

PlugTrack: Multi-Perceptive Motion Analysis for Adaptive Fusion in Multi-Object Tracking

Seungjae Kim, SeungJoon Lee, MyeongAh Cho*

Department of Software Convergence, Kyung Hee University
{tmdwo8814, diplomat3334, maycho}@khu.ac.kr

Abstract

Multi-object tracking (MOT) predominantly follows the tracking-by-detection paradigm, where Kalman filters serve as the standard motion predictor due to computational efficiency but inherently fail on non-linear motion patterns. Conversely, recent data-driven motion predictors capture complex non-linear dynamics but suffer from limited domain generalization and computational overhead. Through extensive analysis, we reveal that even in datasets dominated by non-linear motion, Kalman filter outperforms data-driven predictors in up to 34% of cases, demonstrating that real-world tracking scenarios inherently involve both linear and non-linear patterns. To leverage this complementarity, we propose PlugTrack, a novel framework that adaptively fuses Kalman filter and data-driven motion predictors through multi-perceptive motion understanding. Our approach employs multi-perceptive motion analysis to generate adaptive blending factors. PlugTrack achieves significant performance gains on MOT17/MOT20 and state-of-the-art on DanceTrack without modifying existing motion predictors. To the best of our knowledge, PlugTrack is the first framework to bridge classical and modern motion prediction paradigms through adaptive fusion in MOT.

Code —

<https://github.com/VisualScienceLab-KHU/PlugTrack>

Introduction

Multi-object tracking (MOT) aims to detect and track multiple objects across video sequences while maintaining identity consistency. This fundamental computer vision task enables extensive applications including autonomous driving, surveillance systems, and robotics. Most MOT methods follow the tracking-by-detection paradigm, which decomposes the problem into three subtasks: object detection, motion prediction, and association between detections and predictions. Among these subtasks, motion prediction plays a critical role in maintaining object identities by accurately forecasting object movements during occlusions. Given the importance of real-time performance in MOT applications, tracking-by-detection methods have traditionally adopted the Kalman filter as their motion predictor

*Corresponding author.

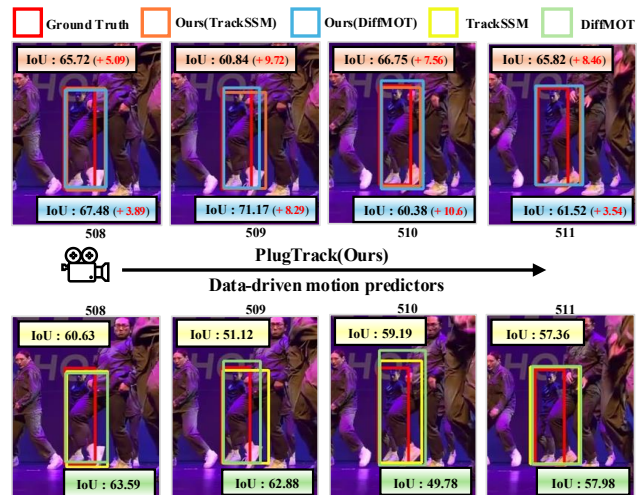


Figure 1: Qualitative comparison of motion prediction on DanceTrack (frames 508-511) dataset. Top row: PlugTrack results using TrackSSM and DiffMOT as base predictors. Bottom row: Standalone TrackSSM and DiffMOT results. PlugTrack adaptively fuses Kalman filter and data-driven predictions to better handle both linear and non-linear motions, achieving up to +10.6 IoU gains. This demonstrates its core novelty: a lightweight, plug-in mechanism that dynamically integrates complementary motion cues to outperform individual predictors in complex scenarios.

due to its parameter-free design and computational efficiency. The Kalman filter’s recursive update mechanism inherently makes linear motion assumptions, proving effective on datasets dominated by linear movements such as MOT17 (Milan et al. 2016) and MOT20 (Dendorfer et al. 2020). However, this same recursive mechanism renders it ineffective on non-linear motion datasets like DanceTrack (Sun et al. 2022) and SportsMOT (Cui et al. 2023), where complex movements and abrupt directional changes violate its fundamental linear assumptions.

To achieve robust tracking under non-linear motion, existing tracking-by-detection methods have pursued two primary approaches. The first approach employs Kalman filter-based methods with rule-based modifications such as OC-

SORT (Cao et al. 2023) and Hybrid-SORT (Yang et al. 2024). These heuristic methods ensure real-time performance but require extensive scenario-specific parameter tuning. This dependency fundamentally limits their domain generalization capabilities. The second approach replaces the Kalman filter with data-driven motion predictors such as DiffMOT (Lv et al. 2024), which leverages diffusion models for probabilistic motion prediction, and TrackSSM (Hu et al. 2024), which employs state-space models to capture temporal dynamics. Unlike the Kalman filter, these learnable architectures demonstrate robust performance in capturing non-linear motion. However, they tend to overfit to their training data distributions, making it challenging to adapt to diverse motion patterns across different domains. These limitations highlight the critical need for developing a motion predictor that achieves both generalization across varied motion patterns and real-time efficiency.

