

HybridTrack: A Hybrid Approach for Robust Multi-Object Tracking

Leandro Di Bella^{1,2} , Yangxintong Lyu^{1,2} , Bruno Cornelis¹ , and Adrian Munteanu^{1,2}  *Member, IEEE*

Abstract—The evolution of Advanced Driver Assistance Systems (ADAS) has increased the need for robust and generalizable algorithms for multi-object tracking. Traditional statistical model-based tracking methods rely on predefined motion models and assumptions about system noise distributions. Although computationally efficient, they often lack adaptability to varying traffic scenarios and require extensive manual design and parameter tuning. To address these issues, we propose a novel 3D multi-object tracking approach for vehicles, HybridTrack, which integrates a data-driven Kalman Filter (KF) within a tracking-by-detection paradigm. In particular, it learns the transition residual and Kalman gain directly from data, which eliminates the need for manual motion and stochastic parameter modeling. Validated on the real-world KITTI dataset, HybridTrack achieves 82.72% HOTA accuracy, significantly outperforming state-of-the-art methods. We also evaluate our method under different configurations, achieving the fastest processing speed of 112 FPS. Consequently, HybridTrack eliminates the dependency on scene-specific designs while improving performance and maintaining real-time efficiency.¹

Index Terms—Multi-Object Tracking, Kalman Filtering, Autonomous Driving

I. INTRODUCTION

MULTI-OBJECT tracking (MOT) involves locating and associating objects across consecutive image and/or LiDAR frames. It allows us to predict object trajectories, handle occlusions, and manage object appearance, making it widely used in Advanced Driver Assistance Systems (ADAS) to enhance traffic safety. The existing MOT techniques can be classified into various paradigms, including **Joint Detection and Tracking** (JDT) [1], [2], [3], [4] and **Tracking-by-Detection** (TBD) [5]. JDT combines detection and tracking in a unified end-to-end deep-learning framework, leveraging temporal and spatial information for robust performance. However, the combination of detection and tracking reduces the modularity of the entire framework, increasing computational complexity. This complexity can limit applicability of the tracking algorithms in real-time traffic situations, particularly

on resource-constrained systems such as smart cameras. Furthermore, the requirements for large-scale tracking data and the extensive training complexity of an end-to-end pipeline, make the use of JDT for real-world MOT potentially inefficient.

In contrast to JDT, TBD separates detection from tracking, enabling greater flexibility and computational efficiency. This approach allows state-of-the-art object detectors [6]–[8] to be easily integrated with tracking algorithms. A notable strength of TBD is its use of recursive frameworks for the tracking component, such as Kalman [9] or Bayesian filtering [10], which continuously update the state of an object based on prior estimates and new observations. This recursive process enhances robustness and adaptability in dynamic environments. However, the reliance of state-of-the-art methods on these traditional model-based approaches presents a significant bottleneck. Specifically, the traditional methods rely on predefined motion models specific to traffic situations e.g., constant turn rate for cars, kinetic bicycle models for bicycles, and constant acceleration for pedestrians [11], [12], [13], [14]. Such specificities limit adaptability to different traffic agents, degrading performance.

Moreover, estimating noise parameters such as process and measurement covariances is challenging with complex sensor data, i.e. LiDAR or image data, requiring extensive manual tuning for specific scenarios. For instance, covariances optimized for highways fail in dense urban settings with frequent occlusions and unpredictable dynamics in traffic. While adaptive techniques, such as adaptive Kalman filtering [15], can adjust parameters dynamically, intensive resources and handcrafted rules are necessary. Therefore, robust, adaptable algorithms that generalize across conditions, while supporting real-time processing, are still an open problem. [16] and [17] replace traditional trackers with fully deep-learning-based alternatives. However, it increases computational complexity, potentially canceling the primary advantage of the model-based TBD framework, namely, real-time processing. As a result, such solutions risk becoming less practical than end-to-end joint detection and tracking approaches.

To address these limitations, we propose HybridTrack, a novel 3D multi-object tracking framework that combines the stability and interpretability of Kalman filters with the flexibility and generalization of deep learning. More specifically, HybridTrack uses a learnable Kalman filter to dynamically adjust stochastic parameters in real-time, eliminating manual tuning. With learnable components integrated in the motion model, Kalman gain, and noise covariances, the system can adapt autonomously to various scenarios without agent-

Manuscript received: December 30, 2024; Revised April 7, 2025; Accepted May 10, 2025.

This paper was recommended for publication by Editor Hyungpil Moon upon evaluation of the Associate Editor and Reviewers' comments.

This work was supported by Innoviris within the research project TORRES.

¹Leandro Di Bella, Yangxintong Lyu, Bruno Cornelis, and Adrian Munteanu are with the Department of Electronics and Informatics, Vrije Universiteit Brussel, Pleinlaan 2, B-1050 Brussels, Belgium.

²Leandro Di Bella, Yangxintong Lyu, and Adrian Munteanu are also with IMEC, Kapeldreef 75, B-3001 Leuven, Belgium.

Digital Object Identifier (DOI): see top of this page.

¹The code is publicly available at: <https://github.com/leandro-svg/HybridTrack>.

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

specific designs and prior knowledge of the scene. To do so, we design the deep-learning-based component with a limited number of parameters, leading to a lightweight framework. Additionally, the inherent recursive property of our learnable Kalman filter ensures data efficiency, while its learnable nature makes it data scalable. Furthermore, we introduce a dynamic scaling mechanism to enhance the stability of our model during noisy initialization. To summarize, our contributions are threefold:

- We propose a new 3D MOT framework, HybridTrack, based on a learnable Kalman filter by integrating novel deep learning modules. It allows us to combine the recursive property of Kalman filtering and the predictive capability of deep learning.
- A novel predictive motion model block and a Kalman Gain block are introduced. By doing so, HybridTrack eliminates the need for manual motion and stochastic parameter design.
- The experimental results demonstrate that the proposed method significantly outperforms the existing model-based tracking techniques on the KITTI dataset with 82.72% HOTA. Furthermore, HybridTrack is able to achieve the fastest processing speed of 112 FPS, which unlocks the potential of real-time traffic applications.

