

DMSORT: An efficient parallel maritime multi-object tracking architecture for unmanned vessel platforms

Shengyu Tang^{a,1}, Zeyuan Lu^{b,1}, Jiazhi Dong^c, Changdong Yu^{d,*}, Xiaoyu Wang^a,
Yaohui Lyu^a and Weihao Xia^e

^aCollege of Information Science and Engineering, Ocean University of China, Qingdao, 266100, Shandong, China

^bMarine Engineering College, Dalian Maritime University, Dalian, 116026, Liaoning, China

^cCSSC Haiwei Tech Co., Ltd, Zhengzhou, 450000, Henan, China

^dCollege of Artificial Intelligence, Dalian Maritime University, Dalian, 116026, Liaoning, China

^eCollege of Marine Electrical Engineering, Dalian Maritime University, Dalian, 116026, Liaoning, China

ARTICLE INFO

Keywords:

Unmanned vessel platforms
Multi-object tracking
Object detection
Transformer
Re-identification
Deep learning

ABSTRACT

Accurate perception of the marine environment through robust multi-object tracking (MOT) is essential for ensuring safe vessel navigation and effective maritime surveillance. However, the complicated maritime environment often causes camera motion and subsequent visual degradation, posing significant challenges to MOT. To address this challenge, we propose an efficient Dual-branch Maritime SORT (DMSORT) method for maritime MOT. The core of the framework is a parallel tracker with affine compensation, which incorporates an object detection and re-identification (ReID) branch, along with a dedicated branch for dynamic camera motion estimation. Specifically, a Reversible Columnar Detection Network (RCDN) is integrated into the detection module to leverage multi-level visual features for robust object detection. Furthermore, a lightweight Transformer-based appearance extractor (Li-TAE) is designed to capture global contextual information and generate robust appearance features. Another branch decouples platform-induced and target-intrinsic motion by constructing a projective transformation, applying platform-motion compensation within the Kalman filter, and thereby stabilizing true object trajectories. Finally, a clustering-optimized feature fusion module effectively combines motion and appearance cues to ensure identity consistency under noise, occlusion, and drift. Extensive evaluations on the Singapore Maritime Dataset demonstrate that DMSORT achieves state-of-the-art performance. Notably, DMSORT attains the fastest runtime among existing ReID-based MOT frameworks while maintaining high identity consistency and robustness to jitter and occlusion. **Code is available at:** <https://github.com/BiscuitsLzy/DMSORT-An-efficient-parallel-maritime-multi-object-tracking-architecture->.

1. Introduction

As the exploitation of marine resources continues to grow, maritime security measures are constantly updated, driving an increasing demand for unmanned surface vessels (USVs) and marine robots to monitor and track the trajectories of ships, fishing vessels, and offshore platforms in predefined waters (Agarwala, 2020; Guo et al., 2024; Hu and Xu, 2025). Maritime multi-object tracking (MOT) plays a pivotal role in various applications such as intelligent surveillance, autonomous navigation, and maritime situational awareness (An et al., 2023; Chen et al., 2025; Cheng et al., 2025; Shao et al., 2024; Su et al., 2025). While satellite remote sensing remains a primary tool for wide-area surveillance, its effectiveness is often hindered by cloud cover and revisit limitations (Langsdale et al., 2025; Soldi et al., 2021). In contrast, vision-equipped USVs can capture real-time imagery

at sea level, overcoming these constraints (Jin et al., 2020). However, the unique challenges of maritime environments—unpredictable camera perspectives, low visibility, heavy occlusions, and drastic target scale variations—make MOT particularly difficult (Tran and Le, 2016). Unreliable maritime communications and poor signal quality further hamper real-time data transfer, rendering cloud-based processing impractical in many scenarios. Moreover, the computational constraints of many onboard systems complicate the deployment of resource-heavy models, highlighting the need for lightweight and efficient solutions (Rangesh and Trivedi, 2019; Tian et al., 2024; Zhang et al., 2021b).

Recent approaches to MOT typically follow one of two paradigms: tracking-by-detection (TBD) or joint detection and tracking (JDT) (Luo et al., 2021). JDT frameworks (Du et al., 2023; Stanojevic and Todorovic, 2024; Wojke et al., 2017; Zhang et al., 2022) aim to unify detection and association within a single network, enabling end-to-end training. However, this approach

 ycd@dlmu.edu.cn (Changdong Yu)
ORCID(s): 0000-0002-5759-4589 (Changdong Yu)

often struggles with learning shared feature representations, particularly in maritime environments where small, low-resolution targets are prevalent. The simultaneous training of detection and appearance feature extraction networks can result in gradient interference, thereby limiting the quality of both object detection and appearance modeling. This conflict undermines the ability to effectively model object motion and appearance, resulting in inaccuracies in both.

In contrast, TBD methods (Du et al., 2023; Stanojevic and Todorovic, 2024; Zhang et al., 2022) decouple the detection and tracking stages, employing robust detectors and modular tracking algorithms such as Kalman filtering. While this design offers better interpretability and flexibility, its sequential nature introduces latency. Moreover, in maritime environments, challenges such as motion blur, low resolution, and visual clutter significantly complicate the extraction of reliable motion and appearance features. These environmental difficulties hinder accurate object tracking and matching, making it challenging for TBD methods to recover correct identity associations. As a result, tracking performance is notably degraded, particularly under long-term occlusions or significant trajectory drift. Ultimately, both JDT and TBD methods face limitations in maritime settings.

From the above analysis, it becomes evident that the limitations of both JDT and TBD approaches arise from their incomplete use of motion and appearance cues. To further clarify this issue, we analyze the role of these two complementary sources of information in maritime MOT. In practice, motion models are typically built upon object detectors combined with Kalman filtering, where the detector provides current observations and the filter predicts target trajectories. While this design ensures short-term continuity, it is easily disrupted by ego-motion, unpredictable platform perturbations from wind and waves, and severe scale variations of maritime targets, often leading to large prediction errors (Aharon et al., 2022; Cai et al., 2022; Lin et al., 2017). On the other hand, appearance cues are essential for long-term identity association, yet are often unreliable due to low resolution, motion blur, frequent occlusions, and cluttered backgrounds. Conventional CNN-based Re-ID models struggle with limited receptive fields (Hermans et al., 2017), while transformer-based approaches offer more robust representations (Graham et al., 2021; Li et al., 2023; Liu et al., 2023; Zhang et al., 2025) but remain computationally heavy under onboard resource constraints.

In addition to these challenges, the crux of maritime MOT lies in how motion and appearance information are fused. Traditional strategies often use simple weighted combinations of the two cues (Wojke et al.,

2017), which fail under drift, noise, or long occlusions. This highlights the need for more adaptive fusion mechanisms that can dynamically balance unreliable cues, ensuring robust identity consistency in the face of occlusion, fragmentation, and adverse visual conditions.

Motivated by the above challenges, we propose a novel maritime multi-object tracking (MOT) framework that jointly enhances detection, tracking, and identity association under resource-constrained conditions. Specifically, we design a Dual-Branch Detection–Tracking Architecture (DDTA) that runs detection and tracking in parallel while isolating platform and target motion, enabling real-time perception without sacrificing accuracy. Building upon this architecture, we develop a Reversible Columnar Detection Network (RCDN), a compact multi-column detector with reversible feature refinement. It optimizes feature transmission to preserve information across layers, thereby achieving high recall and precision under extreme target scale variations while remaining lightweight. To address the limitations of conventional appearance modeling, we further introduce Li-TAE, a lightweight Transformer-based Re-ID module trained with adaptive hard triplets, which captures global contextual features and improves target discrimination under occlusion and low-resolution input. Finally, we propose a Clustering Optimized Feature Fusion (COFF) strategy that adaptively integrates motion and appearance cues using a nonlinear distance metric, unbiased momentum update, and spatial gating, thereby preserving identity consistency under noise, occlusion, and trajectory drift. Together, these components form a coherent solution that effectively tackles the unique difficulties of maritime MOT.

In summary, the main contributions of this work are as follows:

- We design a Dual-Branch Detection–Tracking Architecture (DDTA) that decouples but parallelizes detection and tracking for real-time maritime perception.
- We develop a compact Reversible Columnar Detection Network (RCDN) that preserves shallow features for robust multi-scale detection under resource constraints.
- We introduce Li-TAE, a lightweight Transformer-based Re-ID module with adaptive hard triplet training for enhanced discrimination in occluded and low-resolution conditions.
- We propose a Clustering Optimized Feature Fusion (COFF) mechanism that adaptively balances motion and appearance cues to maintain

identity consistency in challenging maritime environments.

2. Related work

2.1. Traditional Tracking Methods: From Correlation Filters to Deep Learning

Before the advent of deep learning, research in visual tracking mainly focused on the single-object setting (SOT), as traditional detectors and hand-crafted features made large-scale multi-target tracking difficult. In this era, correlation filter-based methods such as MOSSE, CSK, and KCF (Bolme et al., 2010; Danelljan et al., 2014; Henriques et al., 2012, 2015) achieved high efficiency through Fourier domain regression and kernel learning, often running in real time. However, their reliance on shallow features and linear models limited robustness to occlusion, fast motion, and appearance variations, and they lacked mechanisms for identity management.