To understand the strengths of existing motion predictors, we conducted evaluations across both linear and non-linear motion domains. We assessed three motion predictors—Kalman filter, DiffMOT, and TrackSSM—by generating predictions on identical tracklets and measuring their IoU scores against ground truth. Figure 2 presents the number of tracklets where each predictor achieved the highest IoU score. On the linear motion domain MOT17, the Kalman filter demonstrates clear superiority over data-driven motion predictors, achieving the best predictions on 60.3% of all tracklets (12,061 out of 20,000). Surprisingly, even on the non-linear motion domain DanceTrack, the Kalman filter outperforms data-driven motion predictors in 34% of all tracklets (1,700 out of 5,000). This surprising result indicates that linear motion patterns frequently occur even within datasets designed for complex non-linear movements. This observation reveals a crucial insight: **real-world tracking scenarios inherently involve both linear and non-linear motion patterns, regardless of domain-specific characteristics.** This necessitates a unified framework that can adaptively respond to both motion types rather than treating them as mutually exclusive.

Based on these observations and the limitations of existing approaches, we identify two critical requirements for a practical motion prediction system in MOT. (1) **Complementary modeling of linear and non-linear motion:** to handle the full spectrum of real-world motion dynamics, a unified approach must integrate both Kalman filter and data-driven motion predictors. (2) **Real-time capability and generalization:** the framework must be lightweight enough for real-time operation, while maintaining robustness across unseen domains and motion types.

To this end, we propose **PlugTrack**, a framework where a Kalman filter and a data-driven motion predictor work in a complementary and adaptive manner. Our approach consists of two key components: (1) the Contextual Motion Encoder (CME) that performs multi-perceptive motion analysis, and (2) the Adaptive Blending Generator (ABG) that produces adaptive blending factors based on CME’s multi-perceptive understanding. These adaptive blending factors are then used to blend predictions from both the Kalman filter and data-driven motion predictor, achieving refined mo-

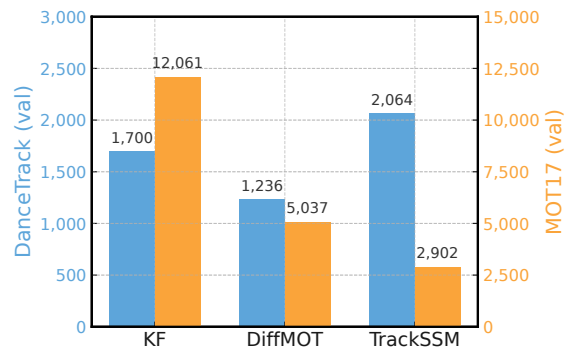


Figure 2: Comparison of motion predictor performance on DanceTrack and MOT17, showing the number of tracklets where each predictor (Kalman filter, DiffMOT, TrackSSM) achieves the highest IoU score with ground truth.

tion prediction. This plug-and-play design enhances existing pretrained data-driven motion predictors, while PlugTrack’s lightweight architecture maintains real-time capability.

Related Work

Kalman Filter Based Motion Prediction Methods

The Kalman filter has dominated motion prediction in MOT due to its computational efficiency. SORT (Bewley et al. 2016) established the tracking-by-detection paradigm, while DeepSORT (Wojke, Bewley, and Paulus 2017) added appearance features. Recent methods attempt to address the Kalman filter’s inability to handle non-linear motion through heuristic modifications: OC-SORT (Cao et al. 2023) introduces observation-centric recovery, C-BIoU (Yang et al. 2023) proposes buffered matching, and Hybrid-SORT (Yang et al. 2024) combines multiple strategies. However, these approaches share critical limitations: (1) they require extensive manual tuning for different scenarios, (2) their rule-based modifications cannot overcome the Kalman filter’s fundamental linear motion assumption, and (3) they lack adaptability to varying motion complexities within sequences. While computationally efficient, these methods have limited domain generalization capabilities and cannot adapt to diverse motion patterns, particularly failing on datasets with complex movements like DanceTrack.

Data-driven Motion Prediction Methods

Data-driven approaches address the limitations of the Kalman filter’s linear motion assumption by capturing complex motion patterns. TrackSSM (Hu et al. 2024) and DiffMOT (Lv et al. 2024) replace the Kalman filter with state-space models and diffusion-based models respectively for motion prediction. Despite their superior performance on non-linear benchmarks, data-driven predictors introduce several challenges: (1) Domain overfitting: Models trained on complex motions struggle with simple linear trajectories, (2) Computational overhead: Methods like MOTIP (Gao, Qi, and Wang 2025) sacrifice real-time capability for accuracy, and (3) Black-box nature: Unlike the interpretable nature of

