

# Deploying Atmospheric and Oceanic AI Models on Chinese Hardware and Framework: Migration Strategies, Performance Optimization and Analysis

Yuze Sun<sup>1,2</sup>, Wentao Luo<sup>2</sup>, Yanfei Xiang<sup>1</sup>, Jiancheng Pan<sup>1,3</sup>,  
Jiahao Li<sup>1</sup>, Quan Zhang<sup>\*1</sup>, Xiaomeng Huang<sup>\*1</sup>

<sup>1</sup>Department of Earth System Science, Ministry of Education Key Laboratory for Earth System Modelling, Institute for Global Change Studies, Tsinghua University, Beijing, China

<sup>2</sup>Huawei Technologies Co., Ltd

<sup>3</sup>INSAIT, Sofia University “St. Kliment Ohridski”  
quanzhang\_zq@163.com, hxm@tsinghua.edu.cn

## Abstract

With the growing role of artificial intelligence in climate and weather research, efficient model training and inference are in high demand. Current models like FourCastNet and AI-GOMS depend heavily on GPUs, limiting hardware independence, especially for Chinese domestic hardware and frameworks. To address this issue, we present a framework for migrating large-scale atmospheric and oceanic models from PyTorch to MindSpore and optimizing for Chinese chips, and evaluating their performance against GPUs. The framework focuses on software-hardware adaptation, memory optimization, and parallelism. Furthermore, the model’s performance is evaluated across multiple metrics, including training speed, inference speed, model accuracy, and energy efficiency, with comparisons against GPU-based implementations. Experimental results demonstrate that the migration and optimization process preserves the models’ original accuracy while significantly reducing system dependencies and improving operational efficiency by leveraging Chinese chips as a viable alternative for scientific computing. This work provides valuable insights and practical guidance for leveraging Chinese domestic chips and frameworks in atmospheric and oceanic AI model development, offering a pathway toward greater technological independence.

## Code —

<https://github.com/zeaccepted/AI-ocean-model-based-on-Mindspore-optimized-for-domestic-chips-in-China>

## Introduction

In recent years, deep learning technology has demonstrated significant potential in atmospheric and oceanic sciences. By learning the spatio-temporal dynamics of climate systems, deep neural networks can efficiently model climate data, enabling high-precision predictions with reduced computational time. In atmospheric science, notable large-scale models include FourCastNet (Pathak et al. 2022), GraphCast (Lam et al. 2023), Fuxi (Chen et al. 2023b), Fengwu (Chen et al. 2023a), and Pangu-Weather (Bi et al. 2023). In oceanography, key models include AI-GOMS (Xiong et al. 2023), Xihe (Wang et al. 2024) Langya (Yang et al.

2024) and Wenhai (Cui et al. 2025). These models leverage atmospheric and oceanic reanalysis data, utilizing high-dimensional feature extraction techniques to generate accurate predictions rapidly. However, efficiently training and running these deep learning models demands substantial computational resources. Currently, Graphics Processing Units (GPUs) dominate as the primary hardware platform due to their powerful parallel computing capabilities, which are essential for deep learning. Yet, GPU-based solutions face challenges, including high costs, supply chain dependencies, and energy inefficiency for certain tasks.

Against this backdrop, Chinese chips and frameworks are gaining prominence. Chinese artificial intelligence (AI) acceleration chips, such as Sugon’s Deep Computing Unit (DCU) and Huawei’s Ascend, are becoming key platforms for deep learning training and inference due to their high cost-effectiveness and maturing software ecosystems. Using Chinese chips for model training and inference in meteorological and oceanic AI can reduce reliance on foreign hardware and advance high-performance computing and autonomous control technologies. Additionally, deep learning frameworks, core tools for large-scale model innovation, play a crucial role in development. Leading Chinese frameworks such as Huawei’s MindSpore (Chen 2021) and Baidu’s PaddlePaddle (Ma et al. 2019) can maximize computational potential and enhance efficiency when optimized for synergy with Chinese chips.

