

EdgeDuet: Tiling Small Object Detection for Edge Assisted Autonomous Mobile Vision

Zheng Yang, *Senior Member, IEEE*, Xu Wang, *Member, IEEE*, Jiahang Wu, *Student Member, IEEE*, Yi Zhao, *Student Member, IEEE*, Qiang Ma, *Member, IEEE*, Xin Miao, *Member, IEEE*, Li Zhang, *Member, IEEE*, Zimu Zhou, *Member, IEEE*

Abstract—Accurate, real-time object detection on resource-constrained devices enables autonomous mobile vision applications such as traffic surveillance, situational awareness, and safety inspection, where it is crucial to detect both small and large objects in crowded scenes. Prior studies either perform object detection locally on-board or offload the task to the edge/cloud. Local object detection yields low accuracy on small objects since it operates on low-resolution videos to fit in mobile memory. Offloaded object detection incurs high latency due to uploading high-resolution videos to the edge/cloud. Rather than either pure local processing or offloading, we propose to detect large objects locally while offloading small object detection to the edge. The key challenge is to reduce the latency of small object detection. Accordingly, we develop EdgeDuet, the first edge-device collaborative framework for enhancing small object detection with tile-level parallelism. It optimizes the offloaded detection pipeline in tiles rather than the entire frame for high accuracy and low latency. Evaluations on drone vision datasets under LTE, WiFi 2.4GHz, WiFi 5GHz show that EdgeDuet outperforms local object detection in small object detection accuracy by 233.0%. It also improves the detection accuracy by 44.7% and latency by 34.2% over the state-of-the-art offloading schemes.



1 INTRODUCTION

Bringing advanced machine vision to mobile devices such as drones and robots enables a wide spectrum of autonomous mobile vision applications. Examples include mobile phones for localization [1] and navigation [2], drones for cost-effective traffic surveillance [3], and robot dogs to enforce social distancing during the COVID-19 pandemic [4]. Crucial in these applications is the capability to detect objects from video inputs. An ideal object detection engine for autonomous mobile vision applications should be *accurate*, *real-time*, and *resource-efficient*. (i) Drones and robots should accurately detect a large number of big and small objects in the scene (e.g., vehicles and pedestrians in an aerial view of a busy street). (ii) Fast object detection (i.e., 30-60 fps) on continuous videos enables decision-making on the go. For instance, a robot may identify the crowd density from live videos and broadcast alerts when moving in a park. (iii) resource-efficient: For portability and mobility, the computation and memory resources in commercial drones are still limited. Object detection algorithms need to be optimized to fit in the resource budgets of mobile devices.

Existing object recognition solutions for resource-limited devices fail to satisfy the accuracy and real-time requirements. For example, in mobile AR/MR scenarios, their systems require to run at 60fps and 2K resolution for the

object detection or object segmentation tasks [5]. For the huge computing resources requirements, the commercial mobile AR solutions such as ARKit [6] and ARCore [7] only track the pose of the camera and fail to track the locations of moving targets in real-time, which results in unnatural virtual effects [8]. (i) One promising approach for fast object detection is to run the model locally on-board. Model compression techniques can dramatically reduce the workload of deep learning models [9]. However, local object detection with compressed models is sub-optimal for autonomous mobile vision because accurate small object detection requires high-resolution input [10], which easily overwhelms mobile memory. (ii) An alternative is to offload object detection to the edge, which utilizes the powerful edge to run large models on high-resolution inputs for accurate detection. Nevertheless, offloading incurs a long delay since it involves the wireless transmission of high-resolution videos to the edge (i.e., a 2000kbps 2k video with 10Mbps wireless network bandwidth means 200ms delay). Long end-to-end detection delay leads to large detection errors as the mobile device's view is constantly changing [11].

Pioneer studies [5], [11] avoid transmitting every frame by using cached detection results of previous frames to track objects in the current frame and only offloading key frames to update the cached results. This “detect+track” strategy supports real-time object detection in case of high bandwidth networks. Its performance tends to deteriorate in the case of low-bandwidth, e.g., outdoors, which autonomous mobile vision applications often target at.