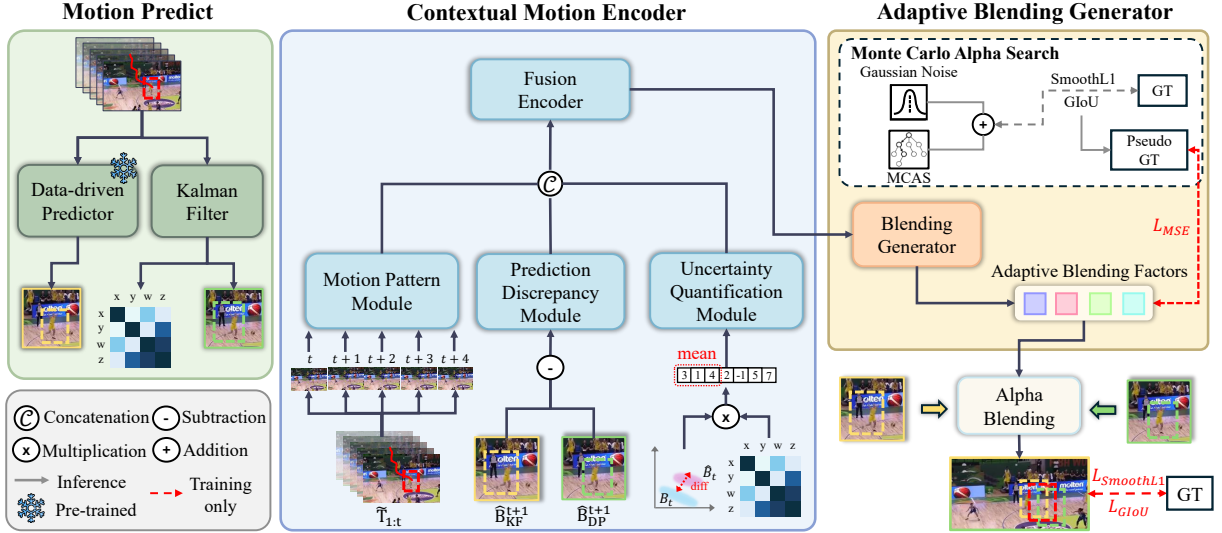


Figure 3: Overview of the PlugTrack architecture. Our framework consists of two main components: (1) Contextual Motion Encoder (CME) that analyzes motion from multi-perceptive through three specialized modules to generate multi-perceptive motion feature. (2) Then Adaptive Blending Generator (ABG) that produces adaptive blending factors for alpha blending. During training, Monte Carlo Alpha Search (MCAS) generates pseudo ground truth blending factors by evaluating multiple candidates with added Gaussian noise. During inference, the learned ABG directly predicts optimal blending factors for real-time adaptive fusion of Kalman filter and data-driven motion predictor outputs.

Kalman filters, neural predictors offer limited transparency regarding their failure modes.

The critical gap is the false dichotomy between Kalman filter and data-driven motion predictors. Current approaches treat them as mutually exclusive, missing their complementary strengths. Real-world tracking scenarios inherently contain both linear and non-linear motion patterns that no single approach can optimally handle. While adaptive methods like RelationTrack (Yu et al. 2022) and StrongSORT (Du et al. 2023) explore fusion strategies, they operate within homogeneous paradigms. None bridge the divide between Kalman filter and data-driven approaches, necessitating an adaptive framework that intelligently blends both based on motion context.

Proposed Method

We present PlugTrack, a novel framework that adaptively fuses the Kalman filter and data-driven motion predictors. As shown in Figure 3, our approach consists of two core components. The Contextual Motion Encoder (CME) extracts multi-perceptive motion features \mathbf{f}_{mult} by analyzing tracklet patterns, prediction discrepancies, and uncertainty measures. This multi-perceptive analysis enables CME to capture nuanced motion characteristics that single predictors might miss. The Adaptive Blending Generator (ABG) then converts \mathbf{f}_{mult} into adaptive blending factors α , which determine the optimal contribution of each predictor based on the current motion context. These factors dynamically weight the predictions from both motion predictors to produce the final bounding box. PlugTrack’s plug-and-play design allows seamless integration with any existing data-driven mo-

tion predictor without architectural modifications.

Contextual Motion Encoder

To formally characterize the input space of our method, we define the following representations. Let $\tilde{\mathbf{T}}_{1:t} = \{\mathbf{B}_1, \dots, \mathbf{B}_f, \dots, \mathbf{B}_t\}$ denote the tracklet up to time t , where each \mathbf{B}_f represents the detection result at frame f . Each bounding box $\mathbf{B}_f = (x_f, y_f, w_f, h_f, \Delta x_f, \Delta y_f, \Delta w_f, \Delta h_f)$ contains both position and velocity information, where Δ denotes the difference from frame $f - 1$ to f . The Kalman filter and data-driven motion predictor generate predictions $\hat{\mathbf{B}}_{t+1}^{KF} \in \mathbb{R}^4$ and $\hat{\mathbf{B}}_{t+1}^{DP} \in \mathbb{R}^4$, respectively. Additionally, we extract motion uncertainty $\sigma_{KF} \in \mathbb{R}^4$ from the Kalman filter’s internal state estimates.

The CME consists of three specialized modules that examine different aspects of motion patterns:

(1) **Motion Pattern Module.** As the first component of multi-perceptive motion analysis, the Motion Pattern Module (MPM) captures temporal dependencies and motion characteristics from the tracklet. We employ a Long Short-Term Memory network (Graves 2012) to encode the sequential motion information:

$$\mathbf{f}_{MPM} = \text{LSTM}(\tilde{\mathbf{T}}_{1:t}) = \mathbf{h}_t \quad (1)$$

where $\mathbf{h}_t \in \mathbb{R}^{128}$ represents the final hidden state of the LSTM. This design enables MPM to learn complex motion patterns such as acceleration, deceleration, and directional changes.

