

Multi-resolution Rescored ByteTrack for Video Object Detection on Ultra-low-power Embedded Systems

*Luca Bompani¹ Manuele Rusci² Daniele Palossi^{3,4} Francesco Conti¹ Luca Benini^{1,4}

¹ Department of Electrical, Electronic and Information Engineering, University of Bologna, Italy

² Department of Electrical Engineering, KU Leuven, Belgium

³Dalle Molle Institute for Artificial Intelligence, USI-SUPSI, Switzerland

⁴Integrated Systems Laboratory, ETH Zürich, Switzerland

*luca.bompani5@unibo.it

Abstract

This paper introduces *Multi-Resolution Rescored ByteTrack (MR2-ByteTrack)*, a novel video object detection framework for ultra-low-power embedded processors. This method reduces the average compute load of an off-the-shelf Deep Neural Network (DNN) based object detector by up to $2.25\times$ by alternating the processing of high-resolution images (320×320 pixels) with multiple down-sized frames (192×192 pixels). To tackle the accuracy degradation due to the reduced image input size, MR2-ByteTrack correlates the output detections over time using the ByteTrack tracker and corrects potential misclassification using a novel probabilistic Rescore algorithm. By interleaving two down-sized images for every high-resolution one as the input of different state-of-the-art DNN object detectors with our MR2-ByteTrack, we demonstrate an average accuracy increase of 2.16% and a latency reduction of 43% on the GAP9 microcontroller compared to a baseline frame-by-frame inference scheme using exclusively full-resolution images. Code available at: https://github.com/Bomps4/Multi_Resolution_Rescored_ByteTrack

1. Introduction

A Video Object Detection (VOD) algorithm identifies objects in a video stream by reporting their categories and the corresponding coordinates in the image space. This function is critical for various cyber-physical systems like surveillance cameras [23] and miniaturized (as big as the palm of a hand) drones [1, 8]. These systems typically integrate a camera with an ultra-low-power embedded processor, i.e., a low-cost Microcontroller Unit (MCU). Compared to other edge/mobile processors (e.g., Nvidia Tegra, Raspberry Pi, Qualcomm Snapdragon, etc.), MCUs feature a 10-100 \times power consumption to enable battery-powered

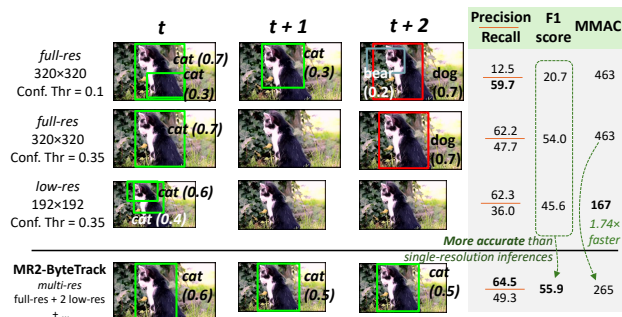


Figure 1. Video object detection using NanoDet-Plus [16] object detector under multiple thresholds and input size settings vs. the proposed MR2-ByteTrack solution.

operations (also referred to as the *extreme edge computing* [15]). On the other side, these devices present a limited on-chip memory (typically a few MB) and low computational power [19]. These severe limitations challenge the design of high-throughput VOD systems using compute-intensive accurate vision algorithms.

Recently, VOD was shown on multi-core MCUs [19] using lightweight (i.e., with a low memory budget) Deep Neural Network-based (DNN) object detectors with a simple frame-by-frame analysis [1, 7, 8]. For every frame of the image stream, the DNN inference task returns a list of output detections with confidence scores between 0 and 1. For example, Fig. 1 shows the output of the recent NanoDet-Plus object detector [16] featuring 1.17 million parameters when applied over three consecutive video frames. The precision-recall tradeoff is adjusted by thresholding the confidence scores, as shown in the figure: a higher threshold value leads to fewer false positives (higher precision) or undesired miss-detections (lower recall). From a computational cost perspective, a common strategy to reduce the number of operations per frame, and thus increasing the

processing throughput, consists of feeding the model with low-resolution (low-res) input images [17, 30], e.g., ~ 296 M Multiply-Accumulate (MAC) using $2.8 \times$ smaller images in our example.

To tackle this problem, this paper proposes *Multi-Resolution Rescored ByteTrack (MR2-ByteTrack)*, a method that transforms an image object detection model into a full VOD algorithmic pipeline, simultaneously improving the detection accuracy and reducing the overall computational cost with respect to a naïve *frame-by-frame* processing pipeline. MR2-ByteTrack combines an off-the-shelf DNN object detector with an extended version of *ByteTrack* [28]. This lightweight Kalman-filter-based tracker updates the output detections based on the object instances already detected by the DNN. Because the original *ByteTrack* [28] lacks recovery mechanisms from misclassification (e.g., a cat is tracked as a “dog” if wrongly predicted at the beginning of the video sequence), we propose a novel *Rescore* algorithm to refine classifications over time, according to a probabilistic logic.

Our approach adopts a *Multi Resolution Inference* scheme to reduce the per-frame workload, inserting low-resolution (low-res) frames between full-resolution (full-res). As shown at the bottom of Fig. 1, our MR2-ByteTrack leverages the temporal correlation of detections in consecutive frames to recover the accuracy drop of low-resolution inferences. When deployed on a multi-core MCU, this method incurs no additional memory cost for parameter storage compared to a single-resolution inference scheme, as the same object detection DNN is applied to both full and low-resolution frames.

