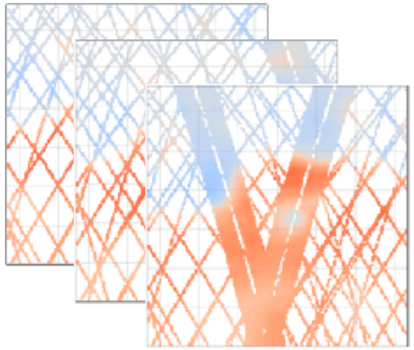

Partial observations y

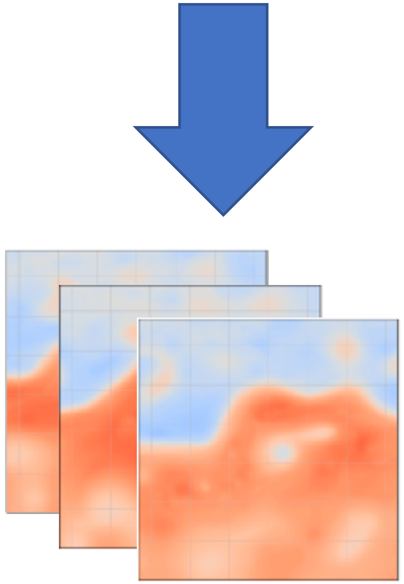

True states x

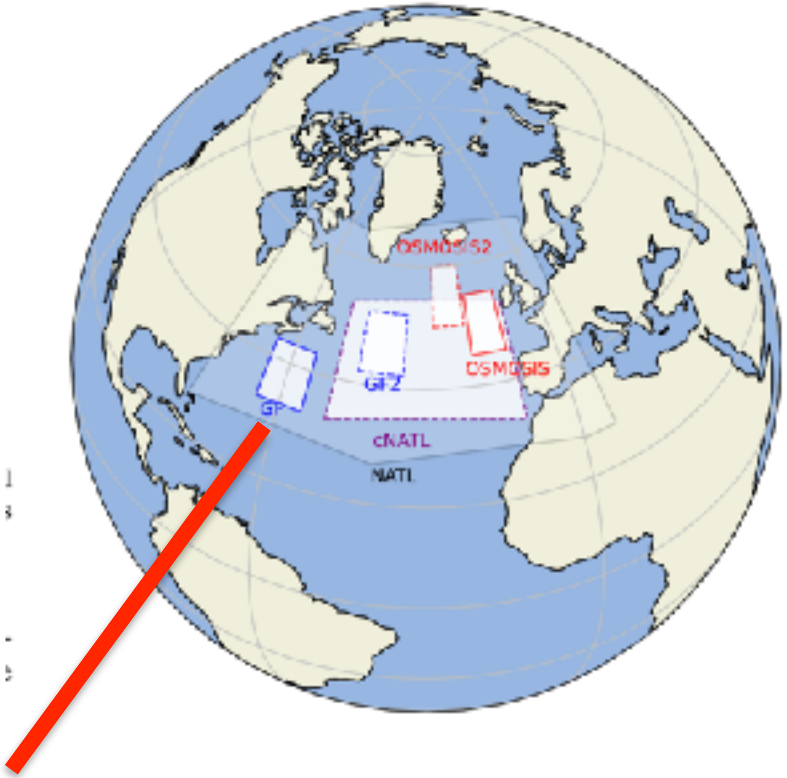

Gulf Stream region

KF (SEEK)

OI

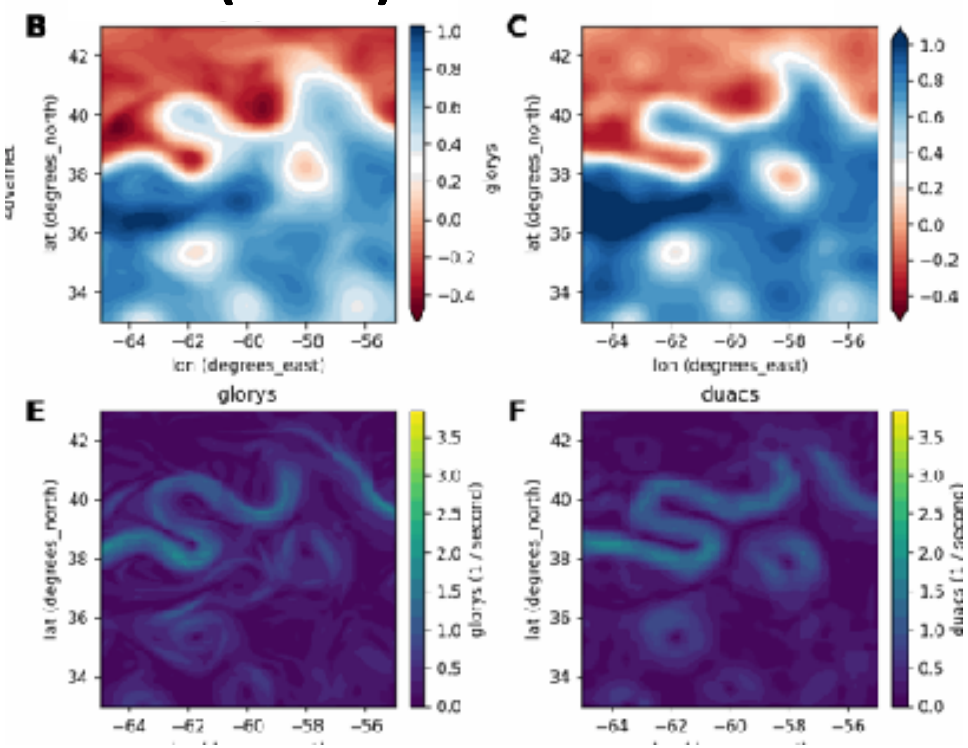

Reconstructed
Sea Surface Height

Deep learning and Data Assimilation (Inverse Problems)