(2) **Prediction Discrepancy Module.** As the second component of multi-perceptive motion analysis, the Prediction Dis-

crepancy Module (PDM) quantifies and interprets the differences between the Kalman filter and data-driven motion predictor predictions. We compute the prediction discrepancy as $\mathbf{B}_{disc} = \hat{\mathbf{B}}_{t+1}^{KF} - \hat{\mathbf{B}}_{t+1}^{DP}$, which is then processed through multiple linear layers:

$$\mathbf{f}_{PDM} = \text{MLP}(\mathbf{B}_{disc}) \quad (2)$$

where $\mathbf{f}_{PDM} \in \mathbb{R}^{32}$. The PDM learns to interpret this discrepancy as indicators of motion complexity and predictor reliability. A large discrepancy often signals transitions between linear and non-linear motion phases. In the subsequent fusion process, \mathbf{f}_{PDM} is fused with \mathbf{f}_{MPM} . This fusion enables the model to contextualize prediction differences within specific motion scenarios, providing valuable cues for determining when to trust each predictor.

(3) Uncertainty Quantification Module. As the third component of multi-perceptive motion analysis, the Uncertainty Quantification Module (UQM) leverages the uncertainty measure extracted from the Kalman filter’s internal state estimation process. We compute dimension-wise Normalized Innovation Squared (NIS) to quantify prediction confidence for each bounding box dimension:

$$\text{NIS}_{t,i} = \frac{(B_{t,i} - \hat{B}_{t,i})^2}{S_{t,ii}} \quad (3)$$

where $B_{t,i} - \hat{B}_{t,i}$ is the innovation (prediction error) for dimension $i \in \{x, y, w, h\}$, and $S_{t,ii}$ is the corresponding diagonal element of the innovation covariance matrix \mathbf{S}_t . Please refer to Appendix for more details.

To obtain a robust uncertainty measure, we aggregate recent NIS values by summing their mean and standard deviation, capturing not only the average prediction error but also its consistency over time. This yields a 4-dimensional uncertainty vector $\boldsymbol{\sigma}_{KF} \in \mathbb{R}^4$ that comprehensively characterizes the Kalman filter’s confidence across all bounding box dimensions.

This uncertainty vector is processed through an MLP to extract motion-relevant features:

$$\mathbf{f}_{UQM} = \text{MLP}(\boldsymbol{\sigma}_{KF}) \quad (4)$$

where $\mathbf{f}_{UQM} \in \mathbb{R}^{32}$. The UQM provides an objective measure of the Kalman filter’s prediction reliability. High uncertainty values indicate that the Kalman filter has low confidence in its predictions, suggesting the presence of non-linear motion patterns that violate its linear assumptions. This uncertainty-based motion characterization enables the model to implicitly distinguish between linear and non-linear motion phases for adaptive blending.

The outputs from the three modules, each providing a distinct component of multi-perceptive analysis, are subsequently fused to create a multi-perceptive motion feature \mathbf{f}_{mult} :

$$\mathbf{f}_{mult} = \text{Encoder}(\text{Concat}(\mathbf{f}_{MPM}, \mathbf{f}_{PDM}, \mathbf{f}_{UQM})) \quad (5)$$

This multi-perceptive analysis enables CME to capture the nuanced relationships between motion patterns, prediction differences, and uncertainty levels, providing rich information for subsequent adaptive blending decisions.

Adaptive Blending Generator

The ABG transforms the multi-perceptive motion feature \mathbf{f}_{mult} into adaptive blending factors $\tilde{\boldsymbol{\alpha}} = \text{MLP}(\mathbf{f}_{mult})$. The resulting $\tilde{\boldsymbol{\alpha}}$ is a 4-dimensional vector containing blending factors for each bounding box coordinate (x, y, w, h). The ABG is implemented as an MLP with sigmoid activation in the final layer, ensuring all blending factors are bounded in range $[0, 1]$.

The final bounding box prediction synthesizes both motion predictions through element-wise weighted blending:

$$\hat{\mathbf{B}}_{ABG} = \tilde{\boldsymbol{\alpha}} \odot \hat{\mathbf{B}}_{KF} + (1 - \tilde{\boldsymbol{\alpha}}) \odot \hat{\mathbf{B}}_{DP} \quad (6)$$

where $\tilde{\boldsymbol{\alpha}} = (\alpha_x, \alpha_y, \alpha_w, \alpha_h)$ contains the adaptive blending factors. This coordinate-specific formulation allows ABG to dynamically adjust the predictions of each predictor based on the motion characteristics identified by CME. For instance, when horizontally linear motion occurs, CME detects stable horizontal patterns through MPM while UQM indicates low uncertainty, then ABG will assign higher weight to the Kalman filter for x-coordinate ($\alpha_x > 0.5$). Conversely, when vertically non-linear motion occurs, PDM reveals large prediction discrepancies during rapid vertical movements, ABG will rely more on the data-driven motion predictor for y-coordinate ($\alpha_y < 0.5$). This adaptive strategy, guided by CME’s multi-perceptive analysis, effectively handles complex motion scenarios without manual intervention.