In contrast, early MOT approaches adopted a tracking-by-detection paradigm, where traditional detectors (e.g., HOG-SVM, DPM) provided candidate boxes, and association was performed across frames. Typical strategies included probabilistic models such as JPDA (Fortmann et al., 1983) and MHT (Reid, 1979), global graph optimization via min-cost flow (Pirsiavash et al., 2011; Zhang et al., 2008) and k-shortest paths (Berclaz et al., 2011), and filtering-based methods using the Kalman filter (Kalman, 1960) or particle filters (Arulampalam et al., 2002). While foundational, these approaches were hindered by the limited discriminative power of hand-crafted features, often resulting in ID switches and fragmented tracks.

The emergence of deep detectors and learned appearance embeddings transformed the field, shifting focus toward deep learning-based MOT. By combining accurate detection with robust identity features, these methods achieve far stronger performance, and can be broadly divided into motion-based trackers and hybrid motion–appearance trackers.

2.2. Deep Learning-Based MOT: Motion and Appearance Fusion

Motion-based MOT prioritizes efficiency, relying on object dynamics and inter-frame geometric consistency to associate detections. SORT (Bewley et al., 2016) remains a widely used baseline, employing Kalman filtering and the Hungarian algorithm to achieve over 250 FPS. ByteTrack (Zhang et al., 2022) improves robustness by incorporating low-confidence detections through a two-stage matching strategy. BOT-SORT (Aharon et al., 2022) introduces

optical flow-based background compensation and extends the Kalman state to better model dynamic camera settings. OCSORT Cao et al. (2023) further refines motion modeling with observation center momentum (OCM) and online smoothing (OOS), enabling virtual motion estimation.

To overcome the limitations of motion-only trackers, hybrid methods fuse motion and appearance cues, typically using deep CNNs for appearance modeling. DeepSORT (Wojke et al., 2017) enhances SORT with deep Re-ID embeddings for identity preservation. StrongSORT (Du et al., 2023) introduces trajectory linking and Gaussian interpolation to improve association robustness. BoT-SORT-ReID (Aharon et al., 2022) and ImprAssoc Stadler and Beyerer (2023) incorporate dynamic feature weighting, while DeepOCSORT (Maggiolino et al., 2023) applies adaptive cue balancing to dynamically adjust motion and appearance contributions.

Despite their improved robustness, most of these trackers adopt a modular, sequential pipeline, which limits real-time performance and struggles under extreme appearance degradation. Their reliance on convolutional appearance extractors also makes them vulnerable to occlusion, low resolution, and environmental interference.

2.3. Domain-Specific Insights: Lessons from Benchmark Studies

Benchmark-driven MOT research further illustrates the domain sensitivity of tracker performance. In DanceTrack (Sun et al., 2022), where inter-target similarity and non-rigid motion dominate, traditional Re-ID pipelines degrade rapidly. SAM2MOT (Jiang et al., 2025) tackles these challenges by employing segmentation-driven association and cross-object interaction mechanisms, effectively mitigating identity switches in such deformable and ambiguous motion settings. In the BEE24 benchmark, bee trajectories exhibit highly nonlinear dynamics, overwhelming classical filters. TOPICTrack (Cao et al., 2025) introduces attention-based appearance reconstruction and feature-level motion-appearance fusion to tackle biologically erratic motion.

These findings reveal two crucial limitations of existing MOT frameworks. First, most methods implicitly assume that targets follow predictable motion patterns with relatively stable appearances—assumptions that break down in scenes with non-rigid motion, crowded interactions, or occlusions. Second, classical association pipelines often lack the adaptability to handle dynamic appearance changes or severe occlusions, leading to fragmented trajectories and frequent identity switches.

These lessons translate directly to maritime environments, which introduce unique challenges: extreme scale variance (from large vessels to distant silhouettes), nonlinear motion due to wave action and platform drift, and frequent visual degradation from fog, sea-glare, and poor illumination. In such settings, objects rarely maintain consistent motion or appearance, and background dynamics further complicate association stability. As a result, conventional motion-based or hybrid trackers often struggle with ID fragmentation, tracking drift, and degraded robustness.

Furthermore, recent studies emphasize that tracking under non-rigid motion requires specialized modeling of both appearance variability and trajectory dynamics. Techniques such as attention-based reconstruction, multi-scale feature fusion, and trajectory smoothing have shown promise in biologically inspired benchmarks like BEE24 and DanceTrack. These insights suggest that maritime tracking systems could benefit from integrating similar strategies, such as adaptive feature learning, context-aware motion modeling, and hierarchical attention mechanisms, to better handle the compounded visual and kinematic distortions inherent to maritime surveillance.

2.4. Maritime Multi-Object Tracking: Current Progress and Limitations

To address these challenges, maritime-specific MOT frameworks have recently emerged. Sea-IoUTracker (Guo et al., 2024) introduces a buffered-IoU strategy and neighboring-frame recovery to stabilize associations under platform jitter. It also integrates a customized detector incorporating ASFF, CBAM, and YOLOv7, optimized for maritime targets. MOMT (Niu et al., 2024) enhances ByteTrack by incorporating a fused-distance metric combining IoU, category, and object scale, alongside Gaussian cascade matching and an augmented Kalman filter with camera motion compensation. These strategies significantly improve maritime tracking robustness while maintaining real-time capability.

Nonetheless, existing maritime approaches remain grounded in terrestrial MOT frameworks, typically extending ByteTrack-like structures with heuristic enhancements. They often lack architectural flexibility and fail to jointly address scale-aware detection, appearance degradation, and dynamic fusion in an integrated fashion. Therefore, there is a growing need for end-to-end MOT systems that unify detection and tracking through parallelizable pipelines, adaptive multi-scale modeling, and transformer-based appearance extraction—capable of maintaining performance under the visual and kinematic distortions inherent to real-world maritime environments.



Figure 1: Schematic diagram of unmanned vessel platform structure.

3. Method

Fig. 1 presents a detailed schematic diagram of the structure and auxiliary devices of unmanned vessel platforms (Wang et al., 2024; Yu et al., 2024). The onboard cameras capture signals from the maritime environment and wirelessly transmit them to a host computer via a video surveillance transceiver. Once received, these video streams are processed by multi-object tracking (MOT) algorithms to identify and follow targets over time.

However, object detection and tracking in maritime environments present significant challenges. The motion of the platform introduces complex and irregular target movement patterns, while the large variation in object sizes—ranging from small buoys to large ships—further complicates reliable multi-object tracking (see Fig. 2) (Prasad et al., 2017). To address these issues, we propose the DMSORT framework, which is specifically designed to enhance tracking performance under such challenging maritime conditions.

3.1. Overview

The overall pipeline of the DMSORT framework is shown in Fig. 3, consisting of online tracking and post-processing. The online module adopts Detection-driven Dual-Task Association (DDTA), which integrates motion compensation and appearance feature extraction for robust tracking.

In DDTA, an RCDN-based detector generates bounding boxes with different confidence levels. For high-confidence detections, a lightweight Transformer encoder (Li-TAE) extracts discriminative appearance features for reliable association. Meanwhile, the observed motion in the video results from the superposition of the platform’s movement and each target’s intrinsic motion. To decouple these components, the motion branch



Figure 2: Illustrative examples of typical maritime scenes. These images showcase diverse oceanic scenarios, including vessels of various sizes, a floating buoy, and dynamic boat movements with wakes. Such scenes reflect the complex environments in maritime surveillance, posing challenges for object detection and tracking.




Figure 3: Overall pipeline of the proposed DMSORT framework. The system includes online tracking, where DDTA combines detection, motion compensation, and appearance encoding for two-stage association and tracklet management, and an optional post-processing step, where tracklet interpolation refines fragmented trajectories into final tracks.

computes the affine transformation between successive frames (Bouguet et al., 2001; FISCHLER AND, 1981; Shi et al., 1994) to estimate the platform’s motion. Once obtained, the affine matrix is incorporated into the

Kalman filter’s prediction step by multiplying it with the predicted state vector and covariance matrix. This correction compensates for camera-induced motion, aligns target trajectories with a stabilized reference




Figure 4: Architecture of the proposed RCDN. The left-bottom part illustrates the reversible column backbone, where each level exchanges features across scales for lossless extraction. The middle modules (SPPF, C3k2, C2PSA) perform multi-scale fusion and attention refinement. The right detection head outputs classification and regression for final object prediction.

frame, and enables more accurate modeling of each target's true trajectory under platform-driven disturbances.

After feature extraction and motion correction, a two-stage data association is performed. The COFF scheme first matches high-confidence detections to predicted tracks by fusing Re-ID features and IoU via the Hungarian algorithm. Remaining tracks are then associated with low-confidence detections using IoU. Finally, the trajectory management module updates motion states and features, initializes new tracks, and removes inactive ones. Overall, DMSORT achieves reliable multi-object tracking under camera jitter, occlusion, and appearance degradation.

3.2. Reversible Columnar Object Detection Network

This study proposes a multi-level feature decoupling object detection framework Reversible Column Detection Network (RCDN). The core innovation of RCDN lies in its lossless feature transformation, which is guided by residual connections and a reversible structure that provides a mathematical guarantee for explicitly reversible approximations.

As illustrated in Fig.4, the network is divided into three major components: the Hierarchical Feature Extraction, Feature Fusion, and Detection Head, which work together to process input data efficiently.

