

Context-Aware Streaming Perception in Dynamic Environments

Gur-Eyal Sela¹, Ionel Gog^{1*}, Justin Wong¹, Kumar Krishna Agrawal¹, Xiangxi Mo¹, Sukrit Kalra¹, Peter Schafhalter¹, Eric Leong¹, Xin Wang², Bharathan Balaji^{3**}, Joseph Gonzalez¹, and Ion Stoica¹

¹ University of California, Berkeley

² Microsoft Research

³ Amazon

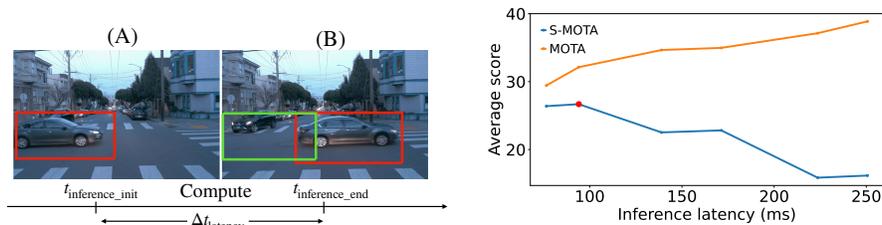
Abstract. Efficient vision works maximize accuracy under a latency budget. These works evaluate accuracy offline, one image at a time. However, real-time vision applications like autonomous driving operate in streaming settings, where ground truth changes between inference start and finish. This results in a significant accuracy drop. Therefore, a recent work proposed to maximize accuracy in streaming settings on average. In this paper, we propose to maximize streaming accuracy for every environment context. We posit that scenario difficulty influences the initial (offline) accuracy difference, while obstacle displacement in the scene affects the subsequent accuracy degradation. Our method, Octopus, uses these scenario properties to select configurations that maximize streaming accuracy at test time. Our method improves tracking performance (S-MOTA) by 7.4% over the conventional static approach. Further, performance improvement using our method comes in addition to, and not instead of, advances in offline accuracy.

1 Introduction

Recent works like EfficientDet [26], YOLO [3], and SSD [16] were designed for real-time computer vision applications that require high accuracy in the presence of latency constraints. However, these solutions are evaluated offline, one image at a time, and do not consider the impact of increase in inference latency on the application performance. In real-time systems such as autonomous vehicles, the models are deployed in an online streaming setting where the ground truth changes during inference time as shown in Figure 1a. To evaluate performance in streaming settings, Li et al. [12] proposed a modified metric that measures the model performance against the ground truth at the end of inference. They evaluated object detection models in a streaming fashion, and found that the average precision of the best performing model drops from 38.0 to 6.2, and picking the model that maximizes streaming average precision reduces the drop to 17.8.

* Now at Google Research.

** Work unrelated to Amazon.



(a) In online streaming settings, the environment changes during inference. Streaming accuracy is computed by evaluating the prediction run on ground truth A against ground truth B.

(b) Offline and streaming accuracy of a tracker of increasing model size. While the offline MOTA (orange, y-axis) of the model increases as inference latency (x-axis) increases, its streaming MOTA (S-MOTA in blue, y-axis) decreases.

Fig. 1: Streaming accuracy deviates from offline accuracy because the ground truth changes during inference.

We confirm the findings of Li et al. [12], and extend their analysis to object tracking. The standard metric for object tracking is MOTA (multiple object tracking accuracy) [18,11], and we refer to its streaming counterpart as S-MOTA. Figure 1b shows that larger models with higher MOTA deteriorate in S-MOTA for the Waymo dataset [28] as the higher latency widens the gap between ground truth between inference start and finish. The MOTA of the largest model (EfficientDet-D7x) is 38.9 while the S-MOTA is 16.2. The model that maximizes S-MOTA is EfficientDet-D4 with MOTA of 32.1 and S-MOTA of 26.7.

Since environment context varies, to further analyze the tradeoffs between latency and accuracy, we identify the model that maximizes the S-MOTA of 1-second video segment scenarios in the Waymo dataset. We observe that the best performing model varies widely from scenario to scenario (Figure 2). Scenarios that are difficult (e.g., sun glare, drops on camera and reflection) and still (e.g., standing cars in intersection) show (Figure 2 center) benefit from stronger perception while incurring marginal penalty from the latency increase. On the other hand, simple and fast scenes with rapid movement (e.g., turning, in Figure 2 left), behave the opposite, where performance degrades sharply with latency. As a result, at the video segment-level the optimization landscape of streaming accuracy looks vastly different than on aggregate (Figure 1b).

In this paper, we propose leveraging contextual cues to optimize S-MOTA dynamically at test time. The object detection model is just one of several choices that we refer to as *metaparameters* to consider in an object tracking system. Our method, called Octopus, optimizes S-MOTA at test time by dynamically tuning the metaparameters. Concretely, we train a light-weight second-order model to switch between the metaparameters using a battery of environment features extracted from video segments, like obstacle movement speed, obstacle proximity, and time of day.

Our contributions in this work can be summarized as:

- We are the first to analyze object tracking in a streaming setting. We show that the models that maximize S-MOTA change per scenario and propose that

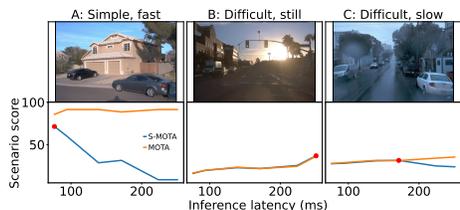


Fig. 2: Offline and streaming accuracies in three different scenarios.