II. RELATED WORK

Tracking-by-Detection (TBD). TBD remains the most widely adopted paradigm in multi-object tracking (MOT) due to its modularity and computational efficiency. In TBD, an object detector [6]–[8] generates detections, which are then processed by a separate tracking module for association and trajectory prediction. The decoupling of detection and tracking allows for flexibility, making TBD a practical choice for a wide range of applications. Traditional TBD methods often rely on statistical techniques such as Kalman filters [9] or, more generally, Bayesian filters [10] for object association and trajectory prediction. These lightweight and interpretable approaches are ideal for real-time applications. However, their reliance on pre-defined motion models, such as constant velocity or constant acceleration assumptions [12], [13], [14], limits their ability to adapt to complex, dynamic scenarios. Extensions to non-linear motion models, such as Extended Kalman filters [18], have not demonstrated significant improvement in the aforementioned challenging conditions. More recent advancements, including Adaptive Kalman filters [15] and probabilistic Bernoulli-based approaches [19] attempt to address these limitations. However intricate parameter tuning and significant domain knowledge are required, reducing these methods’ scalability. Similarly, other works, including UCMCTrack [20] and PMTrack [21], aim to address challenges such as variable frame rates and complex sensor movements by incorporating advanced motion estimation techniques into motion particle filters. To further enhance robustness without abandoning the modular TBD paradigm, EMMS-MOT [22] and BiTrack [23] introduce a multi-modal fusion strategy that corrects 3D detections using 2D cues and jointly models motion across modalities. This approach improves resilience to occlusion and sensor noise,

while retaining the flexibility of classic TBD designs. While real-time performance has always been a critical challenge for trackers generally, Fast-Poly [24] introduces a polyhedral tracking framework that leverages rotated object alignment and dense local computation to significantly boost efficiency. It achieves promising speed-accuracy trade-offs on large-scale datasets with CPU-friendly design.

Data Association. As a critical component of TBD, data association links detections across frames to maintain vehicle identities. Traditional methods like AB3DMOT [11] rely on spatial metrics, Intersection over Union (IoU) and Euclidean distance, respectively. While simple and efficient, these metrics struggle in complex scenarios, leading to identity switches and trajectory fragmentation. To address these issues, advanced approaches incorporate Confidence-guided association [12], [25] and appearance-based features [14]. Recent works also integrate multi-stage association strategies, improving robustness in conditions such as occlusion and low-confidence detections [26]. However, the dependency on empirically defined metrics limits the adaptability to varying environments. Deep learning has further revolutionized TBD by enabling data-driven feature extraction and association. DeepSORT [17] uses neural networks for robust appearance modeling. TripleTrack [27] and AppTracker++ [28] fuse detection and tracking cues at the feature level to enhance robustness. DINO-MOT [29] enhances pedestrian tracking in the TBD framework by integrating a visual foundation model (DINOv2) for pedestrian re-identification. It introduces a visual memory mechanism to reduce ID switches and leverages appearance features from the camera modality. Another promising direction involves the use of graph theory for association. PolarMOT [30] employs spatial and temporal graph-based strategies leveraging 3D geometric cues, while [31] applies graph theory as a post-processing step for trajectory rectification. However, deep learning-based approaches often face challenges related to high data dependency and limited interpretability, making them less generally practical. One can see that most methods leveraging deep learning for multi-object tracking utilize it in either an end-to-end pipeline, where both detection and tracking are integrated, as seen in JDT [1], [2], [3], [4] or within a learnable data association step [17], [27], [28]. In contrast, HybridTrack bridges the gap between these paradigms by embedding the predictive capabilities of deep learning into a traditional Kalman filter framework followed by a one-step greedy-algorithm-based data association. This hybrid approach enables a lightweight and interpretable network that maintains robustness and data efficiency. By combining the adaptability of deep learning with a Kalman filter structure, HybridTrack effectively addresses the limitations of purely traditional or fully deep learning-based methods, providing a scalable and resource-efficient solution for MOT tasks.

III. HYBRIDTRACK

A. Problem Formulation and Architecture Overview

Given a sequence of video/LiDAR point cloud frames from a real-world traffic scene, our goal is to accurately and efficiently identify and track multiple vehicles in 3D space. To

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

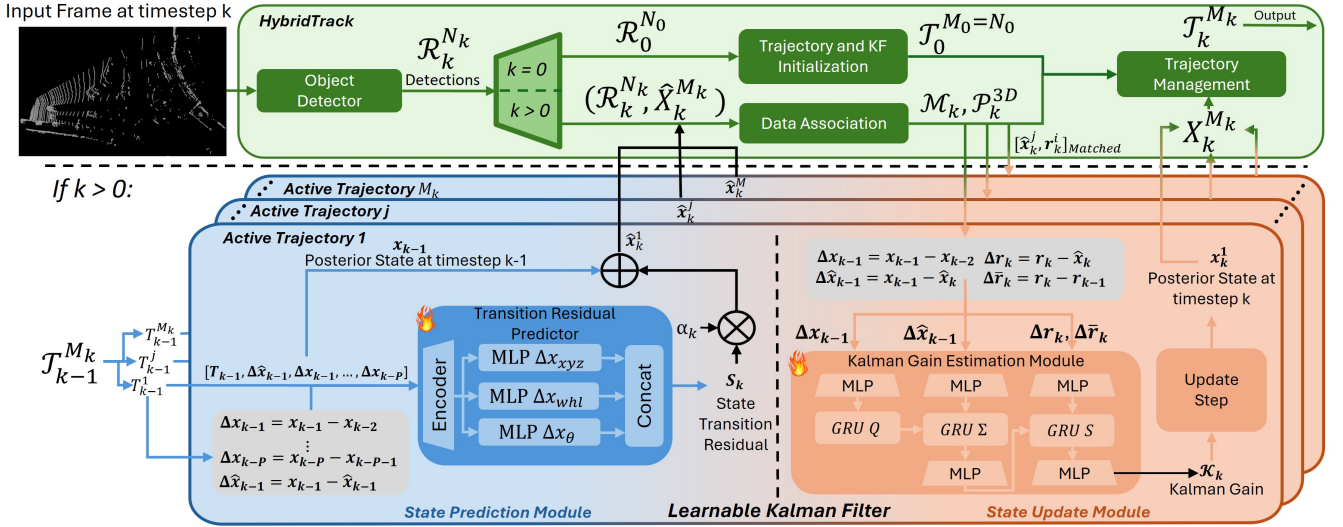


Fig. 1: The proposed HybridTrack. During inference, HybridTrack begins with 3D object detection, localizing vehicles using sensor data. The resulting detections $\mathcal{R}_k^{N_k}$ are sent to the State Prediction Module (SPM), where new detections \mathbf{r}_k^i initialize trajectories $\mathcal{T}_k^{M_k}$. For each trajectory T_k^j , the TRP predicts the next state $\hat{\mathbf{x}}_k^j = \alpha_k S_k^j + \mathbf{x}_{k-1}^j$ forming the set of predicted states $\hat{X}_k^{M_k}$. In the Data Association module, $\hat{X}_k^{M_k}$ is matched with current $\mathcal{R}_k^{N_k}$. The State Update Module (SUM) refines the matched pairs, updating the vehicle states $X_k^{M_k}$. Finally, the Trajectory Management module maintains and updates trajectories.

do so, our tracking-by-detection framework consists of five key modules, shown in Fig. 1: 1) Object Detection: A 3D detection model extracts 3D bounding boxes from sensor data (e.g., video or LiDAR) at each timestep, identifying and localizing vehicles in the 3D world. 2) Trajectory Initialization and State Prediction: Newly detected vehicles initialize a trajectory, and the State Prediction Module (SPM) predicts the next state of each tracked trajectory. 3) Data Association: Detected vehicles from (1) are matched with predicted states from (2). 4) State Update: Vehicle trajectories are refined and updated based on the associations. 5) Trajectory Management: Handles and maintains trajectories for consistent tracking. In the following sections, we elaborate on each component.

