

# UncertaintyTrack: Exploiting Detection and Localization Uncertainty in Multi-Object Tracking

Chang Won Lee and Steven L. Waslander

**Abstract**—Multi-object tracking (MOT) methods have seen a significant boost in performance recently, due to strong interest from the research community and steadily improving object detection methods. The majority of tracking methods, which follow the tracking-by-detection (TBD) paradigm, blindly trust the incoming detections with no sense of their associated localization uncertainty. This lack of uncertainty awareness poses a problem in safety-critical tasks such as autonomous driving where passengers could be put at risk due to erroneous detections that have propagated to downstream tasks, including MOT. While there are existing works in probabilistic object detection that predict the localization uncertainty around the boxes, no work in 2D MOT for autonomous driving has studied whether these estimates are meaningful enough to be leveraged effectively in object tracking. We introduce UncertaintyTrack, a collection of extensions that can be applied to multiple TBD trackers to account for localization uncertainty estimates from probabilistic object detectors. Experiments on the Berkeley Deep Drive MOT dataset show that the combination of our method and informative uncertainty estimates reduces the number of ID switches by around 19% and improves mMOTA by 2-3%. The source code is available at <https://github.com/TRAILab/UncertaintyTrack>

## I. INTRODUCTION

Multi-object tracking (MOT) consists of tracking multiple dynamic objects by associating their detections across multiple frames. MOT for both image (2D) and lidar (3D) is a key component in robust autonomous driving systems as it provides object locations and scales across time for better scene understanding. MOT is also important for applications such as pedestrian counting and video analysis, which operate primarily in 2D. The object trajectories produced by MOT methods are used to predict future object motion, make safe decisions in motion planning for autonomous systems and identify human activities in surveillance settings. While MOT performance has improved significantly due to advances in object association, most state-of-the-art methods rely on accurate detections for strong performance.

The majority of MOT methods follow the tracking-by-detection (TBD) paradigm and separate detection and association into two sequential tasks, meaning the detections are provided as intermediate inputs to the association pipeline [1]. For this reason, TBD trackers are especially affected by detection performance. Unfortunately, despite their high accuracy on in-distribution data, deep object detectors are susceptible to errors with out-of-distribution (OOD) [2], [3]

Chang Won Lee and Steven L. Waslander are with the Institute For Aerospace Studies (UTIAS), University of Toronto, Toronto, Canada.  
Emails: (john.lee@robotics.utias.utoronto.ca; steven.waslander@utoronto.ca)




Fig. 1. Examples of tracking results on the BDD100K dataset [5]. In these images, orange boxes indicate false negatives, red boxes represent false positives, and blue boxes denote identity switches between tracked objects. The sequence of images is in chronological order, arranged from top to bottom. Compared to our baseline (left), UncertaintyTrack (right) effectively reduces the number of ID switches.

examples often seen in deployment for real-world applications. For instance, the sources of error in autonomous driving could be due to complex, unforeseen environments such as night-time, snowy days, and object occlusions. These examples highlight the importance of not only developing robust and reliable autonomous systems to minimize the risks of failure, but also to provide outputs that express the uncertainty associated with object tracks so that OOD events can be approached with caution.

It is therefore critical to capture the uncertainty associated with MOT predictions to enable reliable integration of deep-learning-based systems into autonomy software stacks. By explicitly quantifying the underlying uncertainty of predictive outputs, developers can better understand where the network is uncertain, identify the sources of error, and use this information to make more informed decisions on the design for improvements. Additionally, an ideal robotic system should be able to flag predictions with high uncertainty and seek necessary human intervention autonomously [4].

The introduction of Probabilistic Object Detection [6] has led to several probabilistic object detectors [7], [8], [9], [10], [11] that quantify the semantic and spatial uncertainty of detections. However, there is no prior work in MOT that effectively exploits these uncertainty estimates to develop a more reliable tracking pipeline for autonomous driving. As

current MOT methods are not designed to consider detection uncertainty from probabilistic object detectors when associating detections to tracks, their performance may suffer in the presence of high uncertainty. There are multiple ways trackers could benefit from such uncertainty estimates. For instance, they could filter out detections with high localization uncertainty or prioritize the association between confident predictions. To that end, we put forward the following contributions:

- We present Detection Uncertainty-Aware Kalman Filter, Confidence Ellipse Filtering, Bounding Box Relaxation and Entropy-Based Greedy Matching, which are simple and intuitive but effective extensions to existing 2D trackers to account for detection uncertainty. Our collection of add-ons that we denote *UncertaintyTrack* reduces the number of ID switches by around 19% and increases mMOTA by 2-3% on the Berkeley Deep Drive (BDD100K) dataset, when compared to existing baselines.
- We study how the variation in uncertainty estimates due to the use of different datasets and detectors affects MOT performance.
- We show that the explicit parameterization of localization uncertainty distributions provides meaningful insight for human interpretability and identification of the sources of error in MOT.

The presented methods may also be extended to 3D methods with similar tracking pipelines, as there are no learned modules specific to 2D, such as appearance-based similarities.

## II. RELATED WORKS

### A. Probabilistic Object Detection

The majority of current state-of-the-art object detectors are composed of three primary building blocks: backbone, detection head, and the post-processing pipeline. The detection head maps high-level representations of the input data known as feature maps to object bounding boxes that express the location, orientation and extent of detected objects. Then, it is common to perform Non-Maximum Suppression (NMS) to filter out redundant boxes and return a single tight bounding box for each object. Detectors that follow this standard procedure [12], [13], [14], [15], [16] are said to be deterministic because they output predicted bounding boxes with no estimates of the underlying uncertainty associated with the predictions.