Instead of pure local processing or offloading, we propose to split the object detection task between the mobile device and the edge. Specifically, we offload *small object detection* to the edge. The rationale is intuitive. Commercial mobile devices are now able to accurately and rapidly de-

- Zheng Yang, Xu Wang, Jiahang Wu, Yi Zhao, Qiang Ma, Xin Miao are with the School of Software and TNLIST, Tsinghua University, China. Email: {yangzheng, wangxu2020}@tsinghua.edu.cn, wujx20@mails.thu.edu.cn, zhaoyi.yuan31@gmail.com, {maq, xin}@greenorbs.com
- Li Zhang is with HeFei University of Technology, HeFei, Anhui, China. Email: lizhang@hfut.edu.cn
- Zimu Zhou is with School of Information Systems, Singapore Management University, Singapore. Email: zimuzhou@smu.edu.sg
- Zheng Yang is the corresponding author

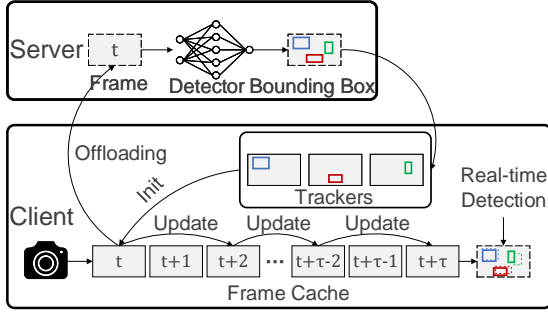


Fig. 1. An illustration of the popular “detect+track” framework for offloaded object detection. The detection results of the current frame are obtained by applying trackers on the cached detection results. The cached results are routinely updated by offloading key frames.

detect large- to medium-sized objects by running compressed models on low-resolution videos [12], [13]. Hence only data relevant to small objects need to be uploaded to the edge in high quality, thus reducing the overall offloading delay and improving detection accuracy. The “small” size is empirically tuned to achieve optimal performance, and we will detail the classification method later.

Realizing the above idea for accurate and real-time object detection needs a systematic design on (i) *what and how to offload to the edge* and (ii) *how to aggregate the detection results*. We base our design upon “detect+track” (Fig. 1), the prevailing framework, to accelerate offloaded object detection [5], [11]. The detection results of the current frame are obtained by adapting cached detection results of prior frames using lightweight trackers [14]. The cached results are routinely updated by offloading key frames for expensive yet highly accurate object detection. In our case, the trackers and the detectors for big objects are lightweight. Hence the bottleneck for real-time detection is the offloaded small object detection. Since the detection results of the current frame rely on the cached results, the bottleneck for accurate detection, especially for small objects, lies in the freshness of the cached results.

We propose EdgeDuet, an accurate, real-time object detection engine, which tiles and offloads small object detection to the edge (Fig. 2). EdgeDuet tackles the aforementioned accuracy and real-time bottlenecks via the following techniques. (i) *Optimizing offloaded small object detection with region-of-interest (RoI) frame encoding and content-prioritized tile offloading*. EdgeDuet applies RoI frame encoding to save network traffic. Only pixel blocks potentially containing small objects are transmitted in high quality, while the rest of the frame is compressed to low quality. EdgeDuet adopts content-prioritized tile offloading to accelerate small object detection at the edge. It processes videos in the unit of tiles rather than the entire frame, so as to improve the parallelism of offloading. It also prioritizes the offloading of tiles containing more objects, so that the cached detection results of more objects are freshly updated. (ii) *Real-time tracking via cache management and adaptive tracker configuration*. EdgeDuet aggregates the detection results from the local and remote object detectors to obtain fresh and consistent cached results via a cache management mechanism. It also applies adaptive tracker configuration to

improve the resource efficiency and real-time performance of the trackers.

We implement EdgeDuet as a *cross-platform* framework and evaluate its performance with mobile phones on VisDrone [15], a public video dataset captured by drone-mounted cameras. Evaluations show that pure local object detection yields a detection accuracy in terms of only 0.096 for small objects, while EdgeDuet achieves an accuracy of 0.319 for small objects.

The main contributions of this work are summarized below.

- EdgeDuet is the first framework that enhances small object detection in crowded scenes via collaboration between the edge and the mobile device.
- We push the state-of-the-art offloaded object detection studies [5], [11] from task-level parallelism to tile-level parallelism, which notably reduces the offloading latency. EdgeDuet is a systematic design that enables accurate, real-time object detection on mobile devices even in the case of low network bandwidth.
- We implement EdgeDuet as a cross-platform framework. Evaluations on VisDrone [15] show that EdgeDuet improves the overall accuracy by 44.7% and the end-to-end latency by 34.2% over the state-of-the-art object detection offloading schemes [5], [11].