The first component of the proposed network, termed Hierarchical Feature Extraction, is detailed in Algorithm 1. In summary, this component consists of N weight-heterogeneous sub-networks that progressively extract multi-scale features through a hierarchical structure. Each sub-network, referred to as a "column", contains four stages (C_1 to C_4), which capture spatial semantics at increasingly higher levels of abstraction. These stages are designed to progressively refine the feature representation as the data passes through, using a residual coupling mechanism that fuses intra- and inter-column information:

$$x_{(m,n)} = \alpha \cdot x_{(m-1,n)} + \phi \left(g \left(x_{(m,n-1)} \right) \oplus h \left(x_{(m-1,n+1)} \right) \right) \quad (1)$$

Here, $g(\cdot)$ and $h(\cdot)$ denote $2\times$ downsampling and $2\times$ upsampling operators, respectively; $\phi(\cdot)$ represents a non-linear transformation using the C_3K_2 module, and α is a learnable scaling factor regulating residual propagation.

This approach facilitates intra-scale residual transmission via $\alpha \cdot x_{(m-1,n)}$, preserving feature continuity

Algorithm 1 Hierarchical Feature Extraction

```

1: Convstem: Conv(k=4, s=4, p=0)
2: Blocklevel: C3K2
3:  $\alpha_{level}$ : learnable scalar
4: function FUSION(down, up, first_col)
5:   down_branch  $\leftarrow$  Conv(k=2, s=2, p=0)(down)
6:   if first_col then
7:     return down_branch
8:   else
9:     up_branch  $\leftarrow$  Upsample(Conv(k=3, s=1, p=1)(up))
10:    return down_branch + up_branch
11:   end if
12: end function
13: function REVCOLFORWARD(x)
14:   x  $\leftarrow$  Convstem(x)  $\triangleright$  Patch embedding
15:   Initialize  $c_0, c_1, c_2, c_3 \leftarrow 0$ 
16:   for  $i = 1$  to  $N$  do
17:     first_col  $\leftarrow (i == 1)$ 
18:     for level = 0 to 3 do
19:       if level == 0 then
20:         down  $\leftarrow x$ 
21:       else
22:         down  $\leftarrow c_{level-1}$ 
23:       end if
24:       if level < 3 and  $\neg first\_col$  then
25:         up  $\leftarrow c_{level+1}$ 
26:       else
27:         up  $\leftarrow$  None
28:       end if
29:       fused  $\leftarrow$  Fusion(down, up, first_col)
30:       out  $\leftarrow$  Blocklevel(fused)
31:        $c_{level} \leftarrow \alpha_{level} \cdot c_{level} + out$ 
32:     end for
33:   end for
34:   return [ $c_0, c_1, c_2, c_3$ ]
35: end function

```

and minimizing semantic degradation as the network progresses through its layers. Simultaneously, cross-scale feature fusion is achieved through bidirectional sampling between neighboring columns, $(x_{(m-1,n+1)}, x_{(m,n-1)})$, which enhances multi-level feature representation.

To ensure lossless feature recovery and enable the reconstruction of previous features during reverse inference, we design the update process as an invertible function. This formulation allows the reconstruction of the earlier features from the final column output using:

$$x_{(m-1,n)} = [x_{(m,n)} - \phi(g(x_{(m,n-1)}) \oplus h(x_{(m-1,n+1)}))] / \alpha \quad (2)$$

The proposed architecture, incorporates a boundary condition ($h(x) = 0$ when $n = n_{\max}$ and $m = 1$) to

ensure the uniqueness of the feature reconstruction process. This formulation enables the explicit recovery of earlier features from final-layer representations through the invertible transformation defined in Eq. 3.2. A key practical advantage of this property is that it eliminates the need to store intermediate activations during the backward pass—features can be recomputed on-the-fly via the inverse operation, leading to substantially reduced memory consumption during training (Cai et al., 2022).

Furthermore, the universal approximation theorem ensures that the network can learn to approximate inverse mappings autonomously. Formally, the learned transformation T_θ satisfies (Cybenko, 1989; Hornik et al., 1989).

$$\left| T\theta^{-1}(x(M, n)) - x_{(m-1,n)} \right| < \epsilon \quad (3)$$

for any small tolerance $\epsilon > 0$, where $x_{(M,n)}$ denotes the output feature and $x_{(m-1,n)}$ the corresponding earlier feature. Thus, even in the absence of an explicitly programmed inversion, the model can discover—through training—efficient strategies to reconstruct or leverage multi-level features directly from the output, potentially surpassing the performance achievable through the hand-designed inverse alone.

To enable accurate inverse transformation fitting during inference, we introduce the Feature Fusion module. This module integrates multi-scale features from the network’s final stages, ensuring feature consistency across resolutions and facilitating reverse mapping during inference.

In the Feature Fusion process, multi-scale features are produced at different resolutions, denoted as $x_{(M,1)}, x_{(M,2)}, x_{(M,3)}, x_{(M,4)}$, with downsampling ratios ranging from 1/8 to 1/64. At the final output feature $x_{(M,4)}$, Spatial Pyramid Pooling Fast (SPPF) is applied. The SPPF layer uses pooling kernels of sizes 5×5 , 9×9 , and 13×13 to capture global contextual information while preserving important fine-grained details across multiple scales.

Following the application of SPPF, the output feature map undergoes attention weighting through the C2PSA module. The attention mechanism helps enhance important features by applying self-attention and feed-forward operations. The C2PSA module first splits the feature map into two parts and processes one part with a series of PSA blocks for attention and transformation. The processed feature is then recombined and passed through a final convolution layer to produce the refined features.

After feature refinement, the aggregated features are fused across scales using a bidirectional feature fusion process. This is achieved through a modified




Figure 5: Architecture of the proposed Li-TAE module. The figure illustrates the overall pipeline and detailed structure of the lightweight Transformer-based appearance encoder, including the use of self-attention with our designed positional embeddings and shrink attention.

PANet architecture. In the top-down pathway, higher-level features are upsampled and fused with shallower representations, while in the bottom-up pathway, lower-level features are downsampled and combined with deeper layers to enhance semantic consistency across scales.

These features are then fed into a YOLOv11’s detection head, where the classification and regression tasks are optimized independently. This design maintains real-time inference capability while improving detection accuracy, particularly in dense and small-object scenes.

While the initial feature extraction provides a performance floor through lossless transmission, the Feature Fusion module surpasses this bound by enhancing multi-scale aggregation and non-linear expressiveness. This improves the model’s ability to capture complex patterns, achieving more accurate detection, especially in challenging object-dense environments.

3.3. Lightweight Transformer-Based Appearance Extractor

To extract discriminative appearance features for maritime multi-object tracking—characterized by low

resolution, occlusion, and background clutter—we propose the Lightweight Transformer-based Appearance Extractor (Li-TAE). As shown in Fig. 5, this module leverages a Vision Transformer (ViT) architecture enhanced with self-attention mechanisms for efficient feature extraction. It overcomes the limitations of CNN-based ReID backbones, which often struggle with limited receptive fields and insufficient global context modeling.

The process begins with an image embedding module that downsamples the input image ($224 \times 224 \times 3$) using four layers of 3×3 convolutions with stride 2, producing a compact $14 \times 14 \times 256$ feature map. These embeddings are then processed by customized Transformer blocks, progressively refining the visual features.

We introduce key architectural modifications to the ViT model. First, the classification head is replaced with a task-specific multi-layer perceptron (MLP) to produce compact identity-preserving embeddings instead of classification logits. Second, the Transformer blocks adopt multi-head self-attention with learnable

relative positional biases, allowing translation equivariance while capturing geometric relationships. Specifically, the attention score between spatial positions (x, y) and (x', y') for attention head h is computed as:

$$A_{(x,y),(x',y')}^h = Q_{(x,y),:} \cdot K_{(x',y'),:} + B_{|x-x'|,|y-y'|}^h \quad (4)$$

This replaces fixed positional encodings with learnable biases conditioned on spatial displacement, improving generalization under viewpoint shifts and target displacements.

As shown in the fig.5, to further enhance efficiency, we adopt shrinking attention blocks that downsample the feature maps by applying strided operations on the queries (Q). This reduces the feature map size, with the output transforming from an input tensor of size (C, H, W) to $(2C, H/2, W/2)$. Lightweight MLP layers—implemented as 1×1 convolutions followed by batch normalization—replace dense layers, with the expansion ratio reduced from 4 to 2, optimizing the trade-off between parameter size and representational power.

The final Transformer output undergoes global average pooling, followed by a linear projection via the output MLP, resulting in a 512-dimensional embedding vector. This vector is designed to serve as a compact descriptor for identity discrimination in the appearance matching pipeline.

In terms of training objectives, we propose an adaptive angular margin loss (Ada-Loss) to replace conventional triplet loss, which is prone to overfitting and gradient instability in deep Transformer models. This loss operates on the hypersphere of L2-normalized features, enforcing angular separation between positive and negative pairs:

$$\theta_1 = \arccos(\vec{a} \cdot \vec{p}) + \theta, \theta_2 = \arccos(\vec{a} \cdot \vec{n}) \quad (5)$$

$$d_i = 1 - \cos \theta_i \quad (i = 1, 2) \quad (6)$$

$$\text{loss} = [\max(d_1 - d_2 + \alpha, 0)]^2 \quad (7)$$

Where θ_1 and θ_2 are the angles between the query vector $\vec{a}$ and the positive vector $\vec{p}$, and the negative vector $\vec{n}$, respectively. The margin α controls the angular separation between positive and negative pairs.

