Reconstructing High-Resolution Ocean Currents via Multistage Technique Applied to a Multiscale Neural Network

Reza Arefidamghani
Department of Applied Mathematics
Federal University of Rio de Janeiro
Rio de Janeiro, Brazil
reza.arefi@matematica.ufrj.br

Hamidreza Anbarlooei
Department of Applied Mathematics
Federal University of Rio de Janeiro
Rio de Janeiro, Brazil
hamidreza@im.ufrj.br

***Abstract*—High-resolution (HR) ocean current data are crucial for resolving submesoscale structures, turbulent energy cascades, and other fine-scale processes that drive oceanic physical and biogeochemical dynamics. However, generating HR outputs using numerical models like the Hybrid Coordinate Ocean Model (HYCOM) is computationally intensive, particularly over large spatial domains or extended time periods. Additionally, satellite and in-situ observations often lack the spatial and temporal resolution necessary to directly capture these features.**

To address these limitations, this study investigates deep learning-based super-resolution (SR) techniques as a cost-effective post-processing solution for reconstructing HR ocean surface velocity fields from low-resolution (LR) numerical simulations. Using synthetically downsampled HYCOM simulations over the Gulf of Mexico as a testbed, we adopt the Downsampled Skip-Connection Multi-Scale (DSC/MS) network as a baseline and benchmark its performance against traditional interpolation methods such as bilinear and bicubic upsampling.

To enhance reconstruction quality, two improvement strategies have been propose: Residual Learning, where the network learns to predict high-frequency residuals that are added to the initial SR output, and Multistage Refinement, where successive SR models iteratively improve the predictions to recover finer-scale structures.

Model performance is assessed using standard image-quality metrics (MSE, PSNR, SSIM, FSIM) and physics-informed diagnostics, including spectral energy distribution and turbulent kinetic energy (TKE) recovery. Results show that the multistage refinement strategy significantly outperforms both interpolation and single-stage SR models, yielding reconstructions that are not only visually sharper but also more physically consistent.

This study highlights the potential of deep super-resolution architectures as efficient tools for enhancing ocean model outputs, enabling more accurate and physically meaningful HR reconstructions without the computational cost of running full-resolution simulations.

***Index Terms*—component, formatting, style, styling, insert**

I. Introduction

A. Motivation

Capturing fine-scale ocean dynamics, such as submesoscale eddies, sharp fronts, and anisotropic shear flows, is critical for improving ocean state estimation, regional forecasting, and climate system modeling. These structures govern key physical and biogeochemical processes, including energy transfer, nutrient transport, and pollutant dispersion. However, resolving such features in numerical simulations or satellite observations remains a major challenge due to computational limitations and sensor resolution constraints [1]–[3].

High-resolution (HR) ocean simulations, such as those produced by the Hybrid Coordinate Ocean Model (HYCOM) [10], [11], are computationally expensive, especially when applied over large spatial domains or extended time periods. As a more affordable alternative, deep learning–based super-resolution (SR) techniques have gained attention as post-processing tools that reconstruct HR fields from low-resolution (LR) inputs with minimal additional cost [4], [5]. These techniques are increasingly applied in Earth system modeling, where surrogate learning methods can enhance spatial resolution without modifying the underlying numerical models.

B. Challenges in Oceanographic Super-Resolution

Despite their potential, SR methods face unique challenges in the context of oceanographic data. Unlike natural images, ocean currents are governed by physical laws that give rise to structured, multiscale variability in both space and energy. Conventional SR techniques often focus on pixel-wise accuracy or visual enhancement, but they may fail to preserve physically meaningful properties such as energy spectra, spatial correlations, or conservation laws [6], [7].

Turbulent flows, in particular, follow well-characterized kinetic energy spectra, such as the $-5/3$ Kolmogorov scaling law in the inertial range [8], [9]. When SR outputs do not reproduce the correct spectral behavior, the reconstructed fields may look plausible yet be dynamically inconsistent, reducing their usefulness in downstream applications like data assimilation, transport prediction, or diagnostics of submesoscale variability.

C. Contributions of This Work

This study presents a deep learning-based framework for super-resolving ocean surface velocity fields using HYCOM data from the Gulf of Mexico, downsampled by a factor of 4.

The goal is to reconstruct high-resolution velocity fields from low-resolution inputs while maintaining both spatial detail and physical realism, particularly in terms of spectral energy content. We adopt the Downsampled Skip-Connection Multi-Scale (DSC/MS) network as a strong baseline architecture and investigate two distinct enhancement strategies:

- **Residual Super-Resolution:** A residual learning framework where the network is trained to predict the residual (i.e., the difference) between the ground truth HR field and the initial output of a first-stage SR model. This approach is more effective than using simple bicubic upsampling as a reference, as it directly targets the components that the first model failed to recover.
- **Multistage Refinement:** A cascaded model architecture where multiple DSC/MS networks are applied sequentially, with each stage refining the output of the previous one. This strategy aims to iteratively reconstruct finer-scale features.

