

UCMCTrack: Multi-Object Tracking with Uniform Camera Motion Compensation

Kefu Yi^{1*}, Kai Luo², Xiaolei Luo², Jianguai Huang², Hao Wu², Rongdong Hu³, Wei Hao^{1*}

¹School of Traffic and Transportation, Changsha University of Science and Technology

²College of Automotive and Mechanical Engineering, Changsha University of Science and Technology

³Changsha Intelligent Driving Institute

{corfyi, haowei}@csust.edu.cn

Abstract

Multi-object tracking (MOT) in video sequences remains a challenging task, especially in scenarios with significant camera movements. This is because targets can drift considerably on the image plane, leading to erroneous tracking outcomes. Addressing such challenges typically requires supplementary appearance cues or Camera Motion Compensation (CMC). While these strategies are effective, they also introduce a considerable computational burden, posing challenges for real-time MOT. In response to this, we introduce UCMCTrack, a novel motion model-based tracker robust to camera movements. Unlike conventional CMC that computes compensation parameters frame-by-frame, UCMCTrack consistently applies the same compensation parameters throughout a video sequence. It employs a Kalman filter on the ground plane and introduces the Mapped Mahalanobis Distance (MMD) as an alternative to the traditional Intersection over Union (IoU) distance measure. By leveraging projected probability distributions on the ground plane, our approach efficiently captures motion patterns and adeptly manages uncertainties introduced by homography projections. Remarkably, UCMCTrack, relying solely on motion cues, achieves state-of-the-art performance across a variety of challenging datasets, including MOT17, MOT20, DanceTrack and KITTI. More details and code are available at <https://github.com/corfyi/UCMCTrack>.

Introduction

At the core of tracking-by-detection paradigm of multi-object tracking (MOT) is the accurate association of detections with tracked objects. Motion cues are widely used due to their effectiveness and simplicity. However, the application of motion model in scenarios with frequent camera movement is highly challenging. This issue is usually addressed by applying additional appearance cues or performing frame-by-frame Camera Motion Compensation (CMC) on video captured by the moving camera. While effective, these additional measures introduce a non-negligible computational burden, posing a obstacle for real-time MOT.

Thus, a pertinent question arises: Is it possible to employ motion cues in MOT that are robust to camera movement

*Corresponding author.

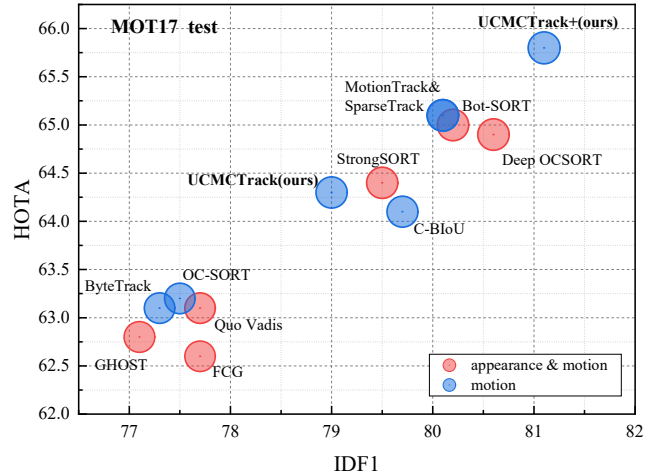


Figure 1: IDF1-HOTA-AssA comparisons of different trackers on the test set of MOT17. The horizontal axis is IDF1, the vertical axis is HOTA, and the radius of circle is AssA. Our UCMCTrack+ achieves 65.8 HOTA, 81.1 IDF1 on MOT17 test, possessing significant competitiveness compared to SOTA trackers. Details are given in Table 1.

without resorting to the cumbersome frame-by-frame CMC? **Our answer is YES.** We have developed a pure motion model-based multi-object tracker that is robust to camera movement. For the same video sequence, it suffices to use the same camera motion compensation parameters, rather than computing the camera motion compensation parameters for every frame as traditional CMC does. We choose to model the target’s motion using a simple Kalman filter on the ground plane, instead of on the imaging plane as most MOT algorithms do, and effectively compensate the motion estimation errors caused by camera movement through the process noise parameters of the Kalman filter. We abandon the commonly used Intersection over Union (IoU), and instead propose the Mapped Mahalanobis Distance (MMD). It computes the projected probability distribution on the ground plane, and utilizes the Mahalanobis distance to calculate the matching costs between targets. It not only effectively leverages the underlying motion patterns of the targets on the ground plane but also efficiently handles the uncertainties

caused by homography projection.

A deep dive into motion-based MOT highlights a significant challenge when employing motion cues in highly dynamic scenes. Historically, IoU has been the favored metric for data association. On the surface, employing IoU on the image plane appears to be a more direct approach. However, its application often leads to inaccurate tracking outcomes, particularly in complex scenes marked by frequent camera movements. Notably, in these settings, detection and tracking boxes might completely fail to overlap, as shown in Figure 3. This observation underscores an imminent necessity: a transition from exclusive reliance on the image plane to harnessing the more robust motion patterns inherent to the ground plane. Embracing such a paradigm shift stands to effectively address challenges spawned by camera movements, setting the stage for superior tracking accuracy. Distinct from the vast majority of trackers relying IoU on the image plane, ground plane-based association can effectively consider camera movement as noise within the motion model. It minimizes the problems induced by camera movements. This methodology is notably more direct, convenient, and efficient than compensating for camera motion frame-by-frame via traditional CMC.

In light of these challenges, we introduce the Uniform Camera Motion Compensation (UCMC) tracker. It is a pure motion-based multi-object tracker that offers a holistic solution robust to camera jitter and motion, without any dependency on IoU-based methodologies.