B. 3D Object Detection

The fundamental concept of the tracking-by-detection approach is to track objects i.e., vehicles in our case, using detections at each timestep k , extracted by a 3D object detector. Formally, the detector \mathcal{F} outputs a set $\mathcal{R}_k^{N_k} = \mathcal{F}(I_k)$, where I_k is the input data frame at timestep k , and $\mathcal{R}_k^{N_k} = \{\mathbf{r}_k^i\}_{i=0}^{N_k} \in \mathbb{R}^{N_k \times F}$ represents the detected vehicles. Here, N_k is the number of detections at timestep k , F is the number of features, and each $\mathbf{r}_k^i \in \mathbb{R}^F$ is the i -th vehicle. In other words, each measurement \mathbf{r}_k^i corresponds to a 3D bounding box. In this paper, measurements are defined as $\mathbf{r}_k^i = [x, y, z, w, l, h, \theta] \in \mathbb{R}^7$, where x, y, z are the 3D centroid coordinates, w, l, h are width, length, and height, and θ is the heading angle.

C. Trajectory Initialization

In HybridTrack, trajectories are classified as active or dead. Active trajectories represent currently tracked vehicles, while dead ones correspond to vehicles undetected for μ_{\max} consecutive timesteps. At $k=0$, each detection $\mathbf{r}_0^i \in \mathcal{R}_0^{N_0}$ initializes an

active trajectory T_0^j , forming a one-to-one mapping. The set of active trajectories is defined as $\mathcal{T}_k^{M_k} = \{T_k^j\}_{j=0}^{M_k}$, where M_k is the number of tracked vehicles at timestep k . Initially, $D_0 = \emptyset$, and the set of dead trajectories D_k is updated when new ones are created. At $k > 0$, each T_k^j includes the current detection \mathbf{r}_k^i , predicted state $\hat{\mathbf{x}}_k^j$, and past posterior states $\mathbf{x}_{k-1}^j, \dots, \mathbf{x}_{k_0}^j$, forming $T_k^j = [\mathbf{r}_k^i, \hat{\mathbf{x}}_k^j, \mathbf{x}_{k-1}^j, \mathbf{x}_{k-2}^j, \dots, \mathbf{x}_{k_0}^j] \in \mathbb{R}^{L \times F}$, where $L = k - k_0 + 2$ and k_0 is the trajectory's start time. At $k=0$, indices i and j are identical. When initializing a trajectory during inference, prior states $\mathbf{x}_{k_0-1}^j, \dots, \mathbf{x}_0^j$ are set to $\mathbf{x}_{k_0}^j - \epsilon$ to avoid duplicates. A learnable Kalman filter (LKF) is also initialized per trajectory, with its prior and posterior estimates.

D. State Prediction Module

At each timestep $k > 0$, the State Prediction Modules (SPMs), integrated in each LKF, predict the next state of the associated tracked trajectories. Rather than computing the prior state directly, we model the Kalman filter's prediction step as a deep-learning problem.

1) *Mathematical Formulation*: Let $\mathbf{x}_k^j \in \mathbb{R}^F$ denotes the posterior state vector of a j -th tracked vehicle at timestep k , which includes the 3D location, 3D dimensions, and orientation of the vehicle. We aim to learn a function $f_\theta(\mathbf{x}_{k-1}^j) : \mathbb{R}^F \rightarrow \mathbb{R}^F$, parameterized by θ , that maps the previous posterior state \mathbf{x}_{k-1}^j to the predicted prior state $\hat{\mathbf{x}}_k^j$. To mitigate the error probability and introduce more stability to the system, instead of directly predicting the prior state $\hat{\mathbf{x}}_k^j$, we aim to predict the state transition residual S_k^j , which captures the state deviation. This enables the model to focus on learning motion dynamics rather than full state estimation. The network is trained to minimize the prediction error relative to the true posterior state at timestep k , formulating the problem as supervised learning with a loss function.

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

2) *Transition Residual Predictor (TRP)*: We parametrize the function $f_\theta(\mathbf{x}_{k-1}^j)$ using an encoder-decoder architecture. It processes the historical trajectory information of a tracked vehicle to predict its transition residual, as illustrated in Fig. 1. The SPM is lightweight and parameter-efficient, as we are leveraging the recursive architecture of the deep learnable Kalman filter to effectively correct small errors and maintain accurate predictions with minimal computational overhead. Each SPM is applied independently to each active trajectory, enabling the TRP module to adaptively generate transition residuals tailored to the specific motion patterns of individual vehicles. To effectively capture the temporal dynamics of vehicle motion, the TRP leverages input that includes the current trajectory T_{k-1}^j and, for any integer $0 \leq P \leq k - k_0 - 1$, the sequence of posterior state differences $[\Delta \mathbf{x}_{k-1}^j, \dots, \Delta \mathbf{x}_{k-P}^j] \in \mathbb{R}^{P \times F}$ each defined as $\Delta \mathbf{x}_{k-i}^j = \mathbf{x}_{k-i}^j - \mathbf{x}_{k-i-1}^j$, for $1 \leq i \leq P$. Additionally, the Kalman gain update term $\Delta \hat{\mathbf{x}}_{k-1}^j$ is included. This term incorporates information on the correction performed by the Kalman filter during the previous timestep. The TRP encodes its input variables, $[T_{k-1}^j, \Delta \hat{\mathbf{x}}_{k-1}^j, \Delta \mathbf{x}_{k-1}^j, \dots, \Delta \mathbf{x}_{k-P}^j]$, using an encoder to extract features. Subsequently, three parallel Multi-Layer Perceptrons (MLPs) compute the residuals for specific vehicle states: (1) Spatial information $\Delta \mathbf{x}_{xyz}^j = [\Delta x, \Delta y, \Delta z]$ which represents the residuals in 3D coordinates of the vehicle, (2) Geometrical properties $\Delta \mathbf{x}_{whl}^j = [\Delta w, \Delta l, \Delta h]$ representing the residuals in dimension of the 3D bounding box, and (3) the residual in vehicle orientation $\Delta \theta$. The final state transition residual S_k^j is computed as: $S_k^j = [\Delta x \ \Delta y \ \Delta z \ \Delta w \ \Delta l \ \Delta h \ \Delta \theta]$. This setup helps capture individual vehicle behavior and maintain stable residual predictions.