The gradient of the loss function, as shown in the following equation:

$$\nabla L = 2(d_1 - d_2 + \alpha)(\nabla d_1 - \nabla d_2)$$

This demonstrates the benefits of dynamic hard sample mining. The gradient magnitude increases with the severity of margin violations, ensuring that hard examples receive larger gradients, which accelerates convergence. Additionally, the loss incorporates adaptive

gradient decay, reducing updates for well-separated samples and addressing potential convergence issues in later training stages. This mechanism can be viewed as actively resampling difficult examples, enhancing the model's robustness and learning efficiency.

Furthermore, the angular formulation aligns with the L2-normalized output embeddings, and the margin parameter α allows fine-grained control over the compactness and discrimination of the learned feature space.

3.4. Clustering Optimized Feature Fusion

To address challenges in multi-object tracking (MOT) within maritime environments, we introduce a clustering optimized feature fusion (COFF) framework. This approach enhances robustness against occlusion, appearance ambiguity, and environmental degradation. Our method incorporates temporal smoothing, nonlinear feature scaling, and spatial gating, ensuring more reliable identity association under complex conditions.

In maritime scenarios, the relatively stable environment leads to low inter-frame motion, enabling the use of long-term appearance features to improve association reliability, even under partial occlusion or adverse weather. To exploit this continuity, we propose an Unbiased Exponential Moving Average (UEMA) for updating appearance representations of tracked targets. The UEMA applies exponential decay to historical features and incorporates a normalization factor to correct bias in early frames. Specifically, given the raw appearance feature f_i^k of target i at frame k , we compute the smoothed embedding e_i^k and its unbiased estimate $\hat{e}_i^k$ as follows:

$$e_i^k = \alpha e_i^{k-1} + (1 - \alpha) f_i^k \quad (8)$$

$$\hat{e}_i^k = \frac{e_i^k}{1 - \alpha^k} \quad (9)$$

Here, α denotes the temporal decay factor and the initial appearance embedding e_i^0 is set to zero. This formulation improves the temporal stability of appearance features without overly relying on noisy early observations, thus preserving identity continuity through prolonged occlusions or partial visual degradation.

Although temporal fusion improves feature consistency, we observe that cosine similarity, commonly used for appearance matching, often lacks sufficient discriminative power. Empirical analysis reveals that intra-class cosine distances average around 10^{-4} , while inter-class distances cluster around 10^{-2} . While this reflects a substantial relative difference, the small absolute gap limits the effectiveness of joint optimization with motion cues, which typically have distances in the




Figure 6: Impact of scaling on distance metrics in multi-object tracking. The left panel displays samples before scaling, where both positive (cyan) and negative (orange) samples are near the center (red star). After scaling (right panel), the distance between samples is magnified, clearly distinguishing positive from negative samples, thus enhancing tracking accuracy.

range of 10^{-1} . This imbalance leads to the dominance of appearance similarity terms, while the influence of motion information is suppressed in the overall association cost.

To address this issue, we propose a nonlinear amplification mechanism to rescale the cosine distance, improving its compatibility with motion-based metrics. Specifically, we multiply the original appearance distance by a scaling factor $\beta = 10^3$ and clip the result to a maximum of 1:

$$d_{i,j}^{\cos} = \min(d_{i,j}^{\cos} \cdot \beta, 1) \quad (10)$$

As illustrated in Fig. 6, this transformation amplifies intra-class appearance distances to approximately 10^{-1} , while inter-class distances move toward 1. After scaling, the similarity space becomes more separable, with positive pairs constrained within $[0, 0.2]$ and negative pairs concentrated near $[1 - \epsilon, 1]$, where ϵ is a small constant. This enhanced contrast significantly improves the ranking reliability of appearance-based associations.

To further suppress false positives caused by abrupt appearance changes or unreliable features, we introduce a spatial gating mechanism based on IoU. When the overlap between a track and a detection falls below a predefined threshold $\theta_{\text{iou}} = 0.3$, their appearance similarity is overridden to the maximum dissimilarity value:

$$d_{i,j}^{\cos} = 1, \quad \text{if } d_{i,j}^{\text{iou}} < \theta_{\text{iou}} \quad (11)$$

This gating strategy effectively filters out visually similar but spatially implausible matches, enhancing the robustness of the association process under occlusion or scene clutter.

Finally, the overall matching cost combines motion and appearance cues through a multiplicative formulation. For each pair of track i and detection j , the final cost $C_{i,j}$ is computed as:

$$C_{i,j} = d_{i,j}^{\cos} \cdot d_{i,j}^{\text{iou}} \quad (12)$$

This formulation ensures that both appearance dissimilarity and motion misalignment contribute to the association cost, making the system sensitive to both modalities. The scaling factor β enhances separability without disrupting the relative order of association costs. Empirical results confirm that across a range of β values (e.g., $\beta \in [200, 2000]$), intra-class costs remain low, while inter-class costs align with motion-based distances, leading to stable matching performance across diverse tracking conditions.

In summary, by combining UEMA-based temporal smoothing, nonlinear distance scaling, and spatial gating, our fusion strategy significantly enhances identity association in maritime MOT, particularly under occlusions, noise, and scale imbalances between motion and appearance cues.

4. Experiments

4.1. Experimental Setup

We evaluate our method on the Singapore Maritime Dataset (Prasad et al., 2017), a benchmark for



Figure 7: Four representative scenes from the Singapore Maritime Dataset illustrating the diversity of maritime conditions. The dataset encompasses a spectrum of scenarios, from night operations and foggy weather to open waters and crowded ports, capturing the complexity of real-world maritime environments.

maritime multi-object tracking (MOT). This dataset includes video sequences captured from static port surveillance cameras and mobile unmanned surface vehicles (USVs), featuring various environmental conditions (day, night, clear weather) at a resolution of 1920×1080 pixels. It introduces domain-specific challenges such as water reflections, dense vessel traffic, occlusions, and low-contrast targets, with annotations for seven maritime object categories: *ferry*, *buoy*, *ship*, *speedboat*, *small boat*, *kayak*, and *sailboat*.

Fig. 7 highlights the unique challenges of maritime environments, including densely packed small targets, motion blur, and severe occlusion. For object detection, we randomly sample temporally independent frames and divide them into training (60%), validation (20%), and testing (20%) subsets, ensuring that temporally adjacent frames are not split across subsets. For Re-Identification (ReID), the first 50% of frames in each video are used for training, while the remaining 50% are reserved for validation.

To comprehensively evaluate tracking performance on the Singapore Maritime Dataset (Capobianco et al., 2021; Choi et al., 2025; Ribeiro et al., 2017), we adopt the CLEAR MOT evaluation framework (Bernardin and Stiefelhagen, 2008), which encompasses a suite of widely adopted tracking metrics. These include Multiple Object Tracking Accuracy (MOTA), which penalizes false positives (FP), false negatives (FN), and identity switches (IDs); IDF1 (Ristani et al., 2016), which

measures identity preservation accuracy; and Higher Order Tracking Accuracy (HOTA) (Luiten et al., 2021), which jointly evaluates detection, association, and localization quality. Additionally, we report Fragmentations (Frag), Localization Accuracy (LocA), Identity Switches (IDs), and runtime efficiency (FPS) to provide a holistic view of tracker performance.

To further validate the generalization ability of our model—especially under dense and non-maritime conditions—we also evaluate it on the widely adopted MOT17 benchmark (Milan et al., 2016), which has become a de facto standard in the MOT community. MOT17 features dense pedestrian crowds with frequent occlusions and complex motion patterns, offering a challenging yet complementary testbed to maritime scenarios.

For fair comparison on MOT17, we follow the community-standard protocol by using the provided public Faster R-CNN detections. Moreover, in alignment with BOXMOT our chosen baseline, which reports only MOTA, IDF1, and HOTA—we restrict our evaluation on MOT17 to these three core metrics. This streamlined approach ensures a consistent and equitable comparison across methods.

The tracking metrics are computed using the following formulations:

$$\text{HOTA} = \sqrt{\text{DetA} \cdot \text{AssA}} = \sqrt{\frac{\sum_{c \in \text{TP}} A(c)}{\text{TP} + \text{FN} + \text{FP}}} \quad (13)$$

$$A(c) = \frac{\text{TPA}(c)}{\text{TPA}(c) + \text{FPA}(c) + \text{FNA}(c)} \quad (14)$$

$$\text{DetA} = \frac{|\text{TP}|}{|\text{TP}| + |\text{FN}| + |\text{FP}|} \quad (15)$$

$$\text{MOTA} = 1 - \frac{\sum_t (\text{FP}_t + \text{FN}_t + \text{IDS}_t)}{\sum_t G_t} \quad (16)$$

$$\text{IDF1} = \frac{2 \cdot \text{IDTP}}{2 \cdot \text{IDTP} + \text{IDFP} + \text{IDFN}} \quad (17)$$

$$\text{LocA} = \frac{1}{N} \sum_{i=1}^N \text{IoU}(GT_i, \text{Det}_i) \quad (18)$$

Here, N denotes the number of true positive detections, and $\text{IoU}(GT_i, \text{Det}_i)$ represents the intersection-over-union between the predicted and the ground-truth bounding boxes.

In summary, higher values of HOTA, MOTA, IDF1, and LocA indicate better tracking performance in terms of detection quality, identity association, and spatial localization. Conversely, lower values of FP, FN, Frag, and IDs correspond to more accurate and stable tracking. FPS serves as an indicator of system efficiency and real-time capability.

4.2. Implementation Details

All experiments are conducted on a workstation equipped with an NVIDIA GeForce RTX 4090D GPU, which provides a consistent and high-performance environment for model training and evaluation. To simulate deployment in resource-constrained maritime scenarios, the inference speed (FPS) in model comparison experiments is additionally measured on a laptop equipped with an NVIDIA GeForce RTX 4060 GPU. This setup demonstrates the real-time capability and efficiency of our method under limited computational resources.