The main contributions of this paper are threefold:

- In the realm of multi-object tracking where IoU is conventionally employed to capitalize on motion cues, our work introduces an innovative non-IoU distance measure, singularly driven by motion cues, and manifests state-of-the-art performance across multiple established datasets, marking a significant departure from traditional tracking techniques.
- In addressing the challenge of camera movements, we propose a method that diverges from conventional camera motion compensation techniques. Instead of computing camera compensation parameters frame-by-frame for video sequences, our approach uniformly applies the same compensation parameters across the entire sequence, substantially reducing the computational burden typically associated with camera motion adjustments.
- We introduce UCMCTrack, a simplistic yet efficacious multi-object tracker that employs a novel, standalone motion-based distance measure. This new measure has the potential to complement commonly-used distance metrics such as IoU and ReID. Remarkably, when provided with detections, UCMCTrack operates at a very fast speed, exceeding 1000 FPS using just a single CPU.

Related Work

Distance Measures

Distance measures play key roles in MOT to associate targets in the current frame with those in previous frames. Currently, most algorithms employ the pixel-based Intersection

over Union (IoU) technique (Du et al. 2023; Liu et al. 2020), which calculates the intersecting area between the detection box and the tracking box for target matching. However, in cases of camera jitter or low sampling rates, the two boxes may not intersect, rendering IoU ineffective. In contrast, Generalized Intersection over Union (GIoU) (Rezatofighi et al. 2019) not only focuses on the overlapping region but also considers the non-overlapping area, thus improving the computation of image overlap. Distance-IoU (DIoU) builds upon GIoU by further incorporating geometric distance into the calculation (Zheng et al. 2020). However, both these methods do not adequately represent the similarity in aspect ratios of the objects. Bidirectional Intersection over Union (BIOU) and the Cascade-BIOU (C-BIOU) proposed to augment the IoU-based approach by introducing a linear average motion estimation model and expanding the search region (Yang et al. 2023). Nevertheless, all these methods operate in the image plane and cannot fully capture the actual motion patterns, leading to faulty tracking during camera motion. Recently, some methods have considered distance measures based on the ground plane. SparseTrack (Liu et al. 2023) goes beyond IoU and incorporates additional estimated pseudo-depth for supplementary metrics. Quo Vadis (Dendorfer et al. 2022) employs homography transformation to calculate the Euclidean distance in the bird’s-eye view and combines it with IoU for target matching. Although these approaches utilize additional depth information, they still rely on IoU and fail to account for the uncertainty in the projection of targets onto the ground plane.

Motion Models

Tracking-by-detection MOT algorithms (Wojke, Bewley, and Paulus 2017; Cao et al. 2023; Maggolino et al. 2023) often favor motion models for their simplicity and effectiveness. Among these, the Constant Velocity (CV) model, which assumes unvarying target motion between frames, is the most favored approach (Bewley et al. 2016; Zhang et al. 2022). Numerous studies have been dedicated to improving motion estimation accuracy, employing methods such as Kalman filtering (Bewley et al. 2016; Zhang et al. 2022; Zhou, Koltun, and Krähenbühl 2020), optical flow (Xiao, Wu, and Wei 2018), and displacement regression (Feichtenhofer, Pinz, and Zisserman 2017; Held, Thrun, and Savarese 2016). However, current MOT algorithms (Du et al. 2021, 2023; Aharon, Orfaig, and Bobrovsky 2022) model the motion of tracking targets directly upon the image plane using detected bounding boxes. This approach fails to reflect the actual motion patterns of the targets on the ground plane, leading to erroneous tracking results during camera motion.

To further leverage the inherent motion patterns of the tracking targets, researchers (Liu, Wang, and Xu 2020; Marinello, Proesmans, and Van Gool 2022) have employed LSTM networks to predict target motion, while others (Babae, Li, and Rigoll 2019) have used RNN networks for similar purposes. Additionally, transformer networks (Yang et al. 2022) have also been utilized to capture object motion patterns. In contrast to employing neural networks to explicitly predict target motion, the tracking-by-query propagation (Zhang, Wang, and Zhang 2023) forces each query




Figure 2: The pipeline of the proposed UCMCTrack.

to recall the same instance across different frames. Alternatively, the approach based on a Graphs framework (Cetintas, Brasó, and Leal-Taixé 2023) is used to model data association. These methods use a learned network to implicitly grasp the dynamics of target motion. While they achieve promising results, their training process can be challenging, requiring a substantial amount of annotated data and computational resources. Moreover, complex network designs may not meet the real-time requirements on end devices.

Camera Motion Compensation

Camera Motion Compensation (CMC) is a prevalent method to address dynamic scenes in the field of MOT (Bergmann, Meinhardt, and Leal-Taixé 2019; Han et al. 2022; Khurana, Dave, and Ramanan 2021). This is often achieved by aligning frames through image registration, leveraging techniques such as Enhanced Correlation Coefficient (ECC) maximization (Evangelidis and Psarakis 2008a), or employing feature matching methods like ORB (Rublee et al. 2011). In BOT-SORT (Aharon, Orfaig, and Bobrovsky 2022), image key-points were extracted frame-by-frame, with sparse optical flow subsequently applied. The affine transformation matrix of background motion is calculated and obtained via RANSAC (Fischler and Bolles 1981), and the affine matrix is used to transform the prediction box from the $(k-1)$ -th frame coordinate system to the k -th frame coordinate system. In (Yu, Kurniaggoro, and Jo 2019), the pyramidal Lucas-Kanade optical flow is implemented to trace grid-based feature points. The affine transformation matrix between two consecutive frames is calculated through matching feature points, and the initial two frames and the background model are aligned with the current frame. In (Yeh et al. 2017), a camera motion compensation framework is proposed with utilization of temporal and spatial structure, which depends on pre-provided background model for background elimination, thereby posing challenges for its adaptation to new scenarios. However, when confronted with high-resolution videos, current CMC techniques impose substantial computational overhead and hinders the implementation of real-time target tracking.

