

MM-Tracker: Motion Mamba for UAV-platform Multiple Object Tracking

Mufeng Yao^{1*}, Jinlong Peng^{2*}, Qingdong He², Bo Peng³, Hao Chen¹, Mingmin Chi^{1†}, Chao Liu^{1†},
Jon Atli Benediktsson^{4†}

¹School of Computer Science, Shanghai Key Laboratory of Data Science, Fudan University

²Tencent YouTu Lab

³Shanghai Ocean University

⁴University of Iceland

mfyao21@m.fudan.edu.cn, jeromepeng@tencent.com, yingcaihe@tencent.com, bpeng@shou.edu.cn,
22210240115@m.fudan.edu.cn, mmchi@fudan.edu.cn, chaoliu20@fudan.edu.cn, benedikt@hi.is

Abstract

Multiple object tracking (MOT) from unmanned aerial vehicle (UAV) platforms requires efficient motion modeling. This is because UAV-MOT faces both local object motion and global camera motion. Motion blur also increases the difficulty of detecting large moving objects. Previous UAV motion modeling approaches either focus only on local motion or ignore motion blurring effects, thus limiting their tracking performance and speed. To address these issues, we propose the Motion Mamba Module, which explores both local and global motion features through cross-correlation and bi-directional Mamba Modules for better motion modeling. To address the detection difficulties caused by motion blur, we also design motion margin loss to effectively improve the detection accuracy of motion blurred objects. Based on the Motion Mamba module and motion margin loss, our proposed MM-Tracker surpasses the state-of-the-art in two widely open-source UAV-MOT datasets.

Code — <https://github.com/YaoMufeng/MMTracker>

Introduction

Multiple object tracking (MOT) aims to identify objects in a given video and is widely used in computer vision (Luo et al. 2021; Kalake, Wan, and Hou 2021), including applications in autonomous driving (Geiger et al. 2013), human-computer interaction (Chandra et al. 2015), and pedestrian tracking (Peng et al. 2020b; Meinhardt et al. 2022). A typical MOT approach consists of two primary stages: detection and association (Shuai et al. 2021). The detection stage identifies all objects in each frame, while the association stage matches objects across consecutive frames to establish complete trajectories for each object (Liu et al. 2022a). Recently, MOT in unmanned aerial vehicle (UAV) platforms has attracted extensive research interest. Compared with conventional MOT (Dendorfer et al. 2020), MOT in UAV view faces more challenges. Firstly, both the ground object motion and the aerial camera motion exist in the UAVMOT

* These authors contributed equally.

† Corresponding authors.

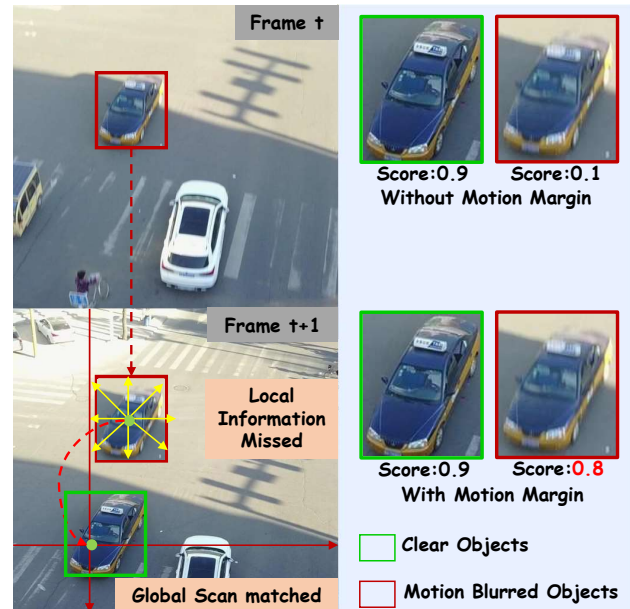


Figure 1: Advantages of our MM-Tracker. Left: a global camera motion was experienced during frame t and $t + 1$, our global scan matched the object that the local correlation missed. Right: with our motion margin loss, the detection score of motion-blurred object is increased.

task, posing a challenge for the tracker to follow the object accurately. Secondly, the severe motion blur (Kurimo et al. 2009) caused by camera motion will significantly increase the difficulty of object detection.

Kalman filtering and its various variants are very commonly used motion modeling schemes (Bewley et al. 2016; Cao et al. 2023; Zhang et al. 2022; Liu et al. 2023). However, the Kalman filtering scheme is a non-learning algorithm, which can only make pre-assumptions about the motion (e.g., assume linear motion), therefore, the accuracy is limited in camera motion scenes. Several studies (Shuai et al. 2021; Zhou, Koltun, and Krähenbühl 2020; Yao et al. 2023) propose learning-based motion modeling, but most of

them are based on local cross-correlation or local convolution and ignore global motion information. Therefore, the lack of global motion modeling limits the tracking accuracy of these trackers in scenes with significant global camera motion. Furthermore, previous studies have ignored the motion long-tailed distributions present in UAV-MOT datasets. Consequently, objects with significant motion, which are difficult to detect, receiving less training compared to easy-to-detect objects that exhibit minimal motion.

To address these issues, we propose MM-Tracker for fast and accurate motion modeling and motion-blur-focused object detector training. Unlike previous work that redundantly extracts object detection features and motion features from raw input images, MM-Tracker estimates object motion from bi-temporal object detection features, which greatly reduces the computational effort of motion modeling. In addition, MM-Tracker improves the detection performance in UAV-MOT scenes by proposing a Motion Margin loss function to impose a larger classification loss for objects with larger motions, thus effectively improves the detection of motion-blurred objects. The main contributions of this paper are summarized as follows:

- We propose the Motion Mamba module, which models object motion by local correlation of detection features and global scan of bi-directional mamba block, reached fast and accurate motion modeling.
- We propose the Motion Margin loss (MMLoss), which imposes a large decision boundary for large motion objects to supervise the object detector, and improves the detection performance on motion blurred objects.
- The MM-Tracker proposed based on the Motion Mamba module and Motion Margin loss reached state-of-the-art on two public UAV-MOT datasets, which promoted the progress of multiple object tracking in UAV scenes.