The object detector is built upon the proposed Reversible Columnar Detection Network (RCDN), which enhances multi-scale feature fusion for robust maritime target detection. Feature maps extracted by RCDN are forwarded to the native decoupled classification head of YOLOv11 for object categorization. For appearance modeling, the Li-TAE module is adopted as the ReID backbone, designed to generate compact and discriminative identity embeddings.

The detector is trained for 300 epochs using the AdamW optimizer with a momentum of 0.937 and a batch size of 32. The initial learning rate is set to 0.01 and gradually reduced to 0.0001 following a cosine annealing schedule. A weight decay of 0.0005 is applied for L2 regularization, and a warm-up phase of 3 epochs is used at the beginning of training. The loss function adopts a weighted scheme, where the box regression loss is weighted at 7.5 and the classification loss at 0.5.

To improve robustness against occlusion and scale variation, data augmentation techniques including mosaic augmentation (probability = 1.0) and random rotation within $\pm 10^\circ$ are applied.

The ReID feature extractor is trained separately for 300 epochs using a hard triplet mining strategy to enhance inter-class separability and intra-class compactness. Input images are resized to 224×224 using bilinear interpolation and augmented with the rand-m9-mstd0.5 policy. Training employs an angular margin loss with a scaling factor $s = 64$ and a margin $m = 0.5$. The batch size is set to 64, and the learning rate is fixed at 10^{-3} , with a 30-epoch warm-up followed by cosine annealing.

During inference, the detection confidence threshold is set to 0.6. For data association, a dual-constraint matching scheme is employed. The linear matching stage uses an appearance similarity threshold of 0.25 and an IoU threshold of 0.75. A two-stage matching strategy is implemented with primary and secondary thresholds set to 0.5 and 0.1, respectively. The appearance distance is scaled by a factor of 800, and the final matching score threshold is 0.85. To handle short-term occlusions, a trajectory buffer of 30 frames is maintained. No post-processing is applied during evaluation to reflect real-world deployment conditions.

4.3. Comparison Experiment

In this study, we benchmark our proposed DMSORT against eight state-of-the-art trackers—BoT-SORT, OCSORT, DeepOCSORT, ImprAssoc, StrongSORT, ByteTrack, BoostTrack, and MOMT—using the Singapore Maritime Dataset (Table 1 and Fig. 8). To ensure a fair comparison, all trackers use the same YOLOv11n detector. We further evaluate an enhanced version of each method (denoted with $\star$) that integrates our proposed RCDN backbone and, where applicable, the Li-TAE appearance encoder.

It is noteworthy that all $\star$ -enhanced methods exhibit noticeable improvements across multiple tracking metrics compared to their vanilla counterparts, demonstrating the general effectiveness and plug-and-play capability of both the RCDN backbone and the Li-TAE module. For instance, DeepOCSORT $\star$ achieves a 3.41% increase in HOTA and 3.87% in MOTA, while BOT-SORT $\star$ improves IDF1 by 3.54% and reduces false negatives significantly. These consistent gains across different tracking paradigms highlight the broad applicability of our proposed components.

Our base DMSORT already demonstrates competitive performance, achieving the highest HOTA (73.44%) and lowest false negatives (16,479) among the original methods. More importantly, the enhanced DMSORT $\star$ consistently outperforms all $\star$ -enhanced

Table 1

MOT results on the Singapore Maritime Dataset. * indicates methods enhanced with RCDN and Li-TAE where Re-ID is available; RCDN only otherwise. Best scores before and after our improvements are bolded.

Method	Re-ID	HOTA↑	LocA↑	FN↓	FP↓	Frag↓	IDF1↑	IDs↓	MOTA↑	FPS↑
BOT-SORT Aharon et al. (2022)	✓	72.13	92.90	19611	8783	25	78.81	162	70.62	20.4
OCSORT (Cao et al., 2023)	–	71.18	92.54	19980	7421	135	78.43	180	70.62	47.4
DeepOCSORT (Maggiolino et al., 2023)	✓	72.48	91.61	17035	8323	114	80.44	186	73.74	25.9
ImprAssoc (Stadler and Beyerer, 2023)	✓	71.56	91.55	16505	8713	293	77.93	494	73.96	25.6
StrongSORT (Du et al., 2023)	✓	72.38	85.22	16819	8559	256	80.55	188	73.88	21.3
ByteTrack (Zhang et al., 2022)	–	72.15	91.72	18055	7972	45	79.81	180	72.77	46.0
BoostTrack (Xu et al., 2024)	✓	71.68	92.45	16611	8654	81	78.23	189	74.25	24.2
MOMT (Niu et al., 2024)	–	72.96	93.38	17319	8264	49	80.93	187	73.53	44.5
DMSORT (ours)	✓	73.44	92.98	16479	8597	53	80.90	172	74.23	28.1
BOT-SORT*	✓	75.47	92.80	14161	11311	52	82.35	179	75.92	24.2
OCSORT*	–	74.88	91.82	13551	10692	98	82.87	201	76.82	46.6
DeepOCSORT*	✓	75.89	91.96	12572	11321	100	82.96	195	77.61	29.6
ImprAssoc*	✓	75.53	91.92	12401	11753	202	82.06	377	77.51	28.3
StrongSORT*	✓	75.75	91.95	12549	11448	173	82.98	193	77.46	24.6
ByteTrack*	–	75.24	91.68	12659	11430	66	83.50	199	77.53	45.3
BoostTrack*	✓	75.20	92.11	13577	11388	91	83.11	191	76.89	27.2
MOMT*	–	75.41	92.24	12468	11642	84	83.61	217	77.67	43.8
DMSORT* (ours)	✓	76.69	93.25	13489	10425	40	83.67	196	77.60	30.6




Figure 8: Comparison of multi-object tracking algorithms on the Singapore Maritime Dataset (SMDS) and MOT17. Left: Results on SMDS using YOLOv11n with its default ReID model. Methods marked with an asterisk (*) are enhanced with RCDN and Li-TAE. Dashed ellipses group each method with its enhanced version for direct comparison; DMSORT denotes our complete method. Right: MOT17 results using detection and ReID models from BoxMOT without enhancement.

counterparts across most key metrics. It attains the best HOTA (76.69%), IDF1 (83.67%), and localization accuracy (LocA = 93.25%), while also yielding the fewest false negatives (13,489) and fragments (40). Although its MOTA (77.60%) is slightly lower than the top performer MOMT* (77.67%), the difference is marginal at only 0.07%.

Regarding localization accuracy (LocA), while RCDN is theoretically designed to enhance localization

through its reversible architecture and feature preservation, some ★-enhanced methods show slight decreases in LocA. This can be attributed to the increased detection sensitivity brought by RCDN, which may introduce more challenging localization cases (e.g., partially occluded or truncated objects) that slightly affect average precision. However, our DMSORT* still achieves the highest LocA (93.25%), demonstrating its superior localization capability.

Table 2

Comparison Experiment of DMSORT on the MOT17 Dataset (Best in Bold)

Tracker	Re-ID	HOTA $\uparrow$	MOTA $\uparrow$	IDF1 $\uparrow$
BoostTrack (Xu et al., 2024)	✓	69.25	75.91	83.21
BOT-SORT (Aharon et al., 2022)	✓	68.50	77.17	80.99
StrongSORT (Du et al., 2023)	✓	68.33	76.35	81.21
ByteTrack (Zhang et al., 2022)	—	66.54	76.91	77.86
ImprAssoc (Stadler and Beyerer, 2023)	✓	64.10	76.51	71.88
DeepOCSORT (Maggiolino et al., 2023)	✓	62.91	74.48	73.46
OCSORT (Cao et al., 2023)	—	65.19	74.82	75.96
DMSORT (ours)	✓	68.92	78.60	81.57

In terms of inference speed, while RCDN has smaller parameter counts and FLOPs than the baseline detector, its deeper architecture may result in slightly slower FPS compared to YOLOv11. However, the lightweight Li-TAE module significantly improves the inference speed for all Re-ID-based methods, as evidenced by the FPS improvements in most $\star$ -enhanced trackers.

For identity switches (IDs) and fragments (Frag), the $\star$ -enhanced versions generally exhibit increases in these metrics, which can be attributed to the substantially improved detection sensitivity of RCDN. While the higher recall successfully reduces false negatives (FN), it also introduces more detected objects—including those that are partially occluded, short-lived, or in close proximity—which inherently increases the complexity of data association. This, in turn, may lead to more frequent track fragmentations and identity switches. Nevertheless, this represents a favorable trade-off: the significant enhancement in recall and the reduction in missed detections far outweigh the moderate increase in IDs and Frag, ultimately contributing to more robust and continuous tracking in challenging maritime environments.

Overall, by combining high-recall detection, the compact yet powerful Li-TAE appearance encoder, and our proposed COFF module, DMSORT $\star$ achieves balanced and state-of-the-art performance in cluttered maritime scenes—excelling in detection, association, localization, and temporal stability.