Method

UCMCTrack follows the tracking-by-detection paradigm, with its pipeline detailed in Figure 2. We introduce significant advancements across crucial dimensions: motion model, distance measure, and process noise compensation.

Together, these improvements bolster UCMCTrack, endowing it with adaptability and efficiency across diverse tracking challenges. For the pseudocode please refer to Appendix A.

Motion Modeling on Ground Plane

We model objects' motion on the ground plane to better capture the fundamental essence of their motion patterns. Selecting the appropriate state vector \mathbf{x} , observation vector \mathbf{z} , and determining the process noise \mathbf{Q}_k and measurement noise \mathbf{R}_k are crucial in establishing the Kalman constant velocity motion model. In order to make observation and calculation more convenient, the choice was made to use the midpoint coordinates of the bottom edge of the bounding box in the image plane, projected onto the ground plane coordinates x and y , as the observation vector. The state vector \mathbf{x} is defined as $\mathbf{x} = [x, \dot{x}, y, \dot{y}]$. According to the linear camera model (Yu and McMillan 2004), the mapping relationship between the ground plane coordinates x and y , and the image plane coordinates u and v , can be expressed as:

$$\begin{bmatrix} u \\ v \\ 1 \end{bmatrix} = \mathbf{A} \frac{1}{\gamma} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix} \quad (1)$$

Where γ is the scale factor, and \mathbf{A} represents the projection matrix which is the product of camera intrinsic and extrinsic parameters. Please refer to Appendix B for more details.

Correlated Measurement Distribution

In general, the measurement errors of detectors on the image plane follow an independent normal distribution, and their covariance matrix \mathbf{R}_k^{uv} can be represented as:

$$\mathbf{R}_k^{uv} = \begin{bmatrix} (\sigma_m w_k)^2 & 0 \\ 0 & (\sigma_m h_k)^2 \end{bmatrix} \quad (2)$$

Where: σ_m represents the detection noise factor as a hyper-parameter, and w_k and h_k denote the detected width and height from the detector.

If we express the inverse matrix of \mathbf{A} in Eq. 1 as:

$$\mathbf{A}^{-1} = \begin{bmatrix} a_{11} & a_{12} & a_{13} \\ a_{21} & a_{22} & a_{23} \\ a_{31} & a_{32} & a_{33} \end{bmatrix} \quad (3)$$

This leads to the covariance matrix of measurement errors in the ground plane as: (Please refer to Appendix B for a detailed derivation)

$$\mathbf{R}_k = \mathbf{C} \mathbf{R}_k^{uv} \mathbf{C}^T \quad (4)$$

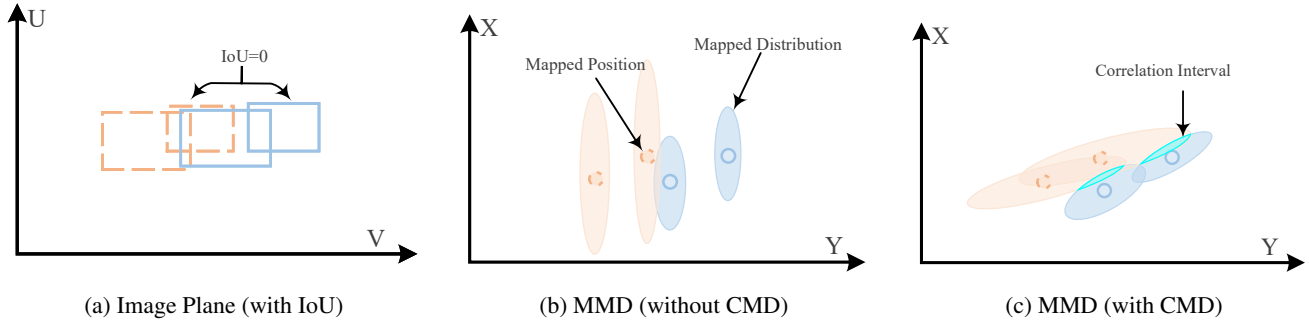


Figure 3: Visualization of distance measures. (a) Visualization of IoU on the image plane. IoU fails as there is no intersection between bounding boxes. (b) Visualization of Mapped Mahalanobis Distance (MMD) without Correlated Measurement Distribution (CMD). Incorrect associations occur due to insufficient utilization of distribution information. (c) Visualization of MMD with CMD. Correct associations after using the correlated probability distribution, undergoing a rotation on the ground plane.

where:

$$\mathbf{C} = \begin{bmatrix} \gamma a_{11} - a_{31}\gamma x & \gamma a_{12} - a_{32}\gamma x \\ \gamma a_{21} - a_{31}\gamma y & \gamma a_{22} - a_{32}\gamma y \end{bmatrix} \quad (5)$$

Thus, we obtain the mapped measurement noise matrix \mathbf{R}_k in the ground plane. It's important to highlight that the mapped distribution exhibits a strong correlation since \mathbf{R}_k is non-diagonal. This allows for more accurate association of targets on the ground plane.