In contrast, probabilistic object detectors are developed to produce reliable uncertainty estimations. Based on the Bayesian framework, predictive uncertainty can be divided into epistemic and aleatoric uncertainties. Epistemic uncertainties arise from model limitations that arise from training on a finite dataset, and aleatoric uncertainties are due to inherent noise in the data, such as far-away objects and occlusions [17]. Most probabilistic object detectors apply MC-Dropout [18] or Deep Ensembles [19] at the detection head to model epistemic uncertainty and Direct Modeling to model aleatoric uncertainty [20]. MC-Dropout involves estimating

the uncertainty by computing sample statistics of detector outputs from multiple forward passes with dropout applied during inference, while ensembles use multiple copies of the network created with different training schemes to generate samples. Direct Modeling consists of using dedicated layers that learn to predict the parameters of output distributions. Direct Modeling relies on the assumption of a certain probability distribution over the predictions; categorical distributions are most common for classification tasks and the Gaussian is the most common for regression [20].

To train the layers to learn bounding box distributions, it is common to use the negative log-likelihood (NLL) for Gaussian distributions [17]. However, it has been shown that minimizing the NLL to obtain the maximum likelihood estimates (MLE) of the network weights results in predicted distributions with high entropy (underconfident predictions). The energy score loss [21] for multivariate Gaussian was introduced as an alternative [22]:

$$ES = \frac{1}{N} \sum_{n=1}^M \left( \frac{1}{M} \sum_{j=1}^M \|z_{n,j} - z_n\| - \frac{1}{2(M-1)} \sum_{j=1}^{M-1} \|z_{n,j} - z_{n,j+1}\| \right) \quad (1)$$

where there are  $N$  ground truth bounding boxes  $z_n$ ,  $M$  i.i.d. samples drawn from  $\mathcal{N}(\hat{\mu}(x_n, w), \Sigma(x_n, w))$  and  $z_{n,j}$  is the  $j^{\text{th}}$  sample.

Finally, most probabilistic object detectors use NMS to follow their deterministic counterparts closely. Some others modify it to consider spatial uncertainty when removing redundancy [20]. BayesOD, the basis for Prob-YOLOX in *UncertaintyTrack*, replaces it with Bayesian inference to consider all redundant outputs instead of discarding them [4]. Covariance Intersection [23] has also been proposed to fuse redundant detections with unknown correlation [24].

### B. Multiple Object Tracking

MOT methods that follow the tracking-by-detection (TBD) paradigm [25], [26], [27], [28], [29], [30] primarily focus on improving the re-identification of objects given their detections, as the object detection problem is already addressed by strong object detectors [31], [12], [32]. For this reason, much of our focus will be on TBD methods to allow the use of bounding box distributions provided by probabilistic object detectors.

The object association task is often set up as a bipartite graph matching problem where two sets of nodes in the graph represent the tracked objects and new observations, respectively. Existing works employ the Hungarian Algorithm [33] to find the optimal assignment between the two sets or use greedy matching [34]. SORT [25] uses a Kalman Filter that assumes constant velocity to predict the trajectories of tracked objects and measures the bounding box overlap with new detections which is used as similarity scores for the assignment problem. DeepSORT proposes to use the Mahalanobis distance between predicted Kalman object states and new detections after realizing that the IoU metric leads to a

high number of ID switches when the motion model’s state uncertainty is high [35].

Furthermore, StrongSORT [30] replaces the vanilla Kalman Filter with the NSA Kalman algorithm [36] that adaptively calculates measurement noise covariance by taking into account the detection confidence score. The authors of OC-SORT [26] argue that the linear motion assumption leads to a large variance in estimated velocity by KF. They propose to decrease the error accumulation in KF by employing a smoothing strategy. Finally, ByteTrack [29] leverages confidence scores to create a secondary set of detections for more potential associations. However, these methods still use the IoU as their geometric distance metric.

In summary, most methods in 2D MOT for autonomous driving, if not all, have only used the state uncertainties from the Kalman Filter or class confidence scores as a measure of uncertainty in the tracking pipeline. Despite the importance of uncertainty, none of the previous works have taken into account localization uncertainty from probabilistic object detectors to show their effectiveness in MOT.

### III. LEVERAGING UNCERTAINTY IN MULTI-OBJECT TRACKING

#### A. Predicting Bounding Box Distributions

1) *Prob-YOLOX Architecture*: We build upon YOLOX [12] to develop Prob-YOLOX that estimates the localization uncertainty of the predicted bounding boxes. To explicitly parameterize the uncertainty distribution, we closely follow [4], [24] where the authors assume a multivariate Gaussian and predict the covariance matrix of the distribution. More specifically, we add a regression head that learns the log-scale diagonal covariance in parallel with the box regression head that predicts the mean. Then, the output of the two heads, the mean and covariance, parameterize a multivariate Gaussian that represents the bounding box distribution.

We emulate MC-Dropout sampling [18] by computing the sample statistics of the bounding box delta samples drawn from the distribution as follows:

$$\begin{aligned}\mu_s &= \frac{1}{N} \sum_{i=1}^N \mathbf{x}_i \\ \Sigma_s &= \frac{1}{N-1} \sum_{i=1}^N (\mathbf{x}_i - \mu_s)(\mathbf{x}_i - \mu_s)^\top\end{aligned}\quad (2)$$

where  $\mathbf{x}_i$  is the  $i$ -th bounding box sample. Then, to cluster bounding boxes together, we perform NMS to determine the cluster centers. However, instead of discarding the other cluster members, which are generally considered redundant detections, we employ the Improved Fast Covariance Intersection [37], a tractable approximation of the Covariance Intersection algorithm [23] to fuse the correlated predictions, but whose correlation is unknown.