Data Assimilation

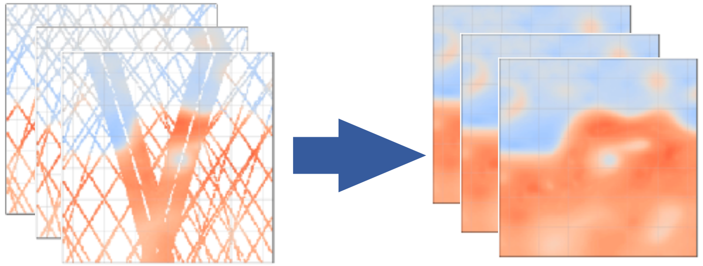

Model-driven vs. Learning-based approaches

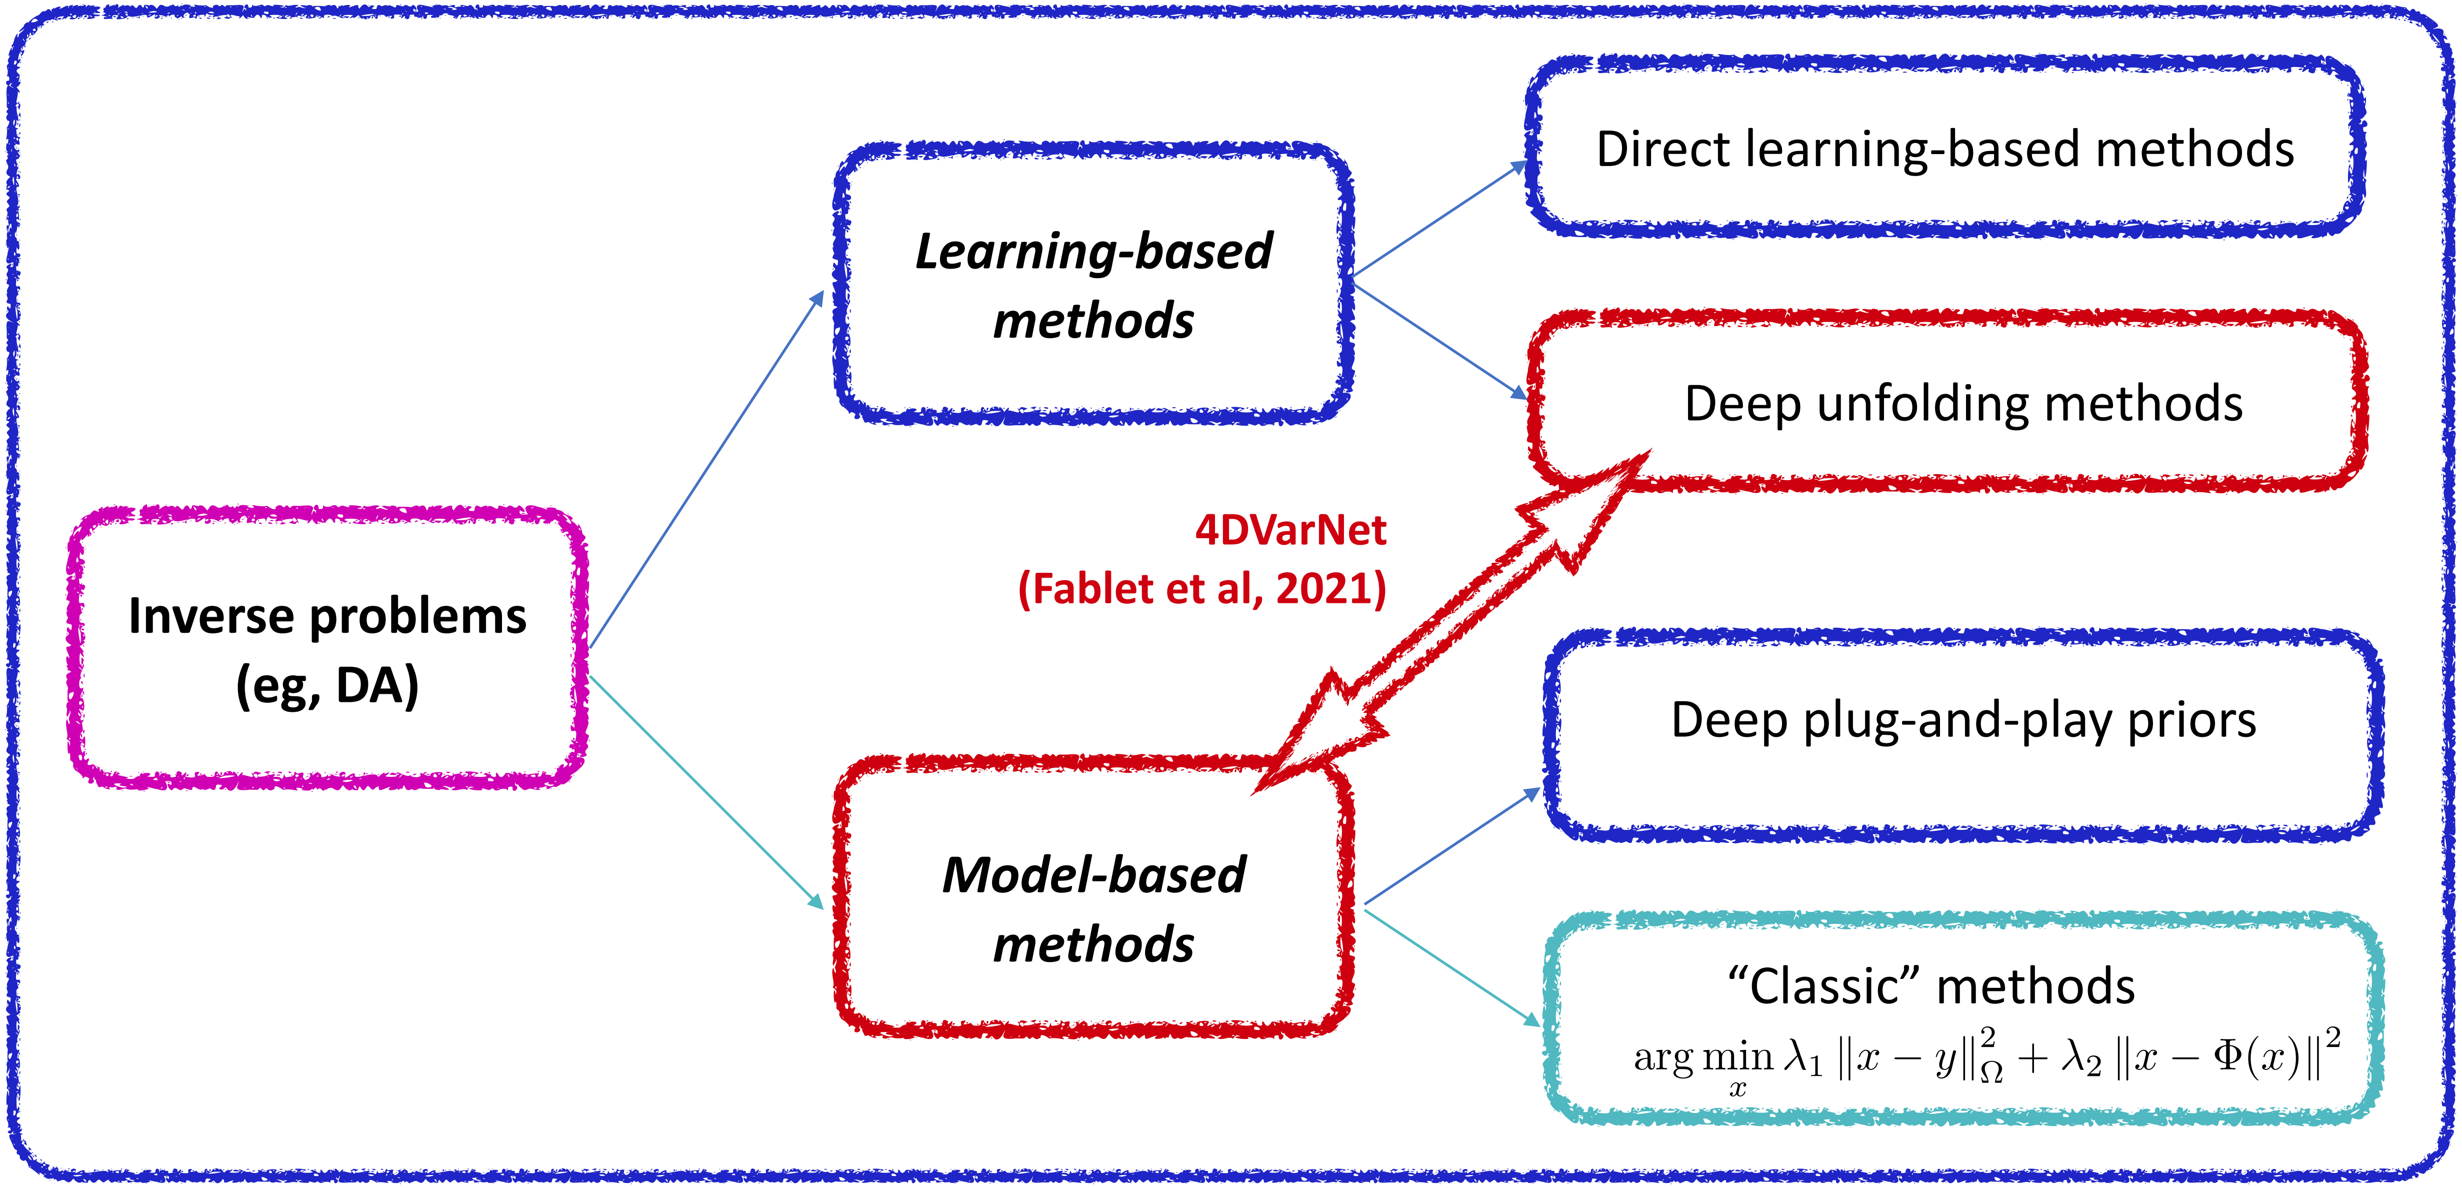

Data assimilation (for a ML scientist)

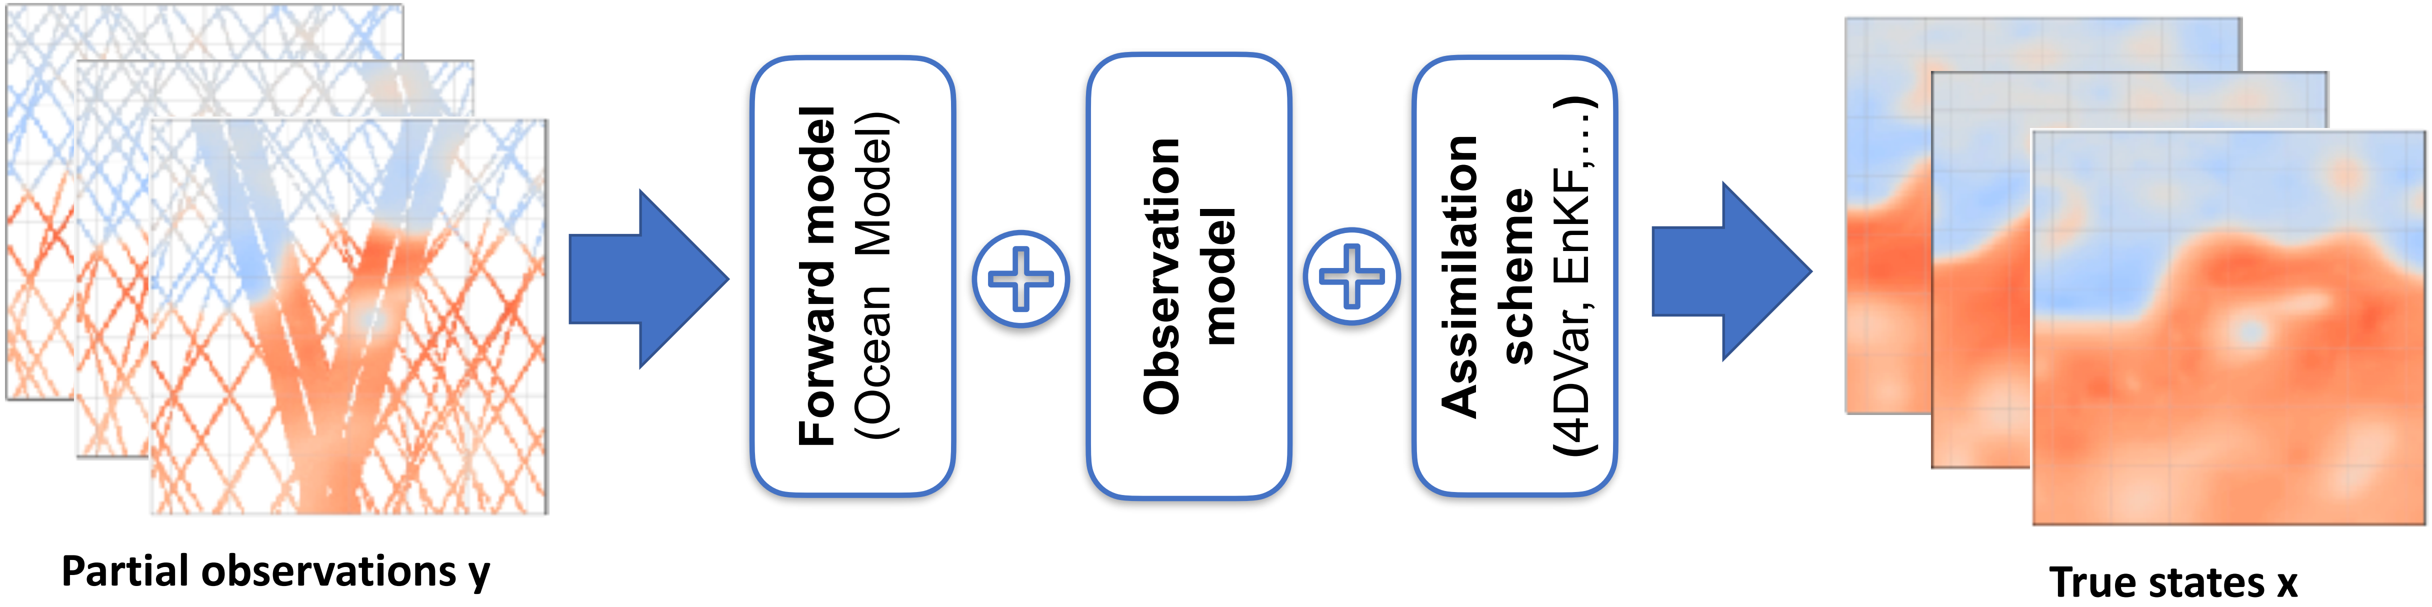

DA scheme regarded as the composition of elementary components based on model-driven principles

End-to-end deep learning schemes

End-to-end architecture

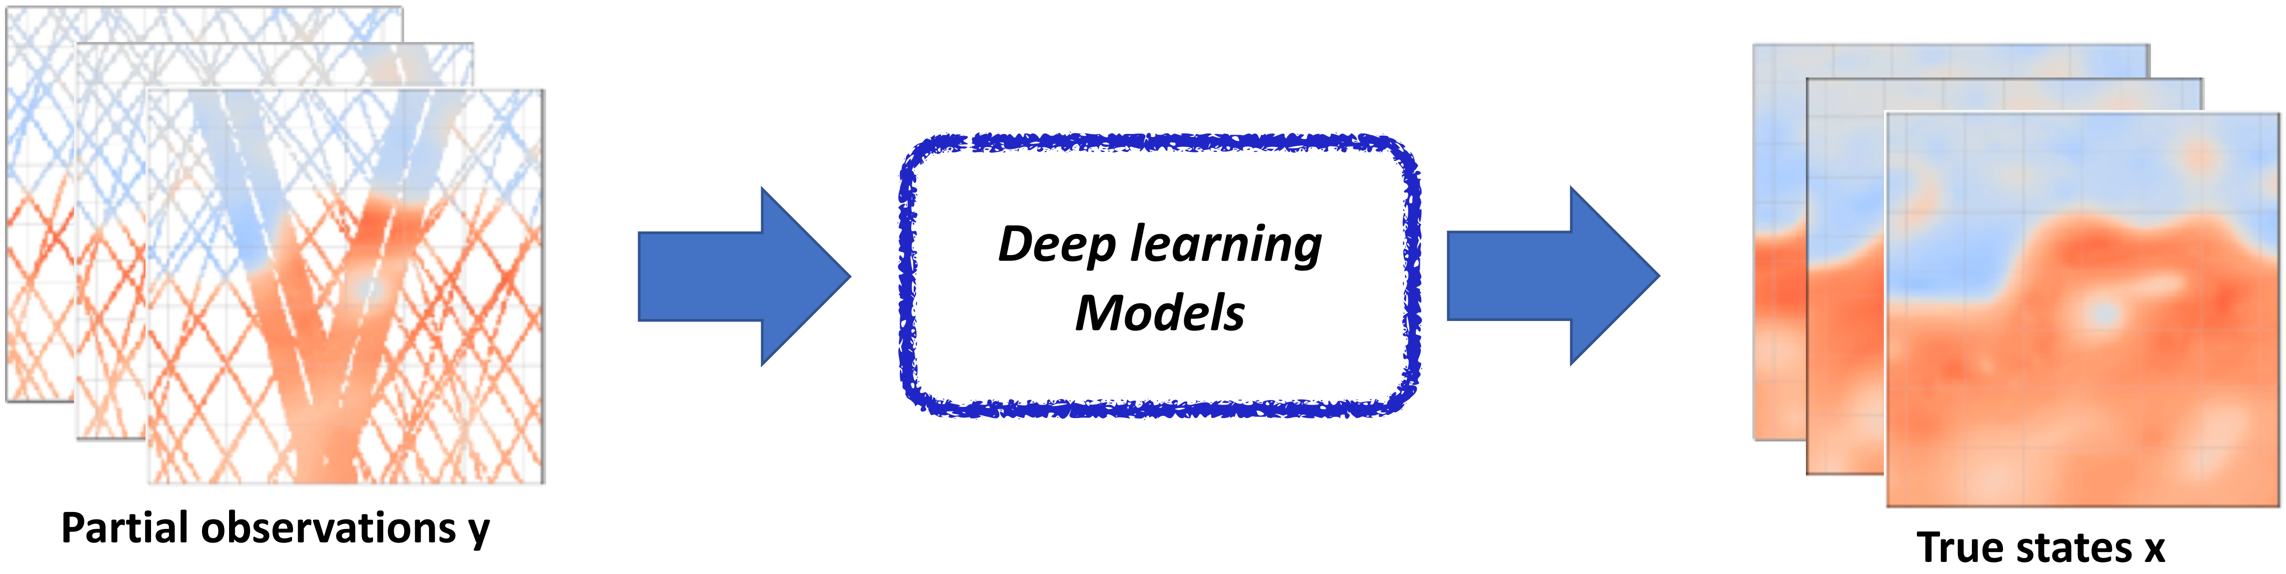

Which neural architectures ?

Which training loss ?