3) *Dynamic Scaling Factor*: One limitation of the LKF is its need for several timesteps to calibrate noise parameters, which complicates prediction stabilization when initial detections are noisy due to sensor errors, occlusions, or environmental factors. This may cause prediction overshooting, with updates that deviate excessively from the true trajectory, leading to missed associations. To mitigate this, we introduce an adaptive scaling factor α that dampens early prediction updates. By reducing the impact of abrupt changes, α limits unstable corrections and allows the Kalman filter to gradually calibrate noise parameters and improve robustness and accuracy. It is defined by the time elapsed since initialization τ_0 :

$$\alpha_k = \frac{1}{T_{\max}} \times (k - \tau_0), \text{ if } (k - \tau_0) < T_{\max}, \quad (1)$$

otherwise, $\alpha_k = 1$. This linear schedule smooths initial predictions and mitigates overshooting. As α approaches 1, the Kalman filter transitions to its standard behavior, relying on calibrated noise parameters for accurate trajectory estimation.

On the other hand, missed detections are an additional challenge for the SPM, particularly in dense traffic scenarios. More specifically, it forces the LKF to rely solely on the process model during prediction, leading to compounding errors and potential divergence. Prior work [12], [25] addresses this by scaling the cost function with confidence scores, expanding the search radius during data association. However, this increases the risk of false positives and ID switches in crowded scenes.

Moreover, such scaling is limited to linear cost functions like Euclidean distance. Instead of modifying the cost function, we propose mitigating the impact of missed detections by directly adjusting the predicted state via a refined scaling factor α . The factor is dynamically adjusted as follows:

$$\alpha_k = \alpha_k \times \left(1 - \frac{\kappa}{\mu_{\max}}\right) \cdot (a_{\max} - a_{\min}), \text{ if } \kappa < \mu_{\max}, \quad (2)$$

Otherwise, $\alpha_k = a_{\min}$. Here, κ is the number of consecutive missed detections, μ_{\max} is the threshold before a trajectory is marked dead, and a_{\max}, a_{\min} are empirically chosen bounds. By reducing α dynamically, our method stabilizes tracking in both early stages and in cases of missed detections, minimizing the impact of noisy predictions while maintaining robust data association.

4) *Prior State Prediction*: Using the predicted transition residual S_k^j and the scaling factor α_k , the Kalman filter prior state can be defined as follows:

$$\begin{aligned} \hat{\mathbf{x}}_k^j &= \alpha_k f_\theta \left(T_{k-1}^j, \Delta \mathbf{x}_{k-1}^j, \Delta \mathbf{x}_{k-2}^j, \dots, \Delta \mathbf{x}_{k-P}^j, \Delta \hat{\mathbf{x}}_{k-1}^j \right) + \mathbf{x}_{k-1}^j, \\ &= \alpha_k S_k^j + \mathbf{x}_{k-1}^j. \end{aligned} \quad (3)$$

with $f_\theta(\cdot) = TRP(\cdot)$. Therefore, for each active trajectory T_{k-1}^j , a predicted prior state is inferred. We define the set of predicted prior states as follows: $\hat{X}_k^{M_k} = \{\hat{\mathbf{x}}_k^j\}_{j=0}^{M_k} \in \mathbb{R}^{M_k \times F}$.

E. Data Association

Data association links current detections $\mathcal{R}_k^{N_k}$ to active tracked trajectories $\mathcal{T}_k^{M_k}$, ensuring consistent vehicle trajectories over time. For this step, we propose a one-stage matching process. First, a cost map is computed, defined as $C_{3D}(\hat{X}_k^{M_k}, \mathcal{R}_k^{N_k}) = 1 - CIoU(\hat{X}_k^{M_k}, \mathcal{R}_k^{N_k})$ where $CIoU$ is the Complete Intersection over Union (CIoU) [32]. The cost map C_{3D} is a $M_k \times N_k$ matrix, with each element $C_{3D}[j, i]$ corresponding to the cost of associating the j -th predicted vehicle with the i -th detected vehicle. By using C_{3D} , we apply a simple yet effective greedy algorithm with a fixed threshold τ_{3D} to associate the newly detected vehicles $\mathbf{r}_k^i \in \mathcal{R}_k^{N_k}$ with the active trajectories T_{k-1}^j and their corresponding prediction states $\hat{\mathbf{x}}_k^j$. The greedy algorithm assigns detections to trajectories by iteratively selecting the lowest-cost $C_{3D}[j, i]$ for the pair $(\hat{\mathbf{x}}_k^j, \mathbf{r}_k^i)$ at each step, τ_{3D} serves as a criterion for deciding whether detection and a trajectory prediction are close enough to be associated. This association step returns the matched indices (j, i) and the set of matched pairs $\mathcal{P}_k^{3D} \in \mathbb{R}^{J_k \times 2 \times F}$ where J_k is the number of matched pairs at timestep k with $J_k \leq M_k$ and one element is represented as $(\hat{\mathbf{x}}_k^j, \mathbf{r}_k^i)$.

F. State Update Module

For every matched pair $(\hat{\mathbf{x}}_k^j, \mathbf{r}_k^i)$ in \mathcal{P}_k^{3D} , we update the previous posterior state \mathbf{x}_{k-1}^j of the corresponding trajectory T_{k-1}^j in the State Update Module (SUM). Our updated step is inspired by KalmanNet [33]. That is, instead of relying on a traditional Kalman filter, *HybridTrack* replaces the Kalman gain $\mathcal{K}_k^j \in \mathbb{R}^{F \times F}$ and the covariances \mathbf{v}_k and \mathbf{w}_k with a Recurrent Neural Network (RNN) [34] based architecture.

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

This architecture is described in Fig. 1 dubbed Kalman Gain Estimation Module. Once the Kalman gain \mathcal{K}_k^j is inferred, we can compute the posterior state as follows:

$$\text{Updated State Estimate} : \mathbf{x}_k^j = \hat{\mathbf{x}}_k^j + \mathcal{K}_k^j (\mathbf{r}_k^i - H\hat{\mathbf{x}}_k^j), \quad (4)$$

where $\hat{\mathbf{x}}_k^j$ is the predicted state estimate, \mathbf{r}_k^i denotes the matched detection/measurement with the associated indexes (j, i) . $H \in \mathbb{R}^{F \times F}$ is the observation model mapping state $\hat{\mathbf{x}}_k^j$ to measurement space with $H : \mathbb{R}^F \rightarrow \mathbb{R}^F$, $\mathbf{r}_k^i \approx H\hat{\mathbf{x}}_k^j$. Here, H is defined as an identity matrix. On the other hand, the Kalman gain \mathcal{K}_k^j effectively maps $(\mathbf{r}_k^i - H\hat{\mathbf{x}}_k^j)$ from the measurement space \mathbb{R}^F back into the state space \mathbb{R}^F . For all updated posterior states \mathbf{x}_k^j of the active trajectories $\mathcal{T}_{k-1}^{M_k}$, we aggregate them into a set of updated detections, $X_k^{M_k} = \{\mathbf{x}_k^j\}_{j=0}^{M_k} \in \mathbb{R}^{M_k \times F}$, which is then passed to the trajectory management module. The main advantage of this approach is that it eliminates the need to manually model state and measurement noise parameters \mathbf{v}_k and \mathbf{w}_k , which are often nonadaptive for diverse and dynamic scenes. Instead, our method employs dynamically changing covariances rather than static, Gaussian-distributed noise parameters.