These methods are compared against standard interpolation techniques (bilinear, bicubic) as well as the single-stage DSC/MS model. Evaluation is conducted using both traditional image-quality metrics (MSE, PSNR, SSIM, FSIM) and physics-informed diagnostics, such as kinetic energy spectra and turbulent kinetic energy (TKE) recovery.

Our results show that the multistage refinement strategy offers superior performance in both spatial accuracy and spectral consistency compared to the other approaches. The findings highlight the potential of tailored SR models to produce physically consistent reconstructions for oceanographic applications.

D. Paper Organization

The remainder of this paper is organized as follows:

- **Section 2** details the data preparation process, including the HYCOM dataset, the generation of synthetic low-resolution inputs, and the evaluation metrics used to assess model performance.
- **Section 3** describes the baseline DSC/MS architecture and the design of the residual and multistage super-resolution models.
- **Section 4** presents the experimental results, offering both visual and quantitative comparisons of the super-resolution strategies.
- **Section 5** discusses the broader implications of the results and outlines potential directions for future research.
- **Section 6** concludes the paper with a summary of the main findings and contributions.

II. Data Sources and Evaluation Metrics

A. HYCOM-TSIS Reanalysis Dataset

This study employs surface ocean velocity fields from the *HYCOM-TSIS Gulf of Mexico Reanalysis Datasets*, generated using the *Hybrid Coordinate Ocean Model (HYCOM)* integrated with the Tendral Statistical Interpolation System (TSIS). Two spatial resolutions are utilized to capture different dynamical regimes: a high-resolution (HR) configuration at **1/100°** (approximately 1 km, GOMb0.01) and a low-resolution (LR) configuration at **1/25°** (approximately 4 km, GOMb0.04) [10], [11].
Both model configurations employ identical bathymetry and adopt a **hybrid vertical coordinate system** that combines isopycnal, terrain-following, and z-level layers. This system consists of 41 vertical levels, ensuring adequate resolution for upper-ocean dynamics.
The models are forced with surface atmospheric reanalyses (CFSR/CFSv2) and incorporate **open boundary conditions** derived from the global HYCOM reanalysis (GOFS 3.1). Additionally, **tidal forcing** is applied using principal constituents (e.g., M2, S2, K1) from the TPXO9 Atlas [11]. Data assimilation is conducted using TSIS, which employs a layer-wise multivariate approach to integrate satellite-derived SLA, gridded SST, and in-situ temperature/salinity profiles.
Further technical details on the assimilation framework, interpolation strategies, and observation coverage are documented in the source references [10], [11], [13], [14].

Available Variables and Study Focus

The reanalysis datasets provide hourly three-dimensional ocean state variables (e.g., temperature, salinity, velocity) as well as surface fields (e.g., sea surface height, barotropic velocities, and derived quantities). This study focuses on the horizontal surface velocity components (u, v), which are essential for evaluating mesoscale and submesoscale circulation features and the effectiveness

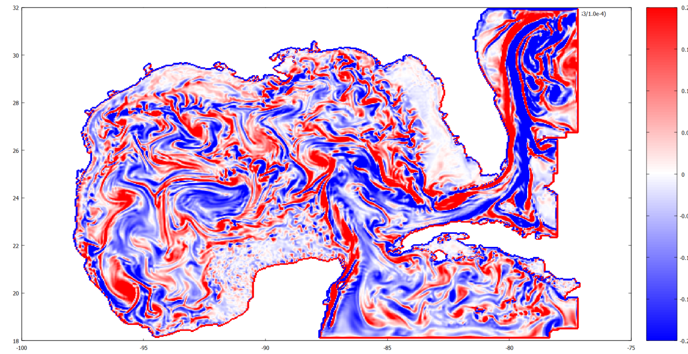

Fig. 1. Snapshot of surface relative vorticity from the HYCOM-TSIS Gulf of Mexico reanalysis, illustrating fine-scale eddy and frontal structures

TABLE I
PRIMARY VARIABLES USED FROM THE HYCOM-TSIS REANALYSIS

Variable	NetCDF Field	Temporal Res	Spatial Res
Zonal Velocity	`u`	Hourly	1/100° or 1/25°
Meridional Velocity	`v`	Hourly	1/100° or 1/25°
Sea Surface Height	`ssh`	Hourly	1/100° or 1/25°
Temperature	`temperature`	Hourly	1/100° or 1/25°
Salinity	`salinity`	Hourly	1/100° or 1/25°
Relative Vorticity	– (diagnostic)	Derived	1/100° or 1/25°

of super-resolution reconstruction. Fig. 1 visualizes fine-scale surface dynamics through the relative vorticity field.