In the rest of this paper, we give an overview of EdgeDuet in Sec. 3 and elaborate on its functional modules in Sec. 4, Sec. 5 and Sec. 6. We present the implementation of EdgeDuet in Sec. 7 and the evaluations in Sec. 8. We review related work in Sec. 2 and finally conclude in Sec. 10.

2 RELATED WORK

Our work is relevant to the following categories of research.

Object Detection Models. Advances in deep learning have resulted in various accurate and fast object detection models such as two-stage models *e.g.*, Faster-RCNN [16] and one-stage models *e.g.*, YOLO [17]. Model compression and acceleration techniques [9], [18]–[22] can substantially reduce the computation workload of deep learning-based models. However, the compressed models suffer from low accuracy on small object detection if the input image/video is low in resolution [10]. For accurate and fast small object detection, customized models [10], [23], [24] have been developed to detect objects on sub-regions of the input image/video. However, these models are computed heavily. Our work also performs object detection on sub-regions. However, rather than design new object detection models, we exploit existing YOLO-family models of different capabilities [17] to process different sub-regions of video frames.

Edge/Cloud Offloading. A popular strategy to enable highly accurate object detection on resource-constrained mobile devices is to offload the compute-intensive object detection to the powerful edge/cloud server [2], [5], [11], [25]–[33]. However, offloading may incur long delays since large amounts of videos need to be uploaded to the server via wireless networks. To enable offloaded object detection on continuous videos, Glimpse [11] proposes to only send trigger frames and proposes the “detect + track” framework

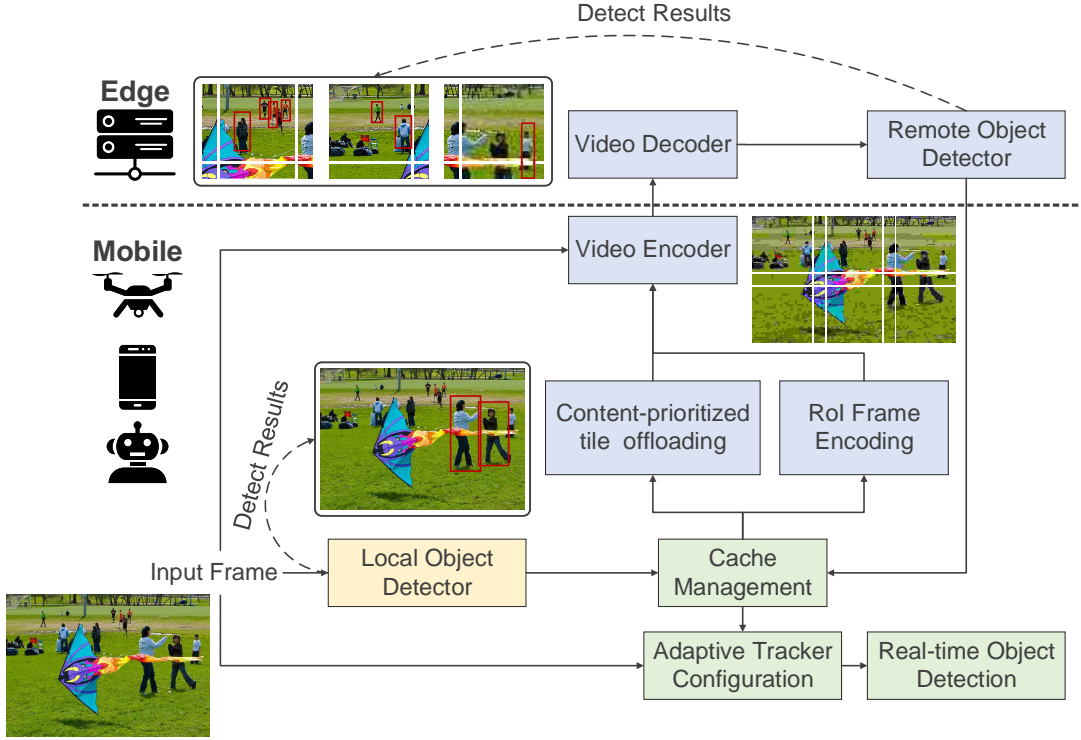


Fig. 2. An overview of EdgeDuet. Rectangles in different colors represent the three functional modules (offloaded small object detection (blue) (Sec. 4), local object detector (yellow) (Sec. 5) and real-time tracking (green) (Sec. 6). EdgeDuet is implemented as a cross-platform framework consisting of both edge-side and device-side modules (Sec. 7)).

for fast object detection. EAAR [5] compresses the uploaded frames via RoI based video encoding and applies parallel streaming and inference techniques to reduce the offloading latency further. Our work is built upon the “detect + track” framework and the pipelined offloading principle, but improves the parallelism of the offloading pipeline to tile-level. Furthermore, these studies do not optimize small object detection. DDS [29] differentiates small and large object detection by first offloading high-resolution, low-quality video frames to detect large objects and locate small objects. Regions containing small objects are then encoded in high quality and offloaded again to detect small objects. The method improves the accuracy of small object detection but doubles the delay for object detection. Unlike DDS, which detects large and small objects sequentially, we run a fast model on low-resolution frames to detect large objects and offload small object detection with high-quality frames at the same time.