WC-4Dvar DA as a deep unfolded scheme

Unfolded architecture
of a 4DVar-WC scheme

$$\widehat{x} = \arg\min_{x} \underbrace{\|y - \mathbf{H}x\|^2 + \lambda\|x - \Phi(x)\|^2}_{\mathcal{J}_{\mathbf{H},\Phi}(x,y)}$$

Iterative update rule

$$x^{(k+1)} = x^{(k)} - \alpha \nabla_x \mathcal{J}_{\mathbf{H},\Phi}\left(x^{(k)}, y\right)$$

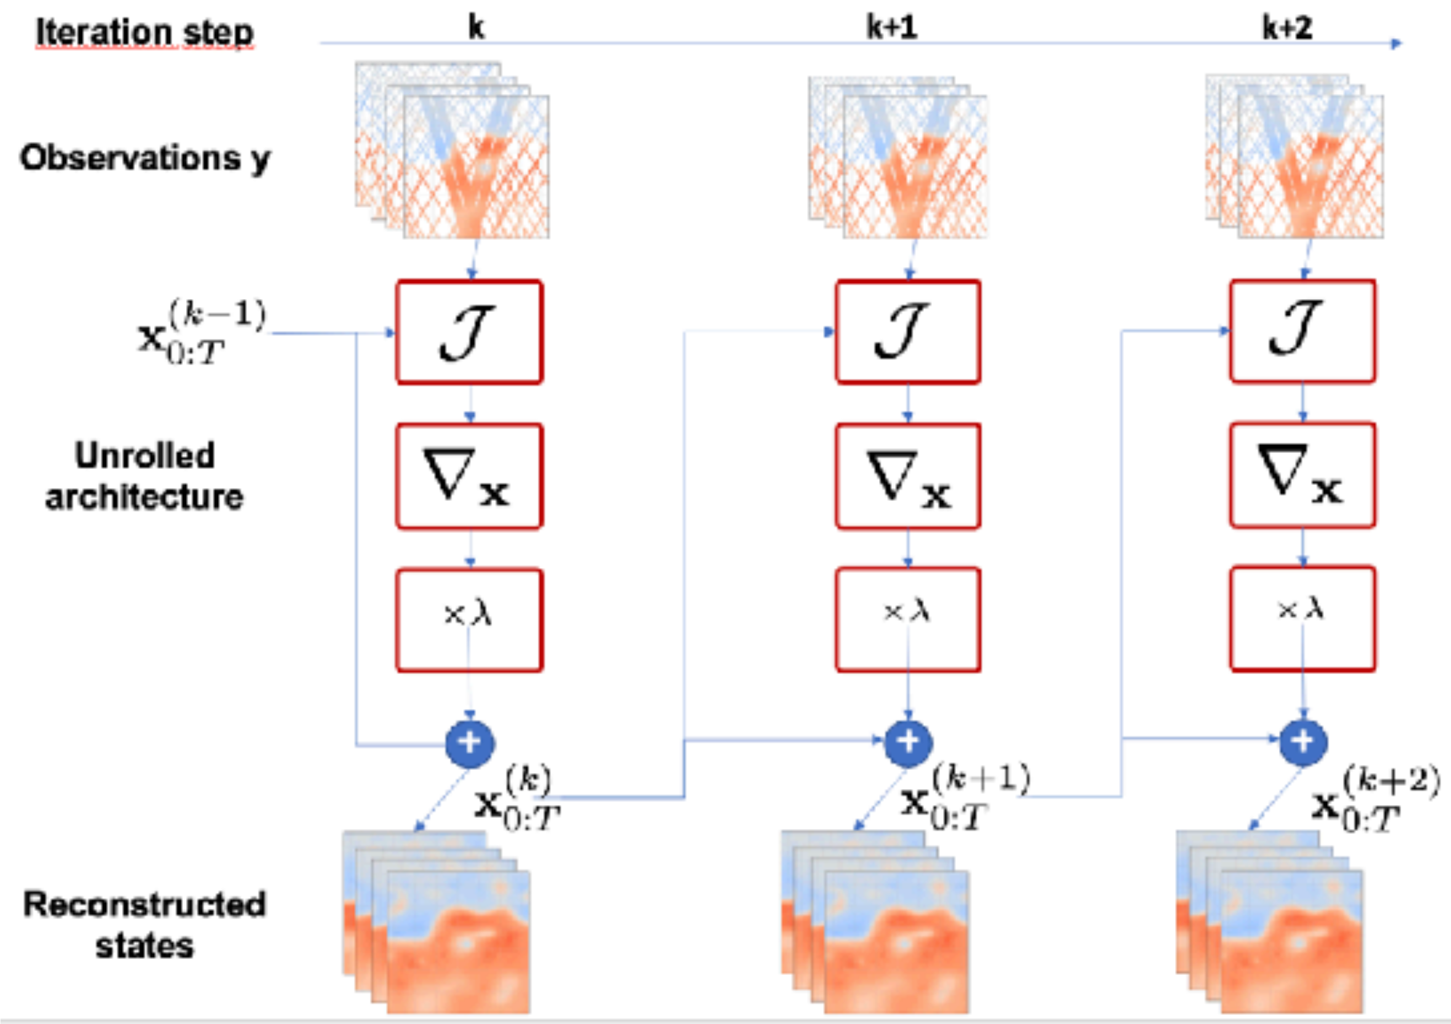

https://doi.org/10.1029/2021MS002572

Folded vs. Unfolded Representations

An example with a ResNet

$$x_{t+1} = x_t + \mathcal{R}\left(x_t\right)$$

Folded Representation

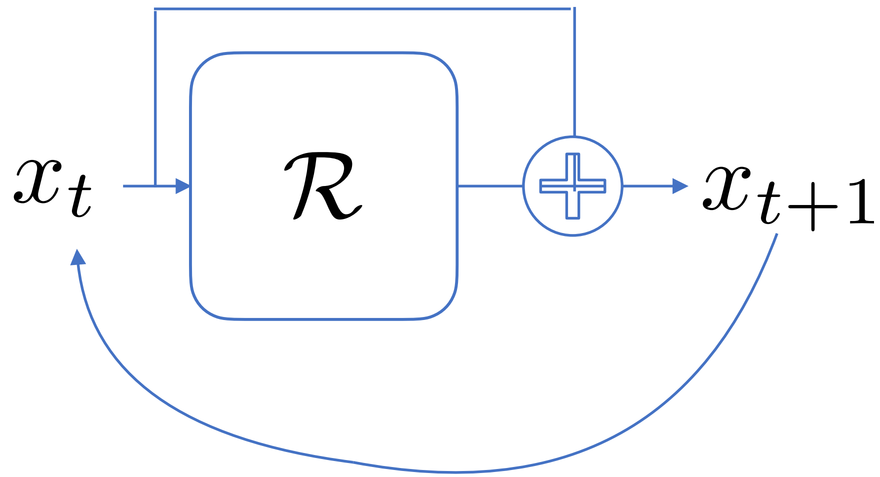

Unfolded Representation

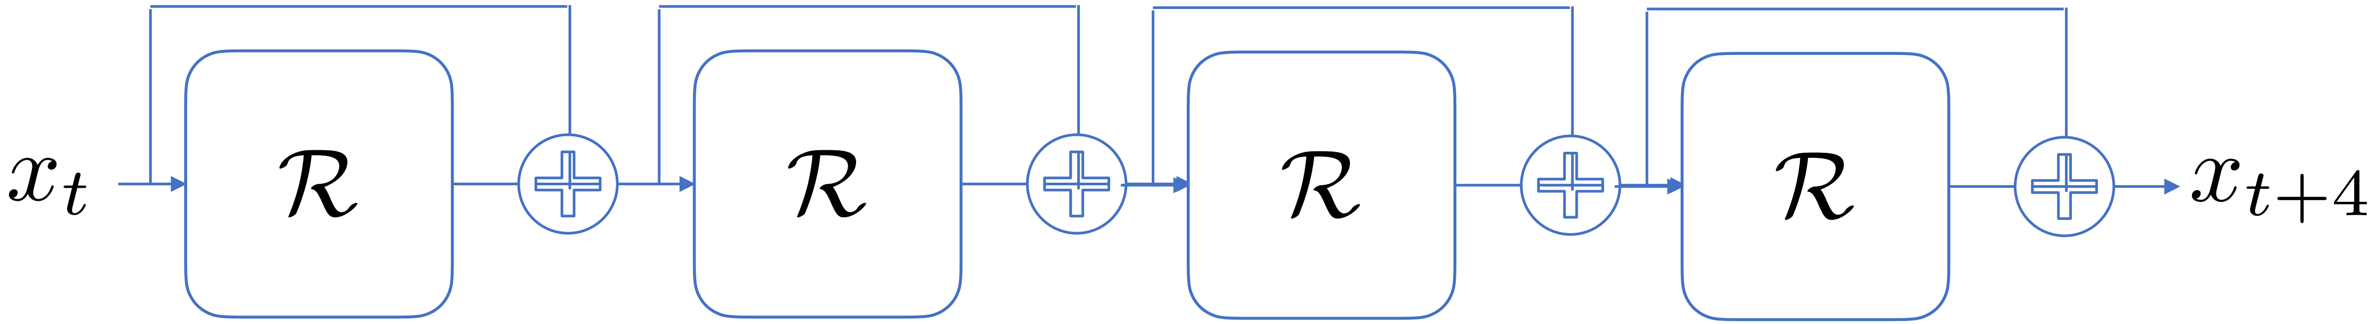

Data Assimilation using Deep unfolding schemes: 4DVarNets

Unfolded architecture of a 4DVarNet scheme

$$\widehat{x} = \arg\min_x \underbrace{\|y - \mathbf{H}x\|^2 + \lambda \|x - \Phi(x)\|^2}_{\mathcal{J}_{\mathbf{H},\Phi}(x,y)}$$

Iterative gradient-based update

$$x^{(k+1)} = x^{(k)} + \mathcal{R}\left[\nabla_x \mathcal{J}_{\mathbf{H},\Phi}\left(x^{(k)}, y\right)\right]$$

