EVALUATING UNROLLED OPTIMIZATION MODEL ARCHITECTURES FOR OCEANOGRAPHIC DATA ASSIMILATION

Chaima Bounadra[1], Paul Denailly[2]
[1] *IMT Atlantique, MEE Department, Brest*

Summary
This study explores unrolled optimization models for oceanographic data assimilation, which are neural network architectures designed to iteratively improve estimates of the ocean state using both observations and prior physical knowledge. We start from 4DVarNet, an unrolled variational architecture that normally updates the estimated state through gradient-based steps of a cost function measuring observation misfit and prior consistency. In this work, we follow the same unrolled philosophy but simplify the mechanism: instead of predicting gradients of the cost function, we learn direct updates on the state itself. This keeps the iterative structure alternating between observation-consistency and prior-based correction while removing the need for adjoint computations. Experiments on real multi-satellite altimetry data evaluated against independent SARAL/AltiKa tracks show that the simplified model achieves 1.5% lower RMSE (0.0373 m vs. 0.0378 m) and 9 km better effective resolution (205 km vs. 214 km) compared to full 4DVarNet, demonstrating that direct state prediction can outperform gradient-based optimization.

BACKGROUND

Mapping sea surface height (SSH) from satellite altimetry is fundamental to resolving mesoscale and submesoscale ocean dynamics (10–100 km), which strongly influence ocean currents and mediate heat and carbon transport relevant for climate projections. Traditional variational data assimilation (4DVar) achieves optimal state estimates through iterative minimization of observation-prior cost functions, but gradient-based solvers impose computational burdens.
Non-unrolled deep learning models such as U-Net Ψ [4] can perform SSH reconstruction by learning a direct mapping from observations $\mathbf{y}$ to the state $\mathbf{x}$: $\mathbf{x}^* = \Psi(\mathbf{y})$. They achieve competitive SSH reconstructions without explicit variational structure, trading physical consistency guarantees for computational efficiency [3].
The 4DVarNet architecture [1] revolutionized this paradigm by embedding trainable neural modules within variational loops, learning both dynamical priors Φ and assimilation operators end-to-end. In its original form, the model starts from an initial state $\mathbf{x}^0$ and performs several unrolled gradient-based updates of the variational cost $J(\mathbf{x}) = J_{\text{obs}}(\mathbf{x}, \mathbf{y}) + J_{\text{prior}}(\mathbf{x}; \Phi)$, where J_{obs} is the observation term measuring the misfit between the predicted state and observations, and J_{prior} is the prior term encoding dynamical consistency through the learned model Φ. This leads to $\mathbf{x}^{(k+1)} = \mathbf{x}^{(k)} - \alpha \nabla_{\mathbf{x}} J(\mathbf{x}^{(k)})$, which requires backpropagation through the dynamical prior and the observation operators at each of typically 5–10 iterations.

AIM

This work pursues two objectives: (1) Develop an unrolled optimization model that replaces the gradient-based solver with a direct state-increment predictor Ψ, applied iteratively to the state:

$$\Delta \mathbf{x}^{(k)} = \Psi(\mathbf{x}^{(k)}, \mathbf{y}_{\text{obs}}), \quad \mathbf{x}^{(k+1)} = \mathbf{x}^{(k)} + \Delta \mathbf{x}^{(k)}, \quad k = 0, \ldots, K-1,$$

where each iteration can be seen as a separate U-Net (or neural module) operating directly on the state and observations. (2) Benchmark this simplified architecture against full 4DVarNet using multi-scale metrics.

METHODS

We build upon the 4DVarNet framework introduced by Fablet et al. [1], relying on its publicly available implementation on GitHub (4dvarnet-global-mapping). The original solver iteratively invokes a lower model (UNet) to predict $\nabla_{\mathbf{x}} J$ for gradient-based updates. Our simplified variant replaces this with a state-increment network Ψ maintaining 4DVarNet's unrolling design but predicting $\Delta \mathbf{x} = \Psi(\mathbf{x}^k, \mathbf{y}_{\text{obs}})$ directly over 5 iterations. The architecture uses a UNet with 128 model channels, bilinear interpolation, and 0.1 dropout.
Training follows the classical OSSE strategy: models are trained on simulated nadir-track altimetry derived from the NATL60 high-resolution ($1/60°$) dataset, aggregated onto a $1/4°$ grid for consistency with the mapping setup. Observations are provided as pseudo-satellite tracks $\mathbf{y}_{\text{obs}}$ sampled on a regular grid, and the task is to reconstruct the full SSH field $\mathbf{x}$. We use the standard Ocean Data Challenge SSH protocol with training on 2010-2017 data and validation on 2018. Training uses 256×256 spatial patches with 15 timesteps, batch size 16, and combines observation MSE loss $\mathcal{L}_{\text{obs}} = \|\mathbf{H}(\mathbf{x}) - \mathbf{y}_{\text{obs}}\|^2$ with spatial gradient regularization weighted at 20:1 ratio.