B. Evaluation Metrics

To rigorously evaluate the reconstructed super-resolution (SR) velocity fields, we adopt a comprehensive set of evaluation metrics that span pixel-level accuracy, perceptual similarity, and physically informed turbulence characteristics. This multifaceted framework ensures both conventional and geophysically relevant assessments:

- **Mean Squared Error (MSE):** Quantifies the average squared difference between the SR prediction and the high-resolution (HR) ground truth. It is a fundamental pixel-wise metric sensitive to large errors.
- **Structural Similarity Index (SSIM):** Measures perceptual similarity by comparing local luminance, contrast, and structure. SSIM is particularly useful for assessing the preservation of spatial features and flow textures.
- **Peak Signal-to-Noise Ratio (PSNR):** Expressed in decibels (dB), PSNR evaluates the ratio of maximum possible signal to the error noise. Higher values indicate greater similarity to the ground truth.
- **Learned Perceptual Image Patch Similarity (LPIPS):** A deep learning-based metric that compares image patches via features extracted from pretrained neural networks. LPIPS correlates well with human visual perception and captures subtle differences in texture and structure.
- **Feature Similarity Index (FSIM):** Assesses perceptual fidelity based on phase congruency and gradient magnitude. FSIM is particularly suited for evaluating edge clarity and detailed flow structures.
- **Energy Preservation Index (EPI):** A physically informed metric designed to quantify how well the SR model preserves the energy spectrum across scales, which is crucial in reproducing turbulent energy cascades.
- **Turbulent Kinetic Energy (TKE):** Compares the total kinetic energy in the reconstructed and reference velocity fields. TKE provides a direct measure of how well the SR model retains physical flow intensity and dynamical consistency.

TABLE II
METRIC SUMMARY AND OPTIMAL VALUES

Metric	Optimal Val	Interpretation
MSE	0	$\downarrow$ better (exact: 0)
SSIM	1	$\uparrow$ better (exact: 1)
PSNR	∞	$\uparrow$ better ($>$ 30 dB good)
LPIPS	0	$\downarrow$ better (perceptual diff.)
FSIM	1	$\uparrow$ better (feature sim.)
EPI	1	$\uparrow$ better (energy preserved)
TKE	1	$\uparrow$ better (turb. energy)

III. Deep Learning Models for Super-Resolution of Ocean Flows

To reconstruct high-resolution (HR) ocean velocity fields from coarse-resolution inputs downscaled by a factor of 4, we employ a deep learning-based approach centered on the Downsampled Skip-Connection Multi-Scale (DSC/MS) network. This architecture is well-suited to the complex, multiscale nature of geophysical flows due to its ability to capture fine-scale features while preserving global spatial coherence. As a baseline, we benchmark against traditional upsampling techniques such as bilinear and bicubic interpolation.

Building on the DSC/MS foundation, we introduce two advanced variants to improve reconstruction fidelity: (1) Res-DSC/MS, which incorporates residual learning to recover high-frequency information, and (2) MS-DSC/MS, a multistage refinement scheme where successive models iteratively enhance reconstruction quality. The following subsections detail the core architecture and the design rationale behind each enhancement.

A. Downsampled Skip-Connection Multi-Scale (DSC/MS) Networks

The Downsampled Skip-Connection Multi-Scale (DSC/MS) network [15]–[17] is a fully convolutional encoder–decoder architecture specifically designed to reconstruct fine-scale structures in turbulent flows from low-resolution inputs. Originally proposed for fluid dynamics applications, the architecture is particularly well-suited for oceanic super-resolution tasks due to its ability to capture and reconstruct multi-scale spatial patterns inherent in geophysical turbulence.

1) Multi-Scale Feature Extraction with Parallel Convolutions: To effectively capture a broad spectrum of spatial scales from mesoscale eddies to submesoscale filaments, the encoder leverages multi-branch parallel convolutional layers, each with a different kernel size: 3×3, 5×5, and 7×7. This design enables the network to extract local, intermediate, and global contextual features in parallel, mimicking the multi-scale nature of turbulent energy cascades. This is in contrast to standard CNNs that may fail to capture scale-adaptive features critical in geophysical flow modeling.

Each convolutional branch is followed by an average pooling layer to downsample the feature maps, compressing spatial dimensions while retaining coherent flow structures. The use of average pooling (as opposed to max pooling) is intentional, preserving the overall energy and smoothness characteristics that are physically meaningful in oceanographic contexts.