In summary this work makes the following novel contributions:

- The MR2-ByteTrack framework, combining an off-the-shelf DNN-based object detector for multi-resolution inference, the *ByteTrack* Kalman-based tracker and the *Rescore* method to refine category assignment of tracked frames, reducing misdetections and misclassifications;
- The first embodiment of a multi-resolution VOD pipeline on a ultra-low-power MCU, Greenwaves Technologies (GWT)’s GAP9, showing a better accuracy-throughput tradeoff than single-resolution inference VOD systems;

We validate our proposed MR2-ByteTrack method using three State-of-the-Art (SoA) tiny object detectors: YOLOX-Nano [4], NanoDet-Plus [16], and EfficientDet-D0 [24], all off-the-shelf pretrained on the COCO dataset [10]. Across the ImageNetVid [20] validation set, MR2-ByteTrack improves mean Average Precision (mAP) scores compared to the object detector baselines by up to 5.17% and the F1 score by 3.58% if applied only on full-res frames. If, instead, we interleave for each full-res frame two low-res frames, the mAP improvement is +2.16% while the average MAC computational costs are reduced by up to 43%.

We deployed our multi-resolution VOD strategy, using the on the GAP9 MCU processor featuring a 9-core compute cluster at a peak power cost of up to 72 mW, also considering the external 32 MB and 8 MB MB FLASH and RAM. Compared to the baseline object detectors [4, 16, 24], MR2-ByteTrack achieves up to $1.76 \times$ lower inference latency with no increase in the DNN’s parameters footprint (2.24 MB for the largest model) and only a modest increase in code size (+186 kB), thus enabling high-accuracy VOD on milliwatt-power extreme-edge devices. Also, when compared to the SoA YOLOV method [22], our MR2-ByteTrack shows the same F1 score but saves the memory and compute overheads (14% and 21%) of the transformer-based front-end that differently, from our *ByteTrack* module, demands an ad-hoc training process.

2. Related Work

2.1. Lightweight Object Detection

DNN-based object detection algorithms can be grouped into two main categories: *two-stage* and *one-stage* detectors. The two-stage methods, e.g., Faster R-CNN [18], generate a set of region proposals that are analyzed during a second inference stage. This dynamic and unpredictable workload, as it depends on the generated proposals, makes these methods unsuitable for resource-constrained embedded devices, where low computation and predictable execution time are crucial to optimally distribute the limited resources. Consequently, many recent works focus on one-stage memory-efficient DNNs, generally composed of a feature extractor and multiple head blocks that predict the bounding box coordinates [4, 16, 21, 24, 27].

SSDLite [21] adds detector heads on top of a MobileNetV2 backbone and achieves an mAP of 22.1 on the COCO dataset [10] with a total memory and computational cost of 4.3 M parameters and 0.8 GMAC, respectively. An instance of this model has been successfully demonstrated on an MCU device with a peak power of 150 mW for object detection aboard nano-drones [8] or license plate detection for smart cameras [7], showing a peak throughput of ~ 1 Frame Per Second (FPS). NanoDet-Plus [16] introduces an “assign guidance module” and a dynamic soft label assigner to improve performance, reaching a mAP of 27 on the COCO dataset with 0.45 GMAC operations and 1.17 M parameters.

YOLOX-Nano [4] starts from the YOLOv3 architecture and introduces depthwise convolutions and an anchor-free decoupled head, which lead to a mAP of 25.4 on COCO with 0.8 GMAC and 0.9 M parameters. Instead, EfficientDet-D0 [24], leveraging a neural architecture search technique, introduces a bidirectional feature pyramid network to aggregate features that, in conjunction with a decoupled head, pushes the mAP to 34.6 at the cost of

2.5 GMAC and 3.9M parameters. Our work builds on these recent frame-by-frame models to construct a novel framework operating on video streams.

2.2. Video Object Detection

A first class of techniques for video object detection uses 3D convolutions instead of 2D ones [25]. This conceptually simple approach is not real-time and demands a large memory footprint to store high-dimensional activation tensors, exceeding the memory available on low-end MCU systems. An alternative line of research aggregates information across frames from region proposals obtained by two-stage detectors: in [31] using optical flow, or convolutional-based trackers in [13, 29], and transformer-based memory layers in [3]. However, the cost of region proposal aggregation exceeds the available budget for MCUs.

A more lightweight group of techniques builds on top of one-stage detectors by modeling the temporal correlation of detections. [5] uses dynamic programming to link objects across frames based on the Intersection over Union (IoU) metric. *Seq-Bbox* [2] and *Motion-based Seq-Bbox* [9] refine the outputs of the underlying object detector using the temporally associated detection information, marking a mAP of 80.9 and 72.7, respectively, on the ImagenetVID dataset (+6.9 and +5.5 mAP points more than frame-by-frame inference). Conversely, our work introduces a novel probabilistic algorithm for aggregating scores over time, while these methods make the average and, differently from our solution, do not achieve any thorough improvement concerning the baseline.

More recently, transformer architectures have also been widely adopted for VOD tasks. The TransVOD Lite [6] uses Deformable DETR object detector [32] and scores an mAP50 of 83.7 on ImagenetVID with a SwinT transformer backbone [12]. The latter features a total of 45.9M parameters and a computational cost of 8.89 GMAC per frame, exceeding the resources of an MCU system. In the best configuration, TransVOD Lite achieves +3.5 mAP score vs. the object detector baseline by processing a batch of 15 frames. This setting requires 272 MB to store the intermediate features ($15 \times$ memory increases compared to single frame execution). If applied on individual frames, the method, instead, reduces the mAP by 3.6 w.r.t. the Deformable DETR detector. On the contrary, our method (i) does not store multiple features across frames but only the bounding boxes (9.5 kB) and (ii) leads to a mAP increase vs. the baseline when running frame-by-frame, i.e., it does not require a batch of images for the inference.

