Efficient and Interpretable On-Device Polynomial Inference for Multi-Variable Ocean Profiling

Vahram A. Martirosyan*
Dept. of Computer and Information Science
Higher Colleges of Technology
Abu Dhabi, United Arab Emirates
vmartirosyan@hct.ac.ae

***Abstract*—Ocean sampling platforms—Argo profiling floats, gliders, and moored nodes—surface only every few days and must run for years on a battery, so reconstructing the temperature, salinity, and oxygen profile on board is appealing, provided it fits a tiny memory and power budget. Neural networks reach high accuracy, yet they are opaque and megabytes in size, which rules out the microcontrollers used in marine electronics. This paper describes an on-board engine built from stagewise-regularized polynomial regression (SRPR). A nine-feature rotation-invariant encoding feeds a deterministic three-way depth partition, and each zone is a single degree-6 polynomial whose coefficients can be read and checked by hand; inference is one sparse-polynomial evaluation. Pooled R^2 on IAP climatologies is 0.973 (temperature), 0.794 (salinity), and 0.895 (oxygen). We also test the model on raw EN4 profiles, under strict spatial and temporal splits, and with sensor noise, and we ablate the zone, degree, and regularization choices. A sparsified kernel runs in $\sim$1.1–4.6 ms on 32-bit microcontrollers and 190 ms on an 8-bit AVR, at $\sim$0.1–14 mJ—small next to an Argo cycle budget. The engine is offered as an auditable edge baseline, not as a claim of physical-law discovery.**

***Index Terms*—marine electronic information, TinyML, edge AI, interpretable machine learning, polynomial regression, ocean profiling, dissolved oxygen, resource-constrained embedded systems.**

I. Introduction

THE ocean's vertical structure of temperature, salinity, and dissolved oxygen (DO) governs heat transport, water-mass formation, primary productivity, and the carbon cycle, and is sampled by a sparse fleet of autonomous platforms that surface only intermittently and must run for years on battery. A central engineering problem is to *reconstruct* the full water column on-board, where computation is nearly free relative to the energy and bandwidth cost of transmitting raw profiles to shore [1], [2]. Deep learning is the dominant approach [3]–[6] but is structurally opaque and multi-megabyte. For an instrument designer two hard constraints dominate: the kernel must fit in tens to hundreds of kilobytes, and the decision rule should be auditable.

*Corresponding author.

We build on stagewise-regularized polynomial regression (SRPR) and make three contributions. First, one engine reconstructs temperature, salinity, and dissolved oxygen from a shared nine-feature Cartesian encoding, split by depth into three zones and fitted at degree six. Second, a two-stage analytic solver—ridge preconditioning, then elastic-net coordinate descent on an active set—produces sparse, readable coefficients without gradient training. Third, we port the result to four real microcontrollers and measure latency, cycle counts, footprint, quantized accuracy, and sparsity, choosing an operating point per board and comparing against iso-footprint baselines. Figure III shows the pipeline.

Elastic net, polynomial regression, and zone-wise models are not new [11], [12]. What we contribute is a deployment-oriented formulation: the depth zones share one encoding, the two-stage solver fits 5004-term bases without gradients, and the artifact carries an explicit fixed-point and footprint budget. Sparse identification and symbolic regression [8]–[10] instead search for governing equations and yield models that are hard to bound for a fixed embedded budget. Our ablations (Section VI) justify each design choice. All results concern monthly climatological means; we do not claim that a single coefficient encodes a dynamical law.

II. Related Work

Deep networks reconstruct subsurface temperature and salinity from satellite or sparse observations with high accuracy [3]–[6], but their parameter counts are megabytes to gigabytes and their rules are opaque. TinyML brings inference to microcontrollers [7], yet is dominated by compressed neural inference. Sparse identification and symbolic regression [8]–[10] recover compact closed-form models, but target the discovery of dynamical equations and produce non-nested models that are hard to bound for a fixed embedded budget. SRPR occupies the intersection: a closed-form, sparse, auditable regressor explicitly sized for a microcontroller.

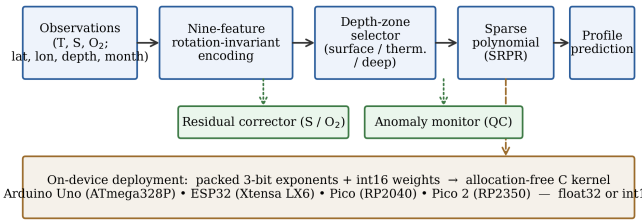


Figure 1. System model of the on-device SRPR pipeline: raw observations are mapped to a nine-feature rotation-invariant encoding, a deterministic depth zone selects one degree-6 sparse polynomial, and the profile prediction is produced on the microcontroller. Optional residual-correction and anomaly-monitoring stages share the same coefficient file.