To further assess the generalization capability of our maritime-specific tracker, we evaluate DMSORT on the widely adopted MOT17 benchmark. In line with BOXMOT, we report the three most critical tracking metrics—MOTA, HOTA, and IDF1—as shown in Table 2 and Fig. 8. Despite being tailored for maritime multi-object tracking, DMSORT demonstrates strong generalization to crowded, occlusion-heavy urban scenes. Specifically, it achieves the highest MOTA score (78.60), outperforming all other state-of-the-art methods, including those explicitly optimized for MOT17. Additionally, it ranks second in both HOTA




Figure 9: Distribution of positive and negative samples post-feature clustering on the Singapore Maritime Dataset. Positive samples (blue) cluster closely below 0.28, indicating high identity similarity. Negative samples (orange) predominantly gather near 1, showing clear separation. This distribution underscores the effectiveness of feature clustering in enhancing tracking accuracy and robustness.

(68.92) and IDF1 (81.57), closely trailing BoostTrack, which attains marginally higher values.

These results confirm that our modular design comprising high-recall detection, appearance modeling via Li-TAE, and robust motion association with COFF remains effective even under drastically different visual domains. The consistent performance across maritime and pedestrian tracking scenarios highlights DMSORT’s strong adaptability and robustness in multi-object tracking tasks.

4.4. Effectiveness Analysis of COFF

Fig. 9 illustrates the distribution of positive and negative samples during the matching process on the Singapore Maritime Dataset. The blue bars represent positive samples (matching the same identity), while the orange bars denote negative samples (non-matching identities). The x-axis represents the distance between samples, ranging from 0 to 1, while the y-axis shows the processed frequency on a logarithmic scale. The distances are divided into 300 intervals, with bins containing fewer than 8 samples excluded to minimize outlier effects.

The plot reveals that after applying feature clustering, negative samples predominantly cluster around a distance of 1, with over 95% falling within the [0.95, 1.0] range, indicating effective separation from the positive samples. In contrast, positive samples are tightly clustered within the [0, 0.28] range, reflecting high feature similarity among the same identity. The intermediate range (0.28, 0.75) contains sparse occurrences, with fewer than 8 samples per bin, suggesting minimal confusion between positive and negative samples. This




Figure 10: Evaluation of the DMSORT tracking method using IDF1 and HOTA metrics: left at different scaling factor thresholds during the matching phase, and right at different IOU thresholds during the detector inference phase with non-maximum suppression.

Table 3
Comparison of Model Performance in Detection and Tracking

Model	Parameters and FLOPs		Detection Performance				Tracking Performance		
	Params	FLOPs	Precision	Recall Rate	mAP 50%	mAP 50-95%	MOTA↑	IDF1↑	HOTA
RCDN (ours)	2.2 M	5.0 G	0.9316	0.8812	0.9332	0.7305	77.60	83.67	76.69
YOLOv11-CAF	3.5 M	10.9 G	0.9369	0.9025	0.9425	0.6900	76.24	83.41	76.68
YOLOv11-tiny	2.6 M	6.5 G	0.9264	0.8503	0.9116	0.6771	75.68	82.00	74.42
YOLOv11-small	9.4 M	21.6 G	0.9570	0.9165	0.9633	0.7698	79.38	85.19	79.38
YOLOv11-medium	20.1 M	68.4 G	0.9692	0.9305	0.9706	0.7608	79.91	85.82	79.20

clear separation between the two classes demonstrates that feature clustering significantly enhances the discriminative power of the tracker, improving both accuracy and robustness in multi-object tracking.

Fig. 10 presents IDF1 and HOTA as the amplification factor β increases from 200 to 2000 (a ten-fold span). Across this range, both metrics remain remarkably stable and consistently exceed the “no feature-scaling” baseline (HOTA = 75.6%, IDF1 = 82.8%). Specifically, HOTA varies only between 76.2% and 76.7%—an absolute gain of approximately 1% ,while IDF1 ranges from 83.0% to 83.7%, peaking at $\beta = 800$. Relative to the unscaled features, our clustering-optimized fusion yields a consistent improvement of about 1.0% in HOTA and 0.9% in IDF1, demonstrating that the proposed scaling is both low-sensitivity and reliably beneficial.

4.5. Detector Analysis

Under the Detection-Before-Tracking (DBT) paradigm, the fidelity of the object detector directly influences the downstream multi-object tracking (MOT) performance. As shown in Table 3, scaling the YOLOv11 family from tiny (6.5G FLOPs) to small (21.6G FLOPs) leads to substantial improvements: MOTA increases by 3.7% , IDF1 by 3.19% , and HOTA by 4.96%. These

consistent gains highlight the benefits of increased detector capacity in enhancing feature quality, which translates into better tracking completeness and identity preservation in complex maritime scenes.

However, further scaling from YOLOv11-small (21.6G FLOPs) to YOLOv11-medium (68.4G FLOPs) yields only marginal improvements: MOTA rises by just 0.53%, IDF1 by 0.63%, and HOTA shows negligible change (from 79.38 to 79.20). This diminishing return suggests that brute-force model enlargement quickly reaches a point of limited utility and highlights the need for architectural innovations that deliver stronger accuracy-efficiency trade-offs.

Although YOLOv11-CAF Hu et al. (2024) and our RCDN share the goal of enhancing information flow and preserving features, RCDN delivers superior multi-object tracking performance. Its reversible architecture ensures lossless feature transformation and yields higher mAP (0.7305 vs 0.6900) while requiring only 5.0 G FLOPs—less than half of YOLOv11-CAF’s 10.9 G—making it better suited for resource-constrained maritime scenarios.

To this end, we propose the Reversible Columnar Detection Network (RCDN), which, with only 5.0G FLOPs, achieves a HOTA of 76.69—just 2.69% below YOLOv11-small yet with less than one-quarter

of its computational cost. Compared to YOLOv11-tiny, RCDN improves HOTA by 2.27% while using fewer FLOPs. These results demonstrate that RCDN effectively bridges the gap between compact and mid-sized detectors. Its reversible columnar architecture and cross-scale feature reuse mechanism enable efficient bi-directional flow and fusion of multi-resolution features, capturing rich spatial semantics with minimal overhead. This lightweight yet performant design is well suited for maritime platforms with constrained resources, such as embedded vessel trackers and autonomous buoys, where both computational efficiency and robust multi-object tracking are essential.

Fig. 10 presents a sensitivity analysis of the model’s performance under varying Intersection-over-Union (IOU) thresholds, illustrating non-monotonic trends in both IDF1 and HOTA scores. At low IOU values (0.1–0.3), IDF1 ranges from 78.5% to 82.2% and HOTA from 71.9% to 75.5%, indicating that overly lenient matching allows more false positives, which degrades target discrimination and tracking accuracy. As the IOU threshold increases into the moderate range (0.4–0.6), both metrics improve markedly—IDF1 peaks at 83.7% and HOTA at 76.7%—showing that this range strikes an optimal balance between acceptance of true matches and rejection of false ones. Beyond 0.7, overly strict criteria begin to exclude valid associations, causing IDF1 to fall back to 81.7%–82.9% and HOTA to 75.3%–76.3%.

Taken together, these results highlight that an IOU threshold of 0.5 yields the best overall performance, achieving the highest IDF1 and HOTA scores while maintaining robust target discrimination and reliable tracking. This underscores the critical importance of selecting an appropriate IOU threshold to optimize multi-object tracking systems.

Through the visualization of layer heatmaps Fig. 11, our proposed RCDN demonstrates clear superiority in dense small-object detection, blurred scenarios, and standard maritime environments. Specifically, RCDN extends its receptive focus across multiple closely spaced targets—outperforming YOLOv11 in robustness and localization accuracy when vessels congregate—while simultaneously exhibiting heightened sensitivity to diminutive objects, such as small boats and fishing vessels, that are easily overlooked in low-contrast or noisy backgrounds. Moreover, RCDN substantially suppresses attention to irrelevant regions (e.g., sea surface and sky), thereby mitigating environmental noise effects and producing sharply concentrated attention zones even amid complex backgrounds like waves and cloud cover. These characteristics underscore RCDN’s enhanced precision and resilience in challenging maritime detection tasks.

4.6. Visualization Results’ Analysis

To further evaluate the robustness of our proposed framework in realistic maritime scenarios, we present a visual comparison of tracking performance across sequential frames in challenging scenes characterized by dense small targets and significant motion blur. As illustrated in Fig. 12, most existing trackers suffer from delayed or incomplete trajectory initialization—many small and visually degraded objects remain undetected or untracked until after frame 60, and some fail to form valid trajectories throughout the entire sequence. This reflects the limitations of their detection sensitivity and long-term association stability under severe maritime degradation. In contrast, our DMSORT successfully establishes stable and continuous trajectories as early as frame 50, including for low-resolution, occluded targets. This improvement stems from the synergy between two key innovations: the Reversible Columnar Detection Network (RCDN), which ensures high recall through progressive multi-scale feature refinement, and the clustering optimized feature fusion (COFF) strategy, which enables resilient data association by adaptively balancing motion and appearance cues through nonlinear fusion and spatial gating. Together, these components allow DMSORT to maintain identity consistency, accelerate trajectory formation, and reduce fragmentation in highly dynamic maritime conditions. This advantage is particularly evident in early frames where competing methods either delay tracking or lose targets altogether, confirming the efficacy of our detection-association design in supporting persistent maritime multi-object tracking.