2) *Losses*: In YOLOX, the regression branch is trained using both the L1 Loss and the IoU Loss. To guide Prob-YOLOX to learn the parameters of the multivariate Gaussian distribution, we replace the L1 Loss with Energy Score Loss

as proposed in [22]. Additionally, we introduce Sample-IoU Loss in place of the standard IoU Loss. This new loss function is designed to work in tandem with Energy Score Loss, helping to learn the distribution parameters while still optimizing the bounding box predictions to minimize the IoU, similar to how the IoU Loss supports the L1 Loss in YOLOX. The Sample-IoU is given by:

$$\begin{aligned}\text{Sample-IoU} &= \frac{1}{N} \sum_{n=1}^N \left( \frac{1}{M} \sum_{j=1}^M \text{IoU}(z_{n,j}^* - z_n^*) \right. \\ &\quad \left. - \frac{1}{2(M-1)} \sum_{j=1}^{M-1} \text{IoU}(z_{n,j}^* - z_{n,j+1}^*) \right)\end{aligned}\quad (3)$$

where there are  $N$  ground truth decoded bounding boxes  $z_n^*$ ,  $M$  i.i.d. decoded samples drawn from  $\mathcal{N}(\hat{\mu}(x_n, w), \Sigma(x_n, w))$  and  $z_{n,j}^*$  is the  $j$ th sample.

#### B. Extending Existing Trackers To Leverage Uncertainty

We propose four possible ways to extend existing trackers to account for localization uncertainty in their association pipeline (see fig. 2 for an overview of UncertaintyTrack). Our work builds upon tracking-by-detection algorithms that use the Kalman Filter as their motion model for two reasons: 1) the detector can be seamlessly replaced with a probabilistic object detector (e.g. Prob-YOLOX) that produces uncertainty estimates, and 2) the inclusion of an explicit motion model makes it more intuitive to incorporate detection uncertainty in the tracking task.

1) *Detection Uncertainty-Aware Kalman Filter*: The Kalman Filter (KF) [38] acts as a motion predictor with a constant velocity assumption that models the object’s location or state as:

$$\mathbf{x} = [x, y, a, h, \dot{x}, \dot{y}, \dot{a}, \dot{h}]^\top, \quad (4)$$

representing the bounding box center position  $(x, y)$ , aspect ratio  $a$ , height  $h$ , and their respective velocities. To measure the geometric similarity between a tracked object and its potential observation, KF predicts the object’s state at the next time step modelled as a multivariate Gaussian distribution (see [39] for more details on KF).

In current state-of-the-art MOT methods, the measurement uncertainty for KF is computed by scaling the aspect ratio with pre-determined coefficients that are inadaptive to changes in detection or track. In fact, this approach produces higher uncertainty for wider objects than taller ones and may cause issues in autonomous driving where both cars and pedestrians are present. Instead, we propose to replace this measurement uncertainty with the predicted covariance matrix from Prob-YOLOX.

For a single detection of an object, Prob-YOLOX outputs  $\mu_{box} = [x_1, y_1, x_2, y_2]^\top$ , where  $(x_1, y_1)$  and  $(x_2, y_2)$  are the top-left and bottom-right corner pixel coordinates, and  $\Sigma_{box} \in \mathbb{R}^{4 \times 4}$  that represents the bounding box covariance. Then, the predicted uncertainty can be incorporated into the KF update after it has been converted to the  $[c_x, c_y, a, h]$  format.

2) *Confidence Ellipse Filtering*: The majority of tracking-by-detection algorithms filter incoming detections based on their class confidence score. By contrast, we suggest

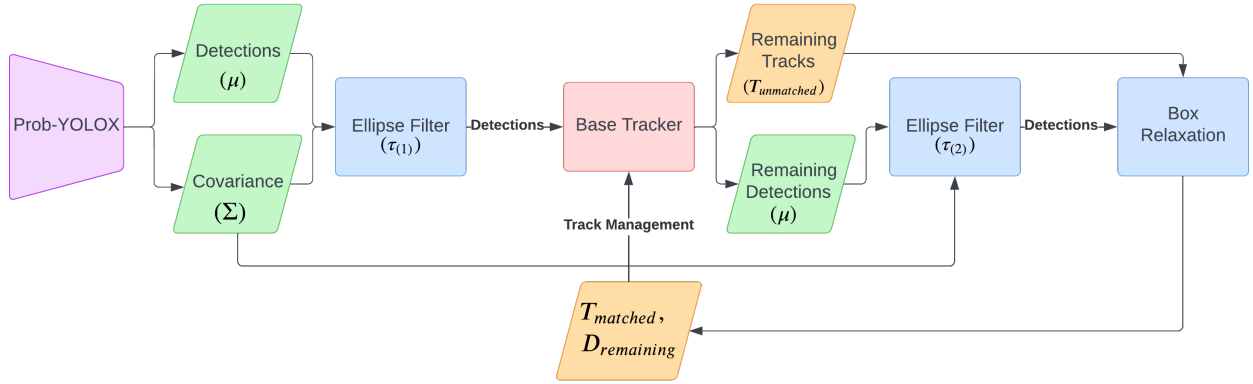


Fig. 2. Overview of UncertaintyTrack. Prob-YOLOX (purple) outputs detections and their covariance matrices (green) which are initially filtered by Confidence Ellipse Filtering (blue). Following the base tracker, UncertaintyTrack applies the filter again before the Box Relaxation step. Matched tracks and remaining detections after this step are passed back to the tracker for track management.