2) Decoder with Skip Connections for Structural Preservation: The decoder stage symmetrically mirrors the encoder, applying progressive upsampling followed by convolutional refinement. A key feature of the DSC/MS architecture is the integration of skip connections between corresponding encoder and decoder layers. These lateral connections enable the direct transfer of high-resolution features, effectively mitigating information loss typically introduced during downsampling.

This mechanism substantially enhances the network's ability to reconstruct sharp gradients and dynamic discontinuities—such as vortex boundaries, frontal jets, and density interfaces—by preserving spatial coherence and fine-scale structures.

The model is trained using a composite loss function that combines pixel-wise accuracy with structural perception:

$$L_{\text{total}} = \alpha L_{\text{MSE}} + \beta(1 - \text{SSIM}), \tag{1}$$

where α and β are tunable weights that balance the contributions of the mean squared error (MSE) and the Structural Similarity Index Measure (SSIM). In this study, we set $\alpha = 0.8$ and $\beta = 0.2$, prioritizing pixel-level accuracy while preserving perceptual structural similarity. Refer to Figure 5 for a schematic of the hybrid DSC/MS architecture, illustrating the multi-scale feature flow and the role of skip connections in structural preservation.

B. Residual Learning with DSC/MS (Res-DSC/MS)

To better capture high-frequency details often lost during coarse gridding, we employ a residual learning framework in conjunction with the DSC/MS architecture. In this two-stage strategy, the network is trained to predict the residual that is, the difference between the ground truth high-resolution (HR) field and the initial output of a first-stage super-resolution (SR) model.

Unlike traditional approaches that rely on bicubic upsampling as a reference, this method directly targets the missing components that the initial model fails to reconstruct, allowing the network to focus on refining unresolved structures such as small-scale eddies, shear layers, and boundary gradients.

The final prediction is computed by adding the network-predicted residual to the output of the first-stage SR model, enhancing detail recovery while simplifying the learning objective. This improves generalization, particularly in dynamically complex flow regions.

The Res-DSC/MS model retains the core architectural elements of the original DSC/MS network, including multi-scale convolutional branches and skip connections. The same hybrid loss function combining mean squared error (MSE) and structural similarity index measure (SSIM), is used to ensure both pixel-wise consistency and structural fidelity.

C. Multistage Refinement with DSC/MS (MS-DSC/MS)

To further enhance reconstruction fidelity and enable the progressive recovery of fine-scale features, we employ a multistage refinement strategy, denoted as MS-DSC/MS. In this framework, a sequence of DSC/MS networks is arranged in a cascaded configuration, where each stage iteratively improves upon the prediction of the previous one.

At each refinement stage s, the model is tasked with learning the residual between the current cumulative estimate and the high-resolution (HR) ground truth:

$$R_s = Y_{\text{HR}} - \sum_{i=1}^{s-1} \hat{Y}^i, \tag{2}$$

where R_s is the residual input and $\hat{Y}^i$ denotes the output of the i-th stage. The stage-s model then predicts a corrective update:

$$\hat{Y}^s = \text{DSC/MS}_s(R_s), \tag{3}$$

and the final reconstruction is obtained by summing the outputs of all stages:

$$\hat{Y}_{\text{final}} = \sum_{s=1}^{S} \hat{Y}^s. \tag{4}$$

This residual-based multistage refinement approach departs from traditional methods by replacing simple bicubic upsampling with a sequence of predictive stages, each tasked with correcting the errors left by its predecessors. By iteratively refining the reconstruction, the model progressively recovers high-frequency details that are often missed in single-stage predictions.

Each stage employs a shared DSC/MS architecture and is optimized using a composite loss function that balances pixel-wise accuracy (MSE) with structural similarity (SSIM). Without relying on explicit spectral constraints, the design promotes spectral coherence as a natural outcome of the refinement process.

This strategy ensures stable training, enhances the recovery of fine-scale structures, and improves interpretability, making it especially effective for reconstructing the complex dynamics of turbulent oceanic flows.

IV. Results

A. Training Configuration

All deep learning models were trained under a unified framework to ensure a fair and consistent comparison. The training setup included the Adam optimizer with an initial learning rate of 1×10^{-4}, which was halved every 20 epochs. A batch size of 32 was used, and each model was trained for 100 epochs. The loss function combined Mean Squared Error (MSE) and Structural Similarity Index Measure (SSIM) in a weighted manner to balance pixel-level accuracy and perceptual quality, thereby enhancing both numerical precision and the visual fidelity of turbulent flow features. Training was conducted on an NVIDIA A100 GPU with 40 GB of VRAM. Hyperparameters were optimized through 5-fold cross-validation using the HYCOM Gulf of Mexico dataset.