III. Method

A. Rotation-invariant feature vector

Each observation (lat, lon, d, M) is encoded into nine features

$$\mathbf{x}^\top = \big(\phi,\ x,\ y,\ z,\ \ln d,\ \sin\tfrac{2\pi M}{12},\ \cos\tfrac{2\pi M}{12},\ |\phi|,\ d^{-1}\big), \quad (1)$$

where ϕ is latitude (rad), (x, y, z) are Cartesian coordinates on the unit sphere, d is depth, and M the month. The Cartesian map removes the polar singularity and dateline discontinuity; $\ln d$ and d^{-1} let low-order terms represent near-surface decay; $|\phi|$ encodes hemispheric symmetry. Features are standardized (training statistics only).

B. Stagewise-regularized polynomial regression

For $D = 9$ and degree $K = 6$ the monomial basis has $\binom{D+K}{K} = 5004$ terms (the constant folded into normalization). We fit

$$\hat{\boldsymbol{\beta}} = \arg\min_{\boldsymbol{\beta}} \ \frac{1}{2n}\sum_{i=1}^{n}\big(y_i - \boldsymbol{\phi}(\mathbf{x}_i)^\top\boldsymbol{\beta}\big)^2 + \lambda_1\|\boldsymbol{\beta}\|_1 + \tfrac{\lambda_2}{2}\|\boldsymbol{\beta}\|_2^2, \quad (2)$$

in two analytic stages: (**1**) a ridge warm start by Cholesky / a memory-bounded streaming Gram accumulation, then (**2**) elastic-net coordinate descent [13] with soft-thresholding over an active set. The result is a sparse, directly readable coefficient vector.

C. Depth zones and on-device kernel

A single global polynomial must spend capacity on all regimes, so we partition the column into surface ($d < 200$ m), thermocline ($200 \le d < 1000$ m), and deep ($d \ge 1000$ m), fitting one degree-6 model per zone; zone selection is a single depth comparison. Inference is one sparse evaluation $\hat{y} = \sum_{j\in\mathcal{A}} \beta_j \phi_j(\mathbf{x})$ with no matrix multiply, tree, or activation, giving a fixed worst-case cost (Algorithm III-C).

Algorithm 1 On-device scalar SRPR inference kernel.

in: ϕ, λ, d, M; coefficients $\{\beta_j\}$/zone
$x \leftarrow \cos\phi\cos\lambda$; $y \leftarrow \cos\phi\sin\lambda$; $z \leftarrow \sin\phi$
$f \leftarrow [\phi, x, y, z, \ln d, \sin(2\pi M/12),$
  $\cos(2\pi M/12), |\phi|, 1/(d+1)]$
normalize f with stored mean/variance
$zone \leftarrow 0$ **if** $d < 200$ **else** (1 **if** $d < 1000$ **else** 2)
$acc \leftarrow 0$
for each active term j in zone $zone$ **do**
  $acc \leftarrow acc + \beta_j \cdot \mathrm{monomial}_j(f)$
end for
return acc (denormalize with stored target statistics)

IV. Data and Validation Protocol

The IAP T/S Climatology 2024 [17] provides monthly temperature and salinity on a 1° grid to 6000 m (534,090 samples, ∼31,596 profiles); the IAP gridded DO product [18] provides 55,204 samples over the same grid. Raw, ungridded EN.4.2.2 profiles [14] (2010–2011, 431,795 QC'd level samples) support raw-profile and temporal-year holdouts; WOA18 [16] is an independent reference only. Table IV states the role, coverage, and scale of each product: DO is far sparser than T/S over the same grid, so DO is reported separately and never pooled with T/S. Folds are stratified by profile; sampling and standardization use training folds only, zone boundaries are fixed a priori, and river/grid correction parameters are fit on training-fold/year residuals only, so the reported gains are not leakage-inflated.

Table 1
Datasets: role, coverage, and scale.

Product	Var.	Sampling	Depth	N
IAP T/S 2024	T,S	1°, monthly	0–5000 m	534,090
IAP DO	DO	1°, monthly	0–5000 m	55,204
EN.4.2.2	T,S	QC stations 2010–11	0–5000 m	431,795
WOA18	T,S,DO	1°, monthly	0–5500 m	grid (ref.)