While offline MOTA (orange, y-axis) of the EfficientDet [26] models increases as inference latency (x-axis) increases, streaming MOTA (S-MOTA in blue, y-axis), responds depending on the scenario context (frames). Compare these plots to Fig. 1b, which shows the same on average over the entire dataset.

the optimization tradeoffs are a result of scene difficulty as well as obstacle displacement.

- We present a novel method of S-MOTA optimization that leverages contextual features to switch the object tracking configuration at test-time.
- Our policy improves tracking performance (S-MOTA) by 7.4% over a static approach by evaluating it on the Waymo dataset [28]. We improve S-MOTA by 3.4% when we apply this approach on the Argoverse dataset [5].

2 Related Works

Latency vs. accuracy. Prior works have recognized the tension between latency and accuracy in perception models [3,7,9,26], and examined facets of the tradeoff between decision speed *vs.* accuracy [24]. However, these works study this tradeoff in offline static settings. Li et al [12] examine this tradeoff in streaming settings, where the ground truth of the world changes continuously. They show that conventional accuracy misrepresents perception performance in such settings, and proposed streaming accuracy. Their results were shown in detection (as well as semantic segmentation [6]), so we first verified that the idea also holds in tracking. Further, while Li et al. maximize streaming accuracy on average, we expose how the environment context plays a crucial role on its behavior. We leverage the context to dynamically maximize streaming accuracy at test time.

Model serving optimization and dynamic test-time adaptation. To reduce the inference latency, techniques such as model pruning [8,13,17] and quantization [23,29,31] have been proposed. While these techniques focus on effectively reducing the size of the models, our focus is on the policy: when and where to do it, in order to optimize accuracy in streaming settings. These techniques could be applied to reduce the latency of the models we use. Most similar to our setting are works that leverage context to dynamically adjust model architecture at test time to reduce resource consumption [10,27] or to improve throughput [25]. This work employs a similar approach in leveraging insight about latency *vs.* accuracy tradeoff in AV (autonomous vehicle) perception.

Inferring configuration performance without running. Several existing works achieve resource savings by modeling from data how a candidate configuration would perform without actually running it. Hyperstar [19] learns to approximate how a candidate hyperparameter configuration would perform on a

dataset without training. Chameleon [10] reduces profiling cost of configurations for surveillance camera detection by assuming temporal locality among profiles. One key differentiating factor with our work is that the prior works were designed for a hybrid setting where some profiling is allowed. In our case, the strict time and compute constraints in the AV setting [14] restrict this approach, requiring an approximation-only method like Octopus.

3 Problem Setup

We first lay out the problem formulation (§3.1). Next, after an introduction of the Octopus dataset (§3.2), we measure the accuracy opportunity gap between the global best policy [12] and the optimal dynamic policy (§3.3). Finally, we perform a breakdown analysis of the components needed in order to optimize streaming accuracy at test time (§3.4).

3.1 Problem formulation

Given real-time video stream as a series of images, we consider the inference of an AV pipeline (obstacle detection and tracking) on this stream. Let \mathbb{S} denote the mean S-MOTA score of the tracking model, which depends on the values of the metaparameters \mathcal{H} , such as object detection model architecture and maximum age of tracked objects. Currently, the metaparameters $h \in \mathcal{H}$ are chosen using offline datasets, and are kept constant during deployment [12]. We refer to this method as the *global best* approach for statically choosing global metaparameters (h_{global}), which are expected to be best across all driving scenarios.

In contrast, in this work we study whether the S-MOTA score \mathbb{S} can be improved by dynamically changing h every $\Delta\tau$ at test time. The metaparameters we choose at each time period $[\tau, \tau + \Delta\tau)$ is h_τ , and the corresponding score-optimal values is h_τ^* .

3.2 The Octopus dataset

To generate the Octopus dataset (\mathcal{D}), we divide each video of a driving dataset into consecutive segments of duration $\Delta\tau$. We run the perception pipeline with a range of values of metaparameters \mathcal{H} , and record the S-MOTA score for each segment s_τ^h . We assign the optimal h_τ^* to the metaparameters that achieve the highest S-MOTA score (i.e., the optimal S-MOTA score s_τ^*). The segment duration $\Delta\tau$ is chosen to be short. This allows more accurately studying the performance potential of dynamic streaming accuracy optimization, because decision-making over smaller intervals generally performs better.

We generate the Octopus dataset by recording metaparameter values h_τ and the corresponding S-MOTA scores s_τ^h of the Pylot AV pipeline [7] for the Argoverse and Waymo datasets [5,28]. We execute Pylot’s perception consisting of a suite of 2D object detection models from the EfficientDet model family [26] followed by the Simple, Online, and Real-Time tracker [2]. For each video scenario, we explore the following metaparameters:

Table 1: Dynamically changing metaparameters creates an accuracy opportunity gap. The *streaming* accuracy (S-MOTA) of the *global best* metaparameters h_{global} is 6.1 points lower on average than that of the *optimal* metaparameters h_{τ}^* . Similarly, there is a 3.2 gap in MOTA (top).