B. Model Variants and Reconstruction Strategies

We evaluated three variants of the DSC/MS architecture, each designed to improve the recovery of fine-scale flow structures. The baseline **DSC/MS** leverages multiscale spatial features through skip connections and parallel convolutional pathways. The **Res-DSC/MS** model incorporates residual learning to better capture high-frequency details that may be overlooked in the baseline predictions. Finally, the **MS-DSC/MS** model employs a multistage refinement approach, where successive networks iteratively correct residual errors, enabling hierarchical reconstruction and enhanced spectral coherence. For comparison, classical bilinear and bicubic interpolation methods were also included as non-learned baselines.

C. Ablation Study: Comparative Evaluation of Methods

An ablation study was conducted to assess the impact of architectural enhancements on reconstruction performance. The following five methods were compared: bilinear interpolation, bicubic interpolation, DSC/MS, Res-DSC/MS, and MS-DSC/MS. Quantitative performance was evaluated on 1000 test samples from the Gulf of Mexico dataset, which were not seen during training. The results, summarized in Table III, reflect the average performance across these samples. Evaluation metrics are defined in subsection II-B and summarized in Table II.

These results clearly demonstrate the superiority of deep learning-based super-resolution methods over classical interpolation techniques. The baseline DSC/MS model substantially improves reconstruction accuracy by capturing multiscale spatial features characteristic of oceanic velocity fields. The Res-DSC/MS variant introduces residual learning, which enhances the recovery of fine-scale structures and sharpens turbulent features. The MS-DSC/MS model, employing multistage iterative refinement, consistently achieves the highest performance across all metrics, including physics-informed indicators such as Turbulent Kinetic Energy (TKE). This confirms its strength in delivering both numerically accurate and dynamically consistent reconstructions. Overall, these findings highlight the critical role of architectural refinement strategies in robustly modeling complex geophysical flows.

TABLE III
AVERAGE QUANTITATIVE PERFORMANCE OF SR MODELS ON 1,000 GULF OF MEXICO TEST SAMPLES

Model	MSE ↓	PSNR ↑	SSIM ↑	FSIM ↑	EDP ↓	TKE ↓
Bilinear	0.0285	27.1	0.671	0.693	0.284	0.266
Bicubic	0.0257	27.8	0.704	0.721	0.261	0.238
DSC/MS	0.0106	32.9	0.864	0.878	0.097	0.084
Res-DSC/MS	0.0089	33.8	0.884	0.894	0.071	0.061
MS-DSC/MS	**0.0081**	**34.5**	**0.902**	**0.913**	**0.054**	**0.043**

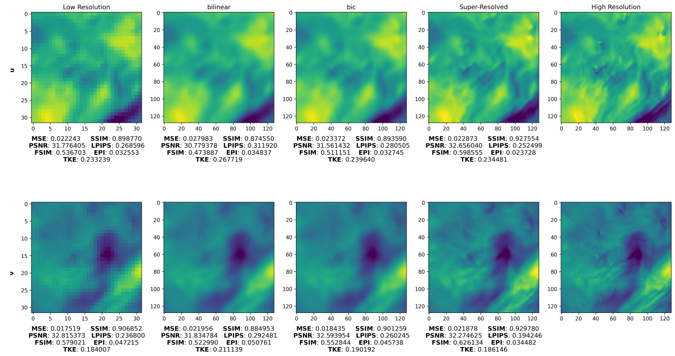

Fig. 2. Qualitative comparison of super-resolved outputs: DSC/MS models vs. interpolative baselines

D. Visual Evaluation

Fig. 2 illustrates qualitative comparisons between the DSC/MS model and classical interpolation methods, including bilinear and bicubic upsampling. Visually, interpolation methods tend to produce overly smooth outputs that blur important flow structures and fail to capture the complex spatial variability characteristic of ocean velocity fields. In contrast, the DSC/MS model demonstrates a clear advantage by preserving fine-scale features such as filaments and eddies, leading to reconstructions that are not only sharper but also more physically consistent with the underlying dynamics.

This enhanced visual fidelity is consistent with the quantitative results, where the DSC/MS model consistently surpasses interpolation baselines across pixel-wise error and structural similarity metrics. These improvements are a direct outcome of the model ability to learn complex spatial correlations and leverage multiscale features during reconstruction. The qualitative evidence reinforces the effectiveness of deep learning-based approaches in delivering high-resolution, physically meaningful reconstructions of ocean dynamics—offering a promising direction for advancing oceanographic observation and modeling capabilities.

Fig. 3 presents a comparative visualization of the DSC/MS and Res-DSC/MS models, emphasizing their ability to reconstruct ocean velocity fields. The second column displays 2D error maps, computed as the pixel-wise difference between the ground truth and the baseline DSC/MS output, while the third column shows the reconstructed velocity fields. These error maps serve as a diagnostic tool to evaluate not only the overall magnitude of reconstruction errors but also their spatial distribution and directional bias—indicating regions of consistent over- or underestimation.