G. Trajectory Management

Active trajectories are updated with the matched updated posterior states $X_k^{M_k}$. For unmatched detections $\mathcal{R}_k^{N_k}$, new trajectories are initialized to account for newly appearing vehicles and added to the set of active trajectories. If a predicted state $\hat{\mathbf{x}}_k^j \in \hat{X}_k^{M_k}$ of an active trajectory $T_{k-1}^j \in \mathcal{T}_{k-1}^{M_k}$ fails to be associated with a detection $\mathbf{r}_k^i \in \mathcal{R}_k^{N_k}$, the trajectory T_{k-1}^j is marked as non-updated. Moreover, the trajectory is classified as dead and added to D_k after μ_{max} consecutive timesteps without updates. Additionally, if a new trajectory is wrongly created due to a detection artifact, it will be designated dead if it remains non-updated for σ timesteps.

H. Optimization

The proposed learnable Kalman filter is trained independently from the tracking pipeline using loss functions that enforce spatial and temporal consistency. That is, the SPM and SUM are trained together in an end-to-end manner. Specifically, HybridTrack is trained by using fixed-length (20 frames) ground-truth tracklets. During training phase, we use ground-truth pose sequences. No data association or occlusion handling is required. The prior and posterior state predictions are optimized by L1 loss. In contrast, no significant impact on performance is observed when augmenting model sequences input with noise. Therefore, we mainly focus on learning accurate motion dynamics by using the original clean data, while the Kalman filter handles uncertainty during inference.

The primary loss, Mean Absolute Error (MAE), minimizes errors in both prior and posterior predicted states, denoted as \mathcal{L}_{states} . To ensure temporal smoothness, a temporal consistency loss $\mathcal{L}_{temporal}$ penalizes large deviations between consecutive states. A direction consistency loss $\mathcal{L}_{direction}$ further stabilizes movement by normalizing speed vectors to avoid abrupt directional changes. Let $\bar{\mathbf{x}}_{i,k}$, $\hat{\mathbf{x}}_{i,k}$ and $\mathbf{x}_{i,k}$ denote the ground

truth, predicted prior and posterior sequences for the single i -th vehicle, respectively, with k indexing the temporal dimension and i the data sample index. The losses are formulated as follows:

$$\mathcal{L}_{states} = \frac{1}{N_D} \sum_{i=1}^{N_D} \frac{1}{T} \sum_{k=1}^T (\|\bar{\mathbf{x}}_{i,k} - \mathbf{x}_{i,k}\|_1 + \|\bar{\mathbf{x}}_{i,k} - \hat{\mathbf{x}}_{i,k}\|_1) \quad (5)$$

$$\mathcal{L}_{temporal} = \frac{1}{N_D} \sum_{i=1}^{N_D} \frac{1}{T-1} \sum_{k=1}^{T-1} \|\mathbf{p}_{i,k+1} - \mathbf{p}_{i,k}\|_1 \quad (6)$$

$$\mathcal{L}_{direction} = \frac{1}{N_D} \sum_{i=1}^{N_D} \frac{1}{T-1} \sum_{k=1}^{T-1} \|\mathbf{d}_{i,k+1} - \mathbf{d}_{i,k}\|_2 \quad (7)$$

where N_D is the dataset size, $\mathbf{p}_{i,k} = [x, y, z]$ is the centroid of the vehicle, $\mathbf{d}_{i,k} = \frac{\mathbf{v}_k}{\|\mathbf{v}_k\| + \varepsilon}$ is the normalized direction vector at time k , and \mathbf{v}_k is the speed vector. $\varepsilon > 0$ is a small constant added to the denominator to avoid division by zero.

Joint loss optimization. The total loss is a weighted sum of these individual losses:

$$\mathcal{L}_{total} = \lambda_{states} \mathcal{L}_{states} + \lambda_{temp} \mathcal{L}_{temporal} + \lambda_{dir} \mathcal{L}_{direction} \quad (8)$$

Where we set $\lambda_\theta = 1$, for all $\theta \in \{states, temp, dir\}$. This loss function ensures the model maintains consistency in the predicted sequences' spatial and temporal dimensions.

IV. EXPERIMENTS

Dataset and Implementations. The KITTI dataset [37] is a widely used dataset for monocular object pose estimation and tracking. The dataset comprises 50 videos, divided into 21 for training and 29 for testing, following the recommended splitting on the official website. Following the standard practice in the KITTI benchmark, the training dataset is further subdivided into 10 sequences for training and 11 sequences for validation. To train *HybridTrack*, we extract the pose predictions from an existing 3D bounding box estimator as \mathbf{r}_k . We use VirConv [7] as a 3D object detector to predict 3D bounding boxes. The network is optimized by Adam Optimizer [38] with a learning rate of 0.001 and a weight decay of 0.00001. We take a batch size of 128 on one Nvidia Titan X (12G). The iteration number for the training process is set to 1,500. During inference, $T_{max} = 8$, $a_{min} = 0.1$, $a_{max} = 0.9$, $\mu_{max} = 22$, and $\sigma = 5$. The IoU threshold for data association is set to 1.20.

Evaluation Metrics. We primarily focus on the Higher Order Tracking Accuracy (HOTA) metric [39], providing a balanced measure of detection, association, and localization accuracy in multi-object tracking. It is defined as $HOTA = \sqrt{\text{DetA} \cdot \text{AssA}}$. DetA is the Detection Accuracy [39] emphasizing the accuracy of object detection, while AssA is the Association Accuracy [39] explicitly measuring the association's effectiveness in maintaining track consistency. We also furnish the Localization Accuracy (LoCA) which describes how accurately the objects' spatial positions are estimated. Additionally, we include the Classification of Events, Activities, and Relationships Multi-Object Tracking (CLEARMOT) [40] metrics, such as MOTA [40] for Multi-Object Tracking Accuracy and IDF1 [41] for identity consistency, to complement the HOTA

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.TABLE I: Comparison of existing methods on the KITTI test set. The best is marked in **bold**, and the second-best in underline. The results are reported in %.