Method	Dataset	MOTA \uparrow	MOTP \uparrow	FP \downarrow	FN \downarrow	ID $_{sw}\downarrow$
Global best	Waymo	37.3	78.1	21515	543193	11615
Optimal	Waymo	40.5	77.6	15738	532678	9137
Global best	Argoverse	63.0	82.1	6721	43376	1284
Optimal	Argoverse	70.8	81.0	5772	34210	828
Method	Dataset	S-MOTA \uparrow	S-MOTP \uparrow	S-FP \downarrow	S-FN \downarrow	S-ID $_{sw}\downarrow$
Global best	Waymo	25.1	72.2	33616	633159	11212
Optimal	Waymo	31.2	71.0	28907	590847	6997
Global best	Argoverse	49.4	75.2	13485	55484	1092
Optimal	Argoverse	57.9	74.1	9562	48354	708

- **Detection model architecture:** selects the model from the EfficientDet family of models, which offers different latency *vs.* accuracy tradeoff points.
- **Tracked obstacles’ maximum age:** limits the duration for which the tracker continues modeling the motion of previously-detected obstacles, under the assumption of temporary occlusion or low detection confidence (flickering).

Other metaparameters had limited effect on performance (Appendix A).

We run every metaparameter configuration in a Cartesian product of selected values for each metaparameter, and record latency metrics, detected objects, and tracked objects. The resulting 18 metaparameter configurations yield \approx 18,000 trials. We make this dataset public (<https://github.com/EyalSel/Contextual-Streaming-Perception>).

3.3 Accuracy opportunity gap

In order to study if the *global best* metaparameters h_{global} offer the best accuracy in all driving scenarios, we split the Octopus dataset (\mathcal{D}) into train (\mathcal{D}_{train}) and test (\mathcal{D}_{test}) sets. Next, we compute global best metaparameters (h_{global}) as the configuration that yields the highest mean S-MOTA score across all video segments in \mathcal{D}_{train} . We denote h_{global} ’s mean S-MOTA scores on the train and test set as s_{train}^{global} and s_{test}^{global} , respectively. Similarly, we denote the mean S-MOTA scores of h_{τ}^* (i.e., optimally changing metaparameters) as s_{train}^* and s_{test}^* .

We define the *S-MOTA opportunity gap* between the optimal dynamic metaparameters and the global best metaparameters as the upper bound of $s_{test}^* - s_{test}^{global}$. We repeat the same calculation for MOTA. In Table 1, we show the opportunity gap for the Argoverse and Waymo datasets [5,28] using $\Delta\tau = 1s$. We conclude that optimally choosing the metaparameters at test time offers a 6.1 (Waymo) and 8.5 (Argoverse) S-MOTA improvement on average, along with reductions in streaming false positives/negatives and streaming ID switches.

Of note, if offline accuracy were to increase uniformly across all configurations and scenarios, the performance improvement of the dynamic approach over the static baseline is expected to persist. This applies to the opportunity gap

shown above (the optimal improvement), as well as for any dynamic policy improvement in this space. This means that performance improvement achieved by dynamic optimization come in addition to, and not instead of, further advances in conventional (offline) tracking.

3.4 Streaming accuracy analysis

We approach dynamic configuration optimization as a ranking problem [15], and solve it by learning to predict the difference in score of configuration pairs in a given scenario context [30]. We decompose this learning task into predicting the difference in (i) MOTA, and (ii) accuracy degradation during inference. To our knowledge we’re the first to perform this analysis.

Decomposition. Streaming accuracy (S-MOTA) is tracking accuracy against ground truth at the end of inference, instead of the beginning (MOTA) [12]. Offline accuracy (MOTA) degrades as a result of change in ground truth during inference. The gap between MOTA and S-MOTA is defined here as the “degradation”. Therefore, S-MOTA is expressed as $S - D$ where S is MOTA and D is the degradation. The difference in S-MOTA between two configurations is $(S_1 - S_2) - (D_1 - D_2)$. This decomposition separates the difference in S-MOTA of two configurations into two parts:

- $(S_1 - S_2)$ is the difference in offline accuracy. This difference originates from the MOTA boost, which is affected by (i) the added modeling capacity influenced by the scene difficulty for detection, a specific case of the more general example difficulty [1], and (ii) the max-age choice which depends on the scene obstacle displacement [32].
- $(D_1 - D_2)$ is the difference in accuracy degradation. This may be derived from:
 - a. the difference configuration latencies and
 - b. scene obstacle displacement.

Predicting both components to optimize streaming accuracy. At test-time, $S_1 - S_2$ and $D_1 - D_2$ are predicted for each scenario. We find that accurately predicting both components per environment context is necessary to realize the *opportunity gap* (see §3.3). First, we show that perfectly predicting $D_1 - D_2$ (ΔD^*) and $S_1 - S_2$ (ΔS^*) (Table 2, top left) yields 31.2 S-MOTA, the same as the optimal policy h_7^* on Waymo (Table 1 bottom panel, row 2). Then, we predict $D_1 - D_2$ and $S_1 - S_2$ on average across all scenarios for each configuration ($\overline{\Delta D}$ and $\overline{\Delta S}$ respectively in Table 2). Combining $\overline{\Delta D}$ and $\overline{\Delta S}$ (Table 2, bottom-right) yields 25.1 S-MOTA, the same as the global best policy h_{global} score on Waymo (Table 1 bottom panel, row 1). The hybrid-optimal policies (Table 2 top-right and bottom-left) achieve 1.8 (26.9 – 25.1) and 2.7 (27.8 – 25.1) of the 6.1 (31.2 – 25.1) optimal policy opportunity gap. Taken together, these results demonstrate the need to accurately predict both the change in MOTA ($S_1 - S_2$) and in degradation ($D_1 - D_2$) per scenario in order to optimize S-MOTA at test time.

	ΔD^*	$\overline{\Delta D}$
ΔS^*	31.2	26.9
$\overline{\Delta S}$	27.8	25.1

Table 2: Both streaming accuracy components must be predicted per scenario in order to realize the full dynamic policy’s opportunity gap. ΔD^* is optimal degradation prediction, ΔS^* is optimal offline gap prediction. $\overline{\Delta D}$ and $\overline{\Delta S}$ are the corresponding global-static policies.