V. Reconstruction Accuracy

All models use 5-fold CV (profile-stratified). Table V summarizes pooled accuracy. Temperature reaches $R^2 = 0.973$ (deep is easiest, 0.967; thermocline hardest, 0.926); salinity is weakest at the surface (0.781) because river plumes and ice melt create sharp freshwater discontinuities a smooth basis represents poorly, while thermocline and deep zones remain strong. Dissolved oxygen reaches $R^2 = 0.895$ with a single global degree-6 model, since the DO field is horizontally smoother.

Strict evaluation. Random-sample CV is optimistic for spatially correlated fields, so we add spatial-block CV ($20° \times 20°$ blocks held out) and repeated 5-fold CV over five seeds, and report the *anomaly* R^2 relative to the (latitude, depth, month) mean. Table V shows the effect the review flags: part of the raw climatological R^2 reflects spatial smoothness rather than transferable skill, and unseen-region salinity is the true weak point.

Raw EN4 profiles and temporal holdout. Table V reports (A) the climatology model applied directly to raw profiles, (B) random 5-fold CV trained on EN4, and (C) a year-to-year holdout. Temperature survives the climatology-to-raw transfer, but salinity does not; training on raw profiles restores skill and generalizes across years. The practical conclusion is that climatology-trained coefficients should not be deployed on

Table 2
POOLED RECONSTRUCTION ACCURACY (DEGREE 6, 5-FOLD CV).

Variable	RMSE	R^2	Macro R^2
Temperature (°C)	1.595	0.973	0.954
Salinity (psu)	0.740	0.794	0.855
Dissolved O_2 (μmol kg^{-1})	30.74	0.895	—

Table 3
STRICT EVALUATION (IAP SPARSE CLIMATOLOGY, DEGREE 6). RAW AND ANOMALY R^2.

Protocol	T R^2	S R^2
Random 5-fold	0.974	0.761
Repeated 5-fold (5 seeds)	0.975±0.002	0.798±0.006
Spatial 20° blocks	0.946	0.626
Anomaly R^2 (random)	0.767	0.581
Anomaly R^2 (spatial)	0.518	0.340

raw sensor data without raw-profile fine-tuning or a residual corrector. A compact 3° grid correction fit on one year of EN4 residuals lifts held-out salinity from 0.11 to 0.833, i.e. the polynomial is the backbone and the learned correction carries raw-data salinity skill.

Sensor noise. Table V perturbs the nine feature channels with Gaussian noise. The degree-6 monomial basis amplifies input noise through high-order terms, so salinity degrades before temperature; a lower-degree model or an input-denoising step is preferable when measurements are noisy. **Cross-product check.** Against an independent DO product (WOA18) the oxygen model reaches $R^2 = 0.707$ but anomaly $R^2 = 0.04$—barely above a climatology-mean baseline—so oxygen, like salinity, requires a correction or fine-tuning stage before raw-data claims. Temperature is deployable as-is.

VI. ABLATION STUDIES

We isolate the main design choices under the same solver and folds (Table VI). *Zones:* the three-zone decomposition improves temperature R^2 from 0.963 to 0.973 and salinity from 0.769 to 0.788 at no extra inference cost. *Degree:* accuracy rises monotonically with degree in every zone with diminishing returns beyond degree 4, while the elastic net retains only a fraction of the 5004 terms (e.g. deep temperature $923/5004$). *Regularization:* increasing λ_1 from 0 to 10 shrinks the active set (deep T $5004 \to 1735$) with R^2 changing by at most 0.004, giving a usable sparsity/accuracy dial.

VII. EMBEDDED DEPLOYMENT AND EFFICIENCY

A sparsified degree-6 kernel was ported to four real microcontrollers and measured against the host (exact float32 agreement); Figure VII and Table VII summarize the results. Table VII then compares, on each board, float32 and int16 fixed-point in wall-clock latency and exact cycles. Cycle counts are hardware counters on the ESP32 (CCOUNT) and the Pico 2 (DWT CYCCNT); the AVR and Pico (RP2040) values are derived from the microsecond timer, since those cores expose no CPU cycle counter.

Table 4
RAW EN4 VALIDATION (2010–2011, DEGREE 4) AND SALINITY CORRECTION.