Visually, the Res-DSC/MS model demonstrates notable improvements over the base DSC/MS network. The residual learning strategy effectively reduces high-frequency noise and sharpens fine-scale ocean features that are

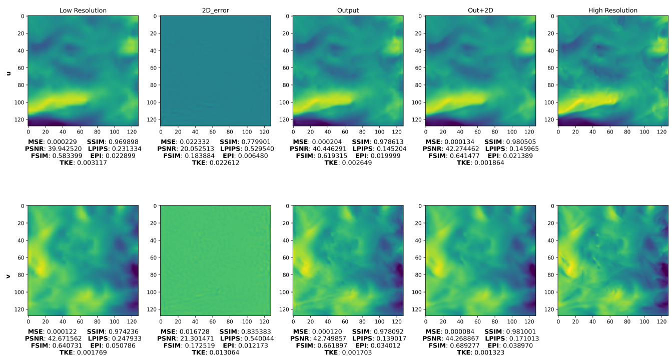

Fig. 3. Qualitative comparison of super-resolved outputs: DSC/MS vs. Res-DSC/MS

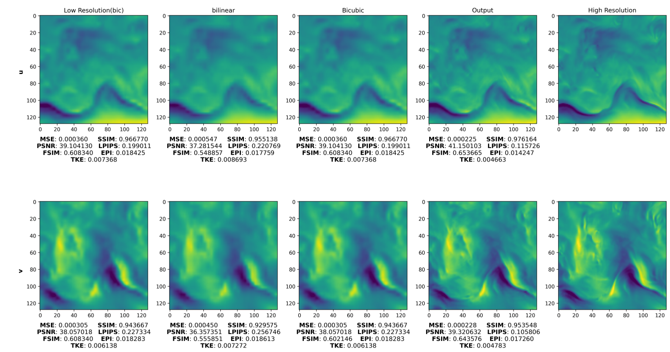

Fig. 4. Qualitative comparison of super-resolved outputs: MS-DSC/MS models vs. interpolative baselines

often attenuated in single-stage networks. Across all quantitative metrics, the Res-DSC/MS model consistently outperforms the baseline DSC/MS architecture. This enhancement is especially valuable for oceanographic applications that depend on the accurate representation of flow patterns, such as eddy tracking and transport diagnostics. Overall, the residual framework strengthens the model (DSC/MS) ability to reconstruct high-resolution ocean flows with improved physical fidelity and predictive accuracy.

Fig. 4 presents a comparative analysis of the proposed MS-DSC/MS architecture against the bicubic interpolation method and the single-stage DSC/MS model. As anticipated, MS-DSC/MS consistently outperforms both baseline models across all evaluated metrics—both pixelwise (e.g., MSE, PSNR) and perceptual (e.g., SSIM, FSIM). The visual reconstructions further confirm that MS-DSC/MS preserves sharp oceanic structures and fine-scale velocity features with higher fidelity and reduced noise. This can be attributed to the multistage refinement strategy, which enables the model to iteratively correct reconstruction errors and progressively enhance the output resolution.

Moreover, MS-DSC/MS also surpasses the performance of the Res-DSC/MS variant. This improvement is well-justified, as the multistage framework not only incorporates the benefits of residual learning but also leverages hierarchical feature correction across multiple scales. This iterative correction mechanism allows the network to better capture complex dynamics and subtle patterns that may be lost in single-pass or residual-only architectures. These promising results demonstrate the robustness and adaptability of MS-DSC/MS, making it a compelling approach for high-fidelity oceanographic data reconstruction and setting a new benchmark in learning-based super-resolution for geophysical applications.

V. Discussion and Future Work

This study demonstrates the strong potential of deep learning-based super-resolution (SR) methods for enhancing the fidelity of oceanographic data. Across both quantitative metrics and visual assessments, the proposed DSC/MS-based models significantly outperform classical interpolation techniques, delivering reconstructions that are not only more accurate but also physically consistent. Among the evaluated architectures, the multistage MS-DSC/MS model emerges as the most effective, consistently achieving superior reconstruction performance by preserving fine-scale oceanic features that are critical for downstream geophysical applications.

Beyond standard pixel-wise accuracy, we assessed the physical realism of model outputs through spectral energy analysis. These evaluations revealed that traditional interpolation methods fail to capture mid-to-high frequency components essential for representing ocean turbulence. In contrast, the DSC/MS architecture demonstrates substantial improvements in recovering these spectral features, with the Res-DSC/MS variant further enhancing fidelity through residual connections. Notably, the MS-DSC/MS model achieves the most accurate spectral reconstruction, closely matching the ground truth across a wide range of spatial frequencies.