4 Octopus: Environment-Driven Perception

We propose leveraging properties of the AV environment context (e.g., ego speed, number of agents, time of day) that can be perceived from sensors in order to dynamically change metaparameters at test time. We first formally present our approach for choosing metaparameters, which uses regression to infer a ranking of metaparameter configurations (§4.1). Then, we describe the environment representation that allows to effectively infer each component of the streaming accuracy (§4.2).

4.1 Configuration Ranking via Regression

In order to find the metaparameters h_τ that maximize the S-MOTA score s_τ^h for each video segment, Octopus first learns a regression model M . The model predicts s_τ^h given the metaparameters h_τ and the representation of the environment e_τ for the period τ . Following, Octopus considers all metaparameter values, and picks the metaparameters that give the highest predicted S-MOTA.

Executing the model M at the beginning of each segment τ requires an up-to-date representation of the environment. However, building a representation requires the output of the perception pipeline (e.g., number of obstacles). Octopus could run the current perception configuration in order to update the environment before executing the model M , but this approach would greatly increase the response time of the AV. Instead, Octopus makes a Markovian assumption, and inputs the environment representation of the previous segment $e_{\tau-1}$ to the model M .

$$M(h_\tau, e_{\tau-1}) = \hat{s}_\tau^h \quad (1)$$

Thus, it is important to limit the segment length $\Delta\tau$ as the longer a segment is, the more challenging it is to accurately predict the scores due to using an older environment representation (see §5.1 for our methodology for choosing $\Delta\tau$).

The model is trained using the mean squared error loss. We choose the metaparameters that the model predicts as the highest score. Let s_τ^{global} denote the S-MOTA score obtained with h^{global} metaparameters, which by definition is a lower bound of the optimal S-MOTA score s_τ^* (i.e., $s_\tau^{global} \leq s_\tau^*$). Thus, in order to pick the best metaparameters, Octopus can only predict s_τ^h relative to s_τ^{global} . As a result, Octopus utilizes the following as the final loss:

$$L = \frac{1}{N} \sum_{i=0}^N \left(\hat{r}_i^h - clip((s_i^h - s_i^{global}), \epsilon) \right)^2 \quad (2)$$

where N is the size of the training data and \hat{r}_i^h is the relative S-MOTA score predicted by the model M , $\hat{r}_i^h = \hat{s}_i^h - s_i^{global}$.

Finally, Octopus clips by lower bounding the predictions by $s_i^{global} - \epsilon$ as the predictions significantly below s_i^{global} are irrelevant to the optimization problem. Moreover, by clipping the predictions, Octopus reduces the dynamic range of the regressor and makes it easier to predict the higher S-MOTA scores.

4.2 Environment Representation

We can represent the environment context e_τ by capturing the characteristics of the video segment. In order to keep the decision-making latency small, we eschew more complex learned representation designs. Instead, we developed hand-engineered features from sensors and the outputs of the object detection model following prior work by Nishi et al. [20], where the features were used to predict human driving behavior. We collect the features per frame and then aggregate by averaging across frames in the video segment. Unless otherwise specified, we use the 10th percentile, mean, and 90th percentile of the following features:

1. **Bounding box speed** is the distance traveled by an object in pixel space across two frames. This feature infers the speed of objects, and thus it is important for capturing obstacle displacement to predict the streaming accuracy degradation.
2. **Bounding box self IoU** is the Intersection-over-Union (IoU) of an object’s bounding box in the current frame relative to the previous frame. Along with the bounding speed, this feature helps isolate the change in size of an obstacle, as it moves towards or away from the ego vehicle.
3. **Number of objects** is measured per frame, and indicates the complexity of a scene as the more objects in a scene, the more likely the AV is to encounter object path crossings and occlusions (scene difficulty). Therefore, this feature signals when to prioritize for high offline accuracy configurations that are robust to object occlusions.
4. **Obstacle longevity** is the number of frames for which an obstacle has been tracked. Lower obstacle longevity implies more occlusion as obstacles enter and leave the scene, making perception more difficult. This feature also guides the choice of tracking metaparameters as low obstacle longevities correlate with lower tracking maximum age giving better performance.
5. **Ego driving speed** is the speed at which the AV is traveling, and indicates the environment in which the AV is driving (e.g., highway vs. city).
6. **Ego turning speed** is the angle change of the AV’s direction between consecutive frames. This feature helps differentiate the source of apparent obstacle displacement between obstacle movement and ego movement.
7. **Time of day** hints if a high-accuracy configuration is required to handle challenging scenes (e.g., night driving, sun glare at sunset).

We compute these feature statistics for different bounding box size ranges in order to reason about obstacle behavior depending on the detection model

strength needed for their accurate detection. For example, higher offline accuracy models do not confer better streaming accuracy if the obstacles in the detectable size range move too rapidly for the longer detection inference time to keep up with.

Note that the choice of metaparameters in each video segment changes the configuration, and hence the outputs of the tracking pipeline. Therefore, the resulting environment representation is no longer independently and identically distributed (i.i.d), making it challenging to use traditional supervision techniques. To keep the data distribution stationary during train time, we use the object detection model as given by the global best metaparameters h_{global} . While this training objective is biased considering that the features may be derived from any configuration, we find that it reduces variance during training, and works better than using features derived from the ground truth.

We concatenate the metaparameters h and the environment representation vector e_τ , and then use a supervised regression model to predict the score, optimizing using the loss given in Eq. (2).