dropping detections with spatial uncertainty above a pre-determined threshold instead. First, we define error ellipses for the top-left and bottom-right corners of each bounding box based on their 95% confidence intervals. Then, we compute the ratio between their axis lengths and the box dimensions and compare it against a percentage threshold. Mathematically, if the box has width  $w_{box}$  and height  $h_{box}$  and the major and minor axis lengths of the error ellipse at the top-left and bottom-right corners are given by  $a_{tl}$ ,  $b_{tl}$ ,  $a_{br}$  and  $b_{br}$ , respectively, only the detections with error ellipses that meet the following conditions are kept:

$$\begin{aligned} \max(a_{tl}, a_{br}) &\leq \tau_{(i)} \times w_{box} \\ &\& \\ \max(b_{tl}, b_{br}) &\leq \tau_{(i)} \times h_{box} \end{aligned} \quad (5)$$

where  $\tau_{(i)}$  is a predefined percentage threshold.

3) *Bounding Box Relaxation*: After plotting the error ellipses along with the bounding boxes, we discovered that many ID switches are caused by an insufficient intersection between the detections and the tracks. Interestingly, for some detections that were incorrectly associated, their enlarged bounding boxes defined by their error ellipses had a significant overlap with the last observed detections of the tracks. Hence, we propose to add a final matching step with bounding boxes that have been enlarged according to the extremities of their error ellipses. More specifically, we compute  $(x_{min}, y_{min})$  of the top-left error ellipse and the  $(x_{max}, y_{max})$  of the bottom-right error ellipse. Then, the enlarged bounding box is defined by  $[x_{min}, y_{min}, x_{max}, y_{max}]$ .

With this expansion, we forcefully increase the probability of the target being in the vicinity of the bounding box. We apply this uncertainty-based enlargement to both new detections and track detections at the previous step and match them using the GIoU [40] distance metric. In addition, to limit the amount of enlargement, we apply the ellipse-based filter a second time before this final matching step.

4) *Entropy-Based Greedy Matching*: Finally, at the final matching step with expanded boxes, we employ greedy matching based on the Gaussian entropy of the predicted

box distributions to prioritize the association of detections with lower uncertainty (more confident predictions).

## IV. EXPERIMENTS AND RESULTS

### A. Experimental Details

1) *Datasets*: BDD100K [5] is used as the primary benchmark for this study. The multi-class dataset contains scenes with diverse environments and varied weather conditions. We also evaluate UncertaintyTrack on the MOT17 pedestrian tracking dataset [41]. MOT17 does not have a separate validation set, so results are generated by taking the second half of the videos in the training set following [34].

### 2) Metrics:

a) *Probabilistic Object Detection Metrics*: NLL and ES are used to evaluate the quality of predictive uncertainty. As proper scoring rules, their minimum score is obtained only when the predicted distribution matches the actual data generating distribution [22]. NLL is also a local scoring rule and penalizes low predictive densities at the exact values of the target. It also prefers predictive distributions that are less informative (high entropy). On the contrary, as a non-local measure, ES penalizes distributions with higher entropy and favors low entropy densities close to the target [22], so the two metrics present complimentary assessments of the uncertainty prediction quality.

b) *MOT Metrics*: We use the CLEAR MOT metrics [42] (e.g. MOTA, FP, FN, IDs), IDF1 [43] and HOTA [44]. MOTA combines the number of false negatives (FN), false positives (FP), and ID switches (IDs) to measure the overall tracking performance, but it is biased toward detection performance. IDF1 computes the F1 score of the object trajectories but focuses on measuring association performance. HOTA seeks a balance between detection and association by taking the geometric mean of their accuracies. The ‘m’ in front of the metrics denotes taking the average across all classes; for MOT17, it is equivalent to the original metrics.

3) *Implementation*: We leverage the open-source libraries MMDetection [45] and MMTracking [46] to develop Prob-YOLOX and UncertaintyTrack. We used the same setup as

Method	Set	mMOTA $\uparrow$	mIDF1 $\uparrow$	mHOTA $\uparrow$	FP $\downarrow$	FN $\downarrow$	IDs $\downarrow$	IDs (%)
ByteTrack*	BDD	32.5	42.1	38.4	21775	120104	46732	0.0%
<b>UncertaintyTrack (Ours)</b>	BDD	<b>35.1</b>	<b>45.5</b>	<b>40.1</b>	27099	<b>113488</b>	<b>37682</b>	<b>-19.4%</b>
ByteTrack*	MOT17	76.0	76.9	66.7	8114	29424	1231	0.0%
<b>UncertaintyTrack (Ours)</b>	MOT17	<b>76.6</b>	<b>78.0</b>	<b>67.2</b>	<b>7218</b>	29894	<b>767</b>	<b>-37.7%</b>
OC-SORT*	BDD	27.0	39.4	36.8	40472	99363	65814	0.0%
<b>UncertaintyTrack (Ours)</b>	BDD	<b>30.1</b>	<b>41.1</b>	<b>37.6</b>	<b>31898</b>	108217	<b>53217</b>	<b>-19.1%</b>
OC-SORT*	MOT17	74.0	77.6	67.0	12099	27060	2946	0.0%
<b>UncertaintyTrack (Ours)</b>	MOT17	<b>75.8</b>	<b>78.9</b>	<b>67.5</b>	<b>7249</b>	30511	<b>1322</b>	<b>-55.1%</b>

TABLE I

COMPARISON OF BYTETRACK\* AND OC-SORT\* WITH UNCERTAINTYTRACK ON BDD100K AND MOT17 VALIDATION SETS. THE BEST RESULTS FROM EACH PAIR OF COMPARISONS ARE SHOWN IN BOLD. \* DENOTES THE BASELINE WITHOUT LINEAR INTERPOLATION.