Related Works

Motion Modeling in MOT

Existing MOT methods typically follow the detection and association scheme (Luo et al. 2021; Fagot-Bouquet et al. 2016; Peng et al. 2020c,a). The association criterion can be classified into two types: positional information (Zhao, Gong, and Medioni 2012; Forsyth et al. 2006; Takala and Pietikainen 2007) and appearance information (Sugimura et al. 2009; Li et al. 2013; Yang and Nevatia 2012; Peng et al. 2018). Association based on positional information needs to model the object motion to predict the position of the object at the next moment (Bewley et al. 2016; Shuai et al. 2021; Zhou, Koltun, and Krähenbühl 2020; Zhang et al. 2022; Bochinsk, Eiselein, and Sikora 2017), while methods based on appearance information need to extract the appearance information of the object to compute the appearance similarity of the object between frames (Wojke, Bewley, and Paulus 2017; Zheng et al. 2016; Wang et al. 2020b; Zhang et al. 2021). The small object size, blurred object appearance, and very similar appearance between ground vehicles make appearance information unreliable for object association in UAV-MOT scenes (Yao et al. 2023). Therefore, this work focuses on motion modeling as association.

The most commonly used motion modeling methods in MOT are Kalman filtering (KF) (Kalman 1960) and its various variants (Bewley et al. 2016; Cao et al. 2023; Zhang et al. 2022). These methods require prior assumptions about object motion patterns and their probability distribution without any learnable parameters, which limits their performance in scenes with human-made global camera motions. Several studies (Zhou, Koltun, and Krähenbühl 2020; Shuai et al. 2021; Yao et al. 2023) introduce learnable neural networks to learn the motion of objects. These methods perform better than those based on preset rules. However, they mainly consist of local convolution and cross-correlation, therefore limited for global motions. In addition, these methods also have the problem of repeated feature extraction, which introduces redundant computations.

Effective Global Information Aggregation

Several studies (Azad et al. 2019; Wang, Yang, and Ning 2018; Huang et al. 2021) combine sequence models (e.g. RNN (Elman 1990), LSTM (Hochreiter and Schmidhuber 1997)) with convolution to extract global image features. However, their recursive inference scheme makes them difficult to train in parallel, resulting in slower training speed. Their stacked sequence architectures are also prone to problems such as gradient explosion or vanishing. By introducing a global attention mechanism, Transformer model (Vaswani et al. 2017) effectively overcomes the problems of low training parallelism and gradient vanishing, therefore achieving success in long sequence text tasks such as language modeling (Subakan et al. 2021). However, for global information modeling of images (Choi et al. 2020; Liu et al. 2021, 2022b), the computational complexity of Transformer is too large, for example $O(n^2)$, which is prone to overfitting and difficult to reach real-time tracking. The emergence of Mamba (Gu and Dao 2023) brings a more efficient global information extraction scheme, which can realize linear-time global attention calculation while also performing parallel training well. Therefore, this paper adopts the bidirectional scanning Mamba structure to compute global information aggregation and thus predict long-range object motion better.

Imbalance Training Loss

The conventional imbalance training (Wang et al. 2020a) refers to the imbalance between different classes that multiple classes only account for a small proportion of the dataset, while a few classes account for a large proportion of the dataset. This imbalance results in classes with few training samples not being adequately trained, affecting the classification accuracy. Many scholars have proposed solutions to this problem, such as resampling (Huang et al. 2016; Li et al. 2020; Ren et al. 2020), Loss reweighting (Lin et al. 2017; Li et al. 2021, 2022), re-margining (Cao et al. 2019; Wang et al. 2018; Menon et al. 2020).

In the UAVMOT scene, the camera's perspective change will cause large movements of objects, and this perspective change is accidental, making this situation account for a small proportion of the dataset. However, the large motion introduces severe motion blur on objects, requiring us