Protocol	T R^2	S R^2
(A) IAP climatology → EN4 raw	0.919	-0.46
(B) EN4 random 5-fold	0.955	0.577
(C) EN4 year 2010→2011	0.939	0.453
(C) EN4 year 2011→2010	0.947	0.501
S + 3° grid (2010→2011)	—	**0.833**

Table 5
SENSOR-NOISE ROBUSTNESS (FEATURE PERTURBATION, RANDOM 5-FOLD).

Noise (σ)	T R^2	S R^2
0.000	0.974	0.761
0.005	0.959	0.606
0.010	0.950	0.496
0.050	0.804	−0.764

The result is core-specific. On the FPU-less 32-bit RP2040, int16 is 2.8× faster (4.64 vs. 1.64 ms) because integer arithmetic beats software float. On the FPU-equipped cores, float32 wins—clearly on the ESP32 (0.88 vs. 1.03 ms) and modestly on the Pico 2 (1.12 vs. 1.21 ms). On the 8-bit AVR both formats are software float and a near-tie (190.6 vs. 197.7 ms). There is therefore no universal "use int16" rule: deploy float32 on FPU cores and int16 on FPU-less 32-bit cores. First-order energy estimates ($E = VIt$ with typical active current) span ∼150× —14.3 mJ (Uno), 96 μJ (Pico 2), 180 μJ (ESP32)—but all remain negligible against an Argo per-cycle (kJ) budget, so on-board inference is not the energy bottleneck.

Quantization accuracy. Int16 is essentially lossless for temperature ($R^2 = 0.968$ unchanged) and salinity (0.731 unchanged) while halving coefficient storage, whereas int8 collapses temperature ($R^2 \to 0.622$); int16 is therefore the recommended fixed-point format.

Iso-footprint baselines. At ∼1.4 KB an int8 MLP(32,32) reaches $R^2 = 0.988$ versus 0.964 for the degree-4 polynomial, and a 0.2 KB int8 MLP(16) matches degree-6 temperature accuracy. We report this as a deliberate boundary result: on a smooth ocean field a tiny quantized network is competitive per kilobyte, so the polynomial's case rests on *auditability and physics-checkability*, not on accuracy or size.

VIII. TRANSPARENCY AND AUDITABILITY

Because every coefficient is explicit, a prediction can be decomposed and audited term-by-term—a property unavailable in tree ensembles or neural networks. The largest temperature weights are absolute latitude $|\phi|$ and log-depth $\ln d$, qualitatively consistent with the dominant solar-latitude and thermocline structure. The basis is highly collinear and features are transformed, so term-level associations are *illustrative, not identifiable* under a single weight. The engineering value is auditability: an operator can see which terms drove an anomalous prediction and confirm that they are physically plausible, without claiming that any coefficient recovers a

Table 6
ABLATIONS: ZONE DECOMPOSITION (DEG 4), DEGREE BUDGET (SURFACE), AND REGULARIZATION (DEEP).

(a) Zone decomposition (global vs. 3-zone, deg 4)				
Model	**T R^2**	**T RMSE**	**S R^2**	**S RMSE**
Global	0.963	1.527	0.769	0.608
3-zone	**0.973**	**1.317**	**0.788**	**0.581**
(b) Degree / active terms (surface zone)				
Deg	**Active/Full**	**R^2 (T)**		**R^2 (S)**
2	52/54	0.897		0.621
4	521/714	0.953		0.737
6	2363/5004	0.966		0.793
(c) Regularization λ_1 (deep zone, deg 6)				
λ_1	**Active**	**R^2 (T)**		**R^2 (S)**
0	5004	0.974		0.915
10	1735	0.970		0.908

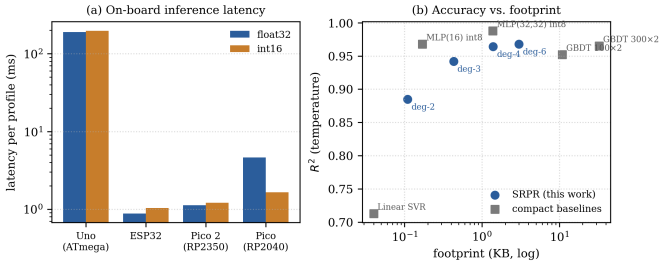


Figure 2. Performance comparison. (a) Per-profile latency of the sparsified degree-6 kernel on four boards, float32 versus int16; the FPU-less Pico (RP2040) favours int16, whereas the FPU cores favour float32. (b) Accuracy versus footprint for the polynomial family and for compact baselines at matched budgets.

dynamical law. Figure VIII shows that residuals are near-Gaussian and centred near zero, with heavy tails at fronts, upwelling, and river/ice zones.