Method	mMOTA $\uparrow$	mIDF1 $\uparrow$	mHOTA $\uparrow$	FP $\downarrow$	FN $\downarrow$	IDs $\downarrow$
SUSHI [27]	40.3	<b>60.0</b>	<b>48.4</b>	-	-	<b>13626</b>
QDTrack [28]	<b>42.4</b>	55.6	-	89376	<b>154797</b>	14282
ByteTrack [29]	40.1	55.8	-	63869	169073	15466
UncertaintyTrack <sub>Byte</sub> (Ours)	33.5	47.6	40.9	<b>52596</b>	202856	73124
UncertaintyTrack <sub>OC</sub> (Ours)	29.0	42.4	38.1	61538	193916	102279

TABLE II

COMPARISON BETWEEN EXISTING MOT METHODS AND UNCERTAINTYTRACK ON BDD100K TEST SET. THE BEST RESULTS ARE SHOWN IN BOLD.

MMTracking’s YOLOX setup for ByteTrack [29] to train Prob-YOLOX on MOT17 and BDD100K.

We use ByteTrack and OC-SORT [26] as our baselines for UncertaintyTrack experiments as they are state-of-the-art methods on multiple MOT benchmarks and have been reproduced by MMTracking for smooth integration with custom object detectors. However, we remove linear interpolation from the post-processing pipeline to keep the methods online and relevant for robotics. We will use an asterisk \* to note that linear interpolation has been removed.

For inference, we use the same hyperparameters as the baselines for the base tracker. For  $\tau_{(1)}$  and  $\tau_{(2)}$ , 0.65 and 0.3 were used for both ByteTrack\* and OC-SORT\* on BDD100K. For MOT17, (0.7, 0.3) and (0.6, 0.3) were used for ByteTrack\* and OC-SORT\*, respectively.

### B. Comparison With Existing Trackers

To address some of the shortcomings in MOT in the presence of detections with high uncertainty, we apply UncertaintyTrack to ByteTrack\* and OC-SORT\* and present the evaluation results on the BDD100K validation set and MOT17 half validation set in table I. Our proposed extensions for localization uncertainty improve the two trackers on mMOTA, mIDF1 and mHOTA and significantly reduce the number of ID switches on both datasets. While our primary objective is to show that detection uncertainty estimates can be helpful for understanding tracking failures and is not purely to develop a state-of-the-art tracker, we include our results on the BDD100K test set for completeness. It is important to note that we do not use learning-based affinity scores for association unlike existing methods, which can significantly boost performance on the BDD100K set.

### C. Relationship Between Detection Uncertainty & Tracking

To understand why our proposed extensions were successful, it is necessary to analyze the causal relationship between




Fig. 3. Examples of ByteTrack\*[29] failures on the BDD100K validation set. Orange, red and blue represent false negatives, false positives and ID switches, respectively.

Detector	Set	mAP (0.5:0.95) $\uparrow$	NLL $\downarrow$	ES $\downarrow$
YOLOX	MOT17	60.5	-	-
YOLOX	BDD	39.1	-	-
Prob-YOLOX	MOT17	60.0	12.6	9.81
Prob-YOLOX	BDD	39.0	639.76	9.61
BayesOD	BDD	24.3	182.19	13.19

TABLE III

EVALUATION OF DETECTION AND LOCALIZATION UNCERTAINTY QUALITY FROM YOLOX [12], PROB-YOLOX AND BAYESOD [4]

inaccurate detections and tracking failures. We plot the 95% confidence ellipses in each top-left and bottom-right corner of the detected bounding boxes that are false positives or lead to ID switches on the BDD100K validation set. No ellipses are drawn for false negative errors. In fig. 3, the two examples on the top show that some MOT errors are due to inaccurate detections of occluded/small objects, which aligns well with the expectations within the object detection community. Interestingly, they also manifest relatively large uncertainty ellipses, indicating the detector has learned to express the difficulty of the occluded/small object detection task well. In addition, other visible failures are caused by insufficient overlap between the detections and ground-truths (bottom-left) or a significant overlap between adjacent detections (bottom-right). Importantly, these errors also have high uncertainty associated with their boxes.

With these few but informative illustrations, it is possible to infer that detections leading to MOT failures tend to have relatively high localization uncertainty compared to correctly associated detections. This highlights the importance of explicit parameterization of the uncertainty distributions as it allows for an in-depth analysis of where the “black-box” deep object detectors are underperforming at inference time and a better understanding of how to adjust MOT association based on the distributions from probabilistic object detectors.