Although Chinese chips and frameworks have made significant progress in AI computing, the current software and hardware ecosystem is still primarily dominated by mainstream GPUs and their supported frameworks (e.g., PyTorch<sup>1</sup>, TensorFlow<sup>2</sup>). The Chinese deep learning ecosystem faces several key challenges compared to traditional frameworks. First, atmospheric and oceanic models are typically developed using frameworks like PyTorch, which are not fully compatible with MindSpore<sup>3</sup> regarding operators. Direct migration may lead to missing operators or reduced computational efficiency, challenging migration. Second, the hardware architecture of Chinese chips significantly

<sup>1</sup><https://pytorch.org/>

<sup>2</sup><https://www.tensorflow.org/>

<sup>3</sup><https://www.mindspore.cn/>

differs from GPU-based ones, necessitating the redesign of many optimization strategies to suit chip-specific features. Finally, the ecosystem of Chinese chips is still immature, lacking a rich library of functions and debugging tools, which adds extra challenges to development and optimization. These issues make it difficult to directly run atmospheric and oceanic models on Chinese chips. Therefore, *how to migrate from PyTorch to Chinese frameworks, optimize operator adaptation and hardware utilization, and ensure efficient model operation on Chinese chips has become a critical research topic.*

Some studies have already analyzed the performance and optimization of MindSpore on Ascend. (Lu et al. 2022) conducted a detailed comparison of PyTorch, TensorFlow, and MindSpore on Ascend Neural Processing Unit (NPU) and GPU. (Wang et al. 2023) performed extensive experiments on Ascend NPUs using MindSpore, analyzing key features related to performance and optimization based on experimental results. (Zhu et al. 2023) conducted comprehensive evaluations at both the model and operator levels on Ascend NPUs, comparing the performance of MindSpore and PyTorch. In addition, some studies (Deng et al. 2024; Mu et al. 2024; Yin, Wu, and Su 2024; Yang and Ha 2025), also focus on the migration performance of framework chips in the fields of computer vision or large language models. However, most of these studies use standard deep learning networks and data sets for testing, such as Convolutional Neural Network (CNN) and Transformer, etc., and there is still no research on migration assessment tests based on large atmospheric and ocean deep learning models.

This study addresses existing challenges by developing a model migration and optimization framework that transfers large-scale atmospheric and oceanic PyTorch-based models to the MindSpore framework. It then systematically evaluates Chinese hardware platforms (DCU and Ascend). In migration and optimization framework, we focus on: (1) Model Migration—adapting PyTorch code to MindSpore while ensuring functional correctness; (2) Performance Optimization—leveraging Chinese chips’ hardware features to refine operators and implement distributed training strategies; and (3) Performance Evaluation—comparing Chinese chips with GPUs across metrics like training/inference speed, accuracy, and energy efficiency to assess their suitability for climate modeling tasks. By enabling large-scale model deployment on Chinese platforms, this work advances the localization of atmospheric and oceanic simulations while validating DCU and Ascend’s capabilities in high-performance scientific computing. The findings expand Chinese chips’ application scope, offer technical guidance for cross-framework model migration, and provide critical data to support the development of Chinese deep learning hardware and software ecosystems. We summarize the main contributions:

- We developed a migration and optimization framework that adapts large-scale atmospheric and oceanic models from PyTorch to MindSpore with functional correctness;
- We performed systematic optimization and evaluation on Chinese chips (DCU and Ascend), comparing them with

GPUs in terms of speed, accuracy, and energy efficiency;

- We advanced the localization of scientific computing, extended the use of Chinese hardware, and provided technical guidance and empirical evidence to support the Chinese AI ecosystem.

## Method

Figure 1 systematically presents the complete process of chip framework migration assessment. The evaluation system includes the migration of large atmospheric and oceanic models to the MindSpore framework, hardware chip adaptation and optimization strategies, and multi-dimensional evaluation index design.

**Model Migration Framework Design.** Due to significant differences in framework design and operator support, direct migration often encounters issues such as model structure incompatibility, missing operators, and insufficient performance optimization. To address these challenges, we designed a step-by-step migration workflow (Figure S2), including key processes such as model structure adjustments, operator adaptation and replacement, and optimizations leveraging MindSpore-specific features, with AI-GOMS as the primary example.

Regarding model structure adjustments, PyTorch employs a dynamic graph mechanism, allowing model structures to be modified dynamically during runtime, whereas MindSpore adheres to a static graph mechanism. This requires a complete redesign of the model’s overall logic during migration. For AI-GOMS, the model architecture had to be redefined in MindSpore using a static computation graph to ensure compatibility and trainability. Additionally, MindSpore’s optimization for static graphs requires developers to declare input dimensions and shapes explicitly. As a result, input tensors needed to be pre-defined during model initialization. Regarding operator adaptation, MindSpore’s operator library does not yet comprehensively cover commonly used PyTorch operators, making operator incompatibility a common issue during migration. For example, AI-GOMS relies on Fourier convolution functions like ‘torch.nn.fft’ in PyTorch.