* Corresponding author. E-mail: chaima.bounadra@imt-atlantique.net

For evaluation, we follow the standard two-step OSE procedure. At inference time, the trained models are applied to real multi-satellite altimetry tracks (six-satellite constellation, year 2019) to produce a full SSH reconstruction. In a second step, these reconstructions are compared against an independent reference track (SARAL/AltiKa), which is withheld during inference. Metrics include: (i) reconstruction RMSE; (ii) spatially-resolved effective resolution computed from spectral coherence; (iii) regional variance score assessment across different ocean regimes (coastal, offshore high/low variability, equatorial band, arctic regions).

RESULTS

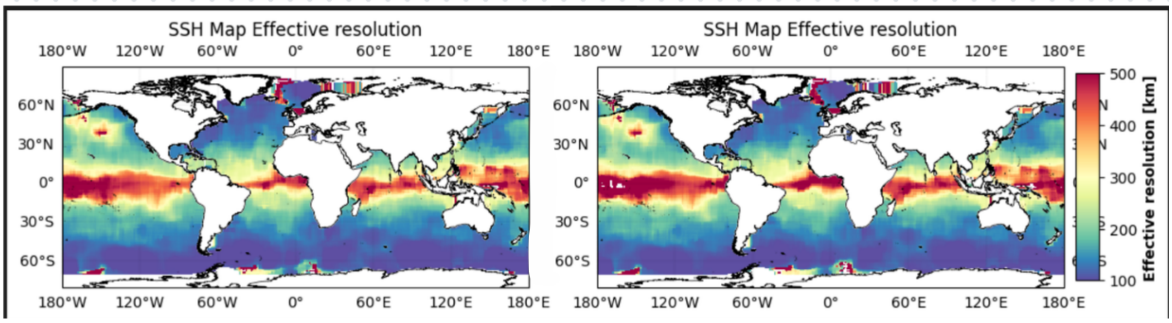

Figure 1. **Spatial distribution of effective resolution.** Left: Simplified 4DVarNet (205 km average). Right: Full 4DVarNet (214 km average). Best performance in mid-latitudes (100–150 km), degraded in tropics and poles (300–500 km).

The simplified 4DVarNet achieves RMSE of 0.037 m compared to 0.038 m for the full one, representing a 1.5% improvement. The effective resolution improves from 214 km to 205 km, indicating enhanced capability to resolve finer-scale oceanographic features (Figure 1). The spatial distribution of effective resolution shows that both methods achieve best performance (100–150 km) in mid-latitude regions where mesoscale dynamics dominate, while tropical and polar regions exhibit degraded resolution (300–500 km) due to observational constraints and complex dynamics.

Region	**Full 4DVarNet**	**Simplified 4DVarNet**
Coastal	0.736	0.744
Offshore high-var	0.948	**0.953**
Offshore low-var	0.807	0.811
Equatorial band	0.748	0.760
Arctic	0.517	0.537

Table 1. Regional variance scores of both architectures. The simplified model shows consistent improvements across all regions.

Regional variance scores (Table 1) show strong performance across diverse ocean regimes: offshore high-variability regions achieve scores exceeding 0.94, offshore low-variability regions maintain 0.81, while coastal (0.74) and equatorial band (0.75) regions show expected degradation due to increased dynamical complexity. Arctic regions exhibit reduced skill (0.52), consistent with sparse satellite coverage.

The simplified architecture's superior performance suggests that direct state prediction better captures the sparse-observation-to-SSH mapping than gradient-based optimization, while eliminating numerical instabilities from adjoint calculations and reducing error accumulation across iterations. These results demonstrate that neural networks can implicitly learn oceanographic variability patterns, enabling hybrid data-physics schemes that preserve dynamical constraints through architectural design rather than explicit optimization. Future work will incorporate temporal modules (ConvLSTM) for eddy propagation modeling and extend applications to climate model parameterizations.

References

[1] Fablet, R., Ouala, S., & Verron, J. (2021). End-to-end learning of variational data assimilation models: Application to ocean dynamics. *Journal of Advances in Modeling Earth Systems*, 13(2), e2021MS002572. https://doi.org/10.1029/2021MS002572

[2] Beauchamp, M., Febvre, Q., Georgenthum, H., & Fablet, R. (2023). 4DVarNet-SSH: end-to-end learning of variational interpolation schemes for nadir and wide-swath satellite altimetry. *Geoscientific Model Development*, 16, 2119–2147. https://doi.org/10.5194/gmd-16-2119-2023

[3] Febvre, Q., Ubelmann, C., Ponte, A., Klein, P., Sommer, J. L., & Fablet, R. (2024). Training neural mapping schemes for satellite altimetry with realistic ocean general circulation model outputs. *Journal of Advances in Modeling Earth Systems*, 16(1), e2023MS003959. https://doi.org/10.1029/2023MS003959

[4] Ronneberger, O., Fischer, P., & Brox, T. (2015). U-Net: Convolutional networks for biomedical image segmentation. In *Medical Image Computing and Computer-Assisted Intervention (MICCAI 2015)*, 234–241. Springer.