Method	Modality	HOTA \uparrow	Detection			Association			LoCA \uparrow
			DetA \uparrow	DetRe \uparrow	DetPr \uparrow	AssA \uparrow	AssRe \uparrow	AssPr \uparrow	
PermaTrack [1] (ICCV'21)	2D	78.03	78.29	81.71	86.54	78.41	81.14	89.49	87.10
PC-TCNN [3] (IJCAI'21)	3D	80.90	78.4	84.22	84.58	84.13	87.46	90.47	87.48
TripletTrack [27] (CVPR'22)	2D	73.58	73.18	76.18	86.81	74.66	77.31	89.55	87.37
RAM [2] (ICML'22)	3D	79.53	78.79	82.54	86.33	80.94	84.21	88.77	87.15
DF-MOT [26] (RAL'22)	2D+3D	75.46	71.59	75.34	85.25	80.05	82.63	89.77	86.07
RethinkingMOT [31] (ICRA'23)	3D	80.39	77.88	<u>84.23</u>	83.57	83.64	87.6	88.90	87.07
OC-SORT [13](CVPR'23)	2D	76.54	77.25	80.64	86.34	76.39	80.33	87.17	87.01
FNC2 [15] (TIV'23)	2D+3D	73.19	73.27	80.98	81.67	73.77	77.05	89.84	87.31
AppTracker [28] (IJCV'24)	2D	75.19	75.55	78.77	86.04	75.36	78.34	88.24	86.59
PNAS-MOT [35] (RAL'24)	2D+3D	67.32	77.69	81.58	85.81	58.99	64.70	80.74	86.94
UCMCTrack [20] (AAAI'24)	2D	77.10	-	-	-	77.20	-	-	-
UG3DMOT [25] (SP'24)	3D	78.60	76.01	80.77	85.44	82.28	85.36	91.37	87.84
MMF-JDT [4] (RAL'24)	2D+3D	79.52	75.83	82.31	83.69	84.01	87.16	90.70	87.65
PMTrack [21] (ACCV'24)	3D	81.36	78.90	82.98	86.76	84.49	87.73	90.18	88.02
MCTrack [36] (ArXiv'24)	3D	82.75	<u>79.40</u>	82.92	87.45	<u>86.85</u>	<u>89.84</u>	<u>91.32</u>	<u>88.07</u>
BiTrack [23] (ArXiv'24)	2D+3D	82.70	80.04	84.53	87.19	86.17	89.11	92.16	88.68
HybridTrack (Ours)	3D	<u>82.72</u>	79.28	82.84	<u>87.39</u>	86.92	89.97	91.24	88.04

evaluation. To further evaluate our tracker, we also analyze MT (Mostly Tracked), PT (Partially Tracked), and ML (Mostly Lost) rates, which highlight the tracker's performance over time. Specifically, MT is the percentage of vehicle tracks that are successfully tracked for more than 80% of their trajectory length, whereas ML is less than 20%.

TABLE II: Further Comparison based on CLEAR metric. ¹ Uses CasA [8] detector. ² Uses VirConv [7] detector. . The results are reported in %.

Methods	MOTA \uparrow	MT \uparrow	PT \downarrow	ML \downarrow	IDFN \downarrow
UGD3MOT ¹ [25]	87.98	79.08	15.54	5.38	1111
HybridTrack ¹	90.16	84.61	9.23	6.15	1350
HybridTrack ²	91.05	85.39	6.31	8.31	534

Comparison with state-of-the-art methods. The performance of our proposed method against the existing methods is shown in Tab. I and II. Our approach achieves near SOTA performance with a HOTA score of 82.72% and a detector accuracy of 79.28%, outperforming the majority of methods that utilize camera-LiDAR data fusion, while our method relies solely on LiDAR data. BiTrack [23], while slightly outperforming in certain metrics, benefits from an intensive post-processing pipeline, including single- and multi-trajectory re-optimization, merging, and smoothing, which significantly contributes to its final scores. Moreover, our one-step association approach achieves state-of-the-art association accuracy of 86.92%, surpassing two-step methods such as MCTrack [36], which achieves 86.85%. It is worth mentioning that in MCTrack, motion models and noise covariances for each individual state variable (e.g., position, size, yaw) and agent types are manually fine-tuned. One can note that in both detection and association, our method prioritizes precision over recall. This approach minimizes false detection and association rates at the cost of introducing missed detections or associations. Additionally, although LocA heavily relies on the performance of the pre-trained 3D object detector, HybridTrack demonstrates strong spatial precision when ve-

hicles are detected or when detections are generated using our LKF. HybridTrack achieves superior MOTA (91.05% with Virconv and 90.16% with CasA) and Mostly Tracked (MT) rate (85.39% and 84.61%) compared to UGD3MOT (87.98%, 79.08%), as shown in Tab. II. This highlights our tracker's ability to maintain object continuity over trajectories, even with occasional misdetections or identity swaps. High MT and low PT rates emphasize long-term continuity, crucial for real-world applications like ADAS. In Fig. 2, we present comparative visualizations demonstrating our method's effectiveness in handling long-term occlusions and distant vehicles, two significant challenges in autonomous driving scenarios. The results show several key improvements over UGD3MOT. For example, when vehicle ID=0 (shown in a,b,c) becomes occluded for more than ten frames, UGD3MOT incorrectly initiates a new trajectory (ID=9 in b). In contrast, our method successfully re-establishes the correct association when the vehicle reappears in frame 24. Additionally, our method effectively tracks distant vehicles (IDs=7,9 in c) despite occlusions. While UGD3MOT doesn't track these vehicles due to noisy initial detections, our method maintains consistent tracking, with the scaling mechanism during early detection stages.

V. ABLATION STUDY

3D Detector impact. Tab. III demonstrates that the HybridTrack tracker's performance (HOTA) strongly correlates with the underlying 3D detector's accuracy, as higher-performing detectors (e.g., VirConv [7], CasA [8]) yield superior tracking results. However, one can note that the tracker is robust to detector performance drops, with more minor losses in HOTA compared to the drop in detection AP. For instance, while Second IoU has a 13.54% lower Moderate AP than VirConv-S, its HOTA decreases only by 7.77%.

Model Performance vs. Dataset Size. One of the main advantages of traditional filtering approaches against deep-learning-based methods is that these model-based methods do not need any data to be trained on. In this experiment, we show that our proposed method, HybridTrack, doesn't

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

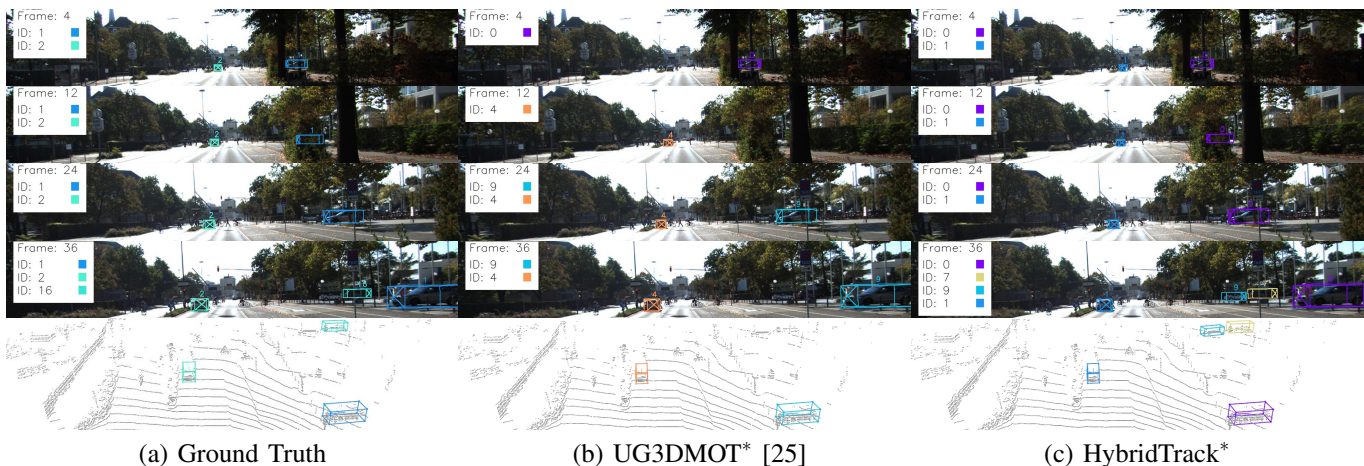


Fig. 2: Qualitative results comparison between ground truth, UG3DMOT and HybridTrack on sequence 15 of the validation set. * uses CasA [8] detector.