Tile-based Video Streaming. Tiling feature [34] in video codecs has provided better quality gains for video streaming [35]–[38]. In tile-based video streaming, the video is first cut into tiles and important tiles are transmitted in high quality, whereas others are transmitted in low quality or not transmitted at all. Since all of them are based on existing video codecs, they encode the whole frame into the bit-streams then process each tile’s bit-stream independently. Differently, we redesign the video encoder such that it outputs the bit-stream of each tile once it is encoded and transmits the bit-stream to the edge server immediately to reduce the offloading latency.

Single Object Tracking. There are many existing accurate tracking algorithms [39], such as optical-flow based Lucas-Kanade tracking [40], correlation-filter based KCF [14], deep learning based Siamese RPN [41]. Most of these techniques require lots of computational resources, and are not efficient for applications on mobile devices with real-time requirements. We choose KCF as our tracking algorithm for its excellent efficiency and accuracy, and limited resource requirements.

3 EDGEDUET OVERVIEW

As shown in Fig. 2, EdgeDuet consists of three functional modules:

- i) An *offloaded small object detection* module which uploads high-resolution frames to the edge to detect small objects. EdgeDuet optimizes offloaded small object detection with region-of-interest (RoI) frame encoding and content-prioritized tile offloading to save network traffic and accelerate small object detection at the edge. EdgeDuet processes videos in the unit of tiles rather than the entire frame, so as to improve the parallelism of offloading.
- ii) A *local object detector* module which detects large objects from low-resolution frames. The local object detector aims to detect medium- to large- sized objects in the video frames locally on the mobile device. We empirically decide the model and input resolution for the local object detector to meet the constraints of resources on the mobile device.
- iii) A *real-time tracking* module which associates the detection results (bounding boxes, a.k.a *bboxes*) from both the edge and the mobile device and tracks each object with

single-object trackers. EdgeDuet adopts multiple single-object trackers for object tracking and it also applies adaptive tracker configuration to improve the resource efficiency and real-time performance of the trackers.

We elaborate on the detailed designs of each functional module in the subsequent sections.

4 OFFLOADED SMALL OBJECT DETECTION

This module aims to (i) reduce the data for transmission to the edge and (ii) accelerate the offloading pipeline for timely updates of the cached detection results on the mobile device. EdgeDuet exploits RoI frame encoding to compress video frames, and content-prioritized tile offloading for highly parallel object detection at the edge.

4.1 RoI Frame Encoding

As mentioned in Sec. 1, accurate small object detection relies on high-resolution, high-quality frames as input. Yet uploading high-quality frames to the edge impairs real-time object detection [23], [24], [26]. The RoI frame encoding module reduces the amount of transmitted data by only keeping the pixel blocks containing small objects in high quality while compressing the rest of the frame to low quality. Although RoI frame encoding has been used in other offloading schemes [5], [29], the definition of RoI (*i.e.*, blocks containing small objects in our case) and the compression level varies and should be tuned for specific applications.

4.1.1 Determining Blocks Containing Small Objects

A pixel block is considered as containing small objects if (i) the local object detector cannot classify the block into a class (or reports low confidence scores); and (ii) the remote object detector can classify the block to a class (or reports high confidence scores). Due to the high temporal correlation between successive frames, we use the detection results of the previous frame to identify blocks potentially containing small objects in the current frame. For simplicity, we decide whether an object is small using a fixed size. The size is empirically tuned such that objects below this size *cannot* be accurately detected by the *local* object detector but *can* be accurately detected by the *remote* object detector. The object size is considered as accurately detected if the recall is above 90%. Experiments show that the optimal size threshold for small objects varies across classes. For example, a size of 2000 pixels results in almost 100% recall for pedestrians but less than 40% recall for cars (see Fig. 3). Hence a different size threshold is set for each targeting class.