IX. CONCLUSION

We presented a hardware-oriented, auditable polynomial inference engine for multi-variable ocean profiling on constrained marine platforms. It reconstructs temperature, salinity, and dissolved oxygen from a shared rotation-invariant encoding (pooled R^2 0.973/0.794/0.895), runs in microseconds on a scalar CPU and $\sim$1–1.5 ms on 32-bit microcontrollers, and exposes every coefficient for operator verification. Strict spatial/temporal and raw-EN4 validation show that temperature transfers directly while salinity and oxygen require a learned correction stage; ablations confirm the zone, degree, and regularization choices; and real-hardware measurements plus iso-footprint baselines delimit the deployment envelope. Limitations: all results concern monthly climatological means, per-coefficient attribution is not identifiable under collinearity, and the energy figures are first-order estimates.

ACKNOWLEDGMENT

The author thanks Laoshan Laboratory and the Institute of Atmospheric Physics for the IAP T/S Climatology 2024 and gridded dissolved-oxygen products, and the Argo community for the observing-system context.

Table 7
REAL-BOARD COMPARATIVE EXECUTION (SPARSE DEGREE-6; K=1500 ON THE 8-BIT AVR, K=1200 ON THE 32-BIT PARTS): MEASURED FLOAT32 VS INT16 LATENCY AND CYCLES.

		float32		**int16**	
Board	**Core (FPU)**	**ms**	**cyc**	**ms**	**cyc**
ESP32	LX6 (yes)	0.88	211,238	1.03	247,479
Pico 2	M33 (yes)	1.12	168,322	1.21	181,430
Pico	M0+ (no)	4.64	927,636	1.64	328,482
Arduino Uno	ATmega (no)	190.6	3,050,086	197.7	3,163,443

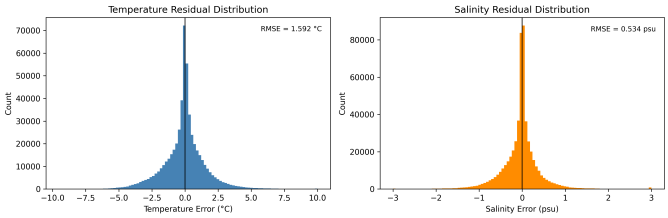


Figure 3. Residual distributions are near-Gaussian and centred near zero; heavy tails correspond to fronts, upwelling, and river/ice-melt zones.

DECLARATION OF AI AND AI-ASSISTED TECHNOLOGIES IN THE WRITING PROCESS

During the preparation of this work the author(s) used OpenCode (DeepSeek-Flash), an AI assistant, to create the initial draft of the manuscript. After using this tool/service, the author(s) reviewed and edited the content as needed and take(s) full responsibility for the content of the publication. All figures were generated from model and system output using Python code and were not created or altered by generative AI.

REFERENCES

[1] G. C. Johnson, S. Hosoda, S. R. Jayne, P. R. Oke, S. C. Riser, D. Roemmich *et al.*, "Argo—Two Decades: Global Oceanography, Revolutionized," *Annual Review of Marine Science*, vol. 14, pp. 379–403, 2022. https://doi.org/10.1146/annurev-marine-022521-102008

[2] A. P. S. Wong, S. E. Wijffels, S. C. Riser, S. Pouliquen, S. Hosoda, D. Roemmich *et al.*, "Argo Data 1999–2019: Two Million Temperature-Salinity Profiles and Subsurface Velocity Observations From a Global Array of Profiling Floats," *Frontiers in Marine Science*, vol. 7, p. 700, 2020. https://doi.org/10.3389/fmars.2020.00700

[3] H. Su, J. Jiang, A. Wang, W. Zhuang and X. H. Yan, "Subsurface Temperature Reconstruction for the Global Ocean from 1993 to 2020 Using Satellite Observations and Deep Learning," *Remote Sensing*, vol. 14, no. 13, pp. 3198, 2022. https://doi.org/10.3390/rs14133198

[4] R. Lguensat, P. Tandeo, P. Ailliot, M. Pulido and R. Fablet, "The Analog Data Assimilation," *Monthly Weather Review*, vol. 145, no. 10, pp. 4093–4107, 2017. https://doi.org/10.1175/MWR-D-16-0441.1

[5] J. Zhang, X. Zhang, X. Wang, P. Ning and A. Zhang, "Reconstructing 3D ocean subsurface salinity (OSS) from T–S mapping via a data-driven deep learning model," *Ocean Modelling*, vol. 184, p. 102232, 2023. https://doi.org/10.1016/j.ocemod.2023.102232