Mapped Mahalanobis Distance

In image plane motion modeling, IoU is the most commonly used distance measure for data association. However, when objects are in high-speed motion or captured at low FPS or in scenes with moving camera, the lack of overlap between detection boxes and tracklets renders IoU ineffective. Conversely, by employing normalized Mahalanobis distance in ground plane modeling, the issue of IoU inefficiency is effectively addressed, as depicted in Figure 3.

The calculation of mapped Mahalanobis distance between track state and measurement involves three steps:

1. Calculate Residual:

$$\epsilon = \mathbf{z} - \mathbf{H}\hat{\mathbf{x}} \quad (6)$$

Here, \mathbf{z} is the mapped measurement on ground plane, $\hat{\mathbf{x}}$ is the predicted track state, and \mathbf{H} is the observation matrix.

2. Compute Residual Covariance Matrix:

$$\mathbf{S} = \mathbf{H}\mathbf{P}\mathbf{H}^T + \mathbf{R}_k \quad (7)$$

Here, \mathbf{P} is the predicted covariance matrix, and \mathbf{R}_k is the mapped measurement noise covariance matrix.

3. Calculate Normalized Mahalanobis Distance:

$$D = \epsilon^T \mathbf{S}^{-1} \epsilon + \ln |\mathbf{S}| \quad (8)$$

Here, $|\mathbf{S}|$ represents the determinant of the matrix \mathbf{S} , and \ln is the natural logarithm.

As seen in Eq. 8, we employed the normalized Mahalanobis distance, incorporating the logarithm of the determinant of the measurement covariance matrix. This ensures

that data association decisions are not solely based on the discrepancies between measurements and predictions, but also holistically consider the accuracy and uncertainty of measurements. Consequently, this will yield more robust and reliable association decisions in object tracking.

Process Noise Compensation

In the context of MOT tasks, many previous works (Bewley et al. 2016; Wojke, Bewley, and Paulus 2017; Zhang et al. 2022; Cao et al. 2023) have treated the target's motion model as a Constant Velocity (CV) model without considering the noise impact caused by camera motion. However, camera motion is quite common in MOT tasks and can introduce significant noise that affects the tracking performance. Assuming that the camera's motion-induced acceleration is the source of noise, we can represent it through the system motion model as follows:

$$\begin{cases} \Delta x = \frac{1}{2} \cdot \sigma \cdot (\Delta t)^2 \\ \Delta v = \sigma \cdot \Delta t \end{cases} \quad (9)$$

where Δx and Δv represent the changes in position and velocity under the influence of noise, respectively. σ denotes the acceleration change due to camera motion, and Δt represents the time interval between two image frames. Expressing Eq. 9 in matrix form yields the matrix:

$$\mathbf{G} = \begin{bmatrix} \frac{\Delta t^2}{2} & 0 \\ \Delta t & 0 \\ 0 & \frac{\Delta t^2}{2} \\ 0 & \Delta t \end{bmatrix} \quad (10)$$

It captures the relationship between the changes in position and velocity caused by noise in each direction. For a two-dimensional CV Model with a Kalman filter, the covariance matrix of the process noise can be represented as follows:

$$\mathbf{Q}_k = \mathbf{G} \cdot \text{diag}(\sigma_x, \sigma_y) \cdot \mathbf{G}^T \quad (11)$$

where σ_x and σ_y denote the process compensation factors along the x and y axes, handling the motion noise cause by camera movements of tilt and rotation respectively.

Tracker	MOT17				MOT20			
	HOTA↑	IDF1↑	MOTA↑	AssA↑	HOTA↑	IDF1↑	MOTA↑	AssA↑
appearance & motion:								
FCG (Girbau, Marqués, and Satoh 2022)	62.6	77.7	76.7	63.4	57.3	69.7	68.0	58.1
Quo Vadis (Dendorfer et al. 2022)	63.1	77.7	80.3	62.1	61.5	75.7	77.8	59.9
GHOST (Seidenschwarz et al. 2023)	62.8	77.1	78.7	-	61.2	75.2	73.7	-
Bot-SORT (Aharon, Orfaig, and Bobrovsky 2022)	65.0	80.2	80.5	65.5	63.3	77.5	77.8	62.9
StrongSORT (Du et al. 2023)	64.4	79.5	79.6	64.4	62.6	77.0	73.8	64.0
Deep OCSORT (Maggiolino et al. 2023)	64.9	80.6	79.4	65.9	63.9	79.2	75.6	65.7
motion only:	HOTA↑	IDF1↑	MOTA↑	AssA↑	HOTA↑	IDF1↑	MOTA↑	AssA↑
ByteTrack (Zhang et al. 2022)	63.1	77.3	80.3	62.0	61.3	75.2	77.8	59.6
C-BIoU (Yang et al. 2023)	64.1	79.7	81.1	63.7	-	-	-	-
MotionTrack (Xiao et al. 2023)	65.1	80.1	81.1	65.1	62.8	76.5	78.0	61.8
SparseTrack (Liu et al. 2023)	65.1	80.1	81.0	65.1	<u>63.4</u>	<u>77.3</u>	78.2	62.8
OCSORT (Cao et al. 2023)	63.2	77.5	78.0	63.4	62.4	76.3	75.7	62.5
UCMCTrack (Ours)	64.3	79.0	79.0	64.6	62.8	<u>77.4</u>	75.5	<u>63.5</u>
UCMCTrack+ (Ours)	65.8	81.1	80.5	66.6	62.8	77.4	75.7	63.4

Table 1: Results on MOT17 & MOT20 test. The detection results were obtained from ByteTrack (Zhang et al. 2022).