The Fig. 13 provides a visual comparison of different detection and tracking algorithms using Kalman filter updates in a bright maritime environment. From this, we can clearly observe that our method outperforms the others in terms of detection, particularly in higher recall rates and more accurate bounding boxes. Specifically, the RCDN’s reversible architecture and precise feature representation enable better detection of small maritime targets, significantly boosting the recall rate. Additionally, our method benefits from robust Kalman filtering, which is facilitated by accurate camera motion estimation and stable detection. The camera estimation effectively reduces errors in the Kalman filter update process, leading to more reliable tracking and minimizing trajectory fragmentation. As a result, the Kalman filter becomes more stable, reducing accumulated uncertainty and improving detection accuracy. In contrast, other methods suffer from issues such as trajectory fragmentation and incorrect associations, which destabilize the Kalman filter process, introducing noise that disrupts target motion estimation.




Figure 11: Comparison of detection heatmaps under various maritime conditions using the proposed Reversible Columnar Detection Network (RCDN) and YOLOv11n. Scenarios include normal weather, fog, and dense small targets.

The Fig. 14 illustrates the attention map produced by the Li-TAE model on ship images extracted from the Singapore Maritime Dataset. These attention maps are generated using multi-head averaging to highlight the focus of the model on key vessel features. After the ships are detected and cropped, they are fed into

the Li-TAE model to extract their re-identification features. The attention map, generated using a multi-head attention strategy, reveals how the model allocates its focus on different regions of the ship. As shown in the figure, our Li-TAE model, after being trained with adaptive hard triplets, effectively captures the vessel's key features with near segmentation-level accuracy. This capability is a result of the global receptive




Figure 12: Visual comparison of different detection and tracking algorithms with Kalman filter updates in a bright environment featuring many objects of varying sizes intertwined. Yellow dashed boxes indicate lost tracking objects.

Table 4
Ablation Study of DMSORT on the Singapore Maritime Dataset

Method	Li-TAE	DDTA	COFF	RCDN	HOTA↑	MOTA↑	IDF1↑	FPS↑
Baseline (BOT-SORT)	—	—	—	—	72.13	70.62	78.81	20.4
Baseline + col. 1	✓	—	—	—	72.96	74.00	80.67	24.6
Baseline + col. 1–2	✓	✓	—	—	72.98	74.00	80.68	31.2
Baseline + col. 1–3	✓	✓	✓	—	74.42	75.68	82.00	31.1
Baseline + col. 1–4	✓	✓	✓	✓	76.69	77.60	83.67	30.6

field facilitated by self-attention mechanisms, which works exceptionally well in maritime environments. The model demonstrates robust feature extraction, ensuring excellent performance even under challenging oceanic conditions.

4.7. Ablation Study

As shown in Table 4, introducing Li-TAE immediately enhances feature discrimination—MOTA jumps by 3.38% and IDF1 by 1.86% thanks to its global receptive field and adaptive angular margin

loss. Adding DDTA further stabilizes long-term track continuity under affine transformations while boosting throughput to 31.2 FPS. With COFF, the fusion of clustered motion cues and appearance embeddings yields an extra +1.44% in HOTA and +1.32% in IDF1, reducing identity switches in crowded scenes. Finally, RCDN’s reversible, columnar detection–tracking co-design elevates detection reliability, driving HOTA up to 76.69% and IDF1 to 83.67%. This incremental ablation confirms that each innovation—appearance encoding, affine-aware tracking, motion–appearance




Figure 13: Visual comparison of different detection and tracking algorithms with Kalman filter updates in a bright environment featuring many objects of varying sizes intertwined. Yellow dashed boxes indicate lost tracking objects.

fusion, and joint optimization—contributes uniquely to our system’s robust, real-time MOT performance.

5. Conclusion

This paper presents DMSORT, a domain-adapted multi-object tracking framework specifically designed for the challenges of maritime environments, including severe occlusion, target scale variation, camera motion, and visual degradation. The proposed system integrates four tightly coupled innovations: a Dual-Branch Detection-Tracking Architecture (DDTA) for real-time motion-aware tracking, a Reversible Columnar Detection Network (RCDN) for multi-scale feature refinement, the Li-TAE Transformer-based appearance encoder for robust identity modeling, and a Clustering Optimized Feature Fusion (COFF) strategy for resilient association under uncertainty.

Through extensive evaluation on the Singapore Maritime Dataset, DMSORT sets a new state-of-the-art with 76.69 HOTA, 83.67 IDF1, and 77.60 MOTA, outperforming existing SOTA trackers in identity consistency, detection quality, and trajectory stability.

Notably, DMSORT also achieves strong cross-domain generalization on MOT17, surpassing urban-focused methods in MOTA (78.60) and performing competitively in HOTA and IDF1, despite its maritime specialization.

These results validate our key design principles: decoupling detection and tracking for parallelism and motion compensation, enhancing detection with reversible multi-scale backbones, leveraging global attention for discriminative Re-ID, and fusing cues through a learned, nonlinear association space. Together, these modules form a modular and efficient tracking solution with broad applicability, from maritime surveillance to general-purpose MOT.

In future work, we plan to explore scalable multi-camera deployment, incorporate temporal modeling for long-term occlusion handling, and extend the framework toward radar-visual multimodal fusion. This integration aims to leverage radar’s robustness in adverse conditions and long-range perception, complementing RGB cues for reliable all-weather maritime tracking.



Figure 14: Visualization of the LiTAE attention map on sample frames from the Singapore Maritime Dataset. For each pair of images, the left one is the original frame, and the right one is the corresponding attention - mapped frame. In the attention - mapped frames, regions with larger attention weights are relatively brighter, while regions with smaller attention weights are darker.

References

- Agarwala, N., 2020. Monitoring the ocean environment using robotic systems: Advancements, trends, and challenges. *Mar. Technol. Soc. J.* 54(5), 42–60.
- Aharon, N., Orfaig, R., Bobrovsky, B.Z., 2022. Bot-sort: Robust associations multi-pedestrian tracking. *arXiv preprint arXiv:2206.14651*.
- An, Y., Wu, J., Cui, Y., Hu, H., 2023. Multi-object tracking based on a novel feature image with multi-modal information. *IEEE Trans. Veh. Technol.* 72(8), 9909–9921.
- Arulampalam, S.S., Maskell, S., Gordon, N., Clapp, T., 2002. A tutorial on particle filters for online nonlinear/non-gaussian bayesian tracking. *IEEE Trans. Signal Process.* 50(2), 174–188.
- Berclaz, J., Fleuret, F., Turetken, E., Fua, P., 2011. Multiple object tracking using k-shortest paths optimization. *IEEE Trans. Pattern Anal. Mach. Intell.* 33(9), 1806–1819.
- Bernardin, K., Stiefelwagen, R., 2008. Evaluating multiple object tracking performance: the clear mot metrics. *EURASIP Journal on Image and Video Processing* 2008, 1–10.
- Bewley, A., Ge, Z., Ott, L., Ramos, F., Upcroft, B., 2016. Simple online and realtime tracking, in: 2016 IEEE international conference on image processing (ICIP), Ieee. pp. 3464–3468.
- Bolme, D.S., Beveridge, J.R., Draper, B.A., Lui, Y.M., 2010. Visual object tracking using adaptive correlation filters, in: *CVPR*.
- Bouguet, J.Y., et al., 2001. Pyramidal implementation of the affine lucas kanade feature tracker description of the algorithm. *Intel Corpor.* 5(1-10), 4.
- Cai, Y., Zhou, Y., Han, Q., Sun, J., Kong, X., Li, J., Zhang, X., 2022. Reversible column networks. *arXiv preprint arXiv:2212.11696*.
- Cao, J., Pang, J., Weng, X., Khiroudkar, R., Kitani, K., 2023. Observation-centric sort: Rethinking sort for robust multi-object tracking, in: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 9686–9696.
- Cao, X., Zheng, Y., Yao, Y., Qin, H., Cao, X., Guo, S., 2025. Topic: A parallel association paradigm for multi-object tracking under complex motions and diverse scenes. *IEEE Trans. Image Process.* 34, 743–758.
- Capobianco, S., Millefiori, L.M., Forti, N., Braca, P., Willett, P., 2021. Deep learning methods for vessel trajectory prediction based on recurrent neural networks. *IEEE Trans. Aerosp. Electron. Syst.* 57(6), 4329–4346.
- Chen, C., Ma, F., Wang, K.L., Liu, H.H., Zeng, D.H., Lu, P., 2025. Shipmot: A robust and reliable cnn-nsa filter framework for marine radar target tracking. *Electronics* 14(8), 1492.
- Cheng, X., Zhao, H., Deng, Y., Shen, S., 2025. Multi-object tracking with predictive information fusion and adaptive measurement noise. *Appl. Sci.* 15(2), 736.
- Choi, J., Cho, D., Lee, G., Kim, H., Yang, G., Kim, J., Cho, Y., 2025. Polaris dataset: A maritime object detection and tracking dataset in pohang canal, in: 2025 IEEE International Conference on Robotics and Automation (ICRA), IEEE. pp. 13626–13632.