Monte Carlo Alpha Search

Training adaptive blending factors through direct optimization presents a fundamental challenge: the ABG tends to converge to dataset-specific biases rather than learning to adaptively respond to varying motion dynamics. Specifically, when trained on MOT17 where linear motion dominates, the ABG consistently assigns high weights to the Kalman filter’s predictions across all scenarios, failing to recognize cases where the data-driven motion predictor would be more appropriate. This *bias collapse* problem prevents the model from discovering the complementary strengths of both predictors.

To address this challenge, we propose Monte Carlo Alpha Search (MCAS), inspired by the success of Monte Carlo methods in navigating complex optimization landscapes. Similar to how Monte Carlo methods in (Hampali et al. 2021) explore discrete proposals in 3D scene understanding by systematically evaluating candidate solutions, MCAS searches through a discrete space of blending factor combinations to find optimal fusion strategies for each training sample.

We define a discrete search space of blending factors:

$$\mathcal{A} = \{0.3, 0.4, 0.5, 0.6, 0.7\}^4 \quad (7)$$

This creates $5^4 = 625$ candidate combinations, constraining the search within $[\lambda_1, \lambda_2] = [0.3, 0.7]$ to ensure meaningful contributions from both predictors.

For each training batch, we introduce stochastic exploration by adding Gaussian noise:

$$\tilde{\mathcal{A}}_b = \text{clamp}(\mathcal{A} + \epsilon_b, 0, 1), \quad \epsilon_b \sim \mathcal{N}(0, 0.1^2) \quad (8)$$

Method	MOT17					MOT20				
	HOTA	IDF1	AssA	DetA	MOTA	HOTA	IDF1	AssA	DetA	MOTA
<i>Kalman Filter Based</i>										
FairMOT	59.3	72.3	58.0	60.9	73.7	54.6	67.3	54.7	54.7	61.8
ByteTrack	63.1	77.3	62.0	64.5	80.3	61.3	75.2	59.6	63.4	77.8
OC-SORT	63.2	77.5	63.4	63.2	78.0	62.4	76.3	62.5	62.4	75.7
C-BIoU	64.1	79.7	63.7	64.8	79.7	/	/	/	/	/
Hybrid-SORT	63.6	78.4	/	/	79.3	62.5	76.2	/	/	76.4
<i>Data-driven Based</i>										
DiffusionTrack	60.8	73.8	58.8	63.2	77.9	55.3	66.3	51.3	59.9	72.8
ETTrack	61.9	75.9	60.5	/	79.0	/	/	/	/	/
TrackSSM	61.4	74.1	59.6	63.6	78.5	59.1	71.1	57.5	60.9	73.9
DiffMOT	64.0	78.9	64.2	64.1	79.1	61.6	74.9	60.5	62.8	76.3
	61.9	75.2	60.3	63.9	78.7	59.7	72.3	58.5	61.3	74.5
Ours(TrackSSM)	(+0.5)	(+1.1)	(+0.7)	(+0.3)	(+0.1)	(+0.6)	(+1.2)	(+1.0)	(+0.4)	(+0.6)
	64.2	79.0	64.4	64.0	79.2	61.8	75.2	60.9	62.9	76.4
Ours(DiffMOT)	(+0.2)	(+0.1)	(+0.2)	(-0.1)	(+0.1)	(+0.2)	(+0.3)	(+0.4)	(+0.1)	(+0.1)

Table 1: Results on MOT17 & MOT20 test. The best results are shown in bold.

where the noise is applied independently to each candidate and batch, encouraging diverse exploration patterns.

Given predictions $\hat{\mathbf{B}}^{KF}$ and $\hat{\mathbf{B}}^{DP}$, we evaluate each candidate $\alpha \in \tilde{\mathcal{A}}_b$ by computing the blended prediction and measuring its accuracy:

$$\mathcal{S}(\alpha) = \text{SmoothL1}(\hat{\mathbf{B}}(\alpha), \mathbf{B}_{GT}) + \text{GIoU}(\hat{\mathbf{B}}(\alpha), \mathbf{B}_{GT}) \quad (9)$$

where $\hat{\mathbf{B}}(\alpha) = \alpha \odot \hat{\mathbf{B}}^{KF} + (1 - \alpha) \odot \hat{\mathbf{B}}^{DP}$.

The optimal blending factors for each sample are selected as:

$$\alpha^* = \arg \min_{\alpha \in \tilde{\mathcal{A}}_b} \mathcal{S}(\alpha) \quad (10)$$

These pseudo ground-truth values α^* serve as supervision signals for training the ABG through an auxiliary loss:

$$\mathcal{L}_{\text{MCAS}} = \text{MSE}(\tilde{\alpha}, \alpha^*) \quad (11)$$

By providing explicit supervision for optimal blending strategies, MCAS guides the ABG to learn context-appropriate fusion rather than dataset-specific biases. During inference, MCAS is not required as the trained ABG directly predicts adaptive blending factors, maintaining real-time efficiency.