Fablet et al., 2021
https://doi.org/10.1029/2021MS002572

Meta-learning review, Hospedales et al., 2020.
arXiv:2004.05439

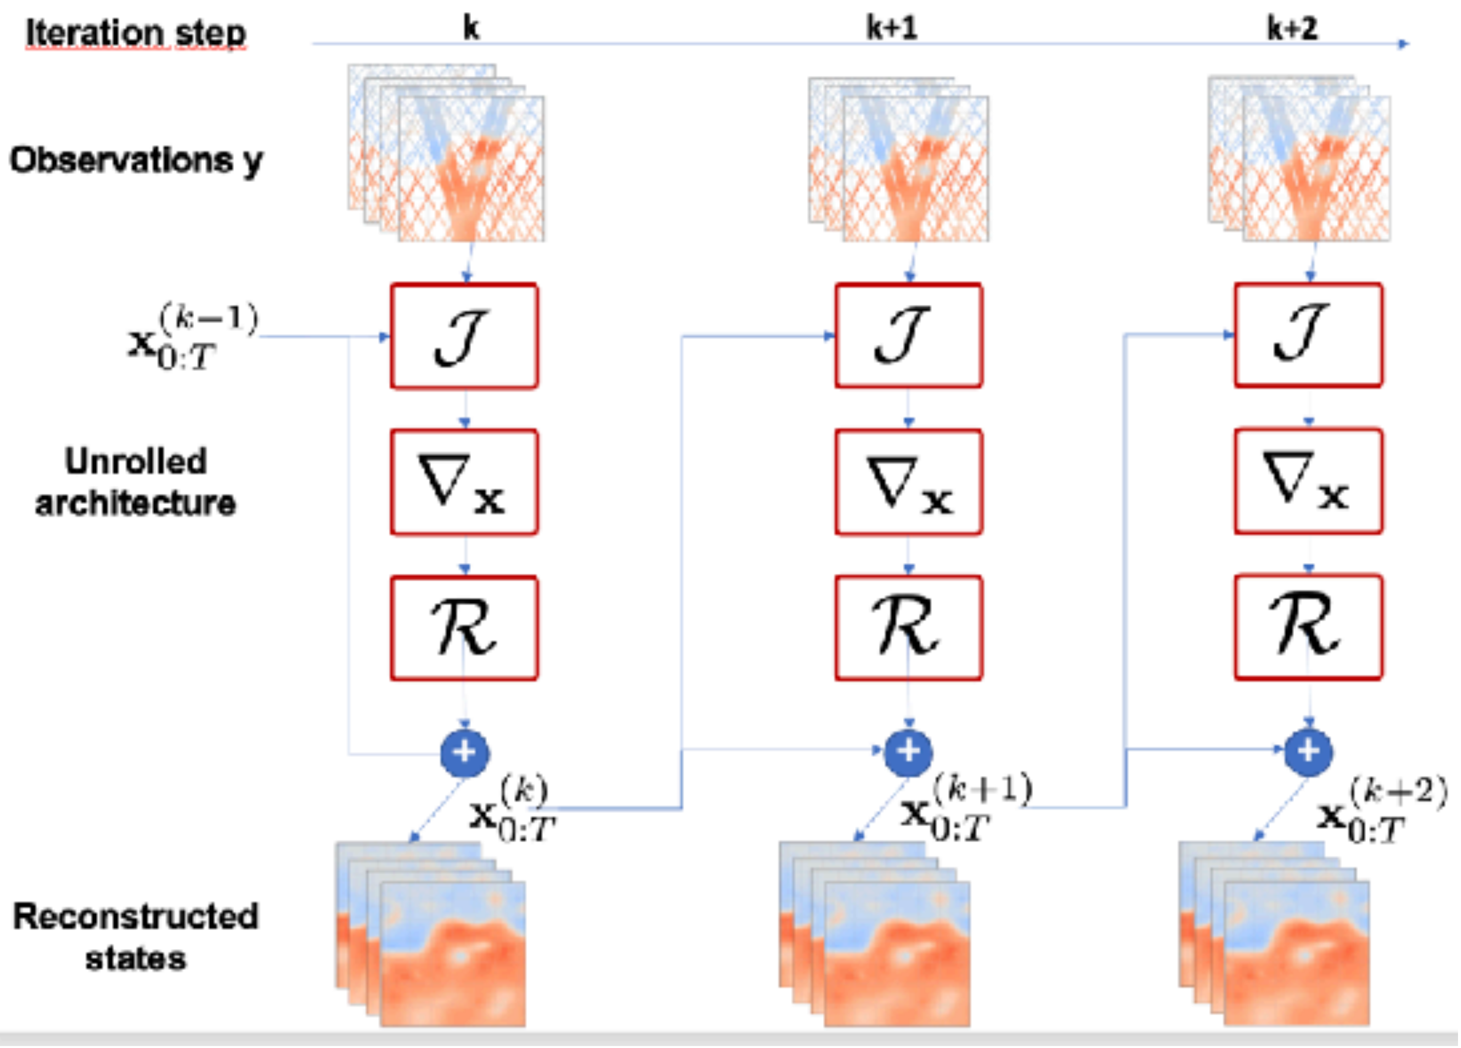

Data Assimilation using Deep unfolding schemes: 4DVarNets

Trainable Variational DA formulation

$$\widehat{x} = \arg\min_x \|y - \mathbf{H}x\|^2 + \lambda\|x - \Phi(x)\|^2$$

Trainable or pre-defined observation model

Trainable or pre-defined prior

Trainable solver

$$x^{(k+1)} = x^{(k)} + \mathcal{R}\left[\nabla_x \mathcal{J}_{\mathbf{H},\Phi}\left(x^{(k)}, y\right)\right]$$

RNN

Automatic differentiation

4DVarNet: end-to-end neural DA

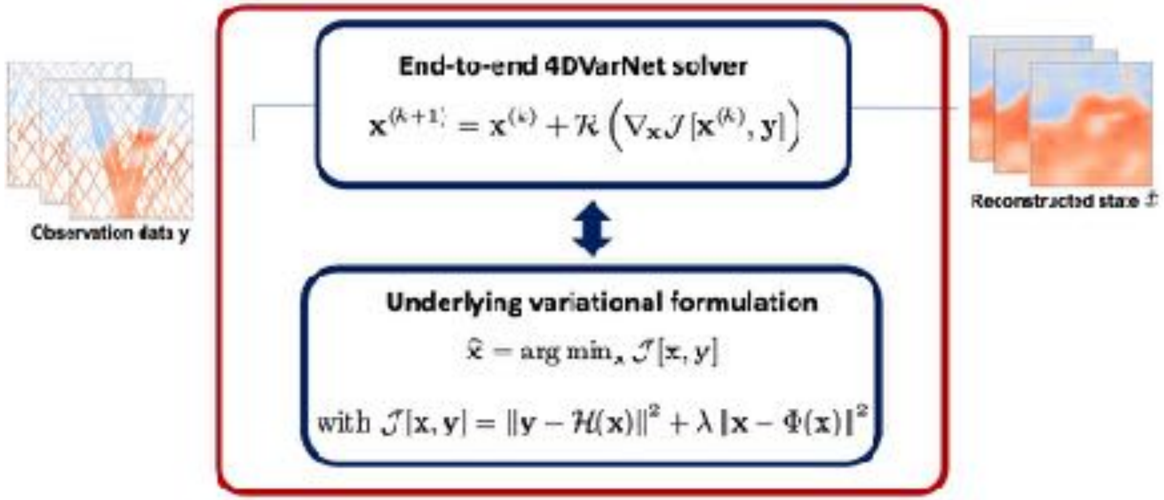

Which training loss?

Supervised vs. Non-supervised?

Paper: Fablet et al., JAMES, 2021. https://doi.org/10.1029/2021MS002572
Code: https://github.com/CIA-Oceanix/4dvarnet-starter

Bi-level formulation for Data Assimilation

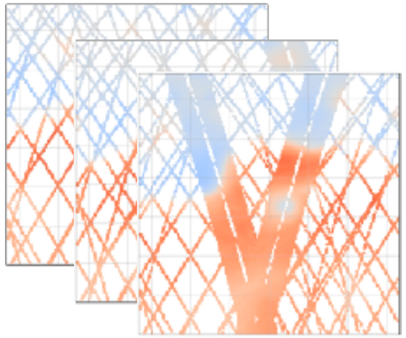

Partial observations y

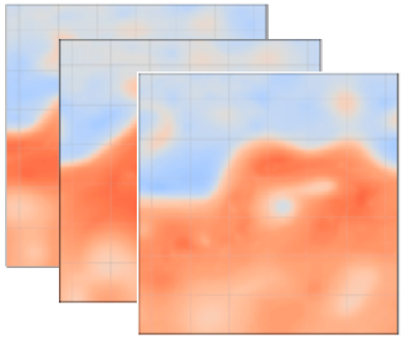

True states x

Model-driven schemes: $\widehat{x} = \arg\min_{x} \|y - \mathbf{H}x\|^2 + \lambda \|x - \Phi(x)\|^2$

End-to-end learning schemes: $\widehat{x} = \Psi(y) \text{ s.t. } \widehat{\Psi} = \arg\min_{\Psi} \mathcal{L}\left(\{x_n, y_n, \widehat{x}_n\}_n\right)$

Training phase as the targeted bi-level optimisation problem

$$\widehat{\Phi} = \arg\min_{\Phi} \mathcal{L}\left(\{x_n, y_n, \widehat{x}_n\}_n\right)$$

$$\text{s.t. } \widehat{x}_n = \arg\min_{x} \|y_n - \mathbf{H}_n x\|^2 + \lambda \|x - \Phi(x)\|^2$$

Relaxed formulation in 4DVarNet

$$\widehat{\Phi}, \widehat{\mathcal{S}} = \arg\min_{\Phi, \mathcal{S}} \mathcal{L}\left(\{x_n, y_n, \widehat{x}_n\}_n\right) \quad \text{s.t. } \widehat{x}_n = \Psi_{\Phi, \mathcal{S}}(y_n, \mathbf{H}_n)$$

Approximate (trainable) solver of the variational minimisation

Ongoing 4DVarNet-related topics (cia-oceanix.github.io)

Implicit layers

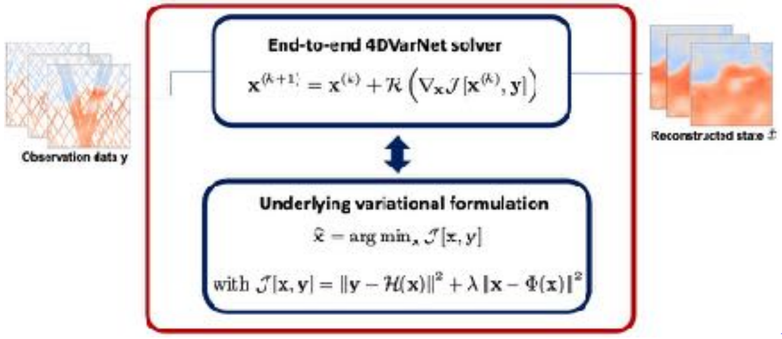

4DVarNet, OI and SPDEs

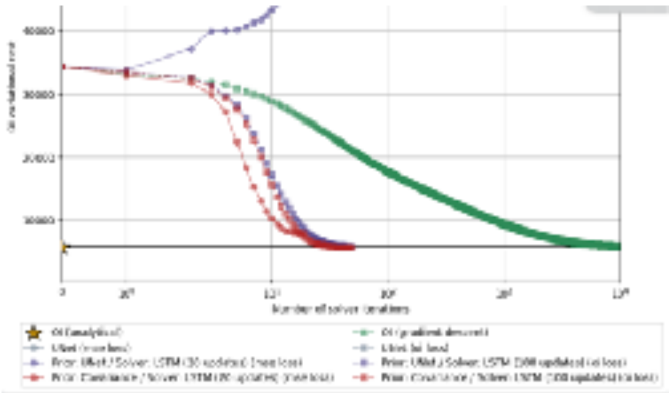

4DVarNet and UQ

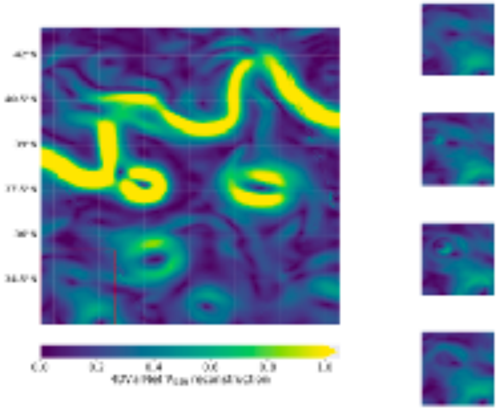

SSH mapping/forecasting

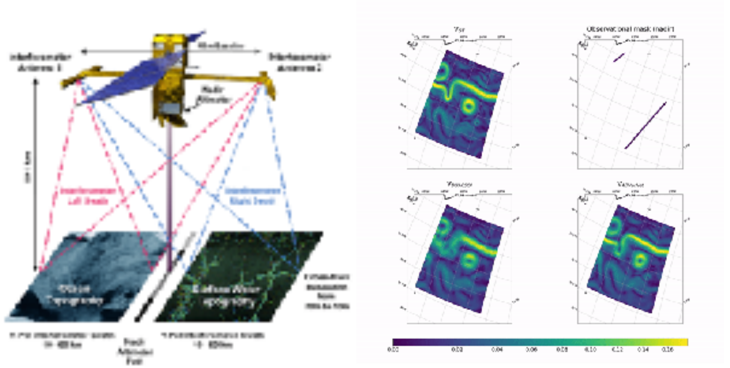

Sea Surface Turbidity

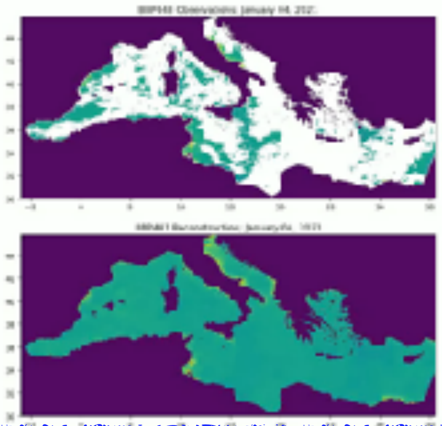

Multimodal synergies

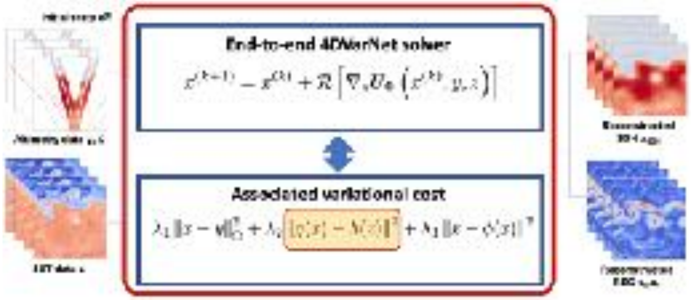

How to benefit from existing knowledge and systems in neural DA schemes?

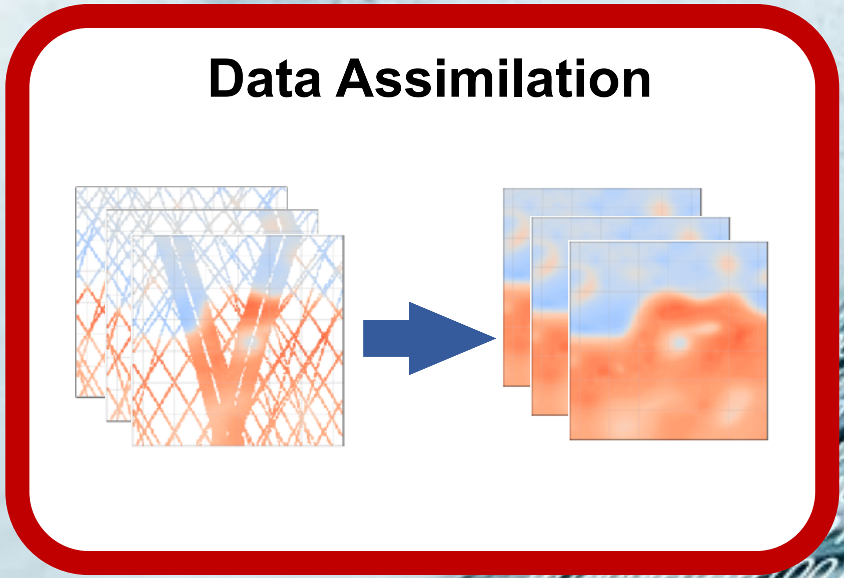

4DVarNet: which projection operator Φ?

