NEURIPS 2026

DC-Ocean: Deep Latent Compression for Global High-Resolution Ocean Forecasting

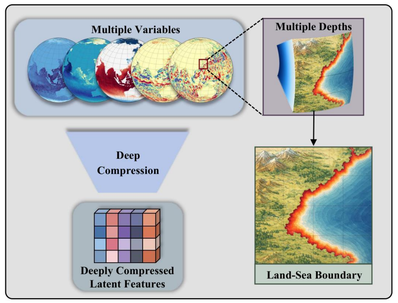
Yuxiang Li* Qiusheng Huang* Hao Li†

[1] Fudan University [2] Shanghai Innovation Institute [3] Shanghai Academy of AI for Science * Equal contribution † Corresponding author

01 Why latent compression?

High-resolution ocean forecasting is limited by data movement — not just GPU compute.

We need a compact, decodable state that represents coupled variables and depths without smoothing sharp land–sea boundaries.



1/12°
global eddy-resolving grid

105 channels
4 variables × 26 depths + SSH

28 TB → 850 GB
latent dataset storage

02 Three design choices

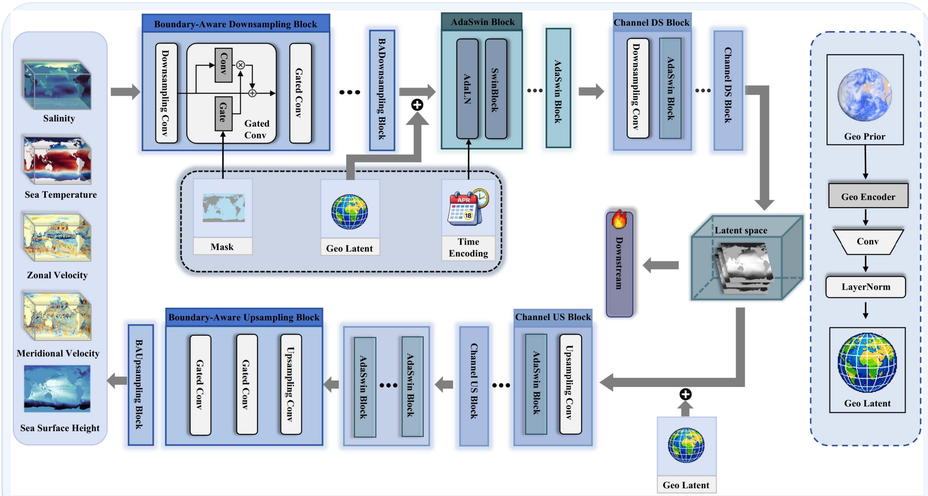
1. **Task-oriented compression**
   A compact, decodable latent for multi-variable, multi-depth ocean states.
2. **Boundary-aware representation**
   Gated convolution and spherical geodesic interpolation suppress coastal artifacts.
3. **Direct latent forecasting**
   Autoregressive rollout reduces I/O while improving long-horizon stability.

03 Experimental setup

DATA	HYCOM-RD · 6-hourly global fields Train 2006–Jun 2014 · validate Jul–Dec 2014 · test 2015
METRICS	Latitude-weighted RMSE, nRMSE and SEDI at extreme quantiles
ROLLOUT	6-hour autoregressive forecasts up to 30 days; GLORYS evaluation

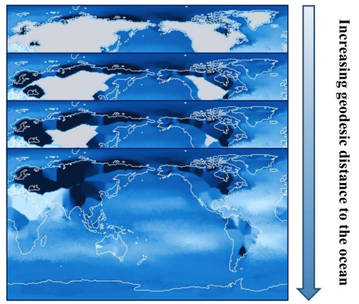
04 DC-Ocean framework

One representation connects reconstruction and downstream prediction.



BOUNDARY-AWARE ENCODER → COMPRESSED LATENT → DECODER / FORECASTER

05 Boundary-aware ocean modeling



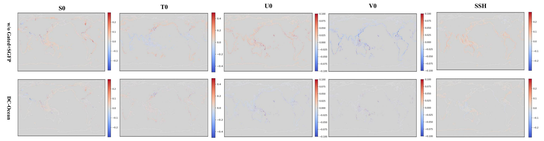
Spherical geodesic interpolation progressively fills land near valid ocean.

Fill land values geodesically
Interpolate outward from valid ocean pixels on the sphere.

Gate features with the mask
Mask-conditioned aggregation limits coastal contamination.

Inject physical context
Geographic and temporal embeddings enter through AdaLN.

BOUNDARY-REGION RECONSTRUCTION ERROR



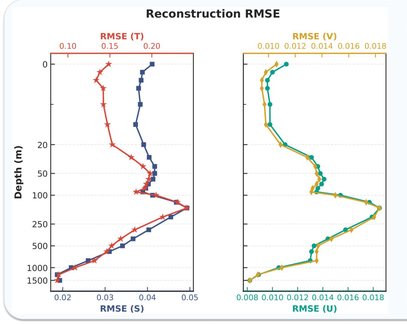
Top: without gated convolution + SGFP · Bottom: DC-Ocean

06 Reconstruction quality

24×
spatial downsampling

≈3 min
annual inference

30 days
latent rollout

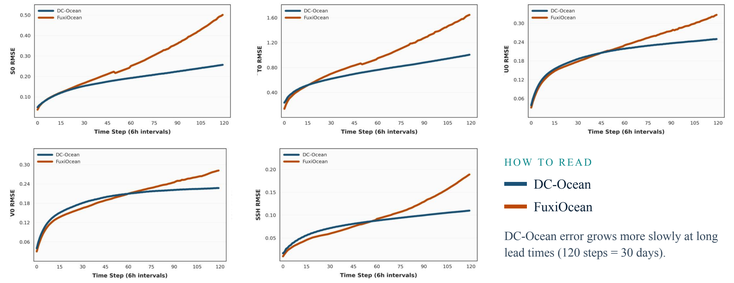


Depth (m)	S	T	U	V	SSH
0	0.0411	0.1486	0.00911	0.0088	0.0084
10	0.0383	0.1423	0.0084	0.0082	–
50	0.0417	0.1979	0.0129	0.0124	–
100	0.0412	0.2051	0.0147	0.0144	–
150	0.0493	0.2418	0.0183	0.0183	–
200	0.0455	0.2130	0.0181	0.0183	–
300	0.0403	0.1797	0.0170	0.0171	–
400	0.0366	0.1629	0.0161	0.0162	–
500	0.0341	0.1521	0.0154	0.0155	–
600	0.0310	0.1463	0.0156	0.0158	–
800	0.0260	0.1317	0.0164	0.0165	–
1000	0.0220	0.1085	0.0144	0.0143	–
1250	0.0187	0.08912	0.0127	0.0125	–
1500	0.0193	0.08724	0.0129	0.0127	–
Average	**0.0371**	**0.1645**	**0.0128**	**0.0128**	**0.0084**

Reconstruction RMSE per depth level (0–1500 m); bottom row averages all levels.

07 Latent-space forecasting

30-day autoregressive rollouts · RMSE versus lead time



HOW TO READ

DC-Ocean
FuxiOcean

DC-Ocean error grows more slowly at long lead times (120 steps = 30 days).

TAKEAWAY

Compact states. Useful forecasts.

Reusable predictive latents reduce data movement while preserving coastal detail and long-rollout stability.

81 h → 36 h
forecast training time

8.0 s → 0.42 s
I/O per snapshot