4.1.2 Determining Compression Levels

Blocks which are determined as containing no small objects cannot be compressed to arbitrarily low quality. This is because the decision is made based on the detection results of the *previous* frame. If a new object appears in the *current* frame, the blocks containing this object may be so heavily compressed that the object cannot be detected by the remote object detector. To avoid missing detection of new objects, we tune a low-quality compression level. Fig. 4 shows the rationality. For low-quality images, the object detection




Fig. 3. An example of the class-dependent size threshold for small objects. Details of datasets, local and remote object detectors see Sec. 7.

model could approximately detect most objects' locations by lowering its confidence score. We choose the compression level such that the remote object detector outputs low confidence scores on the compressed blocks but will not fail to locate objects. These low confidence objects are also returned to the device for offloading their blocks at the next frame, so as to be accurately detected in the next frame. To avoid a cold start, we just offload the whole frame in high quality.

4.1.3 Implementing RoI Frame Encoding

We use the High Efficiency Video Coding (HEVC, a.k.a h.265) codec [42] to encode pixel blocks containing small objects to high quality and compress the rest of the frame to low quality. The Quantization Parameter (QP) is one critical arguments in video compression performance. Larger values mean that there will be higher quantization, more compression, and lower quality. Lower values mean the opposite. In our scheme, we generate a delta QP map describing the delta QP values of each macroblock in the raster order and encode the current frame with the HEVC codec. Fig. 5a and Fig. 5b show an example image before and after RoI frame encoding.

4.2 Content-Prioritized Tile Offloading

This module enables real-time small object detection via fine-grained (tile-level) parallel offloading. It also facilitates timely updates of cached detection results on the mobile device by prioritizing the processing of tiles that contain more small objects. Pipelined offloading proves effective for fast object detection [5], where the offloading process is split into frame encoding, frame upload, frame decoding, object detection, and result downloading. Nevertheless, existing work [5] pipelines the offloading process on a frame basis, which limits the achievable parallelism. In contrast, EdgeDuet breaks a frame into tiles and enables tile-level parallelism, thus allowing faster pipelined and parallel offloading. We explain how to realize tile-level parallelism and content-based priority below.

4.2.1 Enabling Tile-Level Parallelism

A tile is a rectangular region in a frame defined in HEVC [34]. Fig. 5c shows an example of 5x3 tiles. To support tile-level parallelism, we need to modify the frame encoding, frame decoding, and object detection stage, as they are designed to operate on a frame basis. The principle is to eliminate dependencies among tiles for each stage, as described in detail below.




Fig. 4. Detected objects for different image quality. (a) Low quality, low confidence objects. (b) Low quality, high confidence objects. (c) High quality, low confidence objects. (d) High quality, high confidence objects.




Fig. 5. An example of key steps in EdgeDuet. (a) Input frame. (b) Frame after RoI frame encoding, where blocks containing no small objects are compressed to low quality. (c) Tiles. (d) The output of video decoder after enabling tile-level parallelism. (e) Overlap-tiling. (f) Remote object detector results of tiles (red rectangles). (g) Local object detector results of the low-resolution frame (yellow rectangles). (h) Cache management of remote and local object detectors.




Fig. 6. Video encoding parallelism of Kvazaar and EdgeDuet.

- *Frame Encoding.* Existing video encoders [43]–[45] output the encoded bit-stream after processing *all the tiles* in a frame. We redesign the video encoder such that it outputs the bit-stream of *each tile* once it is encoded. Our method is based on Kvazaar [45], which treats the encoding of each tile as an individual task and allows parallel tile encoding via a dynamic task graph. However, Kvazaar outputs bit-streams on a frame basis. Fig. 6(a) shows the task dependencies among tile encoding tasks and frame bit-stream tasks of the current frame and the next frame in Kvazaar. We modify its bit-stream writing module so that the bit-stream tasks operate on a tile basis, as the task dependencies shown in Fig. 6(b). Specifically, we break the bit-stream of a frame into a picture parameter set (PPS) and each tile’s bit-

streams. PPS contains the meta information of each frame on entropy coding mode, slice groups, motion prediction, quantization parameters (QP), and deblocking filters. Consequently, each tile’s bit-stream only depends on PPS and the tile encoding task. Hence the video encoder will first output the PPS, and once one tile is encoded, its bit-stream will be output and sent for offloading. We also introduce a fake bit-stream task to mark the end of the bit-stream tasks in a given frame.