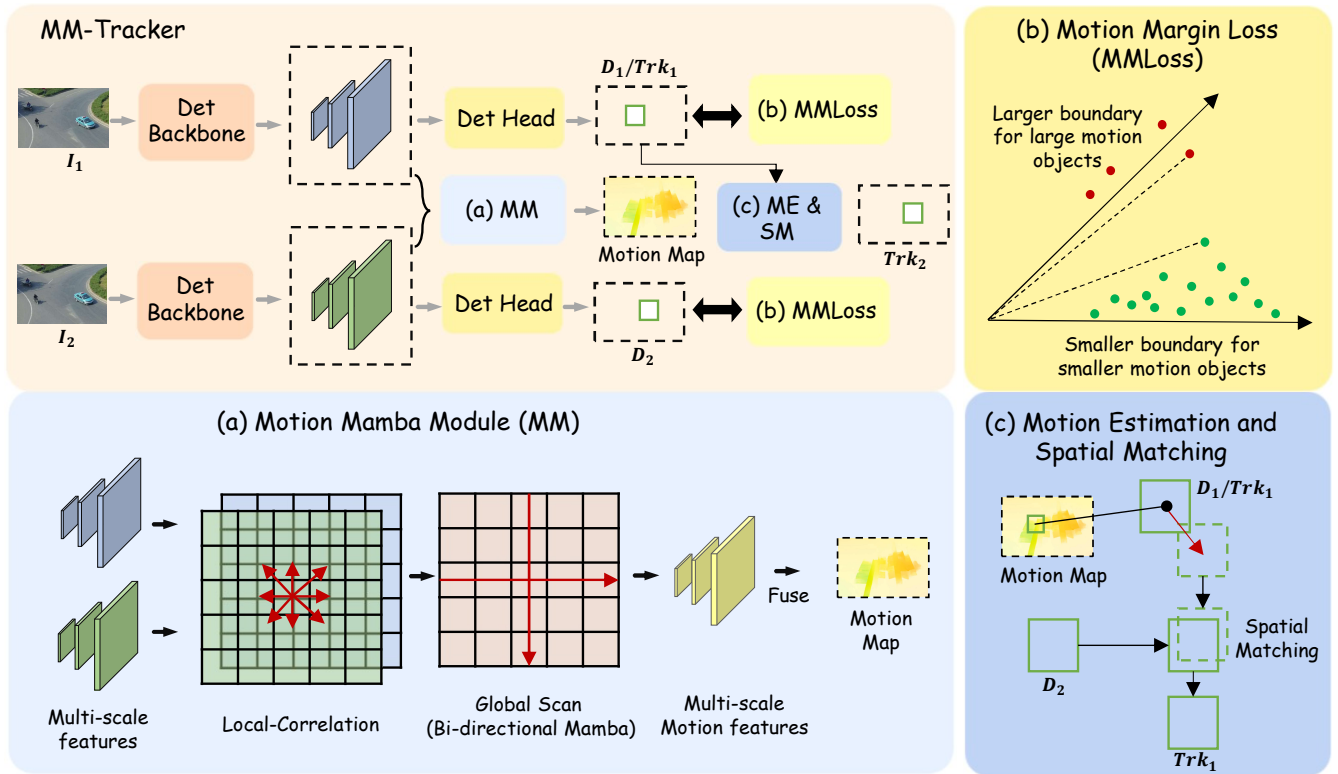


Figure 2: Overall architecture of MM-Tracker.

to focus more on those difficult-to-detect objects, which is missed in previous studies. To this end, we propose Motion Margin loss, which imposes larger classification boundaries for objects with larger motions, thus better solving the problem of less training for large motion objects.

Proposed Method

Overview

Aiming at efficient global motion modeling with a fast tracking speed, we design a Motion Mamba block with a feature reuse scheme. Furthermore, to better detect motion-blurred objects, we propose Motion Margin loss function to supervise the classification branch of the object detector. Fig. 2 shows the overall architecture of our proposed MM-Tracker. The DetBackbone extracts detection features from each frame of the video, which are input to DetHead for object detection and Motion Mamba module for motion feature extraction. The DetBackbone generates three scale feature maps, whose sizes are $1/8$, $1/16$, and $1/32$ of the original image sizes. The DetHead detects objects at these three scales respectively. The classification branch of the detector is supervised by the proposed MMLoss, and the regression branch is supervised by IOU Loss and L1 Loss. The Motion Mamba module extracts motion features from bi-temporal detection features at three scales respectively. The module then upsamples the features from the lowest scale step by step and fuses it with the higher scale, then generates a mo-

tion feature map with a size of $1/8$ of the original image. Finally, this feature map is supervised by L1 Loss with the ground-truth motion map. The generation way of ground-truth motion map is illustrated in Fig. 4. The DetBackbone and DetHead are YOLOX’s backbone and YOLOX’s head respectively.

Motion Mamba Module

The Motion Mamba module is dedicated to efficient and lightweight motion modeling. Therefore, the module extracts motion features from the existing detection features of the previous and next images, which greatly reduces the computational complexity of the model. The Motion Mamba module extracts motion features at three scales. These multi-scale features are then fused step by step from the lowest scale upwards, and finally outputs a motion feature map with a size of $1/8$ of the original image. For each scale, Motion Mamba first uses the cross-correlation of the feature maps at the previous and next moments to extract the local motion information, and then use Motion Mamba block to extract global motion features. As Fig. 3 shows, each Motion Mamba block consists of two branches: vertical state space model (V-SSM) and horizontal state space model (H-SSM). These two branches conduct selective scanning (Gu and Dao 2023) in vertical and horizontal directions respectively. In order to fully scan the entire feature map while reducing the computational complexity, each scan uses the vector at the pixel position as the basic unit. Specifically, for a feature

map with height, width, and channel (H, W, C) respectively, the vertical scan starts from the first row of each column, with a feature vector of size (1, 1, C) as input, and scans to the final row. The scanning of each column is performed separately, and a total of W groups of scans are performed. The horizontal SSM scans a total of H groups, follow the same scheme as the vertical scan. The vertical-scanned feature map and the horizontal-scanned feature map are added to realize global feature interaction. The short-cut connection of the original input to the output is also adopted to accelerate training. The top layer of feature maps outputs a feature map of size (1/8H, 1/8W, 2) through convolution, where the first channel represents the horizontal motion and the second channel represents the vertical motion.




Figure 3: Structures of our Motion Mamba block.

State Space Model The State Space Model (Gu and Dao 2023) used in our Motion Mamba is calculated by an iterative computation of a sequence:

$$\begin{aligned} h_t &= \hat{A}h_{t-1} + \hat{B}x_t, \\ y_t &= Ch_t + Dx_t, \end{aligned} \quad (1)$$

with x_t denote the input at time t , h_t denote model’s hidden state at time t , y_t denote the output at time t , C, D are model parameter matrix, \hat{A}, \hat{B} are defined by:

$$\hat{A} = \exp^{Adt}, \hat{B} = Bdt, dt = MLP(x_t), \quad (2)$$

where A, B are model parameter matrix, dt is depend on our input x_t , MLP denote a multiple layer perceptron. Eq. (2) reveals that the \hat{A} and \hat{B} matrices depend on the model input x_t , which reflects the selective scanning scheme of Mamba. Scanning the entire sequence also ensures that the model can model the long-range dependencies of the feature map, which is better than only focusing on local information.