In addition, we compare our method to a conventional CNN model approach (ignoring its resource and runtime requirements) in Appendix I.

5 Experiments

Next, we analyze our proposed approach. The subsections discuss the following:

1. **Setup (§5.1)**: Description of Datasets and Model Details
2. **Main results (§5.2)**: What is the performance improvement conferred by the dynamic policy over the static baseline?
3. **Explainability (§5.3)**: **I.** Does the learned policy behavior match human understanding of the driving scenario? **II.** Where is the learned policy similar to and different from the optimal policy? **III.** How are scenarios clustered by metaparameter score? **IV.** What is the relative importance of the hand-picked features and of the metaparameters towards S-MOTA optimization?
4. **Ablation study (Appendix H)**: How do various ranking implementation choices affect the final performance?

5.1 Methodology

We evaluate on the Argoverse and Waymo datasets [5,28], covering a variety of environments, traffic, and weather conditions. We do not use the private Waymo test set as it does not support streaming metrics. However, we treat the Waymo validation set as the test set, and we perform cross validation on the training dataset (798 videos for training, and 202 videos for validation). Similarly, we divide the Argoverse videos into 75 videos for training, and 24 for validation. In addition, we follow the methodology from Li et. al [12] in order to create ground truth 2D bounding boxes and tracking IDs, which are not present in the Argoverse dataset. We generate these labels using QDTrack [22] trained on the BDD100k dataset, which is the highest offline accuracy model available.

In our experiments, we run object detectors from the EfficientDet architecture [26] on NVIDIA V100 GPUs, and the SORT tracker [2] on a CPU (simulated evaluation on faster hardware is in Appendix J). We chose to use the EfficientDet models because they are especially optimized for trading off between latency and offline accuracy, and because they are close to the state-of-the-art. The EfficientDet models were further optimized using Tensor RT [21], which both reduces inference latencies to less than 250ms (the latency of the largest model, EfficientDet-D7x) and decreases resource requirements to at most 3 GPUs. We use these efficient models to investigate the effect of optimizing two metaparameters: detection model and tracking max age (see §3.2 for details). We compile 18 configurations of the AV perception pipeline by exploring the Cartesian product of the values of the two metaparameters (see Listing 1.1).

```

1 # Metaparameters
2 detection-model = {EfficientDet: 3, 4, 5, 6, 7, 7x}
3 tracking-maximum-age = {1, 3, 7}

```

Listing 1.1: Values for detection and tracking metaparameters.

We implement Octopus’s policy regression model as a Random Forest [4] using Eq. (2) to choose metaparameters for video segment length $\Delta\tau = 1s$. We describe training details in Appendix B.

Policy Runtime Overhead. Every step $\Delta\tau$, the Octopus policy applies Random Forest regression for all the 18 perception pipeline configurations in order to predict the best one to apply in the next step. Due to the lightweight policy design, inference on all 18 configurations has a latency of at most 6ms using a single Intel Xeon Platinum 8000 core. As a result, the policy decision finishes before the sensor data of the next segment arrives, and thus does not affect latency of the perception pipeline. Moreover, Octopus pre-loads the perception model weights (13.86GB in total in our experiments) and forward pass activations in the GPU memory (32GB), and thus actuates metaparameter changes quickly.

5.2 Main results

We compare the Octopus policy with the optimal policy (h_t^*) and the global best policy (h^{global}) as proposed by Li et al. [12]. In addition, following the discussion in §4.2, we show ablations of the Octopus policy in order to highlight the relative contribution in both the setting of predictive and close-loop metaparameter optimization. Concretely, we include the following setups that Octopus can use to optimize the metaparameters at time t :

- *Ground truth from current segment:* features are derived from the sensor data and the labels at time t .
- *Ground truth from previous segment:* features are derived from the sensor data and the labels at time $t - \Delta\tau$.
- *Closed-loop prediction from previous segment:* features are derived from the sensor data and the output of the perception pipeline for the previous segment (i.e., at time $t - \Delta\tau$).

In Table 3 we show the tracking accuracy results for the Waymo and Argoverse datasets. The results show that the Octopus policy with closed-loop prediction outperforms the global best policy by 1.9 S-MOTA (Waymo) and 1.7 S-MOTA (Argoverse). Both the optimal and Octopus policies achieve further accuracy increases using features derived from the current segment, further illustrating how rapidly configuration score changes over time. The consistent accuracy improvements across the two datasets show that leveraging environment context to dynamically optimize streaming accuracy provides substantial improvements over the state-of-the-art. Moreover, as discussed in §3.3, streaming accuracy improvement over the global best policy will likely persist independently of innovation in offline accuracy of the underlying perception models.

Table 3: Streaming tracking accuracy results on two datasets.