Furthermore, to study why the improvements on BDD100K are more significant than what is observed on MOT17, we evaluate the quality of the predicted distributions of Prob-YOLOX on BDD100K and MOT17 using POD scoring rules and present our findings in table III. Prob-

Method	Set	mMOTA $\uparrow$	mIDF1 $\uparrow$	FP $\downarrow$	FN $\downarrow$	IDs $\downarrow$	IDs (%)
ByteTrack*	BDD	<b>16.6</b>	<b>26.0</b>	<b>20638</b>	236470	25336	0.0%
<b>UncertaintyTrack (Ours)</b>	BDD	12.1	25.7	66140	<b>222100</b>	<b>23219</b>	<b>-8.4%</b>
ByteTrack*	MOT17	<b>58.8</b>	<b>63.9</b>	<b>20470</b>	<b>43890</b>	2205	0.0%
<b>UncertaintyTrack (Ours)</b>	MOT17	51.3	59.6	30350	46473	<b>1847</b>	<b>-16.2%</b>
OC-SORT*	BDD	-11.5	21.7	118483	<b>144665</b>	89154	0.0%
<b>UncertaintyTrack (Ours)</b>	BDD	<b>5.3</b>	<b>22.9</b>	<b>50277</b>	200086	<b>43217</b>	<b>-51.5%</b>
OC-SORT*	MOT17	40.3	54.6	53726	<b>35398</b>	7439	0.0%
<b>UncertaintyTrack (Ours)</b>	MOT17	<b>53.5</b>	<b>60.6</b>	<b>23812</b>	30511	<b>3298</b>	<b>-55.7%</b>

TABLE IV

COMPARISON OF BYTETRACK\* AND OC-SORT\* WITH THE EXTENSIONS ON BDD100K VALIDATION SET AND MOT17 VALIDATION SET USING BAYESOD [4] AS THE DETECTOR. THE BEST RESULTS FROM EACH PAIR OF COMPARISONS ARE SHOWN IN BOLD.

YOLOX has a much higher NLL (local scoring rule) on BDD100K than MOT17, but lower ES (non-local scoring rule) and entropy scores, which signifies that the predicted distributions are informative and confident but centered away from the target values. In contrast, the distributions on MOT17 are closer to the target (lower NLL) while relatively having the same confidence level (similar ES). This observation aligns with the mAP values that show Prob-YOLOX has more accurate detections on MOT17 than on BDD100K. It may also explain UncertaintyTrack’s behaviour on the two datasets. It is sensible that there is more room for improvement on the more complex BDD100K than MOT17 because the performance is saturating on the latter.

#### D. Ablation Studies

1) *BayesOD & UncertaintyTrack*: To test if UncertaintyTrack could improve our baseline trackers with a different probabilistic detector, we repeated the same tracking experiments using BayesOD [4] and present the results in table IV. In contrast to earlier observations with Prob-YOLOX, the improvements are inconsistent among the trackers. While OC-SORT\*’s performance generally increased on most metrics – especially the number of ID switches, ByteTrack\* performed worse with UncertaintyTrack.

Similar to what was done for Prob-YOLOX, we evaluated BayesOD’s uncertainty quality using the same POD metrics and included them in table III. The results not only indicate BayesOD is a weaker detector with a lower mAP score, they show that the uncertainty estimates are less informative. BayesOD has a lower NLL but a higher ES score, meaning that the distributions are closer to the target but are not confident and more spread out. As UncertaintyTrack was designed to improve baseline trackers in the presence of informative and accurate predictive distributions, the limitations of the BayesOD lead to diminished performance in tracking.

The observed discrepancy in UncertaintyTrack’s results demonstrate that improvements in TBD trackers often depend on the detector choice and do not perform in the same way for all methods. It also further solidifies the importance of explicit parameterization of localization uncertainty for in-depth human analysis. Uncertainty estimates can be a complementary source of helpful information in designing a robust tracker as they allow for a better understanding of where detection errors cause MOT to fail.

KF Covariance	Ellipse	Relaxation	Greedy	mMOTA $\uparrow$	mIDF1 $\uparrow$	IDs (%) $\downarrow$
✓				32.4	41.6	-0.6%
	✓			32.6	42.1	-4.1%
	✓	✓		34.8	45.1	-14.3%
	✓	✓	✓	34.9	45.2	-14.3%
✓	✓	✓	✓	<b>35.1</b>	<b>45.5</b>	<b>-19.4%</b>

TABLE V

ANALYSIS OF THE EFFECT OF UNCERTAINTYTRACK’S COMPONENTS ON BYTETRACK\*. ALL VERSIONS ARE EVALUATED ON THE BDD100K VALIDATION SET.

2) *Effect of Each Component*: We study the effect of each extension on ByteTrack\* and present the results on BDD100K validation set in table V. The biggest improvement is observed with uncertainty-aware box relaxation, increasing mMOTA and mIDF1 by 2.2 and 3.0 points, respectively. Box relaxation in particular helps with missed associations between detections and tracks due to insufficient overlap. By increasing the size of the detection and tracklet boxes based on the corner uncertainty, their probability of intersection is essentially increased, leading to more matches.

There is a minimal increase in performance with ellipse-based filtering which hints that detections with high localization uncertainty were well-filtered by ByteTrack’s class confidence-based filter and suggests an underlying correlation between the two. Also, the detections before the box relaxation step are generally far from each other for greedy matching to have sufficient impact.

Interestingly, the use of predicted box covariance as measurement uncertainty in the Kalman Filter complements the other components but decreases performance as a stand-alone extension. We hypothesize that this inconsistency in results may be due to the choice of distribution to model the uncertainty distribution. It is possible that the multivariate Gaussian distribution does not accurately represent the true underlying uncertainty distribution, suggesting that the uncertainty estimates might not be compatible with the Kalman Filter. This remains an area for future investigation.