Motion Ground-truth Generation. As Fig. 4 shows, we first use the pre-trained EMD-Flow (Deng et al. 2023) network to generate the optical flow map of each frame in the dataset. Then we calculate the center offset of each annotated object in the dataset between its previous and next frames and overlap the offset value within the bounding box range of the object on the original optical flow map. This map is used as ground-truth motion map to supervise the Motion Mamba module. The introduced optical flow map helps to accelerate the convergence of the motion model. The usage of the object motion ground-truth also makes full use of the annotation information of the original dataset.

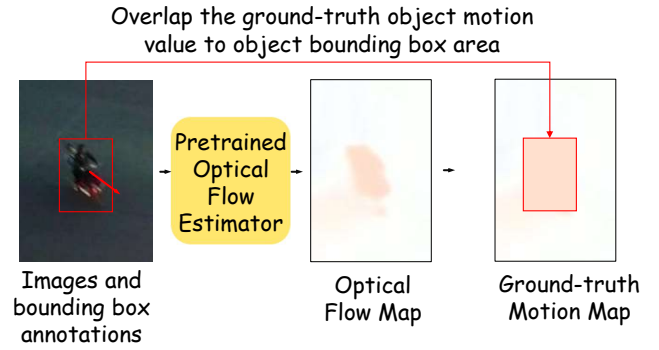


Figure 4: Ground-truth motion map generation procedures.

Motion Prediction and Tracking. For the object detected at time t , we use its center position value on the motion feature map as the motion of the object between time t and $t + 1$. This motion value is used to update the position of the object. The predicted object position at time $t + 1$ is spatially matched with the position of the object detected by the detector at time $t + 1$. This matching gives the final tracking result at time $t + 1$.

Motion Margin Loss

The rotation of the camera view introduces a large motion of the object in the image, resulting in severe motion blur. This motion blur will greatly increase the difficulty of object detection. However, since there are fewer such cases in the dataset, these difficult-to-detect samples have fewer training times than easy samples, which further increases the detection difficulty. For object tracking tasks, even a few frames that cannot be detected will cause tracking interruption, greatly affecting tracking accuracy. To this end, we propose a Motion Margin loss function to assign different decision boundaries according to different object’s motion. We assign larger decision boundaries to objects with larger motion, thereby forcing the model to output higher scores for objects with larger motion during the learning process, so as to effectively detect these objects during inference. Specifically, we calculate the motion margin of each object frame based on the object offset x :

$$D(x) = s * \left(\frac{1}{1 + e^{-(x-5)/s}} \right) - M, \quad (3)$$

where M is defined by:

$$M = s * \left(\frac{1}{1 + e^{-(0-5)/s}} \right) \quad (4)$$

The design of the motion margin function follows these criteria: (1) No margin is set when the object motion is 0. (2) The larger the object’s motion, the larger the motion margin value. (3) The motion margin gradually converges as object motion increases. Fig. 5 shows the motion margin function curves with different s values, which indicate that the function satisfied these three criteria.

Starting from the observation that a motion of about 30 pixels can cause a large motion blur, we select $s=10$ as the

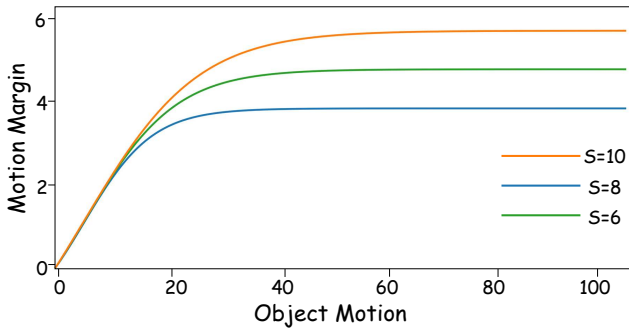


Figure 5: Motion margin function curves of MMLoss.

parameter of MMLoss. Fig. 5 shows that when $s=10$, the motion margin tends to saturation above 30.

For the classification of each object box, the motion margin loss is defined as:

$$\text{MMLoss}(y_i, \hat{y}_i, D_i) = -y_i * \log(\sigma(\hat{y}_i - D_i)) - (1 - y_i) * \log(1 - \sigma(\hat{y}_i)), \quad (5)$$

where \hat{y}_i is the predicted class probability, y_i is the ground-truth label. The function of subtracting D_i from the output of the network classification layer is to assign different decision boundaries to different object boxes based on their motion value.

Experiments

Datasets and Metrics

Datasets. We conduct comparative experiments on the Visdrone (Zhu et al. 2020) and UAVDT (Du et al. 2018) datasets. These two datasets are open-sourced multi-class multi-object tracking datasets, and both are collected from the perspective of UAVs. Therefore, it is suitable to study the tracking objects with global motion and motion blur problems in these two datasets.

The Visdrone dataset consists of a training set (56 sequences), validation set (7 sequences), test-dev set (7 sequences), and test-challenge set (6 sequences). There are 10 categories in the Visdrone dataset: pedestrian, person, car, van, bus, truck, motor, bicycle, awning-tricycle, and tricycle. Each object within the above categories is annotated by a bounding box, category number, and unique identification number. In experiments of Visdrone, we use the full ten categories in training while only using five categories in testing, i.e., car, bus, truck, pedestrian, and van in evaluation, as the evaluation toolkit offered by Visdrone officials only evaluates these five categories.

The UAVDT dataset is a car-tracking dataset based on aerial view. It includes different common scenes, such as squares, arterial streets, and toll stations. UAVDT dataset consists of a training set (30 sequences), and a test set (20 sequences), with three categories: car, truck, and bus. In experiments of UAVDT, all three categories are evaluated using Visdrone’s official evaluation toolkits.