- *Frame Decoding.* Existing video decoders operate on a frame basis. They assume the bit-streams of all the tiles in a frame arrive sequentially and utilize the offset from the first tile in the frame to locate the other tiles. For example, in HEVC, only the location of the first tile is signaled in the slice header. All the

other tiles transmit their bit-stream offsets in the slice header, which introduces dependencies on the first tile. We eliminate such dependencies and enable tile-level parallelism in frame decoding by forcing every tile in a frame as a “first tile”. This is implemented by modifying the bit-stream of each tile in the video encoder (Kvazaar) and the HEVC parser in the video decoder (OpenHevc [46]) accordingly. Fig. 5d shows an example of tile-level frame decoding. Each tile is decoded to its position independent of the other tiles (shown in black).

- **Object Detection.** Performing object detection on each tile separately may miss objects which cross the boundaries of adjacent tiles. We mitigate such dependencies among tiles during object detection via *overlap-tiling*. That is, we split the frame into $M \times N$ tiles, where M and N are odd numbers. We use (i, j) to denote each tile, where $i \in [1, M]$ and $j \in [1, N]$ are the row index and column index, respectively. We classify tiles into two categories. If both i and j are odd numbers, the tile (i, j) is a primary tile. Otherwise, the tile is an overlay tile. Fig. 5e shows an example. Since tile 2 can be denoted as $(1, 3)$, tile 2 is a primary tile. We can conclude that tile 2 and 4 are primary tiles and tile 1, 3, 5, 6, 7, 8 are overlay tiles. We group each primary tile with its surrounding overlay tiles for small object detection. In this example, tile 1, 2, 3, 5, 6, 7 will be grouped. Detecting objects for each tile group reduces the probability of missing objects that exceed the boundary of a primary tile. We only group the surrounding tiles because our remote detector targets at small objects. Only large objects may be present in two primary tiles crossing the overlap tiles, as person A and B in Fig. 5f. The overlap size $S_{overlap}$ (minimal width and height of overlap tiles) is set to the least multiple large coding unit (LCU), which is larger than the maximal size of small objects defined in Sec. 4.1.1 and complies with the tile definition in HEVC. The formula is illustrated as

$$S_{overlap} = \lceil \frac{\max_c \{ \max \{ w_c, h_c \} \}}{LCU_WIDTH} \rceil \times LCU_WIDTH \quad (1)$$

where w_c and h_c are the width and height of the small object threshold for the targeting class c . LCU_WIDTH is the width of LCU. As object size can be arbitrary, the approach of overlap-tiling may not always work well in some extreme cases, we overcome it by simply running object detection on the whole image after the last tile is transmitted. Fig. 5f shows an example of detection results using our method.

4.2.2 Enabling Content-based Priority

Prioritizing tiles containing more objects over those containing fewer objects allows the cached detection results of more objects to stay fresh. Since our implementation of tile-level parallelism (Sec. 4.2.1) ensures tiles are offloaded early to return detection results early, we only need to prioritize tiles at the frame encoding stage.




Fig. 7. An illustration of content-based tile priority.

- **Implementing Tile Priority for Frame Encoding.** We modify the task schedule module in Kvazaar by adding a dynamic priority mapping module to enable the ordering of tiles (see Fig. 7). Specifically, the dynamic priority mapping layer associates a priority value p to each primary tile according to the input content, where $p \in [0, N)$ and N is the number of primary tiles, and each overlap tile calculates its priority p as the maximum priority of its surrounding primary tiles. Then the encoding task of each tile is assigned a priority value of p while its bit-stream task is assigned a priority value of $p + N$. This is to force the bit-stream task to execute once the tile is encoded, which can be before other tiles' encoding tasks.
- **Assigning Tile Priority based on Content.** To determine the priorities (i.e., p) of the N primary tiles, we count the number of small objects of the corresponding tile group. The priority value p of each primary tile is the index in ascending order.

5 LOCAL OBJECT DETECTOR

The local object detector aims to detect medium- to large-sized objects in the video frames locally on the mobile device. Since mobile devices have limited resources compared with the edge, the local object detector should be lightweight and operate on low-resolution frames. We empirically decide the model and input resolution for the local object detector. The local object detector should balance between *offline accuracy* and *latency* to achieve high *online accuracy*. The offline accuracy refers to the accuracy of the object detector, while the online accuracy refers to the accuracy in the “detect+track” framework [5], [11]. Accuracy is measured by metrics such as IoU, as will be defined in Sec. 8.1.5.