[6] T. Tian, L. Cheng, G. Wang, J. Abraham, W. Wei, S. Ren *et al.*, "Reconstructing ocean subsurface salinity at high resolution using a machine learning approach," *Earth System Science Data*, vol. 14, no. 11, pp. 5037–5060, 2022. https://doi.org/10.5194/essd-14-5037-2022

[7] Y. Abadade, A. Temouden, H. Bamoumen, N. Benamar, Y. Chtouki and A. S. Hafid, "A Comprehensive Survey on TinyML," *IEEE Access*, vol. 11, pp. 96892–96922, 2023. https://doi.org/10.1109/ACCESS.2023.3294111

[8] S. L. Brunton, J. L. Proctor and J. N. Kutz, "Discovering governing equations from data by sparse identification of nonlinear dynamical systems," *Proceedings of the National Academy of Sciences*, vol. 113, no. 15, pp. 3932–3937, 2016. https://doi.org/10.1073/pnas.1517384113

[9] A. A. Kaptanoglu, B. M. de Silva, U. Fasel, K. Kaheman, A. J. Goldschmidt, J. Callaham *et al.*, "PySINDy: A comprehensive Python package for robust sparse system identification," *Journal of Open Source Software*, vol. 7, no. 69, p. 3994, 2022. https://doi.org/10.21105/joss.03994

[10] M. Cranmer *et al.*, "Discovering symbolic models from deep learning with inductive biases," in *Proc. NeurIPS*, 2020. https://arxiv.org/abs/2006.11287

[11] H. Zou and T. Hastie, "Regularization and Variable Selection Via the Elastic Net," *Journal of the Royal Statistical Society Series B: Statistical Methodology*, vol. 67, no. 2, pp. 301–320, 2005. https://doi.org/10.1111/j.1467-9868.2005.00503.x

[12] R. Tibshirani, "Regression Shrinkage and Selection Via the Lasso," *Journal of the Royal Statistical Society Series B: Statistical Methodology*, vol. 58, no. 1, pp. 267–288, 1996. https://doi.org/10.1111/j.2517-6161.1996.tb02080.x

[13] J. Friedman, T. Hastie and R. Tibshirani, "Regularization Paths for Generalized Linear Models via Coordinate Descent," *Journal of Statistical Software*, vol. 33, no. 1, 2010. https://doi.org/10.18637/jss.v033.i01

[14] S. A. Good, M. J. Martin and N. A. Rayner, "EN4: Quality controlled ocean temperature and salinity profiles and monthly objective analyses with uncertainty estimates," *Journal of Geophysical Research: Oceans*, vol. 118, no. 12, pp. 6704–6716, 2013. https://doi.org/10.1002/2013JC009067

[15] A. Gholami, S. Kim, Z. Dong, Z. Yao, M. W. Mahoney and K. Keutzer, "A Survey of Quantization Methods for Efficient Neural Network Inference," in *Low-Power Computer Vision*, Chapman and Hall/CRC, 2022, pp. 291–326. https://doi.org/10.1201/9781003162810-13

[16] T. P. Boyer *et al.*, "World Ocean Atlas 2018," NOAA National Centers for Environmental Information. [Online]. Available: https://www.ncei.noaa.gov/products/world-ocean-atlas

[17] L. Cheng *et al.*, "IAP ocean temperature and salinity climatology (2024)," Institute of Atmospheric Physics / Laoshan Laboratory. [Online]. Available: http://www.ocean.iap.ac.cn/

[18] "IAP gridded ocean dissolved oxygen dataset (1960–2025)," Institute of Atmospheric Physics / Laoshan Laboratory, 2025. [Online]. Available: http://www.ocean.iap.ac.cn/

[19] J. D. Sharp, A. J. Fassbender, B. R. Carter, P. D. Lavin and G. C. Johnson, "GOBAI-O2: temporally and spatially resolved fields of ocean interior dissolved oxygen over nearly 2 decades," *Earth System Science Data*, vol. 15, no. 10, pp. 4481–4518, 2023. https://doi.org/10.5194/essd-15-4481-2023

[20] L. Cheng, J. Abraham, K. E. Trenberth, J. Fasullo, T. Boyer, M. E. Mann *et al.*, "Another Year of Record Heat for the Oceans," *Advances in Atmospheric Sciences*, vol. 40, no. 6, pp. 963–974, 2023. https://doi.org/10.1007/s00376-023-2385-2