Tracker	HOTA↑	IDF1↑	MOTA↑	AssA↑
appearance & motion:				
FCG	48.7	46.5	89.9	29.9
GHOST	56.7	57.7	91.3	39.8
StrongSORT	55.6	55.2	91.1	38.6
Deep OCSORT	61.3	61.5	92.3	45.8
motion only:				
ByteTrack	47.3	52.5	89.5	31.4
C-BIoU	60.6	61.6	91.6	45.4
MotionTrack	58.2	58.6	91.3	41.7
SparseTrack	55.5	58.3	91.3	39.1
OCSORT	55.1	54.9	92.2	40.4
UCMCTrack (Ours)	63.4	65.0	88.8	51.1
UCMCTrack+ (Ours)	63.6	65.0	88.9	51.3

Table 2: Results on DanceTrack-test. The detection results were obtained from ByteTrack (Zhang et al. 2022).

Experiments

Setting

Datasets We conducted a fair evaluation of UCMCTrack on multiple publicly available datasets, including MOT17 (Milan et al. 2016), MOT20 (Dendorfer et al. 2020), DanceTrack (Sun et al. 2022), and KITTI (Geiger et al. 2013). Both MOT17 and MOT20 are pedestrian tracking datasets, and their motion is mostly linear. It is worth noting that MOT20 has a significantly higher density of pedestrians, making it a challenging dataset for tracking. The primary task of the DanceTrack (Sun et al. 2022) is to track dancers, who not only have similar appearances but also perform a large number of irregular movements. The KITTI (Geiger et al. 2013) is an autonomous driving dataset, and we only utilized the left color camera images for the visual vehicle and pedestrian tracking task. Compared to other datasets, KITTI has a lower frame rate, only 10 FPS, and the camera’s motion is more intense.

Metrics We employ the CLEAR metrics (Bernardin and Stiefelhagen 2008) which include MOTA, FP, FN, and others, along with IDF1 (Ristani et al. 2016) and TA (Luiten et al. 2021), to evaluate the tracking performance comprehensively in various aspects. MOTA emphasizes the detector’s performance, while IDF1 measures the tracker’s ability to maintain consistent IDs. We also emphasize the use of AssA to evaluate the association performance. On the other hand, HOTA achieves a balance between detection accuracy, association accuracy, and localization accuracy, making it an increasingly important metric for evaluating trackers.

Implementation Details For fair comparison, we directly used the existing baseline object detection method YOLOX (Ge et al. 2021). The weight files for MOT17, MOT20 and DanceTrack were obtained from ByteTrack (Zhang et al. 2022). For KITTI, we used the detection results from PermaTrack (Tokmakov et al. 2021). We applied the Enhanced Correlation Coefficient maximization (ECC) (Evangelidis and Psarakis 2008b) model for camera motion compensation, which is consistent with strongSORT (Du et al. 2023). For MOT17, MOT20, and DanceTrack datasets, we manually estimated the camera parameters since they are not publicly accessible. In contrast, KITTI dataset readily furnishes the requisite camera parameters.

Benchmark Evaluation

Here, we present the benchmark results for multiple datasets. ↑/↓ indicate that higher/lower is better, respectively. The highest scores for each group are highlighted in **bold**, and the highest score for the motion group is marked in underline. "UCMCTrack+" denotes the enhancement of UCMCTrack with the additional incorporation of CMC.

MOT17 and MOT20 Our UCMCTrack results on MOT17 and MOT20 datasets are presented in Table 1, respectively. We used a private detector to generate the detection results and ensured fairness by aligning the detections with OC-SORT (Cao et al. 2023) and ByteTrack

Tracker	Car			Pedestrian		
appearance & motion:	HOTA↑	MOTA↑	AssA↑	HOTA↑	MOTA↑	AssA↑
QD-3DT (Hu et al. 2022)	72.8	85.9	72.2	41.1	51.8	38.8
TuSimple (Choi 2015; He et al. 2016)	71.6	86.3	71.1	45.9	57.6	47.6
StrongSORT (Du et al. 2023)	77.8	90.4	78.2	54.5	67.4	57.3
motion only:	HOTA↑	MOTA↑	AssA↑	HOTA↑	MOTA↑	AssA↑
CenterTrack (Zhou, Koltun, and Krähenbühl 2020)	73.0	88.8	71.2	40.4	53.8	36.9
TrackMPNN (Rangesh et al. 2021)	72.3	87.3	70.6	39.4	52.1	35.5
OCSORT (Cao et al. 2023)	76.5	90.3	76.4	54.7	65.1	59.1
UCMCTrack (Ours)	<u>77.1</u>	90.4	<u>77.2</u>	55.2	67.4	58.0
UCMCTrack+ (Ours)	74.2	90.2	71.7	54.3	67.2	56.3

Table 3: Results on KITTI-test. The detection results were obtained from PermaTrack (Tokmakov et al. 2021).

(Zhang et al. 2022). UCMCTrack+ has attained state-of-the-art (SOTA) performance, notably on the MOT17 dataset, surpassing the SOTA methods by 0.9 in HOTA, 0.5 in IDF1, and 0.7 in AssA. It even surpasses leading algorithms that leverage both motion and appearance features at a considerable margin, highlighting its effective use of motion information to enhance the robustness and efficiency of the tracking.

DanceTrack To demonstrate UCMCTrack’s performance under irregular motion scenarios, we present the test set results on DanceTrack in Table 2. UCMCTrack+ outperforms the SOTA methods with an improvement of 2.3 in HOTA, 3.4 in IDF1, and 5.5 in AssA. This highlights the effectiveness of our tracker in handling targets with irregular motions and further validates its SOTA performance.