Metrics. We select MOTA and IDF1 as our main evaluation metrics.

			Visdrone		UAVDT		
B	MML	MMM	MOTA↑	IDF1↑	MOTA↑	IDF1↑	Time(ms)↓
✓			39.6	50.4	47.5	63.1	15.2
✓	✓		42.1	54.4	49.2	64.1	15.2
✓		✓	40.9	54.5	49.7	66.1	6.9
✓	✓	✓	44.7	58.3	51.4	68.9	6.9

Table 1: Ablation studies on Visdrone and UAVDT test set. B: Baseline model. MML: Motion Margin Loss. MMM: Motion Mamba module.

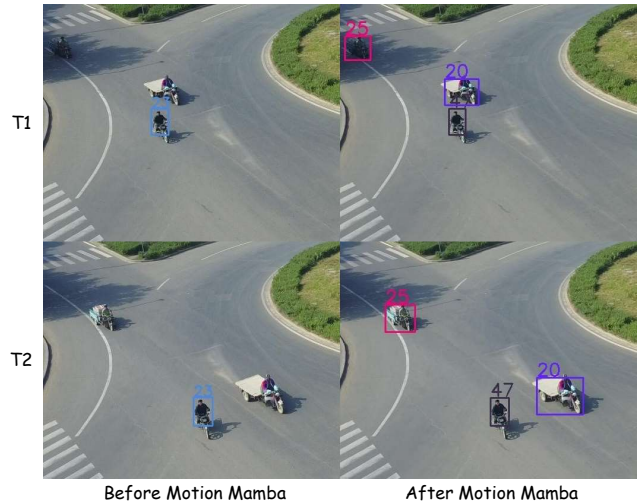


Figure 6: Visualization results for large motion scenes. The model without Motion Mamba missed the fast-moving tricycles on the road, while the model with Motion Mamba tracked them with ID 20, 25, and 47.

Implementation Details

In all experiments, we keep the same train-test split as the official split of the Visdrone and UAVDT datasets. We use the YOLOX-S (Ge et al. 2021) model as the base object detector on both datasets, the input image size is 1088×608 . We use the stochastic gradient descent method (Bottou 2010) to optimize the detector, the learning rate is set to 0.0001, the batch size is set to 8, each dataset is trained for 10 epochs, and the training and testing are completed on a single 2080TI graphic card. The Motion Mamba module and the object regression branch of the detector are trained using the L1 loss function. The object classification branch of the detector is trained using the proposed MMLoss. We select the accuracy evaluation tool officially provided by the Visdrone dataset to complete the comparison of all metrics.

Ablation study

Baseline model. The baseline model we compared with is the model that uses YOLOX-S as the object detector, extended Kalman Filter as motion modeling (Cao et al. 2023), and two-stage spatial matching as association (Zhang et al. 2022).




Figure 7: Visualization results for blurred motion scenes. The model without MMLoss missed the fast-moving bikes on the crossroad while the model with MMLoss tracked it with ID 38.

Motion modeling	MOTA \uparrow	IDF1 \uparrow	Time(ms) \downarrow
OC-SORT	39.6	50.4	15.2
FastFlowNet	40.4	53.1	11
GMA	40.5	53.5	102
EMD-Flow	40.7	54.0	115
Motion Mamba (Ours)	40.9	54.5	6.9

Table 2: Effectiveness of Motion Mamba module.

Motion Mamba Module. We compare Motion Mamba with the OC-SORT tracker, and direct motion prediction using the optical flow network. As shown in Table 2, the tracking accuracy obtained by Motion Mamba is better than that of OC-SORT, and is also better than directly using the optical flow network to predict motion. The results of EMD-Flow are close to those of Motion Mamba, but the inference time of MotionMamba is only 1/16 of that of EMDFlow (6.9 vs 115 in ms), which demonstrates the efficiency of Motion Mamba.

As Fig. 6 shows, after using our Motion Mamba, the tracker successfully tracks the three fast-moving tricycles on the road (with ID1 20, 25, 47), while the model without Motion Mamba only tracked one tricycle (with ID 23).

We also tested the effectiveness of local and global information. As Table 3 shows, both local correlation and global SSM improves the tracking accuracy. Combining local correlation with two SSMs on different directions achieves the best tracking accuracy (40.9 in MOTA and 54.5 in IDF1). In summary, the qualitative and quantitative results above confirm the effectiveness of our proposed Motion Mamba in improving object tracking with large motion.

Motion Margin Loss. We compare the accuracy of the proposed Motion Margin loss (MMLoss) with other imbalance

	MOTA	IDF1
Baseline	39.6	50.4
Local	40.2	52.5
Local & V-SSM	40.5	53.0
Local & H-SSM	40.5	53.5
Local & H-SSM & V-SSM	40.9	54.5

Table 3: Effectiveness of local correlation (Local), vertical-state-space-model (V-SSM), and horizontal-state-space-model (H-SSM).

training loss functions. Table 4 shows that the MMLoss proposed in this paper performs better than the previous category long-tailed loss functions, which proves the effectiveness of our proposed MMLoss.

Loss function	MOTA \uparrow (%)	IDF1 \uparrow (%)
Cross Entropy	39.6	50.4
Focal	36.4	48.3
LDAM	40.3	51.0
MMLoss (Ours)	42.1	54.4

Table 4: Comparisons of different imbalance training loss functions in Visdrone test-dev dataset.

Fig. 7 visualized our ablation results. As shown in Fig. 7, the model using MMLoss successfully tracked the moving bike-rider with blurred appearance (denoted as ID 38), while the baseline model failed to track it. The qualitative and quantitative results above confirm the effectiveness of our proposed loss function.