TABLE III: Ablation Study on 3D Detector Adaptability for KITTI and Corresponding Detection Performance. The results are reported in %.

3D Detector Choice	AP (IoU=0.7) Moderate	AP (IoU=0.7) Easy	HOTA \uparrow
Second IoU [42]	73.66	83.13	78.44
PV-RCNN [6]	81.43	90.25	80.14
CasA [8]	83.06	91.58	83.59
VirConv-S [7]	87.20	92.48	86.35

require a lot of data to be competitive with non-data-driven methods. As illustrated in Fig. 3, our model demonstrates strong data efficiency, achieving 81% HOTA on tracking on the validation set when trained with only 16 sequences (320 timesteps), compared to 86.355% HOTA when trained on 16,000 sequences. Also, one can see that, with a dataset containing only three vehicle trajectories, HybridTrack can already reach 62.92%, as illustrated by the red dot. Moreover, unlike non-deep-learning methods, which tend logically to plateau quickly regardless of the dataset size, our approach is scalable. It can continue to learn from additional data, enabling it to capture specific edge cases and consistently improve its object-tracking performance. One potential use case for our method could be handling accident scenarios, cases that are often lacking in publicly available datasets like KITTI. Unlike classical models that struggle with the unpredictability of accident trajectories, our data-driven approach can learn from accident data and would be able to adapt to these challenging situations, enabling more robust and reliable tracking in real-world applications.




Fig. 3: Performance Metrics vs. Training dataset size.

Data Association Cost Function Choice. Tab. IV evaluates cost map choices and their impact on tracking performance,

highlighting a trade-off between accuracy and computational efficiency. The 3D-CIoU cost map achieves the best tracking performance (HOTA: 86.355%) but is computationally expensive, with the lowest FPS (98.91). In contrast, the simple Euclidian distance-based cost map is the fastest (FPS: 112.83) while maintaining solid performance (HOTA: 85.864%). Combining the L2 error and the size error slightly improves the accuracy. The 2D CIoU method performs worst across metrics, highlighting the importance of 3D domain association for better tracking.

TABLE IV: Ablation Study on Cost Map Choices and Metrics Evaluation. The results are reported in %.

2D CIoU	L2 Error	Size Error	3D CIoU \uparrow	HOTA \uparrow	MOTA \uparrow	IDF1 \uparrow	FPS \uparrow
✓	-	-	-	82.052	86.574	88.742	44.81
-	✓	-	-	85.964	90.345	94.988	112.83
-	-	✓	-	86.037	90.500	95.084	111.60
-	-	-	✓	86.355	91.061	95.512	98.91

Ablation Study on Learning-based Components. We conduct an ablation study to evaluate the impact of our learnable prediction and update steps. As shown in Table V, learning only the update step fails to converge effectively, indicating its dependence on an accurate predicted state. In contrast, learning the prediction step alone already yields strong performance, demonstrating its important role in modeling motion dynamics. Jointly learning both steps achieves the best performance, highlighting the benefit of end-to-end optimization. These results confirm that combining L-PS and L-US significantly improves tracking accuracy compared to a standard Kalman filter with a constant-acceleration model.

TABLE V: Ablation study on learning-based components. 'L-PS' refers to the Learned Prediction Step, 'L-US' to the Learned Update Step. Results are reported in %.

Configuration	Baseline	L-US Only	L-PS Only	L-PS + L-US
HOTA \uparrow	83.927	54.320	85.276	86.355

Computational Cost. HybridTrack achieves high efficiency during inference on a Titan X NVIDIA GPU and an Intel(R) Xeon(R) CPU E5-2630 v4 @ 2.20GHz CPU, respectively. Our method achieves an average of 98.91 FPS across all test

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

sequences on a GPU device. On the other hand, while our method is optimised for GPU devices, it remains reasonably efficient on CPU devices, on average 20.80 frames can be processed per second. In Fig. 4, we visualize the execution time usage percentage of each step in HybridTrack. The primary computational bottleneck remains in the learning-based motion prediction and association steps.




Fig. 4: Share of Cumulative Execution Time for Sequence 1 of the validation set.

VI. CONCLUSION

We introduce HybridTrack, a 3D MOT framework that fuses Kalman filtering with deep learning for robust, efficient, and scalable tracking. HybridTrack eliminates manual tuning by integrating learnable motion models and noise covariances, adapting to diverse traffic scenarios. It achieves state-of-the-art performance and demonstrates strong data efficiency and scalability. Our approach is well-suited for real-world traffic applications, offering reliable tracking even in challenging use cases such as long-term occlusion and distant vehicles.