(a) Waymo					
Method	S-MOTA↑	S-MOTP↑	S-FP↓	S-FN↓	S-ID _{sw} ↓
Global best	25.1	72.2	33616	633159	11212
Optimal	31.2	71.0	28907	590847	6997
Optimal from the prev. segment	27.2	71.2	38631	603777	8092
Octopus with:					
Ground truth from current segment	27.9	72.3	31489	608870	8966
Ground truth from prev. segment	27.3	72.7	31056	612862	9103
Prediction from prev. segment	27.0	72.8	30272	615780	9511
(b) Argoverse					
Method	S-MOTA↑	S-MOTP↑	S-FP↓	S-FN↓	S-ID _{sw} ↓
Global best	49.4	75.2	13485	55484	1092
Optimal	57.9	74.1	9562	48354	708
Optimal from the prev. segment	51.9	74.0	12652	52804	810
Octopus with:					
Ground truth from current segment	53.2	74.9	11447	52638	917
Ground truth from prev. segment	51.6	75.1	11348	54502	1010
Prediction from prev. segment	51.1	74.9	12062	54469	1008

S-MOTA vs. S-MOTP. Table 3 highlights that both Octopus and optimal policy occasionally deteriorate S-MOTP, inversely to the improvement in S-MOTA. This result reflects on a broader pattern where in streaming settings bounding boxes in general lag after the ground truth, even if by a small enough margin to be counted as true positives. S-MOTP, which is weighted by the IOU between correct predictions and the ground truth is especially hurt as a result. S-MOTA, which just counts the number of false positives, is less affected by this. We provide a longer analysis of this tradeoff/pareto-frontier between S-MOTA and S-MOTP in Appendix D.

5.3 Explainability

We survey various aspects of the metaparameter optimization problem, and qualitatively compare the global best (baseline), the optimal, and the Octopus (learned) policies.

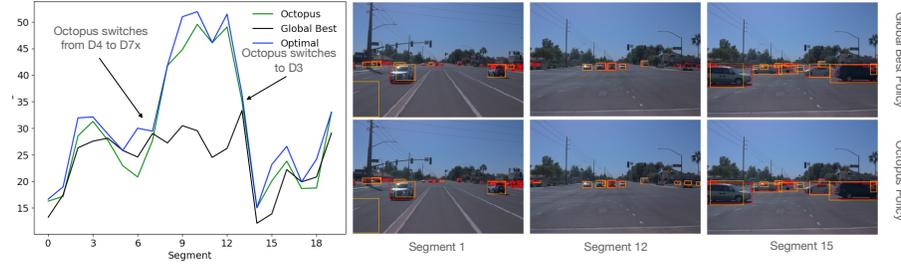


Fig. 3: Busy Intersection Scenario. Left: The S-MOTA score of each $\tau = 1s$ segment for the Octopus (learned), global best, and optimal policies. Right: The front-facing camera feed. Red bounding-boxes represent the ground truth, and orange represent the pipeline’s predictions using the policy’s configuration choice.

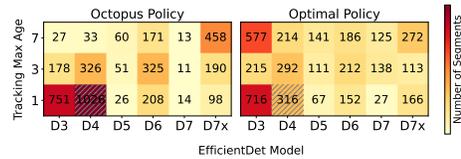


Fig. 4: Policy decision frequency.

The configuration choice frequency (color intensity) of the Octopus policy (left) and of the optimal policy (right). The global best configuration is EfficientDet-D4 with tracking max age of 1, emphasized with gray stripes.

I. Case study: Busy Intersection. Figure 3 shows a scenario where the ego vehicle enters a busy intersection. The scenario is divided into three phases, where the learned (Octopus) policy adaptively tunes the metaparameter to varying road conditions similarly to the optimal policy plan. First, the ego vehicle approaches the intersection with oncoming traffic on the left. The vehicles in the distance cannot be picked up by any of the candidate models. The learned policy chooses D4, the same model as the global best policy. Then, as the light turns yellow and the ego vehicle comes to a stop first in the queue, the learned policy adapts by increasing the strength of the model to D7x. The perception pipeline is now able to detect the smaller vehicles in the opposing lane and adjacent to the road. The optimal policy makes a similar decision while the global best policy remains with the D4 choice, incurring a net performance loss. Finally, as vehicles in the cross-traffic start passing and occluding the vehicles in the background, the learned policy returns to the lower latency D3 model. This scenario illustrates how the learned policy can select the best model in each phase of the scenario, adapting to the change in driving environment.

II. Policy action heatmap. In Figure 4, we visualize the learned policy in comparison to the optimal policy. As we expect, the learned policy often selects the global best configuration (D4-1), but expands out similarly to the optimal policy when performance can be improved by changing configurations. This result illustrates how, contrary to common belief that onboard perception must have low latency (e.g. under 100ms [28]), higher perception latencies are tolerable and even preferable in certain environments. The Octopus policy learns to take advantage of this when it opts for higher offline accuracy models.

III. Scenarios are clustered by metaparameter score.

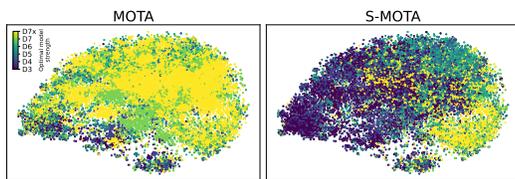


Fig. 5: The impact of the accuracy degradation incurred in online streaming context is scenario dependent. Each point is a 1s driving scenario. Its color is the model that maximizes offline accuracy (left) and streaming accuracy (right).

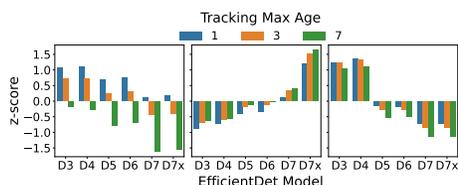


Fig. 6: Three centroids of the configuration score distribution: (i) preference for low max age (left), (ii) preference for offline accuracy (middle), and (iii) preference for low latency (right). The three clusters account for 34% of the dataset.

Here, we evaluate whether scenarios are grouped by common metaparameter score behaviors. To this end, we first visualize the scenario score space (explained below) as a t-SNE plot and then perform more formal centroid analysis.