YOLOV [22] uses a transformer on top of a YOLOX detector to fuse the features extracted from multiple frames and refine the produced detection. This approach reaches a mAP of 85.5 on ImagenetVID. However, tuning the transformer requires a training stage, marking a significant dif-

ference w.r.t. our training-free method, which can be applied to any pre-trained model without additional computation.

2.3. Increasing VOD Throughput

A set of methods proposed data-dependent DNN execution pipelines to skip part of the inference operations based on the information carried by the input. The work of *Zhu et al.* [30] introduces an auxiliary but lightweight network, called dynamic-resolution network (DRNet), to determine the resolution of the input image before running the inference. The policy predicts a higher downsize ratio for images with “easy” objects to recognize. Concerning this work, we adopt a static resize policy avoiding the 4.9M parameters overhead introduced by their DRNet, while achieving similar efficiency gains ($\sim 40\%$).

On the contrary, other works speed up the processing by considering the temporal information and not only the content of a single frame. [26] avoids the computation of the features from specific regions of the current frame by copying the features computed in previous frames. They use an auxiliary lightweight ResNet-8 network to predict the accuracy gain of computing new features vs. copying the previous ones. Compared to this approach, we do not store the activations but only preserve the detected bounding box coordinates from the analysis of previous ones. In the case of the NanoDet-Plus network with a 320×320 px input image, this method would lead to a memory overhead of ~ 50 MB vs. the 9.5 kB of our method. Furthermore, the ResNet-8 auxiliary network requires an extra training procedure using reinforcement learning to learn an efficient policy.

Liu et al. [11] couple a large model with a small one (4.4 M vs. 0.5 M parameters). Similar to our approach, they also interleave the two models’ execution, and to recover from accuracy losses due to the smaller model, they placed an LSTM layer to connect the detection performed across time. A similar idea is also employed in [14], where they interleave different neural networks of varying complexity to determine the pose of a person in front of a drone. These methods deploy the two models on the target platforms, increasing the memory cost compared to the baseline single-model inference. On the contrary, our framework does not require any extra parameters. It is applied to a pre-trained model, saving the cost of training a family of scalable models for the interleaved execution.

3. Multi-resolution Rescored ByteTrack

The MR2-ByteTrack algorithm in Fig. 2 combines three different components: (i) a vision-based Convolutional Neural Network (CNN) object detector that takes in input images of different sizes; (ii) the ByteTrack tracking algorithm [28]; (iii) the novel *Rescore* algorithm to update the detection scores returned by ByteTrack.

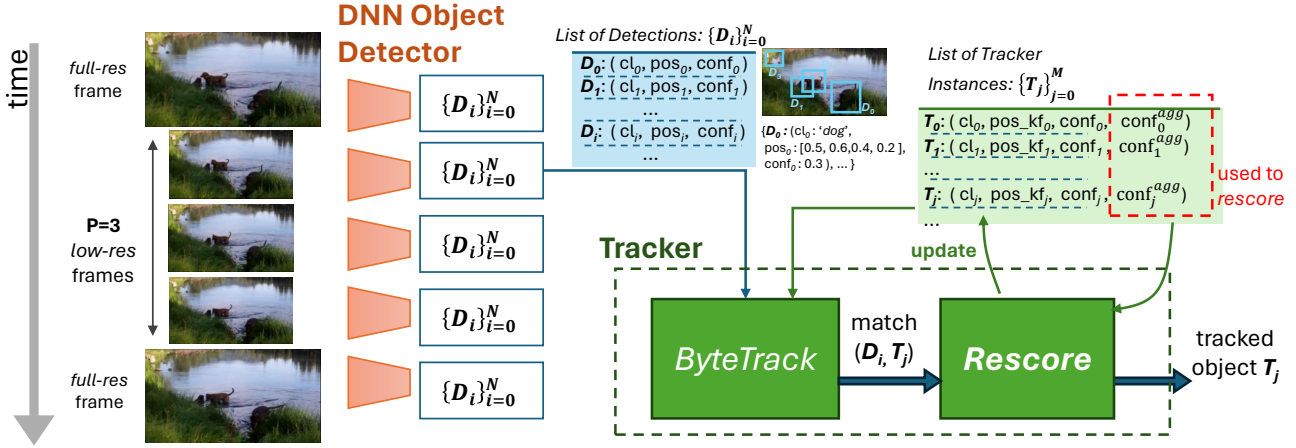


Figure 2. Overview of the proposed Multi-Resolution Rescaled ByteTrack algorithm for video object detection.

3.1. Multi-resolution Object Detector Backbone

We design our framework around existing pre-trained single-shot object detectors. In our formulation, the detection algorithm includes a lightweight CNN model, i.e., the backbone, to analyze the image frames of a video sequence in real time. Given a frame F at time t , we indicate the set of detections the processing task returns as $\{D_i\}_{i=1}^N$, where N is the total number of objects detected within the current frame. Each detection D_i is described by a triplet (bbox $_i$, cl $_i$, conf $_i$): a bounding box bbox, a class index cl and a confidence score conf < 1 . Detections with a confidence score lower than a *low_threshold* are discarded, while the others are retained for further processing.