Parameterization using (trainable) ODE operator

$$\frac{\partial x(t)}{\partial t} = \mathcal{M}_\theta\left(x(t)\right)$$

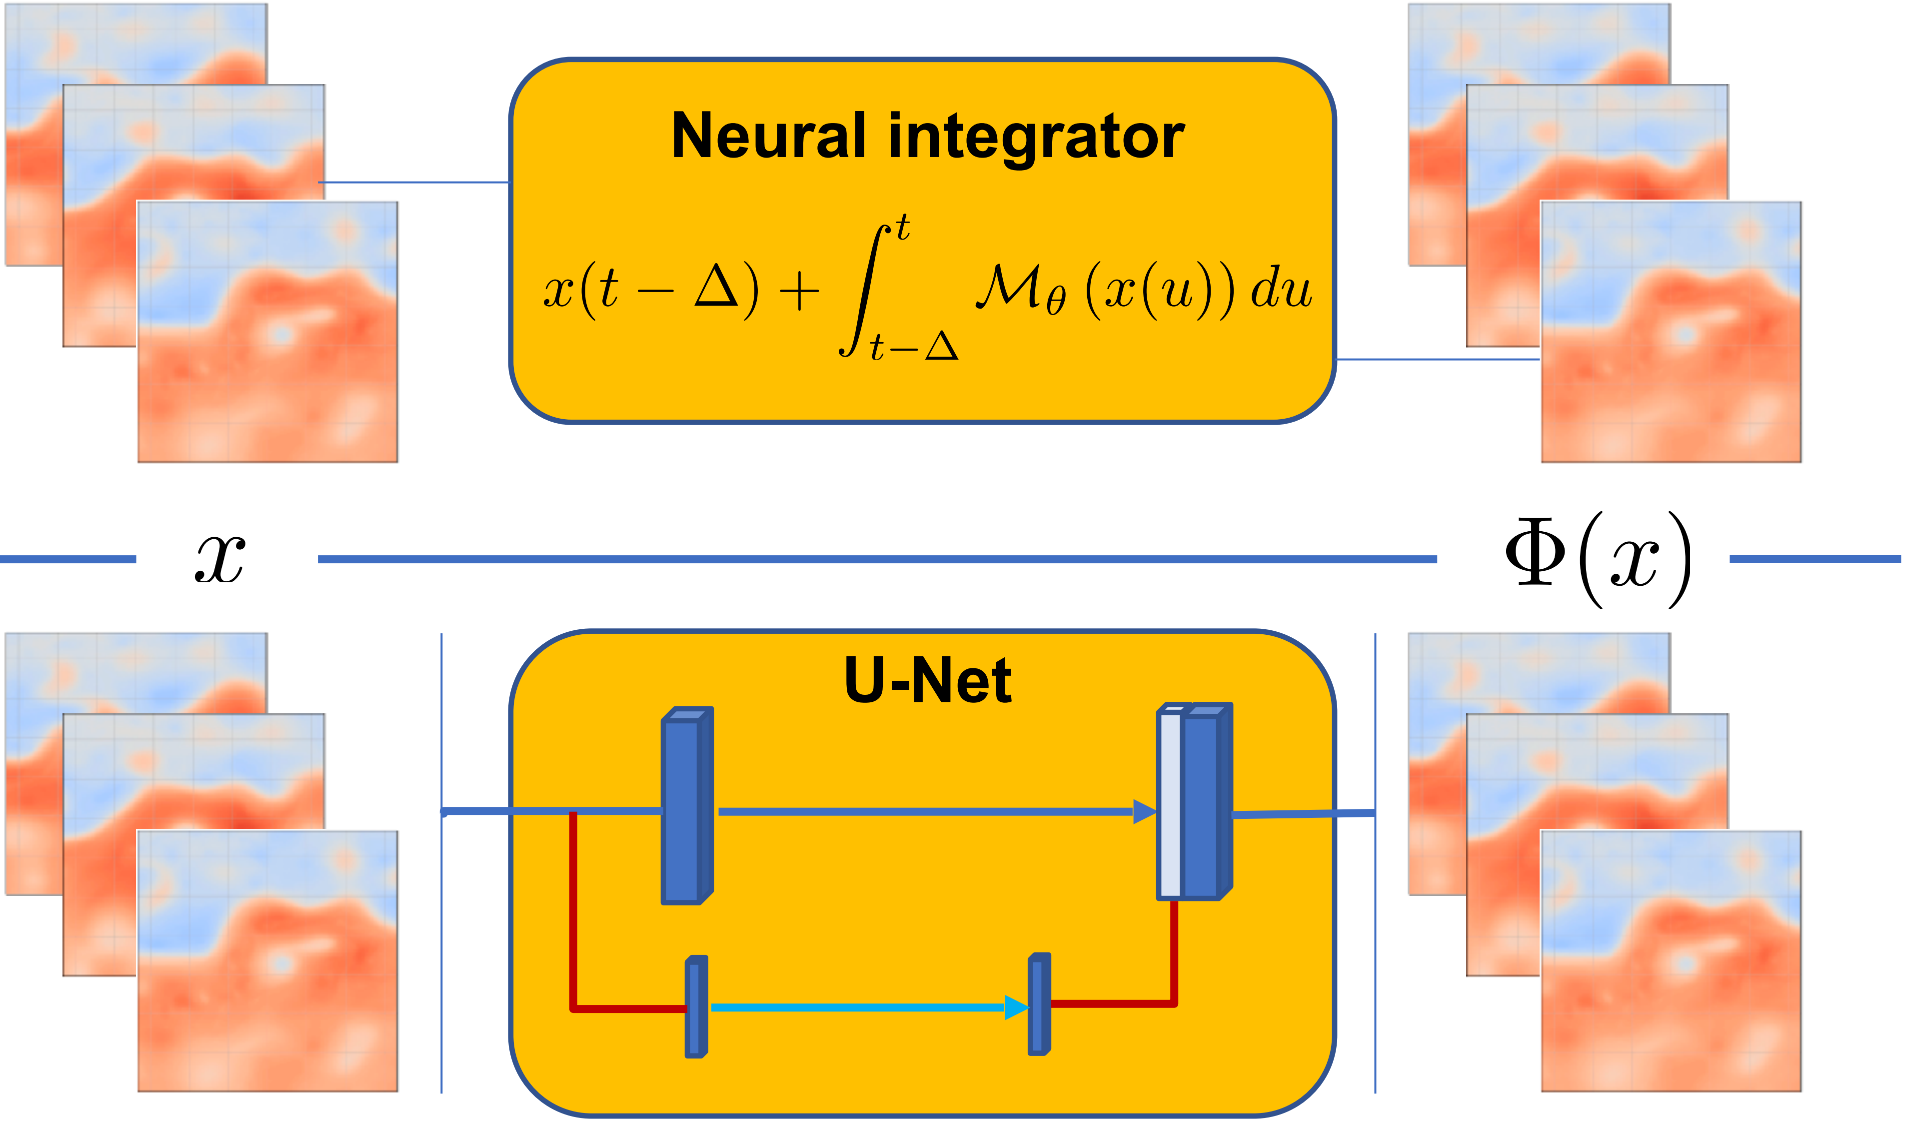

Two-scale U-Net-like Parameterization (similar to Gibbs/ Markov Random Field)

4DVarNet reconstruction of L96 dynamics

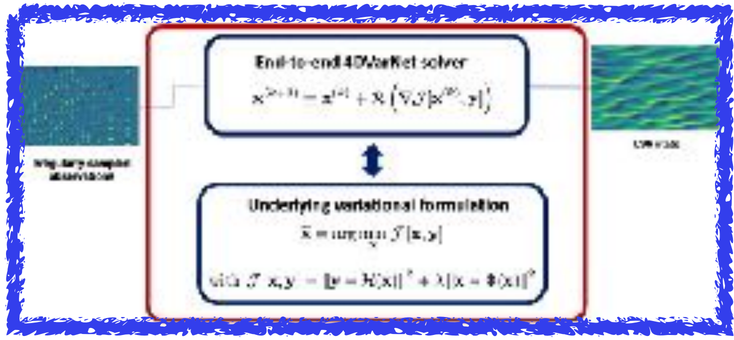

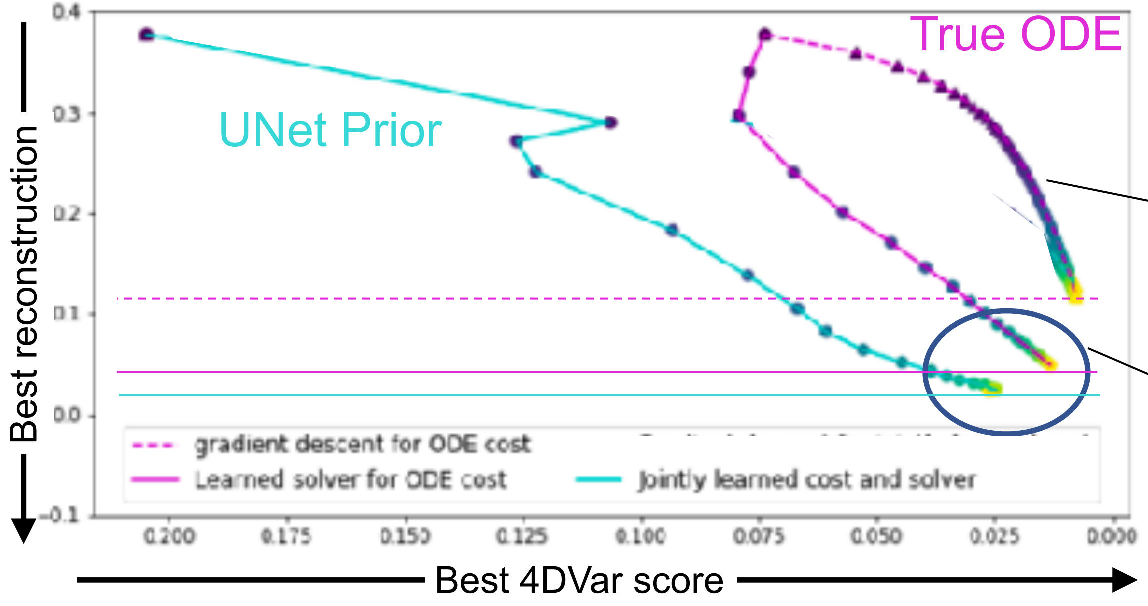

Classic 4DVar minimisation
(Unsupervised setting)

4DVarNets
(Supervised training)

True and observed states

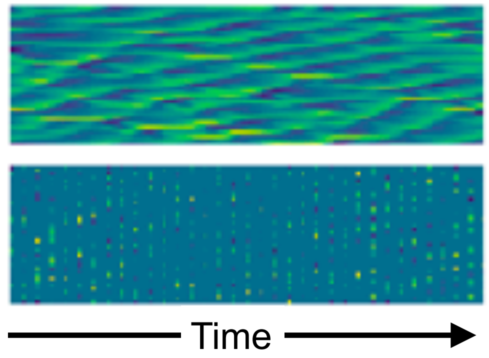

Reconstruction examples and associated error maps

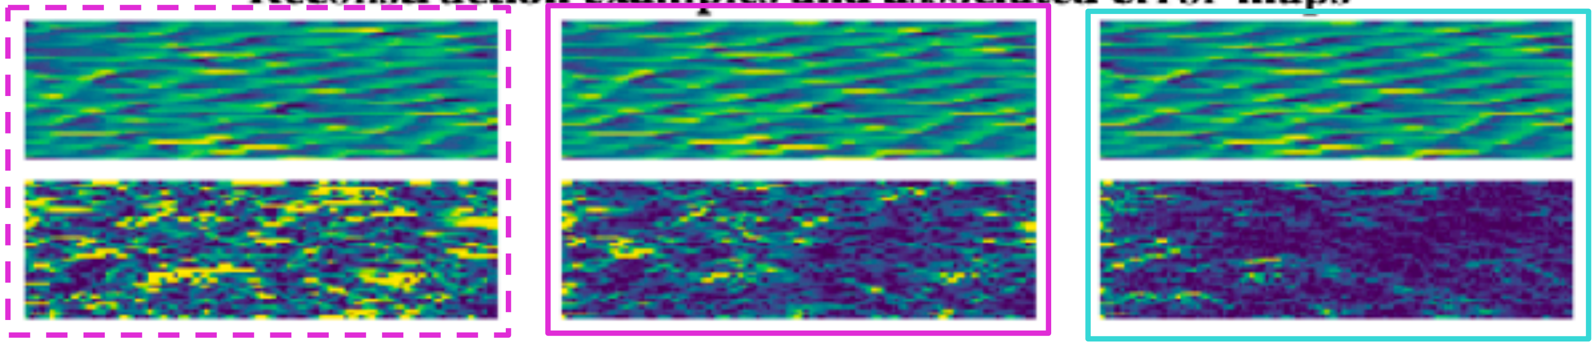

Generic vs. Observation-optimized priors?