KITTI In Table 3, we present the results of UCMCTrack on the KITTI dataset. It’s noteworthy that the addition of CMC on the KITTI dataset did not yield favorable results. We believe that this might be due to the inaccuracies present in the CMC parameters. This observation indicates that UCMC demonstrates stronger generalization capabilities than CMC, particularly in complex scenarios. The performance of UCMCTrack in the KITTI dataset demonstrates its effectiveness in addressing challenges posed by high-speed motion and low frame-rate detections.

Ablation Studies on UCMC

The ablation experiments of UCMCTrack were conducted on the validation sets of MOT17 and DanceTrack. For MOT17, the validation set was split following the prevailing conventions (Zhou, Koltun, and Krähenbühl 2020). The baseline is chosen as ByteTrack (Zhang et al. 2022).

Component Ablation We undertook a comprehensive validation of UCMCTrack’s key components: Mapped Mahalanobis Distance (MMD), Correlated Measurement Distribution (CMD), and Process Noise Compensation (PNC), testing on both MOT17 and DanceTrack datasets. As demonstrated in Table 4, each component of UCMCTrack plays a crucial role in enhancing its overall efficacy. A noteworthy observation is the decline in performance when transitioning from IoU to MMD. This can be attributed to the IoU’s consideration of both position and height informa-

Method	IoU	MMD	CMD	PNC	IDF1↑	HOTA↑
MOT17 Validation Set						
baseline	✓	-	-	-	77.10	68.43
UCMCTrack-v1	-	✓	-	-	75.88	68.09
UCMCTrack-v2	-	✓	✓	-	79.68	70.44
UCMCTrack-v3	-	✓	✓	✓	82.20	71.96
DanceTrack Validation Set						
baseline	✓	-	-	-	47.27	47.93
UCMCTrack-v1	-	✓	-	-	43.76	46.32
UCMCTrack-v2	-	✓	✓	-	53.93	55.06
UCMCTrack-v3	-	✓	✓	✓	62.64	60.42

Table 4: Ablation of UCMCTrack components.

Method	CMC	UCMC	IDF1↑	HOTA↑
MOT17 Validation Set				
baseline	-	-	77.10	68.43
baseline+CMC	✓	-	81.07	70.97
UCMCTrack	-	✓	82.20	71.96
UCMCTrack+	✓	✓	84.05	72.97
DanceTrack Validation Set				
baseline	-	-	47.27	47.93
baseline+CMC	✓	-	46.74	47.55
UCMCTrack	-	✓	62.64	60.42
UCMCTrack+	✓	✓	62.52	59.18

Table 5: CMC ablation.

tion, whereas MMD solely focuses on position, leading to the observed decline in outcomes. However, when CMD is applied post MMD, it better utilizes the distribution information, thus surpassing the performance of the IoU-based baseline. This observation also underscores the potential benefits of integrating height information into MMD for future research. Lastly, the incorporation of PNC effectively mitigates the motion noise introduced by camera movements, elevating the tracking performance to state-of-the-art levels.

CMC Ablation To explore the impact of CMC on UCMCTrack, we conducted ablation experiments on MOT17 and DanceTrack validation datasets, as shown in Table 5. In MOT17, employing UCMC results in a greater performance enhancement compared to using the baseline combined with




Figure 4: In-depth analysis of key parameters and robustness in UCMCTrack.

CMC, underscoring the effective role of UCMC in compensating for camera motion. Furthermore, a subsequent application of CMC to UCMC yields an additional performance boost, consistent with the results observed on the MOT17 test set. Interestingly, upon employing CMC on DanceTrack, the performance of the baseline and UCMC algorithm actually deteriorated. This can be attributed to the fact that most video sequences on DanceTrack don’t exhibit pronounced viewpoint shifts. After employing CMC, minor detection offsets for overlapping targets might result in matching with an incorrect trajectory, culminating in a mild performance setback. In contrast, the use of UCMC resulted in a significant performance boost, suggesting its effectiveness in scenes with only minor camera jitters.

Influence of Process Compensation Factor In order to further investigate the impact of the process compensation factors on the performance of UCMCTrack, we divided the MOT17 validation sequences into dynamic and static scenes. The horizontal axis represents the natural logarithm of the compensation factor, with both σ_x and σ_y set to σ as per Eq. 11, and the vertical axis represents the difference in HOTA compared to the baseline. As illustrated in Figure 4a, the influence of the compensation factor on dynamic and static scenes exhibits two distinct patterns. This divergence underscores the necessity of tailoring the compensation parameters separately for each scene type. Adapting these parameters specific to the scene’s nature ensures that the tracker operates at its optimum performance. This distinction also highlights the role of the compensation factor in mitigating the effects of camera movements. For static scenes, a smaller compensation factor is recommended, while for dynamic scenes, a larger compensation factor is essential to counteract the impact of camera motion, thus enhancing the tracking performance.

Influence of Detector Noise Factor We conducted ablation studies on the hyperparameter σ_m in Eq. 2 to explore its impact on UCMCTrack. As shown in Figure 4b, the tracker’s performance varies with different σ_m . When the σ_m is set to 0.05, the HOTA and IDF1 metrics reach their highest values. It is evident that the influence of the σ_m on HOTA and IDF1 remains relatively minor within the range of 0.04 to 0.1. This

indicates that our UCMCTrack is not sensitive to σ_m .