To address this, we adopted the following strategies: (1) prioritizing the use of equivalent operators already available in MindSpore as replacements for PyTorch operators; (2) developing custom operators to extend MindSpore’s operator library for unsupported operations; and (3) restructuring computational logic or utilizing low-level APIs to implement the functionality of certain complex operators. Regarding optimizations leveraging MindSpore-specific features, we entirely use unique functionalities provided by the MindSpore framework, such as mixed precision training, built-in distributed computation support, and graph mode optimizations. Mixed precision training enables significant reductions in memory usage while maintaining model accuracy, thereby improving computational efficiency. Distributed computation support allowed us to split large models across multiple Ascend chips for training, effectively addressing the limitations of single-chip computing power. Graph mode optimization analyzes the static computation

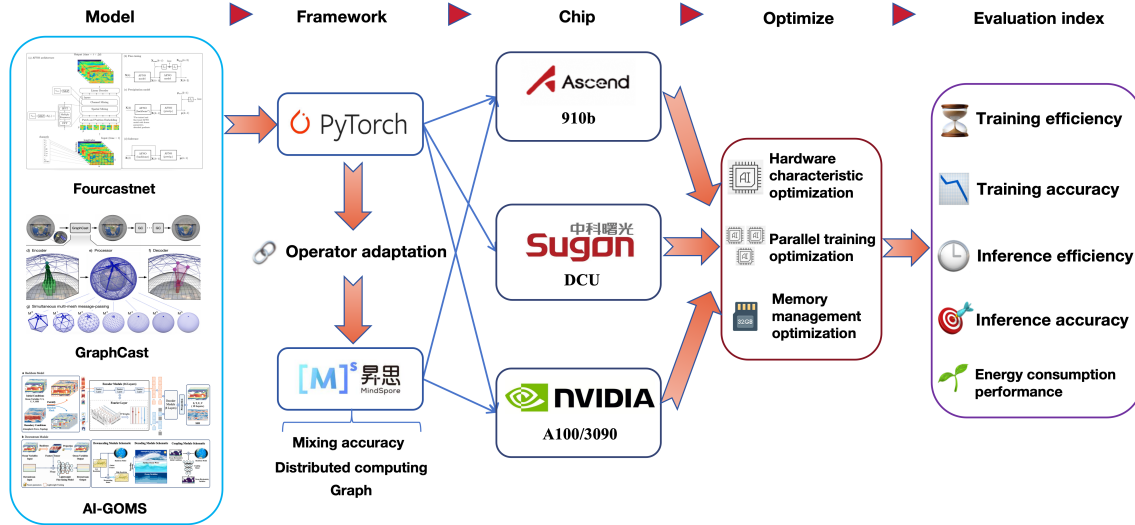


Figure 1: The migration and optimization framework for atmospheric and oceanic models from PyTorch to MindSpore based on different hardware.

graph during the compilation phase, improving training speed by optimizing the computational graph.

**Optimization Strategies for Adapting to Chips.** After migrating the model to MindSpore, we designed a series of targeted optimization strategies to fully exploit the hardware capabilities of Chinese chips (DCU and Ascend). These optimizations include hardware-specific enhancements, distributed training optimizations, and memory management improvements.

The Ascend 910b, as a high-performance AI chip, offers computational power that rivals or surpasses high-end GPUs in specific scenarios. To capitalize on this, we utilize Ascend’s built-in hardware accelerators (such as matrix computation units and AI operator optimization engines) to optimize the model at the operator level. For instance, using Ascend’s matrix acceleration engine significantly improves computational efficiency in the matrix multiplication operations of AI-GOMS. Additionally, we leveraged MindSpore’s built-in operator library for convolution and normalization operations to perform hardware instruction-level optimizations, reducing data transfer and computational latency. During the training of large models, the substantial number of parameters and high computational complexity often make it challenging for a single chip to handle training efficiently. To address this, we utilized MindSpore’s distributed training support to divide the training data into parts processed by multiple Ascend chips in parallel under a model parallelism mode. This implementation of distributed training significantly reduced single-chip training time while enhancing overall training efficiency. In memory management, the memory usage of deep learning models during training and inference is a critical factor affecting

hardware performance. We implemented specific memory allocation and computation scheduling strategies to optimize memory usage on DCU and Ascend. On the one hand, we leveraged MindSpore’s mixed precision training functionality to reduce some model weights from 32-bit floating-point (FP32) to 16-bit floating-point (FP16), significantly lowering memory overhead. On the other hand, we utilized Ascend’s pipelined execution features to schedule the forward and backward propagation stages in a phased manner, avoiding excessive memory usage peaks.