Since we employ a CNN-based backbone, we can interleave one full-res image (e.g., 320×320 px in our setup) with P low-res ones (e.g., 192×192 px). This way, we can trade detection accuracy for fewer MAC operations. Thus, the average number of MAC operations per frame becomes:

$$\text{MAC} = \rho \cdot \text{MAC}^{fr} + (1 - \rho) \cdot \text{MAC}^{lr} \quad (1)$$

Where $\rho = \frac{1}{1+P}$, MAC^{fr} and MAC^{lr} represent the number of MAC operations required to execute the inference on one full-res and one low-res frame, respectively. Because the inference task *utilizes the same network parameters* in both low and full-resolution cases, the weight footprint remains the same as with single-sized frames. In fact, the convolution layers of the CNN models can apply to any scale of input images or feature maps. As an example, a 3x3 filter can slide over a larger feature map to produce a wider output tensor, corresponding, eventually, to a higher number of detections.

3.2. ByteTrack Tracker

Naïve frame-by-frame object detectors can not exploit any temporal correlation between frames, i.e., detections in con-

secutive frames are likely to stay the same or similar. We use the ByteTrack [28] tracker to exploit this correlation, recovering for performance losses due to down-scaling (i.e., low-res images) and possibly improving the precision and recall of a frame-by-frame baseline (see Sec. 5.3). The backbone’s detections $\{D_i\}_{i=1}^N$ are fed to the ByteTrack algorithm based on a Kalman filter for every frame. In its original formulation, at time-step t , ByteTrack operates on a list of M instances: $\{T_j\}_{j=1}^M$, i.e., one instance for each tracked object. Like for detections, each tracker instance (T_j) is described by a triplet (bbox $_{kf_j}$, cl $_j$, conf $_j$). bbox $_{kf_j}$ is the bounding box predicted by a Kalman filter fed with the backbone’s detections. cl $_j$ and conf $_j$ are the class index and confidence scores associated with tracked objects, respectively.

To correlate the detection performed over time, ByteTrack matches the bbox bounding boxes of $\{D_i\}_{i=1}^N$ and the predicted bbox $_{kf}$ ones of $\{T_j\}_{j=1}^M$, using a cost matrix determined by Intersection-over-Union (IoU) scores. An IoU threshold of 0.3, as proposed in the original paper, sets the minimum value to determine a match between a detection and a tracker instance. Unlike the original implementation, we do not adopt any feature-matching criteria to associate tracker instances and detections: the spatial size of the feature maps to match changes over time in our multi-resolution framework. A tracker instance is marked active after a minimum number of consecutive detections has occurred (two in our setting). At the same time, it is removed from the tracker list if it was never updated during the last five time steps. After matching, the detection bbox updates ByteTrack’s Kalman filter. The remaining detections not matched with any existing tracker instance produce a new one if their conf score is higher than a tuned *high_threshold*.

In the original ByteTrack algorithm, the class index cl $_j$ and the confidence score conf $_j$ of active tracker instances

Algorithm 1: Rescore algorithm

```
Inputs: Match
{ $D_i = (cl_i, conf_i), T_j = (cl_j, conf_j, conf_j^{agg})$ }
if  $cl_j == cl_i$  then
|  $conf_j^{agg} \leftarrow 1 - ((1 - conf_i) * (1 - conf_j^{agg}))$ 
else
| if  $conf_j^{agg} < conf_i$  then
| |  $(cl_j, conf_j, conf_j^{agg}) \leftarrow (cl_i, conf_i, conf_i)$ 
| else
| |  $conf_j^{agg} \leftarrow 1 - ((1 - conf_j^{agg}) / (1 - conf_i))$ 
| |  $conf_j^{agg} \leftarrow \max(conf_j^{agg}, 0)$ 
| | if  $conf_j^{agg} < conf_i$  then
| | |  $(cl_j, conf_j, conf_j^{agg}) \leftarrow (cl_i, conf_i, conf_i)$ 
| | end
| end
end
 $conf_j^{agg} \leftarrow \min(conf_j^{agg}, 1 - \epsilon)$ 
return  $T_j = (cl_j, conf_j, conf_j^{agg})$ 
```

are never updated after the match with a new detection. We chose to modify this logic to better cope with the multi-class nature of VOD: specifically, we update cl_j and $conf_j$ if the confidence score of the matched detection is higher than the *high_threshold* value.

3.3. Rescore Algorithm

We augment ByteTrack with the *Rescore* algorithm to improve the estimation of the class index and confidence scores of an active tracker instance by accounting for the history of matched detections. To this aim, we extend the tracker instance T_j with a class status attribute $conf_j^{agg}$, which models the probability of the tracker’s class cl to be correct. The $conf_j^{agg}$ value aggregates the confidence scores $conf$ of the detections assigned to the j -th tracker instance until time $t - 1$. During this aggregation, the $conf$ value is interpreted as the probability that the object’s predicted class is correct in the current-frame detection.

When a new detection D_i is assigned to the tracker instance T_j at time t , the class index cl_j is *rescored* based on $conf_j^{agg}$ as shown in Algo. 1. If the class index cl_i is the same as cl_j , we only update the $conf_j^{agg}$ value, as the match confirms our confidence in the correctness of the classification. We use the product of the opposites of $conf_j^{agg}$ and $conf_i$ as the new value for the status variable. If the class indices do not match, we acknowledge this by decreasing our aggregate confidence, i.e., we decrease the $conf_j^{agg}$ value (the max guarantees a value >0). Then, we check for a class index rescore: if the new detection’s confidence is higher than the $conf_j^{agg}$, we update the tracker class with the new detection: $cl_j \leftarrow cl_i$. Additionally, we impose an upper bound (i.e., $1 - \epsilon$) to the $conf_j^{agg}$ to prevent overconfident detections from blocking future rescors. Finally, the *Rescore* algorithm also outputs the mean confidence score

during the tracking.