- Cybenko, G., 1989. Approximation by superpositions of a sigmoidal function. *Mathematics of control, signals and systems* 2(4), 303–314.
- Danelljan, M., Hager, G., Khan, F.S., Felsberg, M., 2014. Accurate scale estimation for robust visual tracking, in: *British Machine Vision Conference (BMVC)*, pp. 1–11.
- Du, Y., Zhao, Z., Song, Y., Zhao, Y., Su, F., Gong, T., Meng, H., 2023. Strongsort: Make deepsort great again. *IEEE Trans Multimedia* 25, 8725–8737.
- FISCHLER AND, M., 1981. Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. *Commun. ACM* 24(6), 381–395.
- Fortmann, T., Bar-Shalom, Y., Scheffe, M., 1983. Sonar tracking of multiple targets using joint probabilistic data association, in: *Proceedings of the 22nd IEEE Conference on Decision and Control*, IEEE. pp. 807–812.
- Graham, B., El-Nouby, A., Touvron, H., Stock, P., Joulin, A., Jégou, H., Douze, M., 2021. Levit: a vision transformer in convnet's clothing for faster inference, in: *Proceedings of the IEEE/CVF international conference on computer vision*, pp. 12259–12269.
- Guo, Y., Shen, Q., Ai, D., Wang, H., Zhang, S., Wang, X., 2024. Sea-ioutracker: A more stable and reliable maritime target tracking scheme for unmanned vessel platforms. *Ocean Eng.* 299, 117243.
- Henriques, J.F., Caseiro, R., Martins, P., Batista, J., 2012. Exploiting the circulant structure of tracking-by-detection with kernels, in: *ECCV*.
- Henriques, J.F., Caseiro, R., Martins, P., Batista, J., 2015. High-speed tracking with kernelized correlation filters. *IEEE TPAMI*.
- Hermans, A., Beyer, L., Leibe, B., 2017. In defense of the triplet loss for person re-identification. *arXiv preprint arXiv:1703.07737*.
- Hornik, K., Stinchcombe, M., White, H., 1989. Multilayer feedforward networks are universal approximators. *Neural Netw.* 2(5), 359–366.
- Hu, S., Gao, F., Zhou, X., Dong, J., Du, Q., 2024. Hybrid convolutional and attention network for hyperspectral image denoising. *IEEE Geosci. Remote Sens. Lett.* doi:10.1109/LGRS.2024.3393860.
- Hu, Z., Xu, Q., 2025. Usv trajectory optimization based legitimate maritime wireless surveillance, in: *2025 IEEE/CIC International Conference on Communications in China (ICCC)*, IEEE. pp. 1–6.
- Jiang, J., Wang, Z., Zhao, M., Li, Y., Jiang, D., 2025. Sam2mot: A novel paradigm of multi-object tracking by segmentation. *arXiv preprint arXiv:2504.04519*.
- Jin, J., Zhang, J., Liu, D., Shi, J., Wang, D., Li, F., 2020. Vision-based target tracking for unmanned surface vehicle considering its motion features. *Ieee Access* 8, 132655–132664.
- Kalman, R.E., 1960. A new approach to linear filtering and prediction problems. *J. Basic Eng.* 82(1), 35–45.
- Langsdale, M., Verhoelst, T., Povey, A., Schutgens, N., Dowling, T., Lambert, J.C., Compernelle, S., Kern, S., 2025. The challenges and limitations of validating satellite-derived datasets using independent measurements: lessons learned from essential climate variables. *Surv. Geophys.* , 1–38.
- Li, J., Wang, M., Gong, X., 2023. Transformer based multi-grained features for unsupervised person re-identification, in: *Proceedings of the IEEE/CVF winter conference on applications of computer vision(WACV)*, pp. 42–50.
- Lin, T.Y., Dollár, P., Girshick, R., He, K., Hariharan, B., Belongie, S., 2017. Feature pyramid networks for object detection, in: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 2117–2125.
- Liu, Z., Mu, X., Lu, Y., Zhang, T., Tian, Y., 2023. Learning transformer-based attention region with multiple scales for occluded person re-identification. *Comput. Vis. Image Underst.* 229, 103652.
- Luiten, J., Osep, A., Dendorfer, P., Torr, P., Geiger, A., Leal-Taixé, L., Leibe, B., 2021. Hota: A higher order metric for evaluating multi-object tracking. *Int. J. Comput. Vis.* 129, 548–578.
- Luo, W., Xing, J., Milan, A., Zhang, X., Liu, W., Kim, T.K., 2021. Multiple object tracking: A literature review. *Artif. Intell.* 293, 103448.
- Maggiolino, G., Ahmad, A., Cao, J., Kitani, K., 2023. Deep oc-sort: Multi-pedestrian tracking by adaptive re-identification, in: *2023 IEEE International conference on image processing (ICIP)*, IEEE. pp. 3025–3029.
- Milan, A., Leal-Taixé, L., Reid, I., Roth, S., Schindler, K., 2016. Mot16: A benchmark for multi-object tracking. *arXiv preprint arXiv:1603.00831*.
- Niu, L., Fan, Y., Liu, T., Han, Q., 2024. Momt: A maritime real-time visual multi-object tracking algorithm based on unmanned surface vehicles. *IEEE Sens. J.*.
- Pirsiavash, H., Ramanan, D., Fowlkes, C., 2011. Globally-optimal greedy algorithms for tracking a variable number of objects, in: *2011 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, IEEE. pp. 1201–1208.
- Prasad, D.K., Rajan, D., Rachmawati, L., Rajabally, E., Quek, C., 2017. Video processing from electro-optical sensors for object detection and tracking in a maritime environment: A survey. *IEEE Trans. Intell. Transp. Syst.* 18(8), 1993–2016.
- Rangesh, A., Trivedi, M., 2019. No blind spots: Full-surround multi-object tracking for autonomous vehicles using cameras and lidars. *IEEE Trans. Intell. Veh.* 4(4), 588–599.
- Reid, D.B., 1979. An algorithm for tracking multiple targets. *IEEE Trans. Automat. Contr.* 24(6), 843–854.
- Ribeiro, R., Cruz, G., Matos, J., Bernardino, A., 2017. A data set for airborne maritime surveillance environments. *IEEE TCSVT.* 29(9), 2720–2732.
- Ristani, E., Solera, F., Zou, R., Cucchiara, R., Tomasi, C., 2016. Performance measures and a data set for multi-target, multi-camera tracking, in: *European conference on computer vision*, Springer. pp. 17–35.
- Shao, Z., Yin, Y., Lyu, H., Soares, C.G., 2024. A robust method for multi object tracking in autonomous ship navigation systems. *Ocean Eng.* 311, 118560.
- Shi, J., et al., 1994. Good features to track, in: *1994 Proceedings of IEEE conference on computer vision and pattern recognition*, IEEE. pp. 593–600.
- Soldi, G., Gaglione, D., Forti, N., Di Simone, A., Daffinà, F.C., Bottini, G., Quattrocioni, D., Millefiori, L.M., Braca, P., Carniel, S., et al., 2021. Space-based global maritime surveillance. part i: Satellite technologies. *IEEE Aerosp. Electron. Syst. Mag.* 36(9), 8–28.
- Stadler, D., Beyerer, J., 2023. An improved association pipeline for multi-person tracking, in: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 3170–3179.
- Stanojevic, V.D., Todorovic, B.T., 2024. Boosttrack: Boosting the similarity measure and detection confidence for improved multiple object tracking. *Machine Vision Appl.* 35(3), 53.
- Su, H., Arakawa, S., Murata, M., 2025. Cooperative 3d multi-object tracking with cross-agent data association for connected and automated vehicles. *IEEE Trans. Veh. Technol.*.
- Sun, P., Cao, J., Jiang, Y., Yuan, Z., Bai, S., Kitani, K., Luo, P., 2022. Dancetrack: Multi-object tracking in uniform appearance and diverse motion, in: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 20993–21002.
- Tian, S., Duan, M., Deng, J., Luo, H., Hu, Y., 2024. Mf-net: A multimodal fusion model for fast multi-object tracking. *IEEE Trans. Veh. Technol.* 73(8), xxx–xxx.
- Tran, T.H., Le, T.L., 2016. Vision based boat detection for maritime surveillance, in: *2016 International Conference on Electronics,*

- Information, and Communications (ICEIC), IEEE. pp. 1–4.
- Wang, N., Wang, Y., Wei, Y., Han, B., Feng, Y., 2024. Marine vessel detection dataset and benchmark for unmanned surface vehicles. *Appl. Ocean Res.* 142, 103835.
- Wojke, N., Bewley, A., Paulus, D., 2017. Simple online and realtime tracking with a deep association metric, in: 2017 IEEE international conference on image processing (ICIP), IEEE. pp. 3645–3649.
- Xu, Z., Huang, C., Zeng, W., Zhang, Y., Hu, W., 2024. Boosttrack++: using tracklet information to detect more objects in multiple object tracking. *arXiv preprint arXiv:2408.13003*.
- Yu, C., Yin, H., Rong, C., Zhao, J., Liang, X., Li, R., Mo, X., 2024. Yolo-mrs: An efficient deep learning-based maritime object detection method for unmanned surface vehicles. *Appl. Ocean Res.* 153, 104240.
- Zhang, L., Li, Y., Nevatia, R., 2008. Global data association for multi-object tracking using network flows, in: 2008 IEEE Conference on Computer Vision and Pattern Recognition, IEEE. pp. 1–8.
- Zhang, Y., Sun, P., Jiang, Y., Yu, D., Weng, F., Yuan, Z., Luo, P., Liu, W., Wang, X., 2022. Bytetrack: Multi-object tracking by associating every detection box, in: European conference on computer vision, Springer. pp. 1–21.
- Zhang, Y., Wang, C., Wang, X., Zeng, W., Liu, W., 2021b. Fairmot: On the fairness of detection and re-identification in multiple object tracking. *Int. J. Comput. Vis.* 129, 3069–3087.
- Zhang, Y., Wei, L., Zhang, B., Liu, Z., Yi, K., Hu, S., 2025. A lightweight convolution and vision transformer integrated model with multi-scale self-attention mechanism. *Neurocomputing*, 131670.