**Evaluation Benchmarks.** Several key metrics are defined to evaluate model performance after migration. Training efficiency is assessed by recording per-epoch and total training time. Inference efficiency is measured through single-chip inference time, with cross-platform comparisons. Following (Rasp et al. 2020), accuracy is evaluated using RMSE and ACC to analyze atmospheric and oceanic variable predictions over time. Energy performance is monitored during training and inference using hardware tools, with energy efficiency calculated (Bridges, Imam, and Mintz 2016) for comprehensive platform comparisons. The corresponding formulas are shown in the appendix.

## Experimental Setup

**Hardware Configuration.** The experimental hardware platform includes Chinese chips and mainstream GPUs for performance comparison in deep learning tasks. The table 1 details the configurations. The Huawei Ascend 910b chip delivers 256 TFLOPS (FP16), 32GB high bandwidth memory, and a 1.25 GHz clock speed, supporting MindSpore and PyTorch frameworks. Both single-card and multi-card

setups were tested for distributed training and inference. The Sugon DCU Z100L chip, based on the ROCm architecture, emphasizes low power consumption and high performance. It requires precompiled PyTorch packages from Sugon’s platform, as official PyTorch is incompatible. As the DCU platform currently does not support the MindSpore framework, no experiments based on MindSpore framework are conducted on this chip. Experiments cover its 16GB and 32GB memory variants. The GPU platform is based on the CUDA architecture. The GPU platform serves as a control group for comparison of training time, inference speed, and accuracy performance. All hardware platforms were configured with the same version of the operating system and software environment to ensure fair comparisons between platforms.

Hardware	Memory (GB)	FP16 (TFlops)	FP32 (TFlops)	Bandwidth (GB/s)
Ascend 910b	40GB	376	94	392
DCU Z100L	32GB	24.5	12.2	1024
Nvidia 48G memory GPU	40GB	312	19.5	1597
Nvidia 3090	24GB	35.58	35.58	936

Table 1: Hardware Configuration

**Datasets and Models.** We use the ERA5 (Hersbach et al. 2020) and HYCOM (Chassignet et al. 2007) global reanalysis datasets. ERA5 provides 0.25° resolution atmospheric variables (temperature, humidity, wind speed) and oceanic surface data, while HYCOM offers oceanic parameters (salinity, temperature, currents) across depth layers, interpolated from 0.08° to 0.25° resolution. To ensure comparability across platforms, the data is first standardized through normalization and interpolation, then split into training, validation, and testing subsets. This experiment selects three representative models from the atmospheric and oceanic domains, FourCastNet, GraphCast, and AI-GOMS, encompassing diverse deep learning architectures such as Transformer, Graph Neural Network (GNN), and Adaptive Fourier Neural Operator (AFNO), as well as a variety of task scenarios.

FourCastNet (Pathak et al. 2022) is a deep learning weather forecasting model. Based on AFNO (Guibas et al. 2021), it achieves efficient multivariate meteorological predictions through training on large-scale reanalysis datasets. FourCastNet offers superior parallel processing capabilities and computational efficiency, significantly accelerating inference speed compared to traditional numerical weather prediction methods. However, its training process remains computationally intensive - the original study reports requiring 16 hours on 64 NVIDIA 48G memory GPUs for end-to-end training. This model is an effective benchmark for evaluating Chinese chips’ performance handling atmospheric data operations, particularly Fourier transforms.

GraphCast (Lam et al. 2023) is a GNN-based weather forecasting model (Alet et al. 2019) that effectively captures meteorological data’s non-Euclidean characteristics, including flow fields and dynamic spatial relationships. Through sophisticated node-edge representations, it models large-

scale atmospheric interactions with superior performance. However, this comes with significant computational demands - training requires about four weeks on 32 TPUs. We use GraphCast to evaluate the chips’ capability in handling non-Euclidean data structures and computations.

AI-GOMS (Xiong et al. 2023) is the world’s first large-scale oceanic model. It combines Fourier-encoded autoregressive pretraining (He et al. 2022; Guibas et al. 2021) with lightweight fine-tuning to predict key global oceanic variables at 0.25° spatial and 1-day temporal resolution across 15 depth layers for 30-day forecasts. The model significantly improves efficiency and accuracy for operational ocean forecasting and research. Training requires approximately 34.2 hours on 16 NVIDIA 48G memory GPUs. We use AI-GOMS to evaluate Chinese chips’ performance in processing Fourier-encoded oceanic data.