4. Multi-resolution Inference on MCUs

We target ultra-low-power embedded systems, such as tiny MCUs, to deploy our pipeline. Since the end-to-end latency of VOD applications is dominated by DNN inference, we extensively benefit from the throughput optimization of Eq. 1. To demonstrate our approach’s real-time performance and power consumption, we select a GWT GAP9 MCU, which includes a parallel cluster of 9 general-purpose RISC-V cores. The cores can access 4 shared floating point units with vectorized half-precision instructions (FP16 precision). GAP9 features a low-latency scratchpad L1 memory of 128 kB within the cluster and a larger off-cluster L2 memory (1.5 MB).

A Direct Memory Access (DMA) controller is used to copy data between L2 and L1 memories in the background of the compute tasks offloaded to the cores. External volatile and non-volatile memories are interfaced with the GAP9 using the available octoSPI interface: a peripheral DMA is then in charge of read and write operations between the external components and the on-chip L2 memory.

The cluster also includes a convolution accelerator that supports 8-bit operations. However, we only exploit the floating-point multi-core support of the platform to avoid any accuracy degradation of lossy 8-bit quantization, which we leave for future work.

We use GWT’s *GAPflow* toolset to generate the target platform’s inference C code. For every trained model with a specific input size, the tool produces header files with the model parameter and a source file with the graph-level and layer-level routines. Following the automatically generated memory management scheme, weight parameters are stored within an external octoSPI Flash memory. At the same time, the intermediate values of the inference tasks, i.e., the activation tensors, are allocated within the on-chip L2 memory and the external octoSPI RAM. To cope with our multi-resolution scheme, we generate two C implementations, one for the low-res case and a second for the full-res one. The two network functions access the same weight data stored in the Flash memory.

On GAP9, the latency cost to downscale the input image is negligible with respect to the inference time. The overhead for generating a low-res image of 192×192 px from a 320×320 px image is 1.5 ms. This corresponds to only 0.9% of the inference time in the case of NanoDet-Plus (169 ms). The Kalman filter also has low execution time and requires only 9.5 kB of on-chip memory. While our implementation of the MR2-ByteTrack method is specific to the GAP9, we remark on the generality of the method that can be easily applied to other MCUs.

Table 1. Baseline Object Detection Models

Model	Params	full-res (px)	MMAC ^{fr}	low-res (px)	MMAC ^{lr}
YOLOX-Nano	0.9 M	320×320	316	192×192	114
NanoDet-Plus	1.18 M	320×320	463	192×192	167
EfficientDet-D0	3.93 M	384×384	1440	256×256	640

Table 2. Impact of the Rescore Algorithm on ByteTrack using only high-res frames ($P = 0$).

Method	mAP	Precision	Recall	F1 score
YOLOX-Nano	41.30	65.16	44.92	53.18
w/ Naïve-ByteTrack	43.57	63.9	47.96	54.79
w/ MR2-ByteTrack	46.47	64.24	51.23	57.00
NanoDet-Plus	42.70	62.15	47.68	54.00
w/ Naïve-ByteTrack	44.00	58.91	51.40	54.90
w/ MR2-ByteTrack	46.57	64.14	51.50	57.10
EfficientDet-D0	60.70	78.48	64.14	70.30
w/ Naïve-ByteTrack	61.09	78.65	66.42	72.02
w/ MR2-ByteTrack	64.86	80.63	67.32	72.52

5. Results

5.1. Baseline Models and Metrics

We consider three State-of-the-Art object detectors trained on the COCO dataset: YOLOX-Nano, NanoDet-Plus, and EfficientDet-D0. Tab. 1 reports the number of parameters and MMAC operations for the input frames’ considered full-res and low-res sizes. For VOD experiments, we use the ImageNetVID dataset, which shares 16 classes with the COCO dataset. By extracting video sequences that share common classes between the two datasets, we build our validation set of 392 video samples. To distinguish from the entire dataset, we indicate this subset as ImageNetVID^C. To assess the detection performance, we use the mAP50 score, a widely used metric in literature, which we call mAP in the rest of the paper. In addition, we report the precision, recall, and F1 score averaged over the classes. Differently from mAP, which expresses the mean value of the per-class average-precision curves, these three additional metrics show the trade-off between true detections, false positives, and miss-detections. All metrics consider an IoU of 0.5 to mark a correct match between ground truths and predictions.

5.2. Single-resolution Trackers

Tab. 2 compares the results obtained with our MR2-ByteTrack against the baseline object detection models and the Naïve-ByteTrack on the ImageNetVid^C dataset. In this initial experiment, we feed the inference models only with full-res images. The confidence threshold of the baseline models is set to 0.35, 0.3, and 0.4 for, respectively, NanoDet-Plus, YOLOX-Nano, and EfficientDet-D0. These values are optimally tuned to obtain the highest F1 score on the testing dataset and reflected on the *high_threshold*




Figure 3. mAP (blue) and GMAC (red) of MR2-ByteTrack at varying low-res frames vs. the baseline.

value of the ByteTrackers. Using the same procedure, we also set the *low_threshold* of the trackers to 0.3, 0.25, 0.35 for NanoDet-Plus, YOLOX-Nano, and EfficientDet-D0, respectively. By observing the results, we can see that, on average, the Naïve-ByteTrack increases the mAP score measured on the baseline models by 1.55%. This is explained by the capacity of the tracker to positively aggregate low-confident detections, which the baseline models instead discard.