Training loss

The overall training objective of PlugTrack is formulated as:

$$\mathcal{L} = \mathcal{L}_{\text{SmoothL1}} + \mathcal{L}_{\text{GIoU}} + \mathcal{L}_{\text{MCAS}} \quad (12)$$

where $\mathcal{L}_{\text{SmoothL1}}$ and $\mathcal{L}_{\text{GIoU}}$ measure the accuracy of the final blended prediction $\hat{\mathbf{B}}_{ABG}$ against ground truth, while $\mathcal{L}_{\text{MCAS}}$ provides auxiliary supervision for learning optimal blending strategies.

Experiments

Experimental Setup

Datasets. We evaluate our method on three widely-used MOT benchmarks: MOT17 (Milan et al. 2016),

MOT20 (Dendorfer et al. 2020), and DanceTrack (Sun et al. 2022). MOT17 and MOT20 predominantly feature pedestrians with linear motion patterns in surveillance scenarios. DanceTrack contains dancers performing complex non-linear movements with frequent occlusions and similar appearances. For training on linear motion domains, we combine MOT17 and MOT20 training sets to create a mixed dataset, which provides diverse linear motion patterns along with crowded scenes.

Evaluation Metrics. We adopt the standard MOT evaluation metrics for comprehensive assessment. HOTA (Luiten et al. 2021) provides a balanced measure of detection and association performance. IDF1 emphasizes identity preservation across frames. AssA and DetA decompose HOTA into association and detection components respectively. MOTA measures overall tracking accuracy considering false positives, false negatives, and identity switches.

Implementation Details. Following the tracking-by-detection paradigm, we employ YOLOX (Ge et al. 2021) as our detector for all experiments except FairMOT (Zhang et al. 2021) comparisons, ensuring fair comparison across methods. Our framework is trained using Adam optimizer with a learning rate of 0.001. We train for 270 epochs on MIX(MOT17&20) dataset and 220 epochs on DanceTrack without any data augmentation or sampling techniques. The batch size is set to 2,048. We use fixed-length tracklets of 5 frames as the model’s input. The LSTM encoder consists of 2 layers with hidden dimension 128. For the Uncertainty Quantification Module, we use a sliding window size of $w=3$ for computing the averaged NIS values. For MCAS, we set the alpha range bounds $\lambda_1=0.3$ and $\lambda_2=0.7$, with Gaussian noise standard deviation of 0.1 for stochastic exploration. Both TrackSSM and DiffMOT are initialized with their official pre-trained weights. All experiments are conducted on a single NVIDIA GeForce RTX 4090 GPU.

Method	DanceTrack				
	HOTA	IDF1	AssA	DetA	MOTA
Kalman Filter Based					
FairMOT	39.7	40.8	23.8	66.7	82.2
ByteTrack	47.3	52.5	31.4	71.6	89.5
OC-SORT	55.1	54.2	38.0	80.3	89.4
C-BIoU	60.6	61.6	45.4	81.3	91.6
Hybrid-SORT	62.2	63.0	/	/	91.6
Data-driven Based					
DiffusionTrack	52.4	47.5	33.5	82.2	89.3
ETTrack	56.4	57.5	39.1	81.7	92.2
MambaTrack	56.8	57.8	39.8	80.1	90.1
TrackSSM	57.7	57.5	41.0	81.5	92.2
MotionTrack	58.2	58.6	41.7	81.4	91.3
DiffMOT	62.3	63.0	47.2	82.5	92.8
	59.2	59.0	42.9	81.9	92.2
Ours(TrackSSM)	(+1.5)	(+1.5)	(+1.9)	(+0.4)	(0)
	63.3	64.1	48.4	82.5	92.4
Ours(DiffMOT)	(+1.0)	(+1.1)	(+1.2)	(0)	(-0.4)

Table 2: Results on DanceTrack test set. The best results in bold.

Benchmark Evaluation

Performance on Linear Motion Datasets. To evaluate PlugTrack’s ability to enhance data-driven predictors on linear motion, we test on MOT17 and MOT20 datasets. Table 1 shows Kalman filter-based methods achieve state-of-the-art performance across most metrics, consistent with Figure 2 where Kalman filter dominates data-driven predictors in linear scenarios. This explains the strong performance of C-BIoU and Hybrid-SORT, which augment Kalman filter with heuristics. Despite this, PlugTrack improves base predictors: Ours(TrackSSM) gains 0.5/0.6 HOTA on MOT17/MOT20, while Ours(DiffMOT) improves by 0.2/0.2. Notably, Ours(DiffMOT) achieves state-of-the-art HOTA (64.2) and AssA (64.4) on MOT17, demonstrating our adaptive fusion can surpass pure Kalman methods. These results validate that our framework leverages Kalman filter’s linear motion strength while maintaining data-driven flexibility.