**Training and Inference Configuration.** We evaluate model performance across frameworks and hardware platforms by measuring: (1) time efficiency, (2) training accuracy, and (3) energy consumption. To ensure fair comparisons:

- All experiments use identical hyperparameters (learning rate, optimizer, batch size)
- A fixed random seed guarantees reproducibility
- Training duration is determined by convergence (when loss stabilizes)
- Mixed precision training is enabled on GPU/Ascend platforms, with accuracy validated against full precision
- Inference tests use pre-trained weights on single cards to assess speed and accuracy

## Experimental Results

**Training Efficiency Analysis.** Training time is a key indicator for assessing hardware platform performance. We conducted complete training of FourCastNet, GraphCast, and AI-GOMS on different hardware platforms and frameworks, training the models until full convergence. Each model’s average training time per epoch and total training duration were recorded. By comparing the performance of various platforms, we evaluate Chinese chips’ training efficiency and adaptability.

Using the AI-GOMS model as an example, Table 2 shows that in FP16 precision mode, the Ascend 910b platform running on the PyTorch framework achieves training times nearly identical to the 48G memory GPU platform, with a difference of less than 5% in training time per epoch. This result demonstrates that Ascend 910b offers hardware performance comparable to mainstream GPUs. The training time for AI-GOMS on the DCU platform is relatively longer, with the total training time reaching approximately 1.29 times that of the 48G memory GPU platform. This performance difference likely arises from constraints in computing power and memory bandwidth. However, the training time significantly decreases with multi-device parallel training, highlighting its scalability potential. After migrating to the MindSpore framework, Ascend 910b exhibits higher efficiency when leveraging its distributed training features,

Chip	Memory	Number	batch_size	time/epoch(s)	Convergence epoch number	Training time(s)
48G memory GPU	40G	16	2	5400	75	405000
48G memory GPU	40G	16	1	8800	60	528000
3090	16G	16	1	10200	65	663000
910b	40G	16	2	5580	75	418500
910b	40G	16	1	8960	60	537600
DCU	32G	16	1	12600	60	756000
DCU	32G	32	1	6700	75	502500
DCU	32G	64	1	4500	75	337500

Table 2: Comparison of AI-GOMS training efficiency with PyTorch framework

Chip	Memory	Number	batch_size	time/epoch(s)	Convergence epoch number	Training time(s)
48G memory GPU	40G	16	2	5500	75	412500
48G memory GPU	40G	16	1	8960	60	537600
910b	40G	16	2	5400	70	378000
910b	40G	16	1	8800	55	484000

Table 3: Comparison of AI-GOMS training efficiency with MindSpore framework

as shown in Table 3. With the optimization features provided by MindSpore, the total training time on the Ascend 910b platform drops by approximately 10% compared to the PyTorch framework. Additionally, the distributed training mode of MindSpore enhances resource utilization and fully exploits the hardware acceleration advantages of the Ascend platform.

**Training Accuracy Analysis.** During the deep learning model training phase, training accuracy is typically evaluated by monitoring the loss function and the decline in convergence speed. For large-scale models like AI-GOMS, which involve massive datasets and high parameter counts, the convergence speed and the final loss value are critical indicators of model performance after migration.

Using the AI-GOMS model on the PyTorch framework as an example, the training conducted on the Ascend 910b and DCU platforms shows a decline in loss function consistent with GPU platforms. The experiments reveal that the loss function decreases rapidly during the first 20 epochs and gradually stabilizes after 40 epochs. Across different hardware platforms, the overall shape of the loss curves remains similar, indicating that the migrated model maintains consistent training performance on multiple platforms. The training process demonstrates higher stability on the MindSpore framework running on the Ascend 910b platform. The loss curve is smoother than the PyTorch framework, without noticeable fluctuations or instability. This indicates that the migration to the MindSpore framework achieves good compatibility with the optimizer and operators, effectively leveraging the hardware features of the Ascend platform. Further analysis shows that the MindSpore framework, through static graph optimization and mixed precision training, efficiently handles parameter updates and gradient computations for large-scale models, thereby improving training stability and efficiency. The visualization of the training loss is shown in Figure S1.

**Inference Efficiency Analysis.** We evaluate the time consumption for single inference tasks across different models, frameworks, and hardware platforms in the inference time tests. The results are shown in Table 4. Using AI-GOMS as an example, the single-step inference time on the Ascend 910b platform is comparable to that on the GPU platform, while the inference time on the DCU platform is slightly higher than that on the GPU. Notably, when the model is migrated to the MindSpore framework, the single-step inference speed on the Ascend 910b platform improves compared to the PyTorch version, indicating that MindSpore offers particular advantages in optimizing for the Ascend chips. For other models, the inference time for FourCastNet and GraphCast shows more significant differences across hardware platforms. For instance, the inference time of FourCastNet on the Ascend 910b platform is slightly higher than that on the GPU but still outperforms the DCU. In contrast, GraphCast, due to the computational complexity of its graph neural network structure, exhibits higher inference times on both the Ascend 910b and DCU platforms compared to the GPU. These results suggest that while the Ascend and DCU platforms demonstrate competitiveness in certain tasks, further optimization is required for specific types of deep learning models, such as graph neural networks.