Robustness to Camera Parameters Error Our method, UCMCTrack, relies on the camera parameters to project targets from the image plane to the ground plane. However, in scenarios where camera parameters are not provided, we manually estimate them. This means that the estimated camera parameters may not be accurate. We conducted separate ablation experiments for the camera extrinsics errors along the X, Y, and Z axes, and the results are shown in Figure 4c. We observed that errors in the Y axis have a minor impact on performance. On the other hand, errors in the X and Z axes have more significant effects. This is due to the Y-axis corresponding to the yaw direction, variations in this direction have a less impact on the deformation of the estimated ground plane. However, When there are substantial errors in the estimated camera extrinsics along the X and Y axes, the performance of UCMCTrack notably degrades. This can be attributed to significant deformations in the estimated ground plane. Under such circumstances, adjustments to the camera parameters are required to ensure that the estimated ground plane aligns closely with the actual terrain.

Conclusion

In this work, we presented UCMCTrack, a state-of-the-art multi-object tracker demonstrating superior performance across a variety of datasets. UCMCTrack employs a novel distance measure based on normalized mahalanobis distance on mapped ground plane, marking a significant departure from the conventional reliance on IoU. This innovation enables the tracker to adeptly handle the challenges introduced by camera movements, using only a consistent set of compensation parameters across a single video sequence. However, an inherent limitation of UCMCTrack is its assumption that targets reside on a singular ground plane. Looking forward, there is considerable potential to enhance UCMCTrack’s effectiveness by integrating it with well-established distance measures such as IoU and ReID. We believe that this will lay the groundwork for subsequent advancements in the future research of motion-based multi-object tracking.

Acknowledgments

This work was supported by the National Key Research and Development Program of China under Grant 2022YFC3803700, in part by the Natural Science Foundation of China under Grants 52002036, the Changsha Science and Technology Major Project under Grant kh2202002 and kh2301004, the Hunan Provincial Natural Science Foundation of China under Grant 2022JJ30611, the Scientific Research Fund of Hunan Provincial Education Department under Grant 21B0342.

References

- Aharon, N.; Orfaig, R.; and Bobrovsky, B.-Z. 2022. BoT-SORT: Robust associations multi-pedestrian tracking. *arXiv preprint arXiv:2206.14651*.
- Babae, M.; Li, Z.; and Rigoll, G. 2019. A dual CNN-RNN for multiple people tracking. *Neurocomputing*, 368: 69–83.
- Bergmann, P.; Meinhardt, T.; and Leal-Taixe, L. 2019. Tracking Without Bells and Whistles. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*.
- Bernardin, K.; and Stiefelhagen, 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.; and Upcroft, B. 2016. Simple Online and Realtime Tracking. In *2016 IEEE International Conference on Image Processing (ICIP)*, 3464–3468.
- 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 (CVPR)*, 9686–9696.
- Cetintas, O.; Brasó, G.; and Leal-Taixé, L. 2023. Unifying Short and Long-Term Tracking With Graph Hierarchies. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 22877–22887.
- Choi, W. 2015. Near-online multi-target tracking with aggregated local flow descriptor. In *Proceedings of the IEEE international conference on computer vision*, 3029–3037.
- Dendorfer, P.; Rezatofighi, 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*.
- Dendorfer, P.; Yugay, V.; Osep, A.; and Leal-Taixé, L. 2022. Quo Vadis: Is Trajectory Forecasting the Key Towards Long-Term Multi-Object Tracking? *Advances in Neural Information Processing Systems*, 35: 15657–15671.
- Du, Y.; Wan, J.; Zhao, Y.; Zhang, B.; Tong, Z.; and Dong, J. 2021. GIAOTracker: A Comprehensive Framework for MC-MOT With Global Information and Optimizing Strategies in VisDrone 2021. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) Workshops*, 2809–2819.
- Du, Y.; Zhao, Z.; Song, Y.; Zhao, Y.; Su, F.; Gong, T.; and Meng, H. 2023. Strongsort: Make deepsort great again. *IEEE Transactions on Multimedia*.
- Evangelidis, G. D.; and Psarakis, E. Z. 2008a. Parametric image alignment using enhanced correlation coefficient maximization. *IEEE transactions on pattern analysis and machine intelligence*, 30(10): 1858–1865.
- Evangelidis, G. D.; and Psarakis, E. Z. 2008b. Parametric Image Alignment Using Enhanced Correlation Coefficient Maximization. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 30(10): 1858–1865.
- Feichtenhofer, C.; Pinz, A.; and Zisserman, A. 2017. Detect to track and track to detect. In *Proceedings of the IEEE international conference on computer vision*, 3038–3046.
- Fischler, M. A.; and Bolles, R. C. 1981. Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. *Communications of the ACM*, 24(6): 381–395.
- Ge, Z.; Liu, S.; Wang, F.; Li, Z.; and Sun, J. 2021. YOLOX: Exceeding YOLO Series in 2021. *CoRR*, abs/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.
- Girbau, A.; Marqués, F.; and Satoh, S. 2022. Multiple Object Tracking from appearance by hierarchically clustering tracklets. *arXiv preprint arXiv:2210.03355*.
- Han, S.; Huang, P.; Wang, H.; Yu, E.; Liu, D.; and Pan, X. 2022. MAT: Motion-aware multi-object tracking. *Neurocomputing*, 476: 75–86.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep Residual Learning for Image Recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Held, D.; Thrun, S.; and Savarese, S. 2016. Learning to track at 100 fps with deep regression networks. In *Computer Vision—ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part I 14*, 749–765. Springer.
- Hu, H.-N.; Yang, Y.-H.; Fischer, T.; Darrell, T.; Yu, F.; and Sun, M. 2022. Monocular quasi-dense 3d object tracking. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(2): 1992–2008.
- Khurana, T.; Dave, A.; and Ramanan, D. 2021. Detecting Invisible People. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 3174–3184.
- Liu, J.; Wang, Z.; and Xu, M. 2020. DeepMTT: A deep learning maneuvering target-tracking algorithm based on bidirectional LSTM network. *Information Fusion*, 53: 289–304.
- Liu, Q.; Chu, Q.; Liu, B.; and Yu, N. 2020. GSM: Graph Similarity Model for Multi-Object Tracking. In *IJCAI*, 530–536.
- Liu, Z.; Wang, X.; Wang, C.; Liu, W.; and Bai, X. 2023. SparseTrack: Multi-Object Tracking by Performing Scene Decomposition based on Pseudo-Depth. *arXiv:2306.05238*.