Next, the Rescore algorithm of MR2-ByteTrack leads to a further 3.08% improvement by turning false-positive detection into correct classification, as can be observed by the higher Precision, Recall, and F1 scores. Overall, thanks to our MR2-ByteTrack, the smallest NanoDet-Plus and YOLOX-Nano models can achieve up to 46.57 and 46.47 mAP on high-res frames, respectively, +3.9% and +5.2% vs. the baselines.

5.3. Multi-resolution Performance

Fig. 3 plots the measured mAP of MR2-ByteTrack when varying the amount of interleaved low-res images (P from 0 to 10) in the multi-resolution setup, and we benchmark against the baseline models in the same setup. The bars on the bottom show the average GMAC operations required by each configuration (Eq. 1). For baseline models, the mAP starts dropping at $P=1$, while MR2-ByteTrack decreases performance with $P \geq 5$ due to low-res images. This difference between methods arises from the greater robustness of our MR2-ByteTrack approach. This effect is related to the inertia of the trackers, which are kept alive for up to 5 timestamps without new detections in our setup.

Table 3. VOD solutions on the GAP9 MCU: Mr2-ByteTrack vs. frame-by-frame CNN object detectors with full-res or low-res inputs.

Method	Input Size	RAM [MB]	Flash Mem. [MB]	Code [kB]	[FPS]	Energy [mJ]	mAP	Precision	Recall	F1 score
<i>NanoDet-Plus</i>	single full-res	1.6	2.3	186.0	3.3	21.8	42.7	62.2	47.7	54.0
<i>NanoDet-Plus</i>	single low-res	0.6	2.3	186.0	9.8	7.5	27.4	62.9	31.3	41.8
w/ MR2-ByteTrack (our)	multi-res P=2	1.6	2.3	373.0	5.9	12.3	44.9	61.5	49.3	54.7
<i>YOLOX-Nano</i>	single full-res	1.4	1.8	103.0	3.3	19.2	41.3	65.2	44.9	53.2
<i>YOLOX-Nano</i>	single low-res	0.5	1.8	103.0	8.6	7.0	25.7	63.2	28.8	39.6
w/ MR2-ByteTrack (our)	multi-res P=2	1.4	1.8	205.0	5.5	11.1	41.4	64.0	45.7	53.3

Table 4. Comparison between VOD methods on ImagenetNetVID.

Method	Object detector Backbone	Prec.	Recall	F1 score	mAP	Params [M]	GMAC	vs. full-res object detector		
								Δ mAP	Δ params	Δ MAC
YOLOV-S [22]	YOLOX-S[4]	79.9	66.4	72.5	62.5	9.9	12.9	2.7	+14.0%	+21.0%
MR2-ByteTrack		80.8	65.9	72.5	61.8	9.0	10.9	2.0	0	0
MR2-ByteTrack (P=2)		77.5	64.3	70.3	60.3	9.0	6.2	0.5	0	-43.0%
<i>Liu et. al</i> [11]	SSDLite-Mobilenetv2[21]	n/a	n/a	n/a	61.4	4.9	0.2	0.9 ¹	+11.0% ¹	-84.0%
MR2-ByteTrack (P=2)	EfficientDet-D0 ³ [24]	80.0	64.1	71.1	61.2 ²	3.9	0.9	0.5 ²	0	-37.0%

¹ w.r.t. large f_0 non-interleaved model [11], ² measured on ImageNetVID^C, ³ trained on the COCO dataset,

Overall, the mAP scores of the MR2-ByteTrack models using multiple low-res images are always better than the baseline models operating on full-resolution frames. For $P=1$ and $P=2$, on average, our approach marks a mAP improvement of 3% and 0.96%, while the number of MAC operations is reduced by 32% and 43%, respectively, compared with the baseline object detectors. For $P \geq 3$, we observe a progressive reduction of mAP below the original object detection performance, which can be affordable for some use cases. Finally, we chose the configuration $P=2$ for the remainder of our experiments as a good trade-off between accuracy and computational load (i.e., MAC operations).

5.4. State of the Art Comparison

5.4.1 MCU-ready VOD Systems

Tab. 3 compares our MR2-ByteTrack solution with $P=2$ vs. frame-by-frame single-resolution baselines (e.g., used by [7, 8]) with full-res and low-res inputs when running on an MCU system. To this aim, we deploy the different approaches on the GAP9 SoC, coupled with two external memories: a Flash of 64 MB and a RAM of 8 MB. The table reports the code size, the activation and weight memory occupation in MB, denoted as RAM and FLASH Mem, the processing throughput in FPS, and the energy cost, measured when the MCU is clocked at 370 MHz, together with the detection metrics.

We show results for YOLOX-Nano and NanoDet-Plus, while we leave out of this analysis the EfficientDet, as its memory footprint is larger than the available budget. The datatype of weights and activations is cast to half-precision FP16 for a lossless deployment. On GAP9, our

multi-resolution instance features a double code size (up to 373 kB in total) vs. baseline models. The activation memory of the high-res model dominates the total cost of the multi-res deployment; the weight storage is constant among different resolutions. Our solutions can run at 5.9 FPS with NanoDet-Plus and 5.5 FPS with YOLOX-Nano, $1.7 \times$ faster than the high-res version. On the other side, the cost of the ByteTrack algorithm is negligible ($<0.2\%$ than the inference). The latency gains are reflected in the energy costs: the multi-resolution models consume $1.77 \times$ and $1.72 \times$ less than the high-res versions for, respectively, NanoDet-Plus and YOLOX-Nano. Thus, observing the highest mAP and F1 scores, *our MR2-ByteTracker shows the best energy-accuracy trade-off for VOD applications on MCU systems.*