**Inference Accuracy Analysis.** We compare the inference results obtained using the MindSpore framework on the Ascend 910b platform with those obtained using the PyTorch framework on the GPU platform, focusing on the output accuracy of several key variables (e.g.,  $U_0$ ,  $T_0$ ,  $S_0$ ). The results indicate that most variables maintain consistent inference accuracy across the two frameworks, though some variables exhibit slight differences in accuracy.

As shown in Figure 2, the predictions for  $U_0$  demonstrate strong consistency between the MindSpore and PyTorch frameworks, with high inference accuracy, indicating that the model’s dynamic characteristics are well-preserved after migration. Similarly, the predictions for  $S_0$  show high

Model	Framework	Chip	test time/step(s)
Fourcastnet	Pytorch	48G memory GPU	0.353
	Pytorch	910b	0.389
	Pytorch	DCU	0.400
Graphcast	Pytorch	48G memory GPU	1.035
	Pytorch	910b	1.169
AI-GOMS	Pytorch	48G memory GPU	0.675
	Pytorch	910b	0.680
	Pytorch	DCU	0.705
	Mindspore	48G memory GPU	0.598
	Mindspore	910b	0.605

Table 4: Comparison of inference efficiency of different models on different chips and frameworks

similarity, with minimal errors between the two frameworks. However, for  $T_0$ , the predictions from the MindSpore framework exhibit a certain degree of error compared to those from the PyTorch framework, reflecting slightly lower inference accuracy. This discrepancy might be attributed to the numerical characteristics of the temperature field variable, such as its sensitivity to spatial distribution. Additionally, the error could result from numerical approximations in operator implementations or differences in handling floating-point computations.

**Energy Consumption Analysis.** Energy consumption is a critical performance metric for Chinese chips, particularly in scenarios with high energy-saving requirements. In this experiment, hardware monitoring tools were used to record the power consumption of each platform during training and inference tasks, and the relative energy efficiency ratio was calculated, as shown in Table S1. Using AI-GOMS as an example, during the training phase, the Ascend 910b platform demonstrated superior energy efficiency compared to GPUs, with an average power consumption increase of approximately 15% while maintaining comparable performance to GPUs, resulting in a significant improvement in energy efficiency ratio. The DCU platform exhibited even lower energy consumption, but due to its relatively weaker computational power, its overall energy efficiency ratio was comparable to GPUs. In the inference phase, the DCU platform achieved the lowest power consumption, with an energy efficiency ratio 1.3 times higher than GPUs, making it particularly suitable for long-duration inference tasks. The Ascend 910b platform consumed 10% less power than GPUs during inference tasks while maintaining high performance levels. Overall, Chinese chips exhibited outstanding energy efficiency, particularly in inference scenarios. The DCU platform showed a significant advantage in energy savings, while the Ascend 910b platform demonstrated a strong ability to balance energy consumption and performance.

## Conclusion and Discussion

This study demonstrates the feasibility and advantages of migrating large-scale atmospheric and oceanic models from PyTorch to MindSpore for efficient execution on Chinese hardware platforms such as Ascend 910b and DCU. Through systematic experimentation, we confirm that the migrated models retain comparable accuracy to their