Additional validation through turbulent kinetic energy (TKE) analysis confirms the physical integrity of the reconstructions. MS-DSC/MS most faithfully reproduces the spatial distribution and magnitude of TKE, reflecting its ability to preserve the dynamic structures and gradients inherent in ocean flows. These findings highlight the model promise for integration into physical modeling workflows, such as ocean forecasting and coupled data assimilation systems.

These results open several compelling directions for future research. A particularly promising avenue is the incorporation of physics-informed loss functions such as spectral energy constraints or gradient-based regularization to explicitly guide the model toward preserving geophysical structures across spatial scales. Furthermore, embedding these architectures within operational forecasting pipelines could enable real-time enhancement of model outputs or observational data, significantly improving the

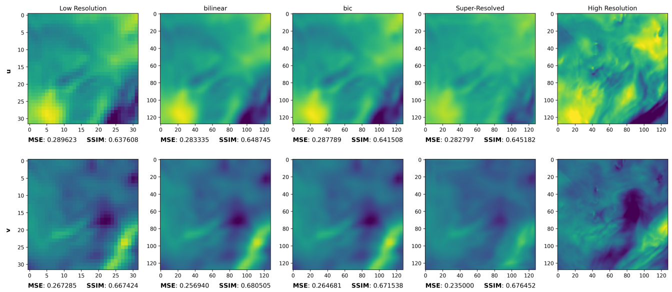

Fig. 5. SR models fail to recover fine-scale structures from unpaired LR-HR

resolution and usability of ocean monitoring products.

We also aim to evaluate the generalization capability of the proposed models across diverse ocean basins, time periods, and variable types. This includes expanding to three-dimensional (3D) volumetric super-resolution, which is critical for fully capturing ocean dynamics. Additionally, integrating domain knowledge such as conservation laws or known flow priors, into the model design may further balance data-driven learning with physical realism, enhancing both interpretability and robustness.

A. Structural Alignment of LR Inputs Using U-Net

In oceanographic applications, low-resolution (LR) and high-resolution (HR) datasets are often unpaired due to differences in spatial grids, data sources, or temporal sampling. This lack of alignment poses a significant challenge for Super-Resolution (SR) models, which typically rely on pixel-wise supervision. As shown in Fig. 5, conventional SR models trained directly on unpaired data often fail to reconstruct fine-scale features, resulting in oversmoothed or physically unrealistic outputs.

To address this issue, we propose an ongoing line of research based on a U-Net-based structural alignment module that acts as a pre-processing stage. This module learns to transform LR inputs to better match the structural characteristics of HR data, effectively generating pseudo-HR targets that make supervised SR training feasible even in the absence of perfectly paired datasets.

Design and Motivation: The alignment model employs a dense U-Net architecture with an encoder-decoder structure and skip connections. This configuration allows the model to learn global spatial mappings while preserving local structures such as eddies and frontal jets. Its capacity to model complex nonlinear transformations makes it ideal for domain adaptation tasks in geophysical datasets. Although computationally intensive, the model is trained and applied only once as a pre-processing step, making it a scalable and practical solution to mitigate data scarcity and misalignment in real-world scenarios.

This alignment strategy opens up opportunities for training SR models on broader and more heterogeneous datasets, enhancing their robustness and transferability across different oceanic regions and observational platforms. It represents a promising step toward practical deployment of deep learning-based SR methods in operational oceanography and climate monitoring.

This study provides a comprehensive assessment of deep learning-based super-resolution (SR) techniques for enhancing oceanographic data resolution, with a focus on capturing physically meaningful fine-scale structures. Through systematic evaluation against classical interpolation methods and among multiple model variants, the proposed DSC/MS-based architectures demonstrate significant advances in both accuracy and physical consistency.

We evaluated three variants of the DSC/MS architecture, each targeting improvements in the reconstruction of fine-scale oceanic features. The baseline DSC/MS model utilizes multiscale spatial feature extraction via skip connections and parallel convolutional branches. The Res-DSC/MS variant enhances this design by incorporating residual learning, which improves the reconstruction of high-frequency details often underrepresented in baseline outputs. Most notably, the MS-DSC/MS model adopts a multistage refinement strategy in which successive networks iteratively correct residual reconstruction errors. This hierarchical reconstruction process leads to progressively more accurate outputs and improved spectral coherence—without altering the base model structure.

Empirical results consistently show that MS-DSC/MS outperforms both DSC/MS and Res-DSC/MS across all evaluated metrics, including turbulent kinetic energy (TKE) recovery and spatial fidelity. The multistage strategy proves particularly effective in capturing complex dynamical patterns, making it the most reliable among the evaluated models for geophysical applications.