5.4.2 VOD Algorithms

Tab. 4 compares our approach with recent real-time methods *not* tailored for MCU deployment given the high number of parameters of the DNN backbones: 9 M for YOLOX-S in the YOLOV work [22] and 4.9 M for SSDLite-MobileNetV2 in the study by *Liu et al.* [11]. To fairly compare vs. YOLOV, which uses a transformer to aggregate the temporal detections of the YOLOX-S backbone, we apply our MR2-ByteTrack to the same CNN object detector pre-trained on the ImageNetVid trainset, considering both the high-res or the multi-resolution ($P=2$) scenarios. For YOLOV, we test the real-time version in which only the detections of the past frames are sent to the transformers. The scores reported in the table adopt confidence threshold values that maximize the F1 score.

Our MR2-ByteTrack reaches a +0.9 higher precision

than the YOLOV, due to the lower number of false positives. At the same time, the lower recall (-0.5) acknowledges the recognition ability of the transformer layer, which increases the memory and computational cost by, respectively, 14 % and 21 % compared to the baseline backbones. Our method presents the same F1 score as the real-time YOLOV, saving the transformer overhead costs. Additionally, with $P=2$, the F1 score is reduced only by 2.2 % with a computational cost reduction of up to 43 %.

On the other side, because the pre-trained SSDLite-Mobilenetv2 model adopted by *Liu et al.* [11] is not openly available, we take the EfficientDet-D0 model pre-trained on COCO for comparison purposes (EfficientDet-D0 features 1 M parameters less than SSDLite-Mobilenetv2). Our MR2-ByteTrack method achieves a substantial latency improvement (37 % vs. 84 % of [11] that uses $P=10$) while slightly increasing the mAP score compared to the baseline. However, our solution does not increase the memory footprint (+11 % in [11]) and does not require any training of the tracker on the video sequence, highlighting the higher flexibility of our approach.

6. Conclusions

This paper proposes the novel MR2-ByteTrack method that (i) improves the accuracy of pre-trained object detection for VOD benchmarks by using the ByteTrack tracker augmented with the novel rescore algorithm and (ii) reduces the computational cost by running multi-resolution inference at no extra memory costs. Leveraging SoA memory-efficient DNNs, we achieved equal or slightly increased mAP score when compared Using SoA memory-efficient DNNs our method delivers up to a $1.76\times$ improvement in inference throughput on the GAP9 MCU, achieving equal or slightly increased mAP score. This result highlights the benefits of the proposed solution for VOD in ultra-low-power embedded devices.

Acknowledgements

This work has been partially supported by the Swiss SNF project 207913 TinyTrainer: On-chip Training for TinyML devices

References

- [1] Eiman AlNuaimi, Elia Cereda, Rafail Psiakis, Suresh Sugumar, Alessandro Giusti, and Daniele Palossi. A deep learning-based face mask detector for autonomous nano-drones (student abstract). In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 12903–12904, 2022. 1
- [2] Hatem Belhassen, Heng Zhang, Virginie Fresse, and El-Bay Bourennane. Improving video object detection by seq-bboxmatching. In *VISIGRAPP(5:VISAPP)*, pages 226–233, 2019. 3
- [3] Yihong Chen, Yue Cao, Han Hu, and Liwei Wang. Memory enhanced global-local aggregation for video object detection. pages 10334–10343, 2020. 3
- [4] Zheng Ge, Songtao Liu, Feng Wang, Zeming Li, and Jian Sun. Yolox: Exceeding yolo series in 2021. *arXivpreprint-arXiv:2107.08430*, 2021. 2, 7
- [5] Wei Han, Pooya Khorrami, Tom Le Paine, Prajit Ramachandran, Mohammad Babaeizadeh, Honghui Shi, Jianan Li, Shuicheng Yan, and Thomas S. Huang. Seq-nms for video object detection. *CoRR*, abs/1602.08465, 2016. 3
- [6] Lu He, Qianyu Zhou, Xiangtai Li, Li Niu, Guangliang Cheng, Xiao Li, Wenxuan Liu, Yunhai Tong, Lizhuang Ma, and Liqing Zhang. End-to-end video object detection with spatial-temporal transformers. In *Proceedings of the 29th ACM International Conference on Multimedia*, page 1507–1516, New York, NY, USA, 2021. Association for Computing Machinery. 3
- [7] Lorenzo Lamberti, Manuele Rusci, Marco Fariselli, Francesco Paci, and Luca Benini. Low-power license plate detection and recognition on a risc-v multi-core mcu-based vision system. In *2021 IEEE International Symposium on Circuits and Systems (ISCAS)*, pages 1–5. IEEE, 2021. 1, 2, 7
- [8] Lorenzo Lamberti, Luca Bompani, Victor Javier Kartsch, Manuele Rusci, Daniele Palossi, and Luca Benini. Bio-inspired autonomous exploration policies with cnn-based object detection on nano-drones. In *2023 Design, Automation & Testin Europe Conference & Exhibition (DATE)*, pages 1–6. IEEE, 2023. 1, 2, 7
- [9] Min Li, Linghan Li, Ruwen Bai, Junxing Ren, Bo Meng, and Yang Yang. A motion-based seq-bbox matching method for video object detection. In *2021 IEEE Symposium on Computers and Communications (ISCC)*, pages 1–7, 2021. 3
- [10] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C. Lawrence Zitnick. Microsoft coco: Common objects in context. In *Computer Vision – ECCV 2014*, pages 740–755, Cham, 2014. Springer International Publishing. 2
- [11] Mason Liu, Menglong Zhu, Marie White, Yinxiao Li, and Dmitry Kalenichenko. Looking fast and slow: Memory-guided mobile video object detection. *arXiv preprint arXiv:1903.10172*, 2019. 3, 7, 8
- [12] Z. Liu, Y. Lin, Y. Cao, H. Hu, Y. Wei, Z. Zhang, S. Lin, and B. Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 9992–10002, Los Alamitos, CA, USA, 2021. IEEE Computer Society. 3
- [13] Ye Lyu, Michael Ying Yang, George Vosselman, and Gui-Song Xia. Video object detection with a convolutional regression tracker. *ISPRS Journal of Photogrammetry and Remote Sensing*, 176:139–150, 2021. 3
- [14] Beatrice Alessandra Motetti, Luca Crupi, Mustafa Omer Mohammed Elamin Elshaigi, Matteo Risso, Daniele Jahier Pagliari, Daniele Palossi, and Alessio Burrello. Adaptive deep learning for efficient visual pose estimation aboard ultra-low-power nano-drones. *ArXiv*, abs/2401.15236, 2024. 3