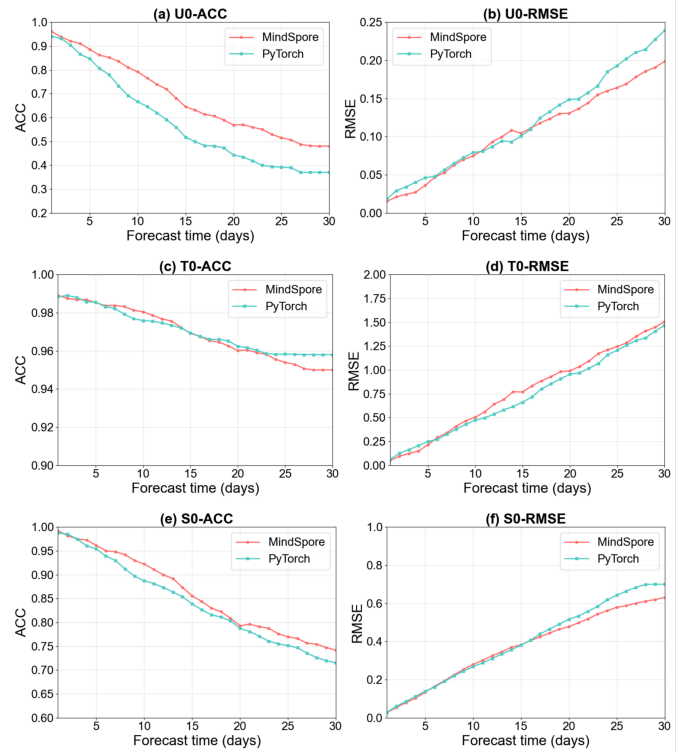


Figure 2: Comparison of Inference Accuracy between MindSpore and PyTorch Frameworks

original PyTorch counterparts, with deviations under 5%, while achieving competitive computational performance. Notably, Ascend 910b exhibits strong capabilities in distributed training scenarios, delivering GPU-level throughput with 10–15% lower energy consumption. In contrast, DCU shows particular promise in energy-efficient, real-time inference tasks due to its superior power efficiency.

However, challenges remain in optimizing operator support for complex model architectures and further bridging the accuracy gap in high-resolution, computationally intensive simulations. Future work should focus on enhancing MindSpore’s operator library, refining distributed training efficiency, and exploring hardware-software co-design strategies tailored to specific meteorological and oceanographic applications. Additionally, fostering a robust ecosystem—including optimized toolchains, open-source collaboration, and industry integration—will be critical to advancing the adoption of Chinese platforms in large-scale scientific deep learning.

Ultimately, this research underscores the significant potential of Chinese chips in climate and ocean modeling, paving the way for further innovations in performance, energy efficiency, and scalability. Continued advancements in framework optimization, hardware specialization, and ecosystem development will be essential to fully realize this potential and strengthen the competitiveness of Chinese solutions in high-performance AI-driven scientific computing.

## Acknowledgments

This work was supported by the following project: National Natural Science Foundation of China (Grant No. U2442206, 42175052, 42125503)