REFERENCES

- [1] P. Tokmakov, J. Li, W. Burgard, and A. Gaidon, "Learning to track with object permanence," in *ICCV*, 2021, pp. 10860–10869.
- [2] P. Tokmakov, A. Jabri, J. Li, and A. Gaidon, "Object permanence emerges in a random walk along memory," *arXiv preprint arXiv:2204.01784*, 2022.
- [3] H. Wu, Q. Li, C. Wen, X. Li, X. Fan, and C. Wang, "Tracklet proposal network for multi-object tracking on point clouds," in *IJCAI*, 2021, pp. 1165–1171.
- [4] X. Wang *et al.*, "A multi-modal fusion-based 3d multi-object tracking framework with joint detection," *IEEE RA-L*, pp. 1–8, 2024.
- [5] D. Ramanan and D. Forsyth, "Finding and tracking people from the bottom up," in *CVPR*, vol. 2, 2003, pp. II–II.
- [6] S. Shi, C. Guo, L. Jiang, Z. Wang, J. Shi, X. Wang, and H. Li, "Pv-rcnn: Point-voxel feature set abstraction for 3d object detection," in *CVPR*, 2020, pp. 10529–10538.
- [7] H. Wu, C. Wen, S. Shi, X. Li, and C. Wang, "Virtual sparse convolution for multimodal 3d object detection," in *CVPR*, 2023, pp. 21653–21662.
- [8] H. Wu, J. Deng, C. Wen, X. Li, C. Wang, and J. Li, "Casa: A cascade attention network for 3-d object detection from lidar point clouds," *IEEE T-GRS*, vol. 60, pp. 1–11, 2022.
- [9] G. Welch, G. Bishop *et al.*, "An introduction to the kalman filter," 1995.
- [10] Z. Chen *et al.*, "Bayesian filtering: From kalman filters to particle filters, and beyond," *Statistics*, vol. 182, no. 1, pp. 1–69, 2003.
- [11] X. Wang, J. Wang, D. Held, and K. Kitani, "3d multi-object tracking: A baseline and new evaluation metrics," in *2020 IEEE/RSS JROS*. IEEE, 2020, pp. 10359–10366.
- [12] H. Wu, W. Han, C. Wen, X. Li, and C. Wang, "3d multi-object tracking in point clouds based on prediction confidence-guided data association," *IEEE T-ITS*, vol. 23, no. 6, pp. 5668–5677, 2021.
- [13] J. Cao, J. Pang, X. Weng, R. Khirodkar, and K. Kitani, "Observation-centric sort: Rethinking sort for robust multi-object tracking," in *CVPR*, 2023, pp. 9686–9696.
- [14] L. Wang *et al.*, "Camo-mot: Combined appearance-motion optimization for 3d multi-object tracking with camera-lidar fusion," *IEEE T-ITS*, 2023.
- [15] C. Jiang, Z. Wang, H. Liang, and Y. Wang, "A novel adaptive noise covariance matrix estimation and filtering method: Application to multi-object tracking," *IEEE T-IV*, vol. 9, no. 1, pp. 626–641, 2023.
- [16] T. Fischer *et al.*, "Qdtrack: Quasi-dense similarity learning for appearance-only multiple object tracking," *IEEE T-PAMI*, 2023.
- [17] N. Wojke, A. Bewley, and D. Paulus, "Simple online and realtime tracking with a deep association metric," in *2017 IEEE ICIP*. IEEE, 2017, pp. 3645–3649.
- [18] L. Di Bella, Y. Lyu, B. Cornelis, and A. Munteanu, "Monokalman: Monocular vehicle pose estimation with kalman filter-based temporal consistency," in *2024 25th IEEE MDM*. IEEE, 2024, pp. 247–250.
- [19] S. Scheidegger, J. Benjaminsson, E. Rosenberg, A. Krishnan, and K. Granström, "Mono-camera 3d multi-object tracking using deep learning detections and pmbm filtering," in *2018 IEEE IV Symp*. IEEE, 2018, pp. 433–440.
- [20] K. Yi *et al.*, "Ucmtrack: Multi-object tracking with uniform camera motion compensation," in *AAAI*, vol. 38, no. 7, 2024, pp. 6702–6710.
- [21] X. Guo, Y. Zheng, and D. Wang, "Pmtrack: Multi-object tracking with motion-aware," in *ACCV*, 2024, pp. 3091–3106.
- [22] X. Xu, W. Ren, X. Chen, H. Fan, Z. Han, and H. Liu, "Exploiting multi-modal synergies for enhancing 3d multi-object tracking," *IEEE RA-L*, 2024.
- [23] K. Huang, M. Zhang, and Q. Hao, "Bitrack: Bidirectional offline 3d multi-object tracking using camera-lidar data," *arXiv preprint arXiv:2406.18414*, 2024.
- [24] X. Li, D. Liu, Y. Wu, X. Wu, L. Zhao, and J. Gao, "Fast-poly: A fast polyhedral framework for 3d multi-object tracking," *arXiv preprint arXiv:2403.13443*, 2024.
- [25] J. He, C. Fu, X. Wang, and J. Wang, "3d multi-object tracking based on informatic divergence-guided data association," *Signal Processing*, vol. 222, p. 109544, 2024.
- [26] X. Wang, C. Fu, Z. Li, Y. Lai, and J. He, "Deepfusionmot: A 3d multi-object tracking framework based on camera-lidar fusion with deep association," *IEEE RA-L*, vol. 7, no. 3, pp. 8260–8267, 2022.
- [27] N. Marinello, M. Proesmans, and L. Van Gool, "Triplettrack: 3d object tracking using triplet embeddings and lstm," in *CVPR*, 2022, pp. 4500–4510.
- [28] T. Zhou, Q. Ye, W. Luo, H. Ran, Z. Shi, and J. Chen, "Aptracker+: Displacement uncertainty for occlusion handling in low-frame-rate multiple object tracking," *IJCV*, pp. 1–26, 2024.
- [29] M. Y. Lee, C. D. W. Lee, J. Li, and M. H. Ang Jr, "Dino-mot: 3d multi-object tracking with visual foundation model for pedestrian re-identification using visual memory mechanism," *IEEE RA-L*, 2024.
- [30] A. Kim, G. Brasó, A. Ošep, and L. Leal-Taixé, "Polarmot: How far can geometric relations take us in 3d multi-object tracking?" in *ECCV*. Springer, 2022, pp. 41–58.
- [31] L. Wang, J. Zhang, P. Cai, and X. Lil, "Towards robust reference system for autonomous driving: Rethinking 3d mot," in *2023 IEEE ICRA*. IEEE, 2023, pp. 8319–8325.
- [32] Z. Zheng *et al.*, "Enhancing geometric factors in model learning and inference for object detection and instance segmentation," *IEEE Trans. Cybern.*, vol. 52, no. 8, pp. 8574–8586, 2021.
- [33] G. Revach, N. Shlezinger, X. Ni, A. L. Escoriza, R. J. Van Sloun, and Y. C. Eldar, "Kalmannet: Neural network aided kalman filtering for partially known dynamics," *IEEE T-SP*, vol. 70, pp. 1532–1547, 2022.
- [34] J. L. Elman, "Finding structure in time," *Cognitive science*, vol. 14, no. 2, pp. 179–211, 1990.
- [35] C. Peng *et al.*, "Pnas-mot: Multi-modal object tracking with pareto neural architecture search," *IEEE RA-L*, 2024.
- [36] X. Wang *et al.*, "Mctrack: A unified 3d multi-object tracking framework for autonomous driving," *arXiv preprint arXiv:2409.16149*, 2024.
- [37] A. Geiger, P. Lenz, and R. Urtasun, "Are we ready for autonomous driving? the kitti vision benchmark suite," in *CVPR*. IEEE, 2012, pp. 3354–3361.
- [38] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *CoRR*, vol. abs/1412.6980, 2014.
- [39] J. Luiten *et al.*, "Hota: A higher order metric for evaluating multi-object tracking," *IJCV*, vol. 129, pp. 548–578, 2021.
- [40] K. Bernardin and R. Stiefelhofen, "Evaluating multiple object tracking performance: The clear mot metrics," *EURASIP J. Image Video Process.*, vol. 2008, pp. 1–10, 2008.
- [41] E. Ristani, F. Solera, R. Zou, R. Cucchiara, and C. Tomasi, "Performance measures and a data set for multi-target, multi-camera tracking," in *ECCV*. Springer, 2016, pp. 17–35.
- [42] Y. Yan, Y. Mao, and B. Li, "Second: Sparsely embedded convolutional detection," *Sensors*, vol. 18, no. 10, p. 3337, 2018.