To this end, each video segment of length $\Delta\tau = 1s$ is vectorized by computing the MOTA score difference ($S_c - S_{h_{global}}$) and degradation difference ($D_c - D_{h_{global}}$) for every metaparameter configuration c and the global best configuration h_{global} (see §3.4). These values are then concatenated and normalized (z-score) across each scenario.

Cluster visualization. We visualize this space in a t-SNE plot in Figure 5. The scenario segment points are colored according to the model that optimizes MOTA (left) and S-MOTA (right). We observe a non-uniform impact of accuracy degradation on the optimal model choice in different video segment regions. This varying effect illustrates that S-MOTA deviates from MOTA setting on an environment context-dependent basis. For case-study analysis of points in the t-SNE please see Appendix F.

Centroid analysis. We now perform a more formal clustering analysis that reveals modes of metaparameter optimality.

We perform K-means clustering, with $k=8$ on the z-score space. Figure 6 shows centroids of three representative scenario clusters (right to left): (i) scenarios that benefit from lower latency detection, (ii) scenarios that benefit from higher accuracy detection, and (iii) scenarios that primarily benefit from a lower tracking max age, and to a lesser extent from lower latency detection. These distinct modes reflect that scenarios are grouped around similar metaparameter behaviors, corroborating with variation in scene difficulty and obstacle displacement. The rest of the scenario cluster visualizations are in Appendix G.

IV. Feature and configuration metaparameter importance.

Feature importance. To study the relative importance of the environment features used in our solution (described in §4.2), we show the feature importance scores derived from the trained Random Forest regressor in Table 4a. The mean bounding box self-IOU and speed together constitute over half of the normalized

Table 4: Importance study
(a) Feature importance (Regression)

Feature	Importance Score	Feature	Importance Score
Mean BBox self IOU	0.304	BBox bin [665, 1024)	0.087
Mean BBox Speed	0.200	BBox bin [1024, 1480)	0.071
Num. BBoxes	0.063	BBox bin [1480, 2000)	0.055
BBox Longevity	0.076	BBox bin [2000, 2565)	0.040
Ego movement	0.096	BBox bin [2565, ∞)	0.417
Time of Day	0.008		

(b) Metaparameter importance (Regression)

Metaparameter	S-MOTA \uparrow	S-MOTP \uparrow	S-FP \downarrow	S-FN \downarrow	S-ID $_{sw}\downarrow$
Global best	25.1	72.2	33616	633159	11212
Detection-model	27.6	72.4	26116	615771	11254
Tracking-max-age	25.6	71.9	39400	625184	8973
Both	27.9	72.3	31489	608870	8966

importance score, as they are predictive of the accuracy degradation that higher-latency detection models would incur. Aggregate bounding box statistics that capture scene difficulty, such as the average number and longevity of the bounding boxes per frame, are also useful for prediction.

Configuration metaparameter importance. In Table 4b, we evaluate the gains attributed to each metaparameter by only optimizing one metaparameter at a time, fixing the other parameter to the value in the global best configuration. We do this to ablate the performance shown in Table 3, where they are optimized together. In both cases, the models were trained on the ground truth and the present. Although the detection model choice is the primary contributor, the additional choice of occlusion tolerance (max age) further contributes to the achieved performance. We also observe how the performance gain when optimizing the tracker’s maximum occlusion tolerance (see §3 for details) does not combine additively with the improvement in the detection model optimization.

6 Conclusions

Streaming accuracy is a much more accurate representation of tracking for real-time vision systems because it uses ground truth at the end of inference to measure performance. In this study we show the varying impact that environment context has on the deviation of streaming accuracy from offline accuracy. We propose a new method, Octopus, to leverage environment context to maximize streaming accuracy at test time. Further, we decompose streaming accuracy into two components: difference in offline accuracy MOTA and the degradation, and show that both must be inferred in every scenario to achieve optimal performance. Octopus improves streaming accuracy over the global best policy in multiple autonomous vehicle datasets.

Acknowledgements. We thank Daniel Rothchild and Horia Mania for helpful discussions.