## References

- Alet, F.; Jeewajee, A. K.; Villalonga, M. B.; Rodriguez, A.; Lozano-Perez, T.; and Kaelbling, L. 2019. Graph element networks: adaptive, structured computation and memory. In *International Conference on Machine Learning*, 212–222. PMLR.
- Bi, K.; Xie, L.; Zhang, H.; Chen, X.; Gu, X.; and Tian, Q. 2023. Accurate medium-range global weather forecasting with 3D neural networks. *Nature*, 619(7970): 533–538.
- Bridges, R. A.; Imam, N.; and Mintz, T. M. 2016. Understanding GPU power: A survey of profiling, modeling, and simulation methods. *ACM Computing Surveys (CSUR)*, 49(3): 1–27.
- Chassignet, E. P.; Hurlburt, H. E.; Smedstad, O. M.; Halliwell, G. R.; Hogan, P. J.; Wallcraft, A. J.; Baraille, R.; and Bleck, R. 2007. The HYCOM (hybrid coordinate ocean model) data assimilative system. *Journal of Marine Systems*, 65(1-4): 60–83.
- Chen, K.; Han, T.; Gong, J.; Bai, L.; Ling, F.; Luo, J.-J.; Chen, X.; Ma, L.; Zhang, T.; Su, R.; et al. 2023a. Fengwu: Pushing the skillful global medium-range weather forecast beyond 10 days lead. *arXiv preprint arXiv:2304.02948*.
- Chen, L. 2021. *Deep learning and practice with mindspore*. Springer Nature.
- Chen, L.; Zhong, X.; Zhang, F.; Cheng, Y.; Xu, Y.; Qi, Y.; and Li, H. 2023b. FuXi: A cascade machine learning forecasting system for 15-day global weather forecast. *npj Climate and Atmospheric Science*, 6(1): 190.
- Cui, Y.; Wu, R.; Zhang, X.; Zhu, Z.; Liu, B.; Shi, J.; Chen, J.; Liu, H.; Zhou, S.; Su, L.; et al. 2025. Forecasting the eddy ocean with a deep neural network. *Nature Communications*, 16(1): 2268.
- Deng, C.; Yang, X.; Qin, X.; Peng, L.; Zhang, J.; Wang, W.; Qian, Q.; and Zhao, J. 2024. Implementation of Person Re-Identification Algorithm Based on Huawei MindSpore and Pytorch Hybrid Architecture. In *2024 IEEE First International Conference on Data Intelligence and Innovative Application (DIIA)*, 1–11. IEEE.
- Guibas, J.; Mardani, M.; Li, Z.; Tao, A.; Anandkumar, A.; and Catanzaro, B. 2021. Adaptive fourier neural operators: Efficient token mixers for transformers. *arXiv preprint arXiv:2111.13587*.
- He, K.; Chen, X.; Xie, S.; Li, Y.; Dollár, P.; and Girshick, R. 2022. Masked autoencoders are scalable vision learners. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 16000–16009.
- Hersbach, H.; Bell, B.; Berrisford, P.; Hirahara, S.; Horányi, A.; Muñoz-Sabater, J.; Nicolas, J.; Peubey, C.; Radu, R.; Schepers, D.; et al. 2020. The ERA5 global reanalysis. *Quarterly Journal of the Royal Meteorological Society*, 146(730): 1999–2049.
- Lam, R.; Sanchez-Gonzalez, A.; Willson, M.; Wirnsberger, P.; Fortunato, M.; Alet, F.; Ravuri, S.; Ewalds, T.; Eaton-Rosen, Z.; Hu, W.; et al. 2023. Learning skillful medium-range global weather forecasting. *Science*, 382(6677): 1416–1421.
- Lu, W.; Zhang, F.; He, Y.; Chen, Y.; Zhai, J.; and Du, X. 2022. Evaluation and optimization for huawei ascend neural network accelerator. *Chinese Journal of Computers*, 45(8): 1618–1637.
- Ma, Y.; Yu, D.; Wu, T.; and Wang, H. 2019. PaddlePaddle: An open-source deep learning platform from industrial practice. *Frontiers of Data and Computing*, 1(1): 105–115.
- Mu, Y.; Zhai, J.; Fang, C.; Chen, X.; Cao, Z.; Yang, P.; Zou, Y.; Zheng, T.; and Chen, Z. 2024. DevMuT: Testing deep learning framework via developer expertise-based mutation. In *Proceedings of the 39th IEEE/ACM International Conference on Automated Software Engineering*, 1533–1544.
- Pathak, J.; Subramanian, S.; Harrington, P.; Raja, S.; Chatopadhyay, A.; Mardani, M.; Kurth, T.; Hall, D.; Li, Z.; Azizzadenesheli, K.; et al. 2022. Fourcastnet: A global data-driven high-resolution weather model using adaptive fourier neural operators. *arXiv preprint arXiv:2202.11214*.
- Rasp, S.; Dueben, P. D.; Scher, S.; Weyn, J. A.; Mouatadid, S.; and Thuerey, N. 2020. WeatherBench: a benchmark data set for data-driven weather forecasting. *Journal of Advances in Modeling Earth Systems*, 12(11): e2020MS002203.
- Wang, B.; Yang, C.; Zhu, R.; Liu, X.; Zhou, M.; and Zheng, N. 2023. Analysis of Performance and Optimization in MindSpore on Ascend NPUs. In *2023 IEEE 29th International Conference on Parallel and Distributed Systems (ICPADS)*, 1701–1708. IEEE.
- Wang, X.; Wang, R.; Hu, N.; Wang, P.; Huo, P.; Wang, G.; Wang, H.; Wang, S.; Zhu, J.; Xu, J.; et al. 2024. Xihe: A data-driven model for global ocean eddy-resolving forecasting. *arXiv preprint arXiv:2402.02995*.
- Xiong, W.; Xiang, Y.; Wu, H.; Zhou, S.; Sun, Y.; Ma, M.; and Huang, X. 2023. Ai-goms: Large ai-driven global ocean modeling system. *arXiv preprint arXiv:2308.03152*.
- Yang, J.; and Ha, W. 2025. Research on Network Migration and Performance Optimization of the YOLOv7 Model Improved Based on MindSpore Framework. In *2025 2nd International Conference on Algorithms, Software Engineering and Network Security (ASENS)*, 211–214. IEEE.
- Yang, N.; Wang, C.; Zhao, M.; Zhao, Z.; Zheng, H.; Zhang, B.; Wang, J.; and Li, X. 2024. Langya: Revolutionizing Cross-Spatiotemporal Ocean Forecasting. *arXiv preprint arXiv:2412.18097*.
- Yin, J.; Wu, F.; and Su, H. 2024. An Image-Retrieval Method Based on Cross-Hardware Platform Features. *Applied System Innovation*, 7(4): 64.
- Zhu, Z.; Wang, B.; Yang, C.; Zhu, R.; Zhou, M.; and Zheng, N. 2023. Performance Evaluation of MindSpore and PyTorch Based on Ascend NPU. In *2023 IEEE 29th International Conference on Parallel and Distributed Systems (ICPADS)*, 1826–1832. IEEE.