Performance on Non-linear Motion Dataset. We further assess whether adaptive fusion improves tracking on complex non-linear motion. Table 2 presents DanceTrack results, where PlugTrack demonstrates substantial improvements: HOTA gains of 1.5 and IDF1 of 1.5 for TrackSSM, and 1.0 and 1.1 for DiffMOT. The significant AssA gains indicate effective identity maintenance during challenging motion transitions. PlugTrack(DiffMOT) achieves state-of-the-art performance with 63.3 HOTA and 48.4 AssA, surpassing specialized methods like MotionTrack and MambaTrack. These results confirm that intelligently combining Kalman filter and data-driven predictors yields superior performance on non-linear motion-dominant datasets.

Cross-domain Generalization

To investigate PlugTrack’s generalization across fundamentally different motion domains, we conduct bidirectional transfer experiments. Table 3 shows cross-domain results using DanceTrack validation and MOT20 test sets.

Model	HOTA	IDF1	AssA
DanceTrack → MOT20			
TrackSSM	47.7	56.0	41.7
PlugTrack(TrackSSM)	54.2	65.2	50.4
MOT20 → DanceTrack			
TrackSSM	52.5	49.0	34.8
PlugTrack(TrackSSM)	54.3	52.3	37.1

Table 3: Cross-domain Performance Comparison

When trained on DanceTrack and tested on MOT20, PlugTrack(TrackSSM) achieves remarkable improvements: HOTA +6.5, IDF1 +9.2, and AssA +8.7. This large gain occurs because MOT20 contains many linear motion patterns where Kalman filter naturally excels, as shown in Figure 2. Our CME correctly identifies these linear patterns, and ABG assigns higher weights to the Kalman filter’s predictions. Conversely, when trained on MOT20 and tested on DanceTrack, PlugTrack still improves performance: HOTA +1.8, IDF1 +3.3, and AssA +2.3. These gains are smaller but significant—both the Kalman filter and linear-domain-trained predictor face extreme challenges in non-linear domains. That our framework still finds optimal blending strategies in this adverse setting demonstrates the strength of multi-perceptive motion analysis, proving PlugTrack can extract meaningful improvements even when both predictors are mismatched to the target domain.

Efficiency Analysis

Real-time performance is critical for practical MOT applications. Table 4 demonstrates that PlugTrack maintains real-time capability while achieving substantial tracking improvements. Our lightweight design adds only 0.54M parameters to both base predictors, representing a 22% increase for TrackSSM and 4.7% for DiffMOT. The resulting processing speeds remain practical: PlugTrack(TrackSSM) operates at 34.2 FPS (down from 37.2 FPS) and PlugTrack(DiffMOT) at 24.7 FPS (down from 25.9 FPS), both exceeding the 20 FPS real-time threshold. This efficiency stems from our compact architecture—CME uses a 128-dimensional LSTM with lightweight MLPs, while ABG consists of simple fully-connected layers. These results validate that our multi-perceptive motion analysis introduces minimal computational overhead, successfully balancing enhanced tracking accuracy with real-time constraints.

Model	Params (M)	FPS
TrackSSM	2.5	37.2
PlugTrack(TrackSSM)	3.04	34.2
DiffMOT	11.52	25.9
PlugTrack(DiffMOT)	12.06	24.7

Table 4: Efficiency of PlugTrack

Ablation Study

Component Analysis. To understand each component’s contribution to multi-perceptive analysis, we conduct systematic ablations. Table 5 analyzes the contribution of each

MPM	PDM	UQM	HOTA	AssA	IDF1
Ours (DiffMOT)					
			59.2	44.5	59.7
✓			60.2	45.8	61.2
✓	✓		60.4	46.0	61.8
✓		✓	60.3	46.1	61.4
✓	✓	✓	60.8	46.6	61.7

Table 5: Ablation study of CME Analyzers on DanceTrack validation set. The best results are shown in bold.

CME module on DanceTrack validation set. The baseline configuration (first row) directly feeds the tracklet $\tilde{\mathbf{T}}_{1:t}$ to ABG without specialized modules, achieving 59.2 HOTA. Adding only Motion Pattern Module (MPM) improves HOTA by 1.0 to 60.2, demonstrating temporal pattern analysis value. Combining MPM with PDM increases performance to 60.4 HOTA, while MPM with UQM achieves 60.3 HOTA. The full configuration with all three modules achieves the best performance of 60.8 HOTA, with AssA improving from 44.5 to 46.6. This demonstrates that multi-perceptive motion analysis creates synergistic interactions for enhanced tracking accuracy.

Alpha Range Analysis. Table 6 investigates the impact of alpha range bounds in MCAS. We evaluate four range settings: [0.1, 0.9], [0.2, 0.8], [0.3, 0.7], and [0.4, 0.6]. For PlugTrack(TrackSSM), the optimal range [0.3, 0.7] achieves 54.9 HOTA, while PlugTrack(DiffMOT) performs best with [0.3, 0.7] reaching 61.7 HOTA. Wider ranges like [0.1, 0.9] allow extreme blending values that may completely ignore one predictor, leading to suboptimal performance. Conversely, narrow ranges like [0.4, 0.6] restrict the adaptive capability. The consistent optimal range across different base predictors suggests that [0.3, 0.7] provides sufficient flexibility for adaptive blending while ensuring meaningful contributions from both predictors.