## V. CONCLUSION

This paper studies the quality of learned bounding box uncertainty from a probabilistic object detector and presents UncertaintyTrack, a collection of extensions applicable to several state-of-the-art tracking-by-detection MOT methods to leverage the predicted uncertainty. Experiments using UncertaintyTrack show that probabilistic object detectors produce meaningful localization uncertainty estimates which can be used in trackers to improve overall performance under complex scenarios. This work provided concrete examples that show the usefulness of probabilistic object detection for downstream tasks such as MOT and paves the path for more innovative probabilistic tracking techniques as well as more applications of probabilistic object detection. Potential future work includes extending UncertaintyTrack to 3D trackers and combining learned detection uncertainty with feature representation uncertainty in trackers that use learned similarities for association.

## REFERENCES

- [1] D. Ramanan and D. A. Forsyth, "Finding and tracking people from the bottom up," in *2003 IEEE Computer Society Conference on Computer Vision and Pattern Recognition, 2003. Proceedings.*, vol. 2. IEEE, 2003, pp. II–II.
- [2] D. Hendrycks and K. Gimpel, "A baseline for detecting misclassified and out-of-distribution examples in neural networks," *arXiv preprint arXiv:1610.02136*, 2016.
- [3] B. Lakshminarayanan, A. Pritzel, and C. Blundell, "Simple and scalable predictive uncertainty estimation using deep ensembles," *Advances in neural information processing systems*, vol. 30, 2017.
- [4] A. Harakeh, M. Smart, and S. L. Waslander, "Bayesod: A bayesian approach for uncertainty estimation in deep object detectors," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2020, pp. 87–93.
- [5] F. Yu, H. Chen, X. Wang, W. Xian, Y. Chen, F. Liu, V. Madhavan, and T. Darrell, "Bdd100k: A diverse driving dataset for heterogeneous multitask learning," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 6 2020.
- [6] D. Hall, F. Dayoub, J. Skinner, H. Zhang, D. Miller, P. Corke, G. Carneiro, A. Angelova, and N. Sünderhauf, "Probabilistic object detection: Definition and evaluation," in *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 2020, pp. 1031–1040.
- [7] F. O. Catak, T. Yue, and S. Ali, "Prediction surface uncertainty quantification in object detection models for autonomous driving," in *2021 IEEE International Conference on Artificial Intelligence Testing (AITest)*. IEEE, 2021, pp. 93–100.
- [8] D. Miller, N. Sünderhauf, H. Zhang, D. Hall, and F. Dayoub, "Benchmarking sampling-based probabilistic object detectors." in *CVPR Workshops*, vol. 3, 2019, p. 6.
- [9] D. Feng, X. Wei, L. Rosenbaum, A. Maki, and K. Dietmayer, "Deep active learning for efficient training of a lidar 3d object detector," in *2019 IEEE Intelligent Vehicles Symposium (IV)*. IEEE, 2019, pp. 667–674.
- [10] M. T. Le, F. Diehl, T. Brunner, and A. Knol, "Uncertainty estimation for deep neural object detectors in safety-critical applications," in *2018 21st International Conference on Intelligent Transportation Systems (ITSC)*. IEEE, 2018, pp. 3873–3878.
- [11] F. Kraus and K. Dietmayer, "Uncertainty estimation in one-stage object detection," in *2019 IEEE intelligent transportation systems conference (itsc)*. IEEE, 2019, pp. 53–60.
- [12] Z. Ge, S. Liu, F. Wang, Z. Li, and J. Sun, "Yolox: Exceeding yolo series in 2021," *arXiv preprint arXiv:2107.08430*, 2021.
- [13] S. Ren, K. He, R. Girshick, and J. Sun, "Faster r-cnn: Towards real-time object detection with region proposal networks," *Advances in neural information processing systems*, vol. 28, 2015.
- [14] T.-Y. Lin, P. Goyal, R. Girshick, K. He, and P. Dollár, "Focal loss for dense object detection," in *Proceedings of the IEEE international conference on computer vision*, 2017, pp. 2980–2988.
- [15] K. Duan, S. Bai, L. Xie, H. Qi, Q. Huang, and Q. Tian, "Centernet: Keypoint triplets for object detection," in *Proceedings of the IEEE/CVF international conference on computer vision*, 2019, pp. 6569–6578.
- [16] J. Redmon and A. Farhadi, "Yolov3: An incremental improvement," *arXiv preprint arXiv:1804.02767*, 2018.
- [17] A. Kendall and Y. Gal, "What uncertainties do we need in bayesian deep learning for computer vision?" *Advances in neural information processing systems*, vol. 30, 2017.
- [18] Y. Gal and Z. Ghahramani, "Dropout as a bayesian approximation: Representing model uncertainty in deep learning," in *International Conference on Machine Learning*. PMLR, 2016, pp. 1050–1059.
- [19] B. Lakshminarayanan, A. Pritzel, and C. Blundell, "Simple and scalable predictive uncertainty estimation using deep ensembles," *Advances in neural information processing systems*, vol. 30, 2017.
- [20] D. Feng, A. Harakeh, S. L. Waslander, and K. Dietmayer, "A review and comparative study on probabilistic object detection in autonomous driving," *IEEE Transactions on Intelligent Transportation Systems*, 2021.