- [15] Jorge Portilla, Gabriel Mujica, Jin-Shyan Lee, and Teresa Riesgo. The extreme edge at the bottom of the internet of things: A review. *IEEE Sensors Journal*, PP:1–1, 2019. 1
- [16] RangiLyu. Nanodet-plus superfast and high accuracy lightweight anchor-free object detection model. 2021. 1, 2
- [17] Joseph Redmon and Ali Farhadi. Yolov3: An incremental improvement, 2018. cite arxiv:1804.02767Comment: Tech Report. 2
- [18] S. Ren, K. He, R. Girshick, and J. Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. pages 1137–1149, Los Alamitos, CA, USA, 2017. IEEE Computer Society. 2
- [19] Davide Rossi, Francesco Conti, Manuel Eggiman, Alfio Di Mauro, Giuseppe Tagliavini, Stefan Mach, Marco Guermandi, Antonio Pullini, Igor Loi, Jie Chen, Eric Flamand, and Luca Benini. Vega: A ten-core soc for iot endnodes with dnn acceleration and cognitive wake-up from mram-based state-retentive sleep mode. *IEEE Journal of Solid-State Circuits*, 57(1):127–139, 2022. 1
- [20] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C. Berg, and Li Fei-Fei. ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision (IJCV)*, 115(3):211–252, 2015. 2
- [21] Mark Sandler, Andrew G. Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. Mobilenetv2: Inverted residuals and linear bottlenecks. *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4510–4520, 2018. 2, 7
- [22] Yuheng Shi, Naiyan Wang, and Xiaojie Guo. Yolov: Making still image object detectors great at video object detection. *Proceedings of the AAAI Conference on Artificial Intelligence*, 37(2):2254–2262, 2023. 2, 3, 7
- [23] Ernesto Sola-Thomas and Masudul Haider Imtiaz. An ultra-low-power design of smart wearable stereo camera. In *South-eastCon 2021*, pages 1–8, 2021. 1
- [24] M. Tan, R. Pang, and Q. V. Le. Efficientdet: Scalable and efficient object detection. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10778–10787, Los Alamitos, CA, USA, 2020. IEEE Computer Society. 2, 7
- [25] D. Tran, L. Bourdev, R. Fergus, L. Torresani, and M. Paluri. Learning spatiotemporal features with 3d convolutional networks. In *2015 IEEE International Conference on Computer Vision (ICCV)*, pages 4489–4497, Los Alamitos, CA, USA, 2015. IEEE Computer Society. 3
- [26] Thomas Verelst and Tinne Tuytelaars. Blockcopy: High-resolution video processing with block-sparse feature propagation and online policies. In *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 5138–5147, 2021. 3
- [27] Yunyang Xiong, Hanxiao Liu, Suyog Gupta, Berkin Akin, Gabriel Bender, Yongzhe Wang, Pieter-Jan Kindermans, Mingxing Tan, Vikas Singh, and Bo Chen. Mobiledets: searching for object detection architectures for mobile accelerators. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3825–3834, 2021. 2
- [28] Yifu Zhang, Pei Sun, Yi Jiang, Dongdong Yu, Zehuan Yuan, Ping Luo, Wenyu Liu, and Xinggang Wang. Bytetrack: Multi-object tracking by associating every detection box. In *European Conference on Computer Vision*, 2021. 2, 3, 4
- [29] Zheng Zhang, Dazhi Cheng, Xizhou Zhu and Stephen Lin, and Jifeng Dai. Integrated object detection and tracking with tracklet-conditioned detection. *ArXiv*, abs/1811.11167, 2018. 3
- [30] Mingjian Zhu, Kai Han, Enhua Wu, Qiulin Zhang, Ying Nie, Zhenzhong Lan, and Yunhe Wang. Dynamic resolution network. In *Neural Information Processing Systems*, 2021. 2, 3
- [31] Xizhou Zhu, Yujie Wang, Jifeng Dai, Lu Yuan, and Yichen Wei. Flow-guided feature aggregation for video object detection. In *2017 IEEE International Conference on Computer Vision (ICCV)*, pages 408–417, 2017. 3
- [32] Xizhou Zhu, Weijie Su, Lewei Lu, Bin Li, Xiaogang Wang, and Jifeng Dai. Deformable {detr}: Deformable transformers for end-to-end object detection. In *International Conference on Learning Representations*, 2021. 3