We further investigate the effectiveness of MMLoss to improve the detection of fast-moving objects. As Fig. 9 shows, introducing MMLoss effectively increases the average object scores of large motion objects (objects with larger relative velocity). When using MMLoss, the average score of large motion objects is higher than 0.5, which is more reliable than the model without MMLoss.

Combination. We also evaluate the effectiveness of our MMLoss and Motion Mamba which are used in combination. As Table 1 shows, both Motion Mamba and MMLoss can effectively improve the tracking performance on the Visdrone dataset.

With the MMLoss and Motion Mamba commonly used, the best tracking performance is obtained (44.7% in MOTA and 58.3% in IDF1 on Visdrone, 51.4% in MOTA and 68.9% in IDF1 on UAVDT), and the motion-inference time is lower than the baseline model (15.1 to 6.9 in ms).

Comparison with state-of-the-art

As shown in Table 5, the MM-Tracker proposed in this paper exceeds the current state-of-the-art (SOTA) methods in MOTA and IDF1. We also visualize the comparison results of the latest best methods (OC-SORT and FOLT) and our MM-Tracker in the Visdrone dataset and the UAVDT dataset. As the left part of Fig. 8 shows, our MM-Tracker




Figure 8: Visual comparisons on the Visdrone and UAVDT datasets. The same number denotes the same object in different frames.




Figure 9: Visualization of object's detection scores with different training loss. With MMLoss introduced, the average detection score of large motion objects is higher than 0.5 (denoted as the red dash line), which is more reliable than the model without MMLoss.

successfully tracks the small and fast-moving bus and truck in both $T1$ and $T2$ frames, while the OC-SORT failed to track these three vehicles, and the FOLT failed to track the half-occluded truck with ID 12. As the right part of Fig. 8 shows, the camera and the object are both in motion, introducing large and irregular motion and causing the object feature to blur. In this difficult situation, our MM-Tracker successfully tracks all five cars in frame $T1$, while the OC-SORT missed the left-bottom car and the FOLT has duplicated tracking of the same car with ID 72,61 at $T1$ time and 72,75 at $T2$ time. The FOLT also missed the right-top car which is severely blurred due to large and irregular motion.

Conclusion

Aiming at efficient and fast global motion modeling, we propose the Motion Mamba module. The Motion Mamba estimates a motion map based on features extracted by the object detector, which greatly reduces the computational cost of motion modeling. The introduced bi-directional selective scan module extracts global motion features effectively, therefore outperforms previous motion modeling methods in UAVMOT scenes. Aiming at addressing the motion imbal-

Method	Visdrone		UAVDT		FPS \uparrow
	MOTA \uparrow	IDF1 \uparrow	MOTA \uparrow	IDF1 \uparrow	
SiamMOT (Shuai et al. 2021)	31.9	48.3	39.4	61.4	11.2
FairMOT (Zhang et al. 2021)	34.3	46.1	41.5	59.2	17.2
ByteTrack (Zhang et al. 2022)	35.7	37	41.6	59.1	27.0
UAVMOT (Liu et al. 2022a)	36.1	51	46.4	67.3	12.0
OC-SORT (Cao et al. 2023)	39.6	50.4	47.5	64.9	26.4
FOLT (Yao et al. 2023)	42.1	56.9	48.5	68.3	29.4
U2MOT (Liu et al. 2023)	42.8	53.9	47.1	65.2	24.1
TrackSSM (Hu et al. 2024)	41.9	55.3	48.1	65.4	24.1
Ours	44.7	58.3	51.4	68.9	31.1

Table 5: Comparisons of the proposed MM-Tracker with state-of-the-art methods on Visdrone and UAVDT test sets. Our proposed MM-Tracker surpasses state-of-the-art in both tracking accuracy and inference speed.

ance training problems in current UAV-MOT datasets, we propose Motion Margin loss (MMLoss), which supervises the object detector according to the object's motion, and improves the detection performance on fast-moving objects. Experiments show that both MMLoss and Motion Mamba can improve the accuracy of multiple object tracking in UAV view. Comparison with state-of-the-art shows that our MM-Tracker achieves the best results on the two public UAV-MOT datasets in both tracking accuracy and inference speed.

Acknowledgements

This work was supported in part by Natural Science Foundation of China under contract 62171139, and in part by Zhongshan science and technology development project under contract 2020AG016, and in part by the Shanghai Science and Technology Project, grant number 24YF2716900.

References

Azad, R.; Asadi-Aghbolaghi, M.; Fathy, M.; and Escalera, S. 2019. Bi-directional ConvLSTM U-Net with densley connected convolutions. In *Proceedings of the IEEE/CVF international conference on computer vision workshops*, 0–0.