Table 1 shows the performance of different combinations of object detectors and input resolutions evaluated on the VisDrone dataset with an iPhone 11. For resource efficiency, the models are quantized to `float16`. Since the “real-time” detector is not a must with the “detect+track” framework, we choose YOLOv3 (640x640) as the local object detector based on the analysis. Fig. 5g shows an example of detection results of the local object detector. Note that we only aim to show the feasibility of running an object detector locally for accurate and real-time medium- to large-sized objects. An exhaustive search on the optimal local object detector is out of the scope of this paper.

6 REAL-TIME TRACKING

This module aggregates the offloaded and the local detection results into the cache and adjusts the cached results via

TABLE 1
Performance of local detector on iPhone 11.

Model	IoU (offline)	Latency	IoU (online)
YOLOv3-tiny (320x320)	0.015	12.4ms	0.012
YOLOv3-tiny (640x640)	0.078	19.5ms	0.052
YOLOv3-tiny (960x960)	0.140	38.9ms	0.090
YOLOv3 (320x320)	0.176	23.8ms	0.092
YOLOv3 (640x640)	0.361	62.5ms	0.193
YOLOv3 (960x960)	0.522	178.7ms	0.161

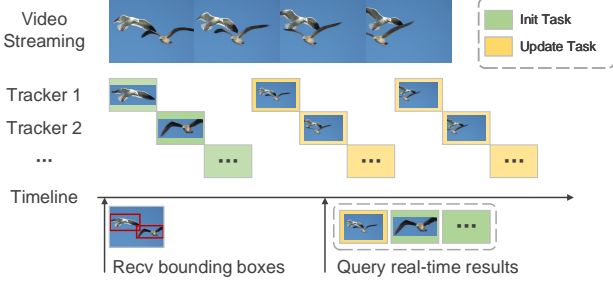


Fig. 8. General workflow using multiple single-object trackers.

object trackers to output the bounding boxes for the current frame. EdgeDuet adopts multiple *single-object trackers* for object tracking, as in [5], [11]. Fig. 8 shows the general workflow. Trackers return latest updated bounding boxes so as to query as the same fps as the video input. Since we target at video streams with high frame rates (30/60/120 fps) and the cached results come from two object detectors, the general workflow needs to be optimized for EdgeDuet, as we describe below.

6.0.1 Cache Management

We cache the detection results received from the local or remote object detector and discard the old results upon receiving new ones. One issue in our cache management is that the local and the remote detector may introduce duplicated detection results of the same object. We drop the results of the local detector for small objects and those of the remote detector for medium- to large-sized objects in case of duplicated results. Fig. 5h shows an example of merging local detection results and remote detection results.

6.0.2 Adaptive Tracker Configuration

To optimize the tracking performance on mobile devices, we consider the following.

- *Choice of Single-object Tracker.* We empirically choose KCF [14] as our single-object tracker since it is both faster and more accurate than the optical flow based tracker in [11] and has higher accuracy than the motion vector based tracker in [5].
- *Priority-based Tracker Scheduling.* To execute multiple single-object trackers on resource-constrained devices, we adaptively update the tracking results based on the speed of the objects, because it is unnecessary to frequently update the tracking results of objects that are static or moving slowly. Specifically, we estimate the object's speed by the object's move distance in continuously tracked frames. Then we set a different weight for each speed range, and the




Fig. 9. Internal Architecture of Apple A13 SoCs.

priority of each tracker is updated to the product of its weight value and the default priority (distance between the current frame and last tracked frame in sequential task scheduling). We schedule the tracker with high priority to track first to ensure high-speed objects frequently updated. To accelerate the tracking module, we use a thread pool to execute multiple trackers parallelly.

7 IMPLEMENTATION

This section presents the implementation of EdgeDuet on the device-side and the edge-side.

7.1 Implementation of Core Device-Side Modules

We implement the device-side of EdgeDuet on a popular mobile device, iPhone 11, which equips the A13 Bionic chip. As shown in Fig. 9, this SoC contains 4 processors: two high-performance Lightning CPU cores running at 2.65 GHz, four energy-efficient Thunder cores running at 1.80 GHz, four GPU cores with 0.69 TFLOPS for FP32 floating point computing and eight Neural Engine cores with 5.5 TOPS for AI computing [47]. Benefiting from the excellent performance, the Soc makes it efficient for the tasks of video compression, object detection and real-time object tracking.