Which robustness to the sampling pattern?

Reduced-order representation of AMOC dynamics

$$\begin{cases} \dot{\omega}(t) = & -\lambda\omega(t) - \epsilon\beta S_{NS}(t) \\ \dot{S}_{BT}(t) = & (\Omega_0 + \omega(t))\, S_{NS}(t) - K\, S_{BT}(t) + \frac{F_0 S_0}{h}, \\ \dot{S}_{NS}(t) = & -(\Omega_0 + \omega(t))\, S_{BT}(t) - K\, S_{NS}(t) \end{cases}$$

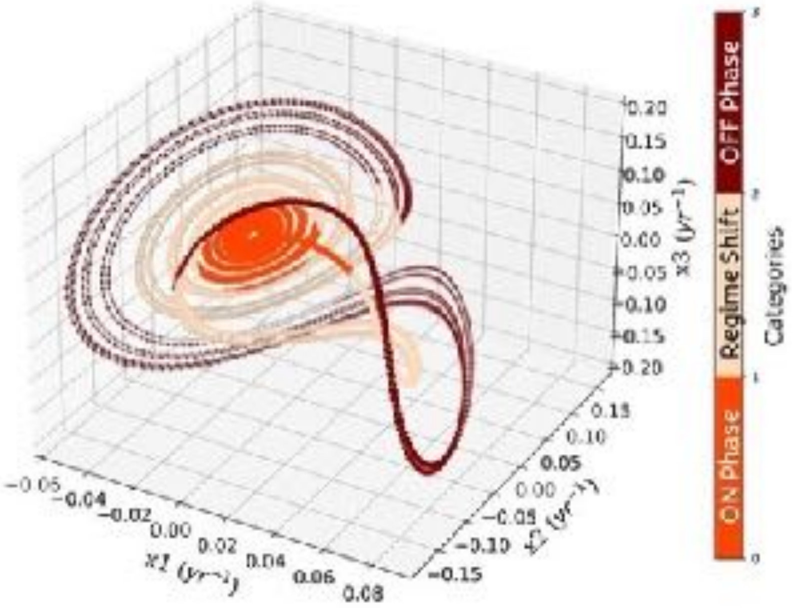

Bauchot et al. JAMES, 2025. [10.1029/2024MS004462](https://doi.org/10.1029/2024MS004462)

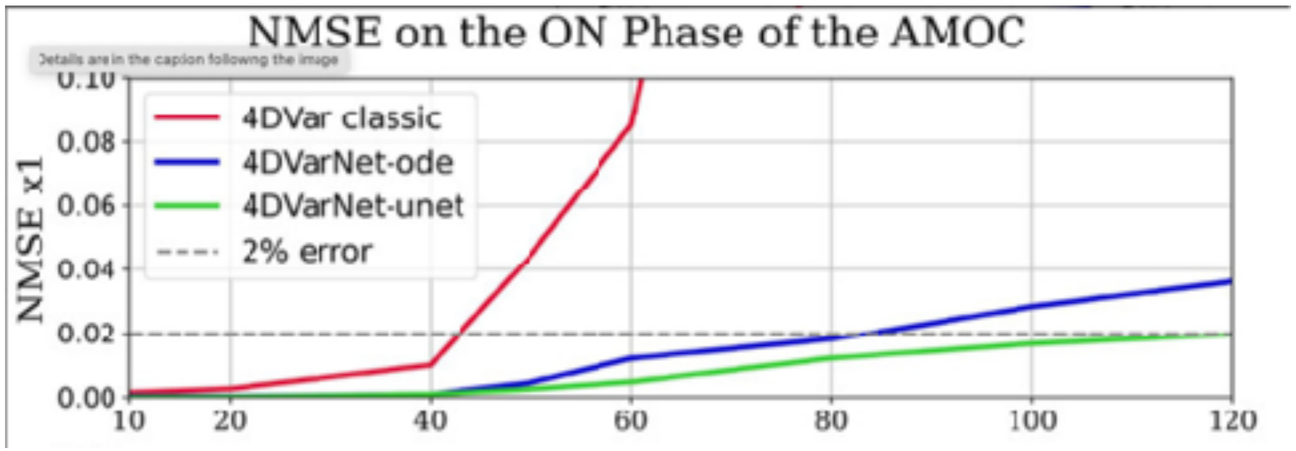

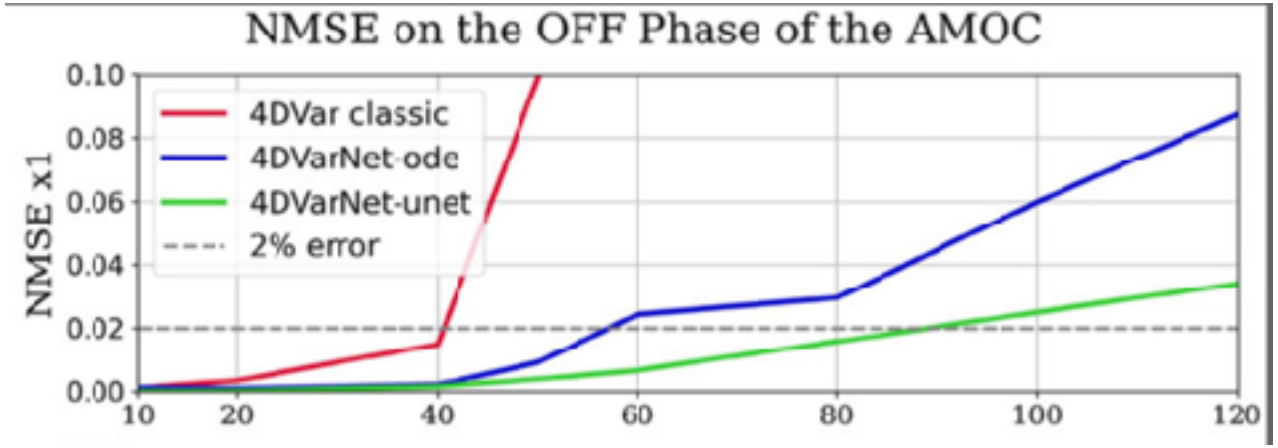

The sparser the sampling, the more the difference (beyond Shannon**)**

Can we benefit from a first guess from operational DA systems?

4DVarNet scheme using an I.C.

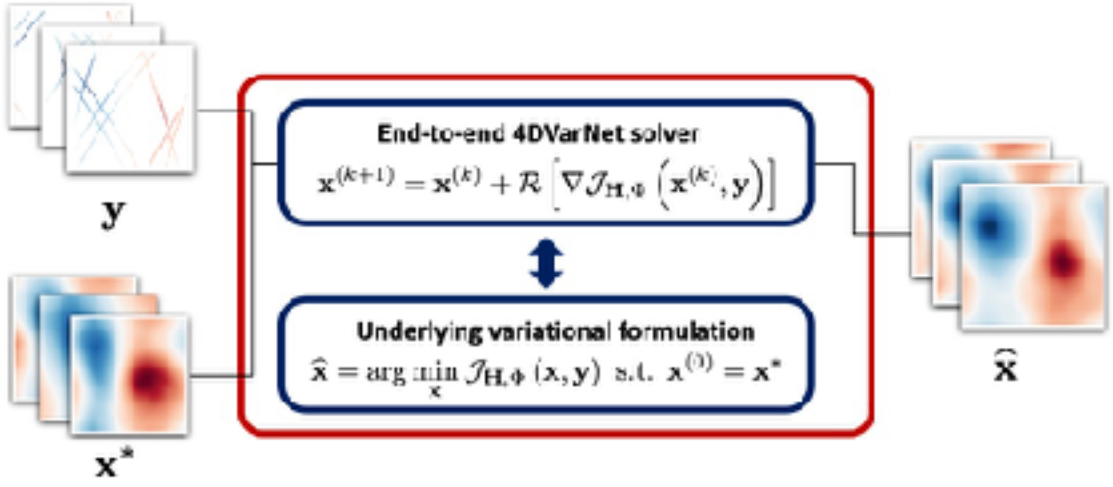

Random perturbation of true states

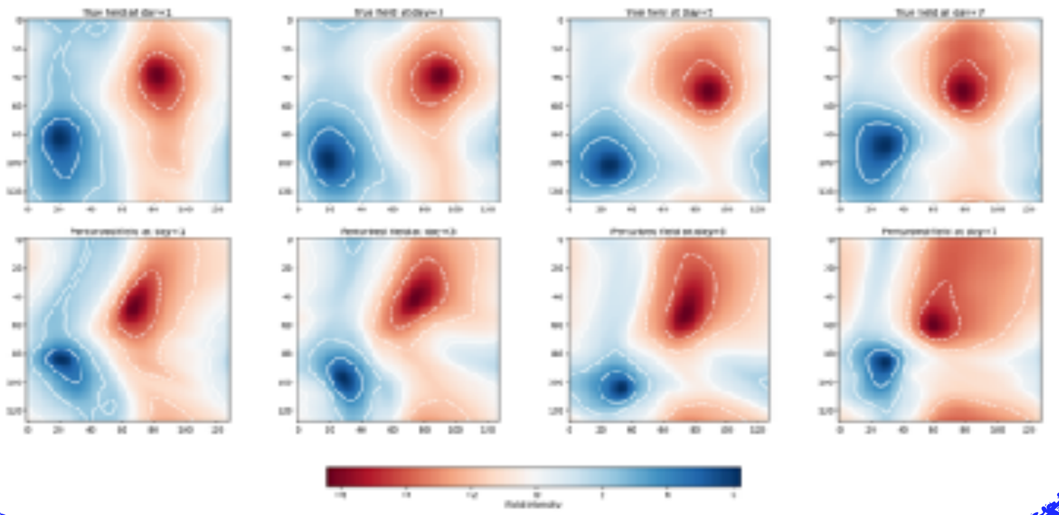

Impact of the uncertainty level of the DA baseline

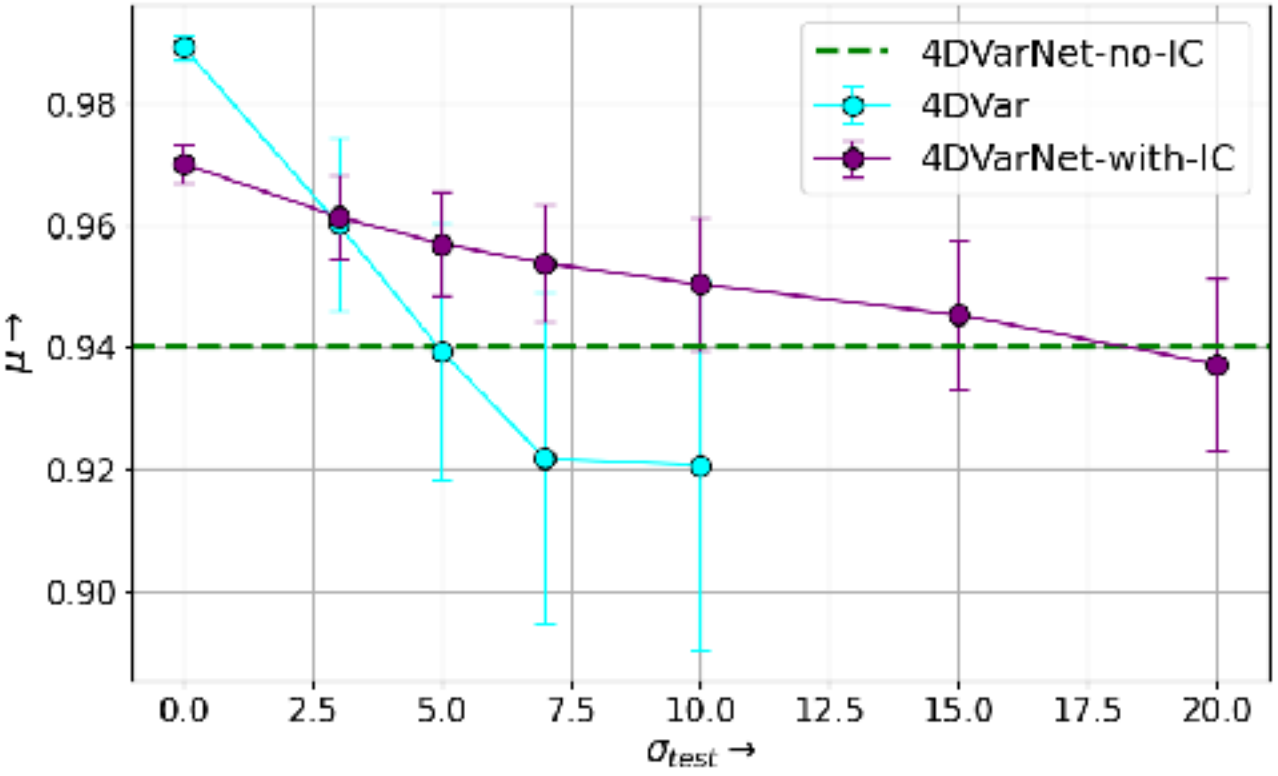

Uncertainty levels seen during training

Uncertainty levels not seen during training

Roy et al. EGU, 2025.