In addition to reconstruction accuracy, we introduce a structural alignment strategy based on a U-Net module to address the frequent issue of unpaired low- and high-resolution datasets. This pre-processing step enables more effective supervised learning under realistic data constraints, enhancing the model generalizability and applicability.

In summary, the MS-DSC/MS architecture offers a robust and physically consistent solution for super-resolving oceanographic fields. Its ability to accurately capture both fine-scale structures and broader spatial dynamics makes it a promising tool for data assimilation, forecasting, and operational oceanography. These contributions lay a strong foundation for future research in applying SR techniques across diverse geoscientific and climate domains.

References

[1] D. Roemmich and J. Gilson, “The 2004–2008 mean and annual cycle of temperature, salinity, and steric height in the global ocean from the Argo program,” *Progress in oceanography*, vol. 82, no. 2, pp. 81–100, 2009.

[2] P. Klein and G. Lapeyre, “The oceanic vertical pump induced by mesoscale and submesoscale turbulence,” *Annual review of marine science*, vol. 1, no. 1, pp. 351–375, 2009.

[3] D. Roemmich, G. C. Johnson, S. Riser, R. Davis, J. Gilson, W. B. Owens, S. L. Garzoli, C. Schmid, and M. Ignaszewski, “The Argo Program: Observing the global ocean with profiling floats,” *Oceanography*, vol. 22, no. 2, pp. 34–43, 2009.

[4] C. Dong, C. C. Loy, K. He, and X. Tang, “Image super-resolution using deep convolutional networks,” *IEEE transactions on pattern analysis and machine intelligence*, vol. 38, no. 2, pp. 295–307, 2015.

[5] C. Ledig, L. Theis, F. Huszár, J. Caballero, A. Cunningham, A. Acosta, A. Aitken, A. Tejani, J. Totz, Z. Wang, and W. Shi, “Photo-realistic single image super-resolution using a generative adversarial network,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 4681–4690, 2017.

[6] T. Vandal, E. Kodra, S. Ganguly, A. Michaelis, R. Nemani, and A. R. Ganguly, “Deepsd: Generating high resolution climate change projections through single image super-resolution,” *Proceedings of the 23rd acm sigkdd international conference on knowledge discovery and data mining*, 2020.

[7] K. Stengel, A. Glaws, D. Hettinger, and R. N. King, “Adversarial super-resolution of climatological wind and solar data,” *Proceedings of the National Academy of Sciences*, pp. 1663-1672. 2017.

[8] P. Bartello, “Geostrophic adjustment and inverse cascades in rotating stratified turbulence,” *Journal of the atmospheric sciences*, vol. 52, no. 24, pp. 4410–4428, 1995.

[9] A. N. Kolmogorov, “The local structure of turbulence in incompressible viscous fluid for very large Reynolds numbers,” *Dokl. Akad. Nauk SSSR*, vol. 30, pp. 301–305, 1941. [English trans.: *Proceedings of the Royal Society of London. Series A: Mathematical and Physical Sciences* , vol. 434, pp. 9–13, 1991.]

[10] HYCOM Consortium, “HYCOM reanalysis data,” 2020. [Online]. Available: https://www.hycom.org/

[11] J. Xie *et al.*, “T-SIS: Tendral Statistical Interpolation System for HYCOM,” 2020. [Online]. Available: https://www.hycom.org/data/goml0pt04/expt-90pt2

[12] G. D. Egbert and S. Y. Erofeeva, “Efficient inverse modeling of barotropic ocean tides,” *Journal of Atmospheric and Oceanic technology*, vol. 19, pp. 183–204, 2002.

[13] Bureau of Ocean Energy Management (BOEM), 2020. [Online]. Available: https://www.boem.gov

[14] E. Chassignet and J. Xie, “GOMb0.01 and GOMb0.04: HYCOM Gulf of Mexico reanalysis datasets,” 2020. [Online]. Available: https://www.hycom.org/data/goml0pt04/expt-90pt2

[15] M. Bode, M. Gauding, , Z. Lian, D. Denker, M. Davidovic, K. Kleinheinz, J. Jitsev, and H. Pitsch, “Using physics-informed enhanced super-resolution generative adversarial networks for subfilter modeling in turbulent reactive flows. ,” *Proceedings of the Combustion Institute*, vol. 38, pp. 2617–2625, 2021.

[16] K. Fukami, K. Fukagata, and K. Taira, “Super-resolution reconstruction of turbulent flows with machine learning,” *Journal of Fluid Mechanics*, vol. 870, pp. 106–120, 2019.

[17] Z. Li, N. Kovachki, K. Azizzadenesheli, B. Liu, , K. Bhattacharya, A. Stuart, and A. Anandkumar, “Fourier neural operator for parametric partial differential equations,” *arXiv preprint arXiv:2010.08895*.