- [21] T. Gneiting, L. I. Stanberry, E. P. Gritmit, L. Held, and N. A. Johnson, "Assessing probabilistic forecasts of multivariate quantities, with an application to ensemble predictions of surface winds," *Test*, vol. 17, pp. 211–235, 2008.
- [22] A. Harakeh and S. L. Waslander, "Estimating and evaluating regression predictive uncertainty in deep object detectors," in *International Conference on Learning Representations*, 2021. [Online]. Available: <https://openreview.net/forum?id=YLeWtnvKgR7>
- [23] D. Smith and S. Singh, "Approaches to multisensor data fusion in target tracking: A survey," *IEEE transactions on knowledge and data engineering*, vol. 18, no. 12, pp. 1696–1710, 2006.
- [24] A. Harakeh, "Estimating and evaluating predictive uncertainty in deep object detectors," Ph.D. dissertation, University of Toronto (Canada), 2021.
- [25] A. Bewley, Z. Ge, L. Ott, F. Ramos, and B. Uppcroft, "Simple online and realtime tracking," in *2016 IEEE international conference on image processing (ICIP)*. IEEE, 2016, pp. 3464–3468.
- [26] J. Cao, J. Pang, X. Weng, R. Khirodkar, and K. Kitani, "Observation-centric sort: Rethinking sort for robust multi-object tracking," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 9686–9696.
- [27] O. Cetintas, G. Brasó, and L. Leal-Taixé, "Unifying short and long-term tracking with graph hierarchies," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 22 877–22 887.
- [28] T. Fischer, T. E. Huang, J. Pang, L. Qiu, H. Chen, T. Darrell, and F. Yu, "Qdtrack: Quasi-dense similarity learning for appearance-only multiple object tracking," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2023.
- [29] Y. Zhang, P. Sun, Y. Jiang, D. Yu, F. Weng, Z. Yuan, P. Luo, W. Liu, and X. Wang, "Bytetrack: Multi-object tracking by associating every detection box," in *European Conference on Computer Vision*. Springer, 2022, pp. 1–21.
- [30] Y. Du, Y. Song, B. Yang, and Y. Zhao, "Strongsort: Make deepsort great again," *arXiv preprint arXiv:2202.13514*, 2022.
- [31] Z. Cai and N. Vasconcelos, "Cascade r-cnn: Delving into high quality object detection," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 6154–6162.
- [32] P. Sun, R. Zhang, Y. Jiang, T. Kong, C. Xu, W. Zhan, M. Tomizuka, L. Li, Z. Yuan, C. Wang, *et al.*, "Sparse r-cnn: End-to-end object detection with learnable proposals," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2021, pp. 14 454–14 463.
- [33] H. W. Kuhn, "The hungarian method for the assignment problem," *Naval research logistics quarterly*, vol. 2, no. 1-2, pp. 83–97, 1955.
- [34] X. Zhou, V. Koltun, and P. Krähenbühl, "Tracking objects as points," in *European conference on computer vision*. Springer, 2020, pp. 474–490.
- [35] N. Wojke, A. Bewley, and D. Paulus, "Simple online and realtime tracking with a deep association metric," in *2017 IEEE international conference on image processing (ICIP)*. IEEE, 2017, pp. 3645–3649.
- [36] Y. Du, J. Wan, Y. Zhao, B. Zhang, Z. Tong, and J. Dong, "GiaoTracker: A comprehensive framework for mcmt with global information and optimizing strategies in visdrone 2021," in *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) Workshops*, 10 2021, pp. 2809–2819.
- [37] D. Franken and A. Hupper, "Improved fast covariance intersection for distributed data fusion," in *2005 7th International Conference on Information Fusion*, vol. 1. IEEE, 2005, pp. 7–pp.
- [38] R. E. Kalman, "A new approach to linear filtering and prediction problems," 1960.
- [39] T. D. Barfoot, *State estimation for robotics*. Cambridge University Press, 2017.
- [40] H. Rezatofighi, N. Tsoi, J. Gwak, A. Sadeghian, I. Reid, and S. Savarese, "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*, 2019, pp. 658–666.
- [41] A. Milan, L. Leal-Taixé, I. Reid, S. Roth, and K. Schindler, "MOT16: A benchmark for multi-object tracking," *arXiv:1603.00831 [cs]*, Mar. 2016, arXiv: 1603.00831. [Online]. Available: <http://arxiv.org/abs/1603.00831>
- [42] K. Bernardin and R. Stiefelhagen, "Evaluating multiple object tracking performance: the clear mot metrics," *EURASIP Journal on Image and Video Processing*, vol. 2008, pp. 1–10, 2008.
- [43] E. Ristani, F. Solera, R. Zou, R. Cucchiara, and C. Tomasi, "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.* Springer, 2016, pp. 17–35.

- [44] J. Luiten, A. Osep, P. Dendorfer, P. Torr, A. Geiger, L. Leal-Taixé, and B. Leibe, “Hota: A higher order metric for evaluating multi-object tracking,” *International journal of computer vision*, vol. 129, pp. 548–578, 2021.
- [45] K. Chen, J. Wang, J. Pang, Y. Cao, Y. Xiong, X. Li, S. Sun, W. Feng, Z. Liu, J. Xu, Z. Zhang, D. Cheng, C. Zhu, T. Cheng, Q. Zhao, B. Li, X. Lu, R. Zhu, Y. Wu, J. Dai, J. Wang, J. Shi, W. Ouyang, C. C. Loy, and D. Lin, “MMDetection: Open mmlab detection toolbox and benchmark,” *arXiv preprint arXiv:1906.07155*, 2019.
- [46] M. Contributors, “MMTracking: OpenMMLab video perception toolbox and benchmark,” <https://github.com/open-mmlab/mtracking>, 2020.