From deterministic neural DA to generative neural DA?

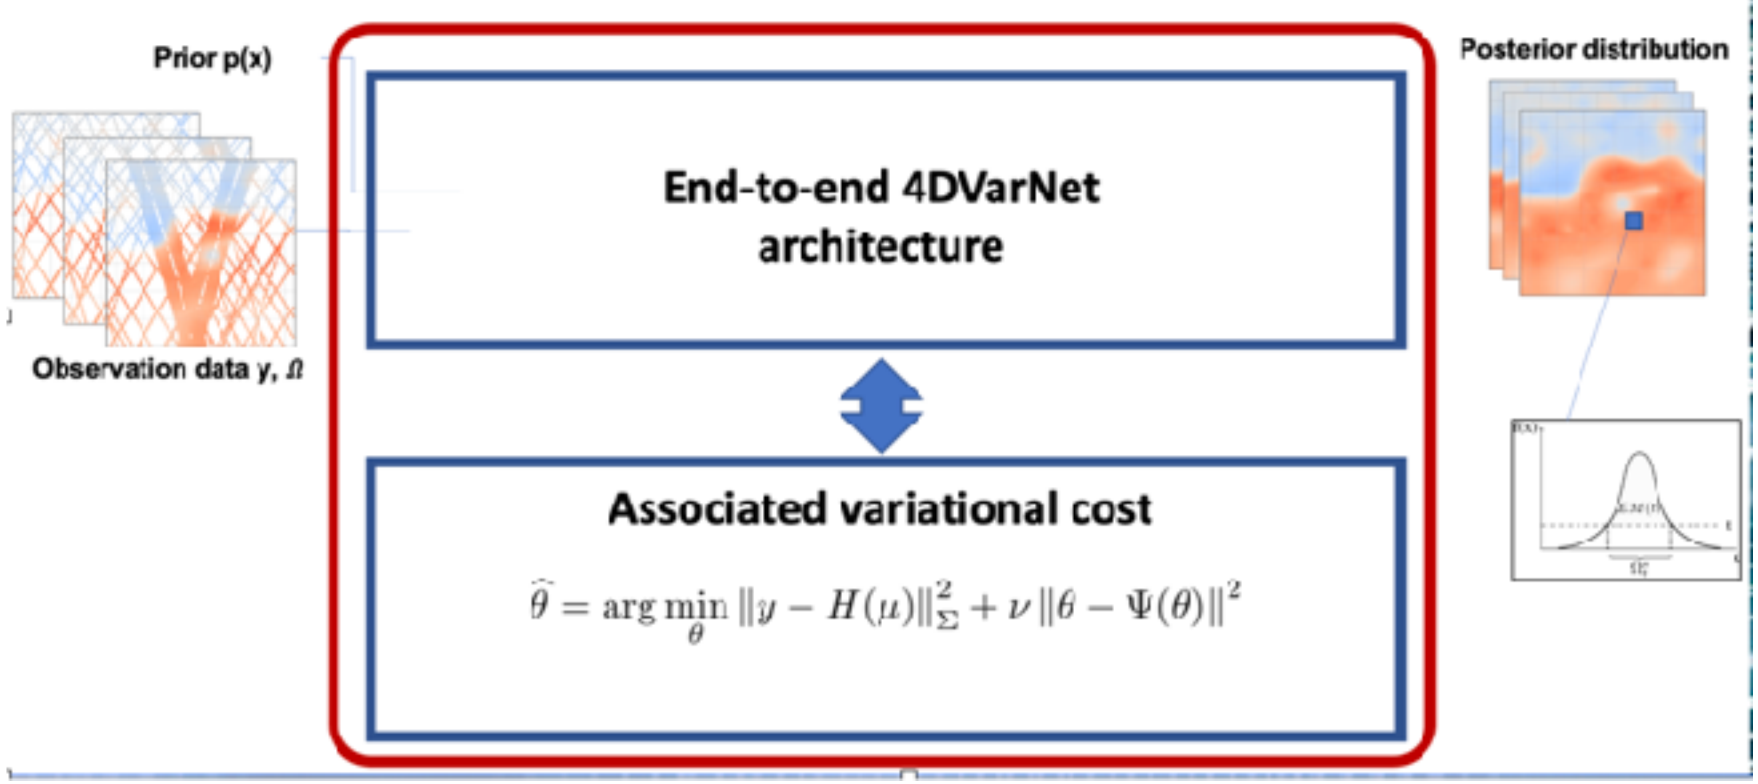

4DVarNet and Uncertainty Quantification

Ensemble of 4DVarNets

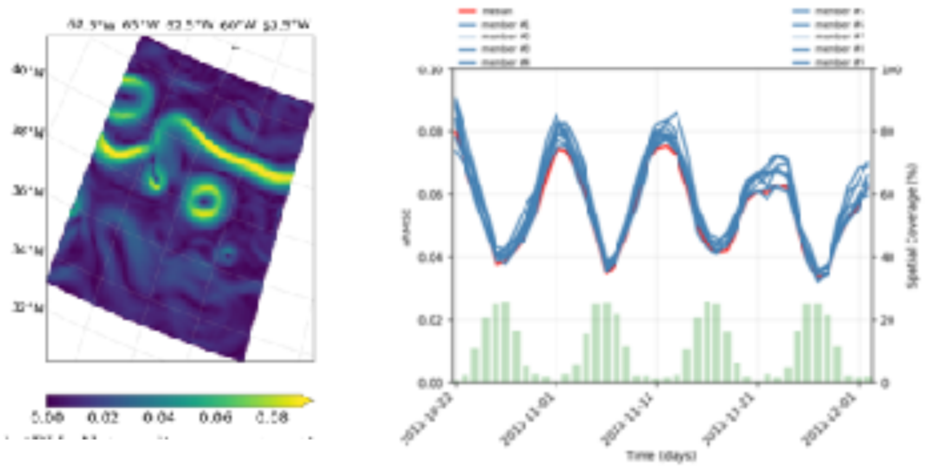

Beauchamp et al., 2023. doi.org/10.5194/gmd-16-2119-2023

Random Perturbation and neural DA

Basic idea: sampling error signals using reconstruction from random perturbation of the interpolated field

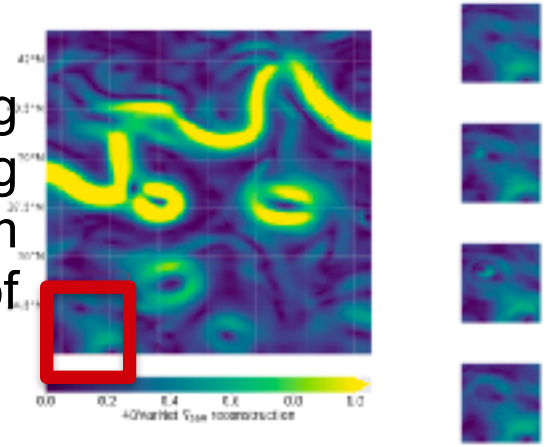

Roy et al., in prep.

Neural DA & Flow Methods

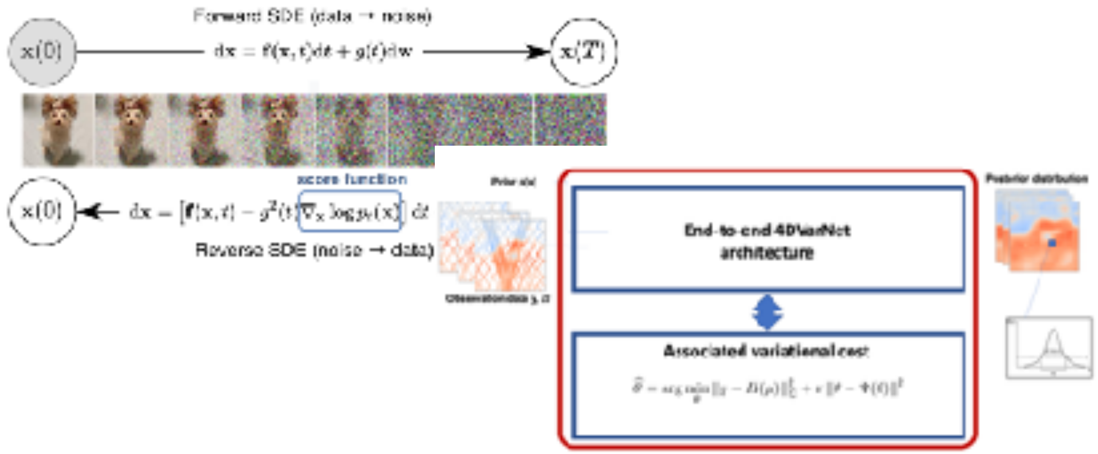

4DVarNets with SPDE priors

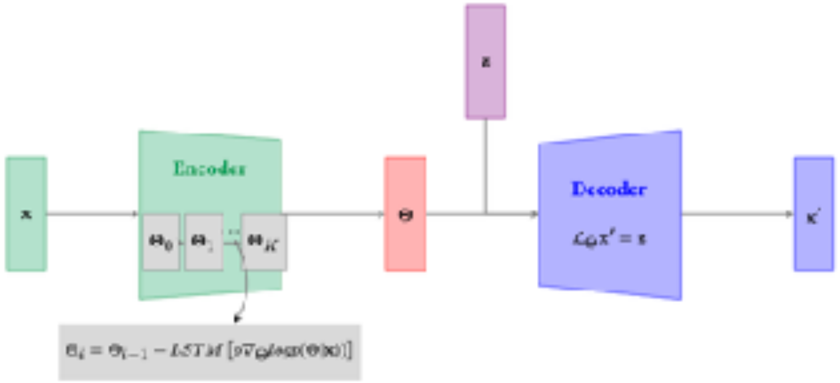

Beauchamp et al., 2025. https://arxiv.org/abs/2311.01783

SPDE priors with 4DVarNets

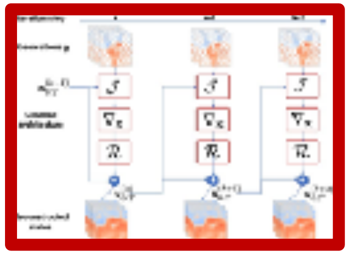

SPDE-4DVarNet in inference mode

Step I: estimation of the augmented state from observation data using the trained 4DVarNet

$$\widetilde{\mathbf{X}} = [\mathbf{X} \quad \theta]^T$$

Step II: sampling the SPDE prior with estimated parameter θ

Step III: computing the error samples from simulated SPDE states using the same mask as the real observations

Step II: sampling “posterior” states as the combination of the neural interpolation of the real observations and of an error sample

$$\mathbf{x}^{\star,i}(\mathbf{s},t) = \mathbf{x}^{\star}(\mathbf{s},t) + (\mathbf{x}_s^{i}(\mathbf{s},t) - \mathbf{x}_s^{\star,i}(\mathbf{s},t))$$

SPDE priors with 4DVarNets for Satellite Altimetry

Performance wrt 4DVarNet baseline (nadir+SWOT OSSE)