Qualitative Analysis

Finally, we perform a case analysis by comparing PlugTrack with individual motion predictors on a challenging DanceTrack sequence. Figure 4 demonstrates PlugTrack’s effectiveness in maintaining identity consistency through adaptive fusion. The visualization compares tracking results for the same object across frames 475-490, where the Kalman filter (top) successfully maintains ID, DiffMOT (middle) suffers from ID switching at frame 485, and PlugTrack (bottom) prevents this failure. Even on this non-linear motion dataset, certain scenarios favor the Kalman filter’s predictions over data-driven methods. At frame 485, our adaptive blending factors ($\alpha_x = 0.874$, $\alpha_y = 0.413$, $\alpha_w = 0.721$, $\alpha_h = 0.912$) show PlugTrack assigns higher weights to the Kalman filter for most coordinates, particularly x-dimension and box size, with only y-coordinate receiving lower weight (0.413) due to non-linear vertical movement. This coordinate-specific adaptation demonstrates that CME correctly identifies motion contexts where the Kalman filter excels, enabling ABG to prevent ID switches through intelligent fusion.

Model	Range	HOTA	IDF1	AssA
Ours (TrackSSM)	(0.1 0.9)	54.1	37.2	53.8
	(0.2 0.8)	54.7	38.1	54.8
	(0.3 0.7)	54.9	38.4	54.7
	(0.4 0.6)	54.8	38.2	54.7
Ours (DiffMOT)	(0.1 0.9)	59.8	45.2	60.7
	(0.2 0.8)	60.0	45.5	60.9
	(0.3 0.7)	60.7	46.5	61.7
	(0.4 0.6)	59.5	44.6	60.5

Table 6: Ablation study of different α ranges on DanceTrack validation set. The best results are shown in bold.



Figure 4: Qualitative comparison on DanceTrack dataset showing tracking results across frames 475-490. (a) Kalman filter maintains ID consistency. (b) DiffMOT suffers from ID switching at frame 485. (c) PlugTrack (Ours) successfully maintains ID consistency through adaptive blending ($\alpha_x=0.874$, $\alpha_y=0.413$) Boxes of the same color indicate the same tracked identity.

Conclusion

In this paper, we presented PlugTrack, a novel framework that adaptively fuses Kalman filter and data-driven motion predictors for robust multi-object tracking. Our key insight—that Kalman filters outperform data-driven predictors in up to 40% of cases even on non-linear motion datasets—reveals the inherent coexistence of diverse motion patterns in real-world scenarios. Through the Contextual Motion Encoder’s multi-perceptive analysis and the Adaptive Blending Generator’s coordinate-specific fusion, PlugTrack achieves consistent performance improvements across MOT17, MOT20, and DanceTrack while maintaining real-time capability. The framework’s plug-and-play design opens promising future directions: beyond integrating data-driven predictors, PlugTrack can potentially accommodate Kalman filter-based heuristic methods like OC-SORT and Hybrid-SORT, enabling adaptive fusion among multiple motion prediction paradigms. We believe this principle of intelligently combining classical and modern approaches provides a valuable paradigm for advancing multi-object tracking and other vision tasks where complementary methods exist.

Acknowledgments

This work was supported by the National Research Foundation of Korea (NRF) grant funded by the Korea government(MSIT)(RS-2024-00456589)

References

- 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. Ieee.
- 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.
- Cui, Y.; Zeng, C.; Zhao, X.; Yang, Y.; Wu, G.; and Wang, L. 2023. Sportsmot: A large multi-object tracking dataset in multiple sports scenes. In *Proceedings of the IEEE/CVF international conference on computer vision*, 9921–9931.
- 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*.
- 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*, 25: 8725–8737.
- Gao, R.; Qi, J.; and Wang, L. 2025. Multiple object tracking as id prediction. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 27883–27893.
- Ge, Z.; Liu, S.; Wang, F.; Li, Z.; and Sun, J. 2021. Yolox: Exceeding yolo series in 2021. *arXiv preprint arXiv:2107.08430*.
- Graves, A. 2012. Long short-term memory. *Supervised sequence labelling with recurrent neural networks*, 37–45.
- Hampali, S.; Stekovic, S.; Sarkar, S. D.; Kumar, C. S.; Fraundorfer, F.; and Lepetit, V. 2021. Monte carlo scene search for 3d scene understanding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 13804–13813.
- 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*.
- 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(2): 548–578.
- Lv, W.; Huang, Y.; Zhang, N.; Lin, R.-S.; Han, M.; and Zeng, D. 2024. Diffmot: A real-time diffusion-based multiple object tracker with non-linear prediction. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 19321–19330.
- 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*.
- 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*, 20993–21002.
- 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. IEEE.
- 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, M.; Han, G.; Yan, B.; Zhang, W.; Qi, J.; Lu, H.; and Wang, D. 2024. Hybrid-sort: Weak cues matter for online multi-object tracking. In *Proceedings of the AAAI conference on artificial intelligence*, volume 38, 6504–6512.
- Yu, E.; Li, Z.; Han, S.; and Wang, H. 2022. Relationtrack: Relation-aware multiple object tracking with decoupled representation. *IEEE Transactions on Multimedia*, 25: 2686–2697.
- 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.