- Luiten, J.; Osep, A.; Dendorfer, P.; Torr, P.; Geiger, A.; Leal-Taixé, L.; and Leibe, B. 2021. Hota: A higher order metric for evaluating multi-object tracking. *International journal of computer vision*, 129: 548–578.
- Maggiolino, G.; Ahmad, A.; Cao, J.; and Kitani, K. 2023. Deep oc-sort: Multi-pedestrian tracking by adaptive re-identification. *arXiv preprint arXiv:2302.11813*.
- Marinello, N.; Proesmans, M.; and Van Gool, L. 2022. TripletTrack: 3D Object Tracking Using Triplet Embeddings and LSTM. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, 4500–4510.
- Milan, A.; Leal-Taixé, L.; Reid, I.; Roth, S.; and Schindler, K. 2016. MOT16: A benchmark for multi-object tracking. *arXiv preprint arXiv:1603.00831*.
- Rangesh, A.; Maheshwari, P.; Gebre, M.; Mhatre, S.; Ramezani, V.; and Trivedi, M. M. 2021. TrackMPNN: A message passing graph neural architecture for multi-object tracking. *arXiv preprint arXiv:2101.04206*.
- Rezatofighi, H.; Tsoi, N.; Gwak, J.; Sadeghian, A.; Reid, I.; and Savarese, S. 2019. Generalized intersection over union: A metric and a loss for bounding box regression. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 658–666.
- Ristani, E.; Solera, F.; Zou, R.; Cucchiara, R.; and Tomasi, C. 2016. Performance measures and a data set for multi-target, multi-camera tracking. In *Computer Vision—ECCV 2016 Workshops: Amsterdam, The Netherlands, October 8–10 and 15–16, 2016, Proceedings, Part II*, 17–35. Springer.
- Rublee, E.; Rabaud, V.; Konolige, K.; and Bradski, G. 2011. ORB: An efficient alternative to SIFT or SURF. In *2011 International conference on computer vision*, 2564–2571. Ieee.
- Seidenschwarz, J.; Brasó, G.; Serrano, V. C.; Elezi, I.; and Leal-Taixé, L. 2023. Simple Cues Lead to a Strong Multi-Object Tracker. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 13813–13823.
- Sun, P.; Cao, J.; Jiang, Y.; Yuan, Z.; Bai, S.; Kitani, K.; and 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 (CVPR)*, 20993–21002.
- Tokmakov, P.; Li, J.; Burgard, W.; and Gaidon, A. 2021. Learning To Track With Object Permanence. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 10860–10869.
- Wojke, N.; Bewley, A.; and Paulus, D. 2017. Simple Online and Realtime Tracking with a Deep Association Metric. In *2017 IEEE International Conference on Image Processing (ICIP)*, 3645–3649.
- Xiao, B.; Wu, H.; and Wei, Y. 2018. Simple baselines for human pose estimation and tracking. In *Proceedings of the European conference on computer vision (ECCV)*, 466–481.
- Xiao, C.; Cao, Q.; Zhong, Y.; Lan, L.; Zhang, X.; Cai, H.; Luo, Z.; and Tao, D. 2023. MotionTrack: Learning Motion Predictor for Multiple Object Tracking. *arXiv:2306.02585*.
- Yang, F.; Odashima, S.; Masui, S.; and Jiang, S. 2023. Hard to track objects with irregular motions and similar appearances? make it easier by buffering the matching space. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 4799–4808.
- Yang, J.; Ge, H.; Su, S.; and Liu, G. 2022. Transformer-based two-source motion model for multi-object tracking. *Applied Intelligence*, 1–13.
- Yeh, C.-H.; Lin, C.-Y.; Muchtar, K.; Lai, H.-E.; and Sun, M.-T. 2017. Three-pronged compensation and hysteresis thresholding for moving object detection in real-time video surveillance. *IEEE Transactions on Industrial Electronics*, 64(6): 4945–4955.
- Yu, J.; and McMillan, L. 2004. General linear cameras. In *Computer Vision—ECCV 2004: 8th European Conference on Computer Vision, Prague, Czech Republic, May 11–14, 2004. Proceedings, Part II* 8, 14–27. Springer.
- Yu, Y.; Kurnianggoro, L.; and Jo, K.-H. 2019. Moving object detection for a moving camera based on global motion compensation and adaptive background model. *International Journal of Control, Automation and Systems*, 17: 1866–1874.
- 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 *Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXII*, 1–21. Springer.
- Zhang, Y.; Wang, T.; and Zhang, X. 2023. MOTRv2: Bootstrapping End-to-End Multi-Object Tracking by Pretrained Object Detectors. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 22056–22065.
- Zheng, Z.; Wang, P.; Liu, W.; Li, J.; Ye, R.; and Ren, D. 2020. Distance-IoU loss: Faster and better learning for bounding box regression. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, 12993–13000.
- Zhou, X.; Koltun, V.; and Krähenbühl, P. 2020. Tracking objects as points. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part IV*, 474–490. Springer.