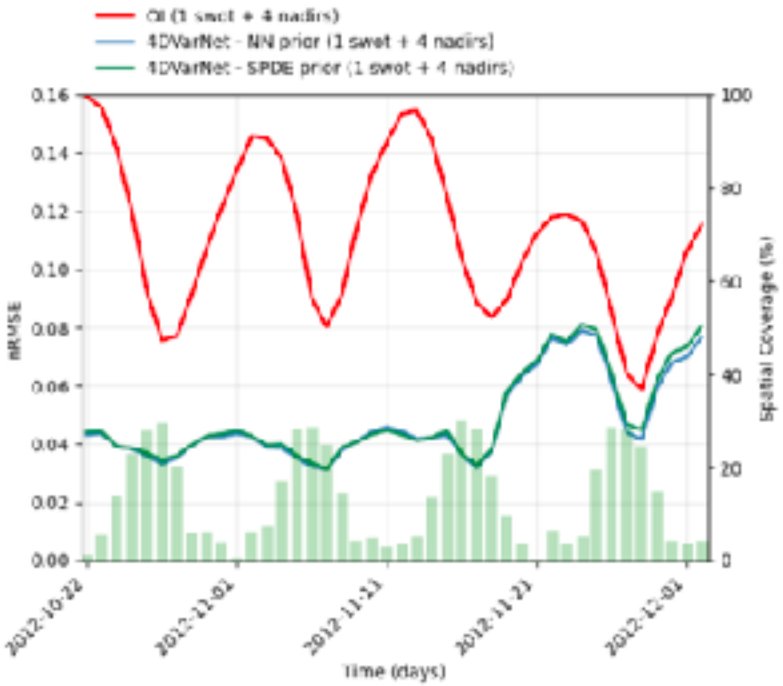

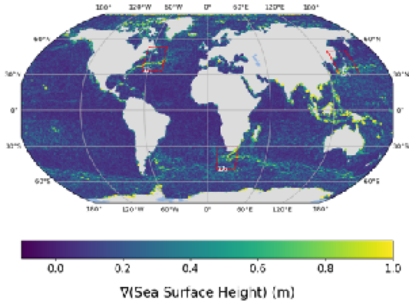

Generated “posterior” samples

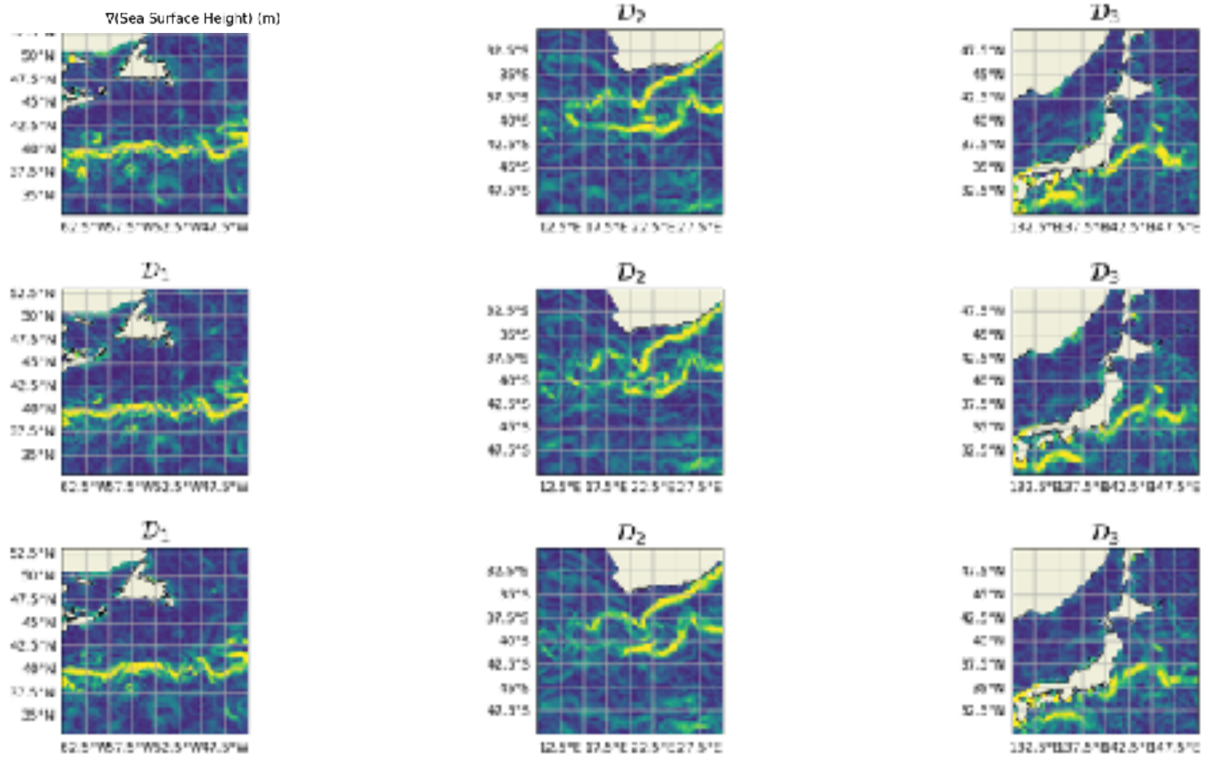

Lessons learnt

Take-Home messages

- *Growing literature bridging DA, UQ, DL*
- *Learning from simulations for applications to real data*
- *Moving towards conditional generative approaches*
- *State-of-the-art performance for the reconstruction and forecasting of ocean surface dynamics*
- *Importance of collaborative interdisciplinary frameworks*

Fostering interdisciplinary and open science

opening : architectures, workflows, datasets and codes

Key for developing robust, well-proven evolutive solutions

New French initiative to design ocean benchmarks: https://github.com/ppr-ocean-ia/data-challenges-info

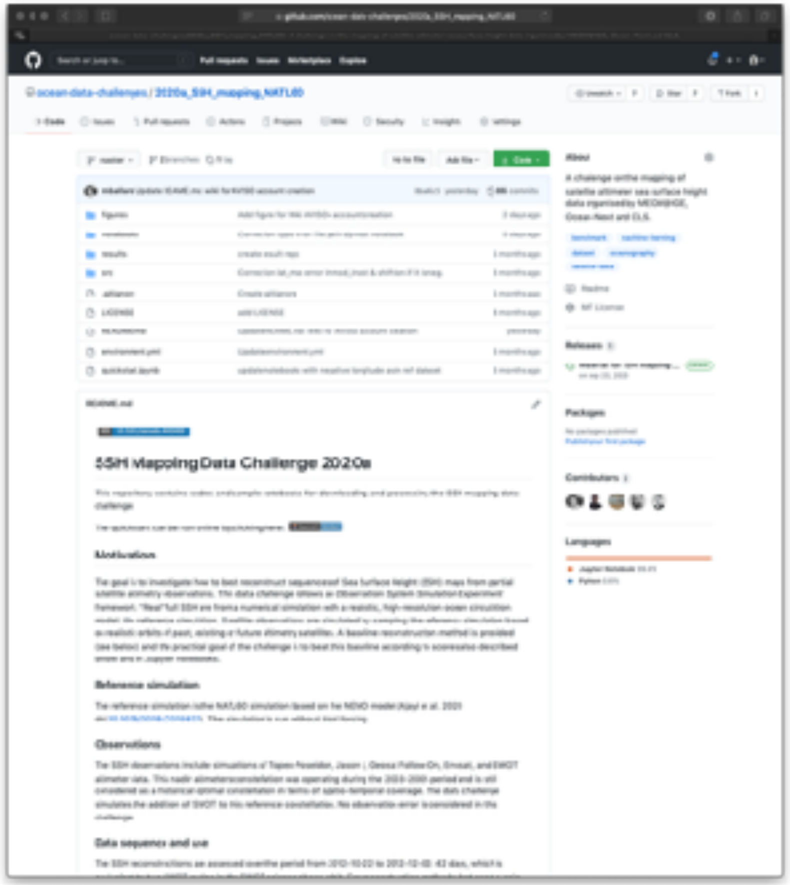

e.g. : collaborative data challenges

Initiative “AI & Ocean”: Ocean Data Challenges

https://github.com/ppr-ocean-ia/data-challenges-info

Scientific coordination: R. Fablet, IMT Atlantique

Technological coordination: Q. Febvre, J.F. Pill (Ifremer), K. Ait Mohand, G. Cossio (IMT Atlantique), Q. Gaudel (Mercator Ocean Intl)

Selected Benchmarks

- ***DC1. Emulation of Global Ocean Reanalyses***
  - Scientific lead: J. Le Sommer (IGE), A El Aouni (Moi), S. Ouala (Odyssey), A. Storto (CNR)
- ***DC2. Probabilistic short-term forecasting of global ocean dynamics***
  - Scientific lead: L. Drumetz, F. Sévellec (Odyssey), J. Le Sommer (IGE), S. Metref (Dallas)
- ***DC3. Arctic Sea Ice Forecasting***
  - Scientific lead: S. Metref (Datlas), P. Rampal (IGE), J. Brajard (NERSC)
- ***DC4. High-resolution monitoring and forecasting of tropical cyclones***
  - Scientific lead: Q. Febvre, A. Mouche, B. Chapron (Odyssey)
- ***DC5. Marine Biodiversity Prediction***
  - Scientific lead: J.O. Irisson (LOV), L. Guidi (LOV). D. Mouillot (Marbec), A. Joly (LIRMM)

Distribution Date: September 2025 (before Kick-Off workshop)

Git organisation for the development phase: ***https://github.com/ppr-ocean-ia***

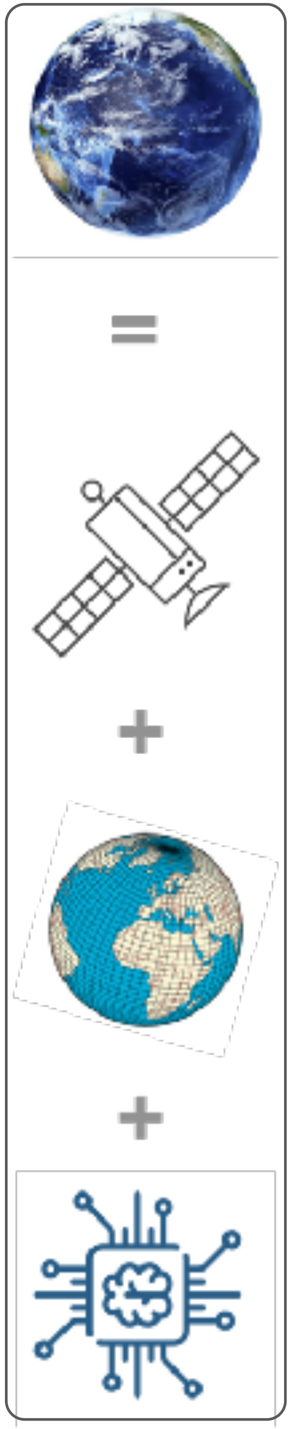

- **Model calibration/identification**
  - Frezat et al. Physics-informed neural networks for sub-grid scale modeling in filtered turbulence. PRF, 2021. link
  - Frezat et al. A posteriori learning for quasi-geostrophic turbulence parametrization. arXiv, 2022. link
  - Ouala et al. Learning Runge-Kutta Integration Schemes for ODE Simulation and Identification. arXiv, 2021. link.
  - Ouala et al. Extending the extended dynamic mode decomposition with latent observables: the latent EDMD framework. MLST, 2023. link

- **Data assimilation and Forecasting**
  - Ouala et al. Learning Latent Dynamics for Partially-Observed Chaotic Systems. Chaos, 2020. link
  - Fablet et al. Learning Variational Data Assimilation Models and Solvers. JAMES, 2021. link
  - Lafon et al. Uncertainty quantification when learning dynamical models and solvers with variational methods. Preprint, 2022. link
  - Ouala et al. Bounded nonlinear forecasts of partially-observed geophysical systems (….) Physica D, 2022. link
  - Beauchamp et al. Learning Neural Optimal Interpolation Models and Solvers. MLDADS, LNCS, 2023. link
  - Cheng et al. Machine learning with data assimilation and UQ for dynamical systems: a review. IEEE/CAA JAS, 2023. link

- **Trajectory data analysis and modeling**
  - Nguyen et al. GeoTrackNet-A Maritime Anomaly Detector using Probabilistic NN Representation of AIS Tracks and A Contrario Detection. IEEE TITS, 2020. link
  - Roy et al. Deep Learning and Trajectory Representation for the Prediction of Seabird Diving Behaviour. MEE, 2020. link
  - Nguyen et al. TrAISformer-A generative transformer for AIS trajectory prediction, arXiv, 2021. link

- **Reconstruction and forecasting of ocean dynamics**
  - Pauthenet et al. 4D temperature, salinity and mixed-layer depth in the Gulf Stream, reconstructed from remote-sensing and in situ observations with neural networks, Ocean Sci., 18, 1221–1244. link
  - Fablet et al. Multimodal 4DVarNets for the reconstruction of sea surface dynamics from SST-SSH synergies. IEEE TGRS, 2023. link
  - Fablet et al. Inversion of sea surface currents from satellite-derived SST-SSH synergies with 4DVarNets. arXiv, 2022. link
  - Beauchamp et al. 4DVarNet-SSH: end-to-end learning of variational interpolation schemes for nadir and wide-swath satellite altimetry. GMD, 2023. link
  - Roussillon et al. A Multi-Mode Convolutional Neural Network to reconstruct satellite-derived chlorophyll-a time series in the global ocean from physical drivers. Front. Mar. Sc. link

AI Chair OceaniX

Physics-informed AI for Observation-Driven Ocean AnalytiX

Thank you.

R. Fablet, Prof. IMT Atlantique, Brest
ronan.fablet@imt-atlantique.fr
Web: *https://cia-oceanix.github.io/*

Visiting scholarships, PhD and postdoc opportunities

AI4PEX