References

1. Baldock, R., Maennel, H., Neyshabur, B.: Deep learning through the lens of example difficulty. *Advances in Neural Information Processing Systems* **34**, 10876–10889 (2021) [6](#)
2. Bewley, A., Ge, Z., Ott, L., Ramos, F., Upcroft, B.: Simple Online and Realtime Tracking. In: *Proceedings of the 23th IEEE International Conference on Image Processing (ICIP)*. pp. 3464–3468 (2016) [4](#), [10](#)
3. Bochkovskiy, A., Wang, C.Y., Liao, H.Y.M.: YOLOv4: Optimal Speed and Accuracy of Object Detection. *arXiv preprint arXiv:2004.10934* (2020) [1](#), [3](#)
4. Breiman, L.: Random forests. *Machine learning* **45**(1), 5–32 (2001) [10](#)
5. Chang, M.F., Lambert, J., Sangkloy, P., Singh, J., Bak, S., Hartnett, A., Wang, D., Carr, P., Lucey, S., Ramanan, D., et al.: Argoverse: 3D Tracking and Forecasting with Rich Maps (2019) [3](#), [4](#), [5](#), [9](#)
6. Courdier, E., Fleuret, F.: Real-time segmentation networks should be latency aware. In: *Proceedings of the Asian Conference on Computer Vision* (2020) [3](#)
7. Gog, I., Kalra, S., Schafhalter, P., Wright, M.A., Gonzalez, J.E., Stoica, I.: Pylot: A Modular Platform for Exploring Latency-Accuracy Tradeoffs in Autonomous Vehicles. In: *Proceedings of IEEE International Conference on Robotics and Automation (ICRA)* (2021) [3](#), [4](#)
8. Han, S., Pool, J., Tran, J., Dally, W.J.: Learning both Weights and Connections for Efficient Neural Networks. In: *Proceedings of the 28th International Conference on Neural Information Processing (NeurIPS)*. pp. 1135–1143 (2015) [3](#)
9. Huang, J., Rathod, V., Sun, C., Zhu, M., Korattikara, A., Fathi, A., Fischer, I., Wojna, Z., Song, Y., Guadarrama, S., Murphy, K.: Speed/Accuracy Trade-Offs for Modern Convolutional Object Detectors. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2017) [3](#)
10. Jiang, J., Ananthanarayanan, G., Bodik, P., Sen, S., Stoica, I.: Chameleon: Scalable Adaptation of Video Analytics. In: *Proceedings of the 2018 Conference of the ACM Special Interest Group on Data Communication (SIGCOMM)*. pp. 253–266 (2018) [3](#), [4](#)
11. Leal-Taixé, L., Milan, A., Schindler, K., Cremers, D., Reid, I., Roth, S.: Tracking the trackers: an analysis of the state of the art in multiple object tracking. *arXiv preprint arXiv:1704.02781* (2017) [2](#)
12. Li, M., Wang, Y., Ramanan, D.: Towards Streaming Perception. In: *Proceedings of the European Conference on Computer Vision (ECCV)* (2020) [1](#), [2](#), [3](#), [4](#), [6](#), [9](#), [10](#)
13. Lin, J., Rao, Y., Lu, J., Zhou, J.: Runtime Neural Pruning. In: *Proceedings of the 31st International Conference on Neural Information Processing Systems (NeurIPS)*. pp. 2178–2188 (2017) [3](#)
14. Lin, S.C., Zhang, Y., Hsu, C.H., Skach, M., Haque, M.E., Tang, L., Mars, J.: The Architectural Implications of Autonomous Driving: Constraints and Acceleration. In: *Proceedings of the 23rd International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*. pp. 751–766 (2018) [4](#)
15. Liu, T.Y., et al.: Learning to rank for information retrieval. *Foundations and Trends® in Information Retrieval* **3**(3), 225–331 (2009) [6](#)
16. Liu, W., Anguelov, D., Erhan, D., Szegedy, C., Reed, S., Fu, C.Y., Berg, A.C.: SSD: Single shot multibox detector. In: *Proceedings of the European Conference on Computer Vision (ECCV)*. pp. 21–37. Springer (2016) [1](#)
17. Luo, J.H., Wu, J., Lin, W.: Thinet: A Filter Level Pruning Method for Deep Neural Network Compression. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 5058–5066 (2017) [3](#)

18. Milan, A., Leal-Taixé, L., Reid, I., Roth, S., Schindler, K.: Mot16: A benchmark for multi-object tracking. arXiv preprint arXiv:1603.00831 (2016) [2](#)
19. Mittal, G., Liu, C., Karianakis, N., Fragoso, V., Chen, M., Fu, Y.: HyperSTAR: Task-Aware Hyperparameters for Deep Networks. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 8736–8745 (2020) [3](#)
20. Nishi, K., Shimosaka, M.: Fine-Grained Driving Behavior Prediction via Context-Aware Multi-Task Inverse Reinforcement Learning. In: 2020 IEEE International Conference on Robotics and Automation (ICRA). pp. 2281–2287. IEEE (2020) [8](#)
21. NVIDIA: Tensor RT. <https://developer.nvidia.com/tensorrt> [10](#)
22. Pang, J., Qiu, L., Li, X., Chen, H., Li, Q., Darrell, T., Yu, F.: Quasi-dense similarity learning for multiple object tracking. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 164–173 (2021) [9](#)
23. Park, E., Ahn, J., Yoo, S.: Weighted-entropy-based Quantization for Deep Neural Networks. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 5456–5464 (2017) [3](#)
24. Pleskac, T.J., Busemeyer, J.R.: Two-stage dynamic signal detection: a theory of choice, decision time, and confidence. *Psychological review* **117**(3), 864 (2010) [3](#)
25. Shen, H., Han, S., Philipose, M., Krishnamurthy, A.: Fast Video Classification via Adaptive Cascading of Deep Models. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 3646–3654 (2017) [3](#)
26. Tan, M., Pang, R., Le, Q.V.: EfficientDet: Scalable and Efficient Object Detection. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2020) [1](#), [3](#), [4](#), [10](#)
27. Wang, X., Yu, F., Dou, Z.Y., Darrell, T., Gonzalez, J.E.: Skipnet: Learning Dynamic Routing in Convolutional Networks. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 409–424 (2018) [3](#)
28. Waymo Inc.: Waymo Open Dataset. <https://waymo.com/open/> [2](#), [3](#), [4](#), [5](#), [9](#), [12](#)
29. Xu, Y., Wang, Y., Zhou, A., Lin, W., Xiong, H.: Deep Neural Network Compression with Single and Multiple Level Quantization. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 32 (2018) [3](#)
30. Yogatama, D., Mann, G.: Efficient transfer learning method for automatic hyperparameter tuning. In: Artificial intelligence and statistics. pp. 1077–1085. PMLR (2014) [6](#)
31. Zhao, R., Hu, Y., Dotzel, J., De Sa, C., Zhang, Z.: Improving Neural Network Quantization without Retraining using Outlier Channel Splitting. In: International Conference on Machine Learning (ICML). pp. 7543–7552 (2019) [3](#)
32. Zhou, X., Koltun, V., Krähenbühl, P.: Tracking objects as points. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 474–490. Springer (2020) [6](#)