- Bewley, A.; Ge, Z.; Ott, L.; Ramos, F.; and Upcroft, B. 2016. Simple online and realtime tracking. In *ICIP*, 3464–3468. IEEE.
- Bochinsk, E.; Eiselein, V.; and Sikora, T. 2017. High-speed tracking-by-detection without using image information. In *AVSS2017*, 1–6. IEEE.
- Bottou, L. 2010. Large-scale machine learning with stochastic gradient descent. In *Proceedings of COMPSTAT'2010: 19th International Conference on Computational Statistics-Paris France, August 22-27, 2010 Keynote, Invited and Contributed Papers*, 177–186. Springer.
- Cao, J.; Pang, J.; Weng, X.; Khirodkar, R.; and 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*, 9686–9696.
- Cao, K.; Wei, C.; Gaidon, A.; Arechiga, N.; and Ma, T. 2019. Learning imbalanced datasets with label-distribution-aware margin loss. *Advances in neural information processing systems*, 32.
- Chandra, S.; Sharma, G.; Malhotra, S.; Jha, D.; and Mittal, A. P. 2015. Eye tracking based human computer interaction: Applications and their uses. In *2015 International Conference on Man and Machine Interfacing (MAMI)*, 1–5. IEEE.
- Choi, M.; Kim, H.; Han, B.; Xu, N.; and Lee, K. M. 2020. Channel attention is all you need for video frame interpolation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, 10663–10671.
- Dendorfer, P.; Rezatofghi, H.; Milan, A.; Shi, J.; Cremers, D.; Reid, I.; Roth, S.; Schindler, K.; and Leal-Taixé, L. 2020. Mot20: A benchmark for multi object tracking in crowded scenes. *arXiv preprint arXiv:2003.09003*.
- Deng, C.; Luo, A.; Huang, H.; Ma, S.; Liu, J.; and Liu, S. 2023. Explicit motion disentangling for efficient optical flow estimation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 9521–9530.
- Du, D.; Qi, Y.; Yu, H.; Yang, Y.; Duan, K.; Li, G.; Zhang, W.; Huang, Q.; and Tian, Q. 2018. The unmanned aerial vehicle benchmark: Object detection and tracking. In *ECCV*, 370–386.
- Elman, J. L. 1990. Finding structure in time. *Cognitive science*, 14(2): 179–211.
- Fagot-Bouquet, L.; Audigier, R.; Dhome, Y.; and Lerasle, F. 2016. Improving multi-frame data association with sparse representations for robust near-online multi-object tracking. In *European Conference on Computer Vision*, 774–790. Springer.
- Forsyth, D. A.; Arikan, O.; Ikemoto, L.; O’Brien, J.; Ramanan, D.; et al. 2006. Computational studies of human motion: Part 1, tracking and motion synthesis. *Foundations and Trends® in Computer Graphics and Vision*, 1(2–3): 77–254.
- Ge, Z.; Liu, S.; Wang, F.; Li, Z.; and Sun, J. 2021. Yolox: Exceeding yolo series in 2021. *arXiv preprint arXiv:2107.08430*.
- Geiger, A.; Lenz, P.; Stiller, C.; and Urtasun, R. 2013. Vision meets robotics: The kitti dataset. *The International Journal of Robotics Research*, 32(11): 1231–1237.
- Gu, A.; and Dao, T. 2023. Mamba: Linear-time sequence modeling with selective state spaces. *arXiv preprint arXiv:2312.00752*.
- Hochreiter, S.; and Schmidhuber, J. 1997. Long short-term memory. *Neural computation*, 9(8): 1735–1780.
- Hu, B.; Luo, R.; Liu, Z.; Wang, C.; and Liu, W. 2024. TrackSSM: A General Motion Predictor by State-Space Model. *arXiv preprint arXiv:2409.00487*.
- Huang, C.; Li, Y.; Loy, C. C.; and Tang, X. 2016. Learning deep representation for imbalanced classification. In *CVPR*, 5375–5384.
- Huang, H.; Zeng, Z.; Yao, D.; Pei, X.; and Zhang, Y. 2021. Spatial-temporal ConvLSTM for vehicle driving intention prediction. *Tsinghua Science and Technology*, 27(3): 599–609.
- Kalake, L.; Wan, W.; and Hou, L. 2021. Analysis Based on Recent Deep Learning Approaches Applied in Real-Time Multi-Object Tracking: A Review. *IEEE Access*, PP(99): 1–1.
- Kalman, R. E. 1960. A new approach to linear filtering and prediction problems. *Journal of Basic Engineering*.
- Kurimo, E.; Lepistö, L.; Nikkanen, J.; Grén, J.; Kunttu, I.; and Laaksonen, J. 2009. The effect of motion blur and signal noise on image quality in low light imaging. In *Image Analysis: 16th Scandinavian Conference, SCIA 2009, Oslo, Norway, June 15-18, 2009. Proceedings 16*, 81–90. Springer.
- Li, B.; Yao, Y.; Tan, J.; Zhang, G.; Yu, F.; Lu, J.; and Luo, Y. 2022. Equalized focal loss for dense long-tailed object detection. In *CVPR*, 6990–6999.
- Li, X.; Hu, W.; Shen, C.; Zhang, Z.; Dick, A.; and Hengel, A. V. D. 2013. A survey of appearance models in visual object tracking. *ACM transactions on Intelligent Systems and Technology (TIST)*, 4(4): 1–48.
- Li, X.; Wang, W.; Hu, X.; Li, J.; Tang, J.; and Yang, J. 2021. Generalized focal loss v2: Learning reliable localization quality estimation for dense object detection. In *CVPR*, 11632–11641.
- Li, Y.; Wang, T.; Kang, B.; Tang, S.; Wang, C.; Li, J.; and Feng, J. 2020. Overcoming classifier imbalance for long-tail object detection with balanced group softmax. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 10991–11000.
- Lin, T.-Y.; Goyal, P.; Girshick, R.; He, K.; and Dollár, P. 2017. Focal loss for dense object detection. In *Proceedings of the IEEE international conference on computer vision*, 2980–2988.
- Liu, K.; Jin, S.; Fu, Z.; Chen, Z.; Jiang, R.; and Ye, J. 2023. Uncertainty-aware unsupervised multi-object tracking. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 9996–10005.
- Liu, S.; Li, X.; Lu, H.; and He, Y. 2022a. Multi-Object Tracking Meets Moving UAV. In *Proceedings of the*

- IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 8876–8885.
- Liu, Z.; Lin, Y.; Cao, Y.; Hu, H.; Wei, Y.; Zhang, Z.; Lin, S.; and Guo, B. 2021. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF international conference on computer vision*, 10012–10022.
- Liu, Z.; Ning, J.; Cao, Y.; Wei, Y.; Zhang, Z.; Lin, S.; and Hu, H. 2022b. Video swin transformer. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 3202–3211.
- Luo, W.; Xing, J.; Milan, A.; Zhang, X.; Liu, W.; and Kim, T.-K. 2021. Multiple object tracking: A literature review. *Artificial Intelligence*, 293: 103448.
- Meinhardt, T.; Kirillov, A.; Leal-Taixe, L.; and Feichtenhofer, C. 2022. Trackformer: Multi-object tracking with transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 8844–8854.
- Menon, A. K.; Jayasumana, S.; Rawat, A. S.; Jain, H.; Veit, A.; and Kumar, S. 2020. Long-tail learning via logit adjustment. *arXiv preprint arXiv:2007.07314*.
- Peng, J.; Gu, Y.; Wang, Y.; Wang, C.; Li, J.; and Huang, F. 2020a. Dense scene multiple object tracking with box-plane matching. In *ACM International Conference on Multimedia*, 4615–4619.
- Peng, J.; Qiu, F.; See, J.; Guo, Q.; Huang, S.; Duan, L.-Y.; and Lin, W. 2018. Tracklet siamese network with constrained clustering for multiple object tracking. In *VCIP*, 1–4. IEEE.
- Peng, J.; Wang, C.; Wan, F.; Wu, Y.; Wang, Y.; Tai, Y.; Wang, C.; Li, J.; Huang, F.; and Fu, Y. 2020b. Chained-tracker: Chaining paired attentive regression results for end-to-end joint multiple-object detection and tracking. In *ECCV*, 145–161. Springer.
- Peng, J.; Wang, T.; Lin, W.; Wang, J.; See, J.; Wen, S.; and Ding, E. 2020c. TPM: Multiple object tracking with tracklet-plane matching. *Pattern Recognition*, 107: 107480.
- Ren, J.; Yu, C.; Ma, X.; Zhao, H.; Yi, S.; et al. 2020. Balanced meta-softmax for long-tailed visual recognition. *Advances in neural information processing systems*, 33: 4175–4186.
- Shuai, B.; Berneshawi, A.; Li, X.; Modolo, D.; and Tighe, J. 2021. SiamMOT: Siamese Multi-Object Tracking. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 12372–12382.
- Subakan, C.; Ravanelli, M.; Cornell, S.; Bronzi, M.; and Zhong, J. 2021. Attention is all you need in speech separation. In *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 21–25. IEEE.
- Sugimura, D.; Kitani, K. M.; Okabe, T.; Sato, Y.; and Sugimoto, A. 2009. Using individuality to track individuals: Clustering individual trajectories in crowds using local appearance and frequency trait. In *ICCV*, 1467–1474. IEEE.
- Takala, V.; and Pietikainen, M. 2007. Multi-object tracking using color, texture and motion. In *2007 IEEE Conference on Computer Vision and Pattern Recognition*, 1–7. IEEE.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.
- Wang, D.; Yang, Y.; and Ning, S. 2018. DeepSTCL: A deep spatio-temporal ConvLSTM for travel demand prediction. In *2018 international joint conference on neural networks (IJCNN)*, 1–8. IEEE.
- Wang, F.; Cheng, J.; Liu, W.; and Liu, H. 2018. Additive margin softmax for face verification. *IEEE Signal Processing Letters*, 25(7): 926–930.
- Wang, T.; Li, Y.; Kang, B.; Li, J.; Liew, J.; Tang, S.; Hoi, S.; and Feng, J. 2020a. The devil is in classification: A simple framework for long-tail instance segmentation. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XIV 16*, 728–744. Springer.
- Wang, Z.; Zheng, L.; Liu, Y.; Li, Y.; and Wang, S. 2020b. Towards real-time multi-object tracking. In *European Conference on Computer Vision*, 107–122. Springer.
- Wojke, N.; Bewley, A.; and Paulus, D. 2017. Simple online and realtime tracking with a deep association metric. In *ICIP*, 3645–3649. IEEE.
- Yang, B.; and Nevatia, R. 2012. Multi-target tracking by online learning of non-linear motion patterns and robust appearance models. In *CVPR*, 1918–1925. IEEE.
- Yao, M.; Wang, J.; Peng, J.; Chi, M.; and Liu, C. 2023. FOLT: Fast Multiple Object Tracking from UAV-captured Videos Based on Optical Flow. In *Proceedings of the 31st ACM International Conference on Multimedia*, 3375–3383.
- Zhang, Y.; Sun, P.; Jiang, Y.; Yu, D.; Weng, F.; Yuan, Z.; Luo, P.; Liu, W.; and Wang, X. 2022. Bytetrack: Multi-object tracking by associating every detection box. In *ECCV*, 1–21. Springer.
- Zhang, Y.; Wang, C.; Wang, X.; Zeng, W.; and Liu, W. 2021. Fairmot: On the fairness of detection and re-identification in multiple object tracking. *International Journal of Computer Vision*, 129(11): 3069–3087.
- Zhao, X.; Gong, D.; and Medioni, G. 2012. Tracking using motion patterns for very crowded scenes. In *ECCV*, 315–328. Springer.
- Zheng, L.; Bie, Z.; Sun, Y.; Wang, J.; Su, C.; Wang, S.; and Tian, Q. 2016. Mars: A video benchmark for large-scale person re-identification. In *European conference on computer vision*, 868–884. Springer.
- Zhou, X.; Koltun, V.; and Krähenbühl, P. 2020. Tracking objects as points. In *European Conference on Computer Vision*, 474–490. Springer.
- Zhu, P.; Wen, L.; Du, D.; Bian, X.; Fan, H.; Hu, Q.; and Ling, H. 2020. Detection and tracking meet drones challenge. *arXiv preprint arXiv:2001.06303*.