The device-side modules of EdgeDuet consist of a video streamer, a video encoder, a local object detector and an object tracker. Each module has a unique thread and communicates with other modules by blocking queues. The video streamer simulates the video camera streaming process and feeds raw frames to the video encoder, the local object detector and the object tracker with different frame rates. The video encoder encodes raw frames to bit-stream and sends network packets to the server. The local object detector runs a light-weight object detection model and updates the frame cache for object tracking. The object tracker tracks all objects detected from both the local object detector and the remote object detector. Most modules on the device side are implemented in C++ 17 [48] for easy deployment on different platforms such as iOS Frameworks [49], Android NDK [50] and Nvidia Jetson [51]. We next explain more detail for each module.

7.1.1 Video Streamer

This module is used for simulating the video camera streaming process using standard video datasets. The streamer loads a video file and feeds video frames into EdgeDuet at 30/60/120 fps. The module is implemented in C++ using the VideoCapture module in OpenCV [52] to read RGB images from a video file and convert the image to I420 format for video encoding.

7.1.2 Video Encoder

This module is implemented in C++ based on `Kvazaar` [53], an open-source HEVC encoder. We modify the library to support tile-based parallel encoding and priority, as described in Sec. 4.2.1 and Sec. 4.2.2. The modified library is open-sourced at <https://github.com/xu-wang11/kvazaar>. We empirically encode and offload frames at a fixed frame rate (e.g., 5 fps).

7.1.3 Local Object Detector

This module is implemented in Objective-C [54] with `Core ML` [55], which optimizes on-device performance by jointly leveraging the CPU, GPU, and Neural Engine. We use the pre-trained compressed YOLOv3 model YOLOv3FP16 (640x640). for medium- to large-sized object detection, as explained in Sec. 5. We empirically run the local object detector at a fixed frame rate (e.g., 10).

7.1.4 Object Tracker

This module is implemented in C++ with the `KCF` [14] Tracker and `ThreadPool` [56] to schedule multiple object tracking. We use the implementation of `KCFcpp` [57] without the HOG features [58] for fast object tracking, as described in Sec. 6.0.2.

7.2 Implementation of Core Edge-Side Modules

We implement the edge-side modules of EdgeDuet on a CentOS 7.0 server. It is equipped with two 8-core Intel Xeon CPU E5-2560 v4 CPUs, two GTX 2080ti GPUs and 256GB memory. The edge-side modules of EdgeDuet consist of a video decoder and a remote object detector. The video decoder receives packets from the device-side and converts them to image tiles. The image tiles will be queued in the remote object detector for object detection.

7.2.1 Video Decoder

The video decoder is implemented in C++ based on the `OpenHEVC` library [46]. We modify the library to support tile-based parallel decoding as in Sec. 4.2.1. We use `OpenCV-Python` bindings [59] for Python to run the decoder and access the decoding results from memory.

7.2.2 Remote Object Detector

The detector for small objects on the edge is implemented as a pre-trained full-precision YOLOv3-spp [60] model in `PyTorch` [61]. We run the model in multiple processes for parallel inference on tiles. The GPU is set to `CUDA Multi-Process Service` mode [62] to reduce GPU context switching.

7.3 Implementation of Auxiliary Modules

For the edge- and device-side modules to work in synergy, we also implement the necessary functions for network communication, inter-thread communication, and data logging.

7.3.1 Network Communication

The communication between the edge and the mobile device is via TCP [63]. We use the `sockpp` [64] library for network programming. The up-link traffic is transmitted as bit-streams, and the down-link traffic (bboxes from the edge) is transmitted in the format of JSON. The device will parse it with the `JSON` [65] library.

7.3.2 Inter-thread Communication

In our implementation, each module works as an independent thread. We use the `concurrentqueue` [66] library to block the thread when no task is put into the queue.

7.3.3 Data Logging

We use `spdlog` [67] to log the timestamps of each processing step for tracing the entire workflow.

Similarly, we use “queue” for message communication, “socket” for network communication, “logging” for logging the work flow. All of them are in the python standard library.

8 EVALUATION

This section presents the evaluations of EdgeDuet.

8.1 Experiment Setup

8.1.1 Datasets

To evaluate the performance for small object detection, we compare different methods on VisDrone, a dataset of videos captured by drone-mounted cameras, which contains lots of small objects. We filter out the low-resolution videos and only keep six 2K videos (2560x1440) captured along a street. The count of frames is 1886. We upsample the origin 30fps videos to 60/120 fps with `Super-SloMo` [14] to evaluate the performance with the video frame rate.

Table 2 gives the basic statistics of the dataset. Fig. 10 shows certain vital statistics of the VisDrone dataset. From Fig. 10a, the average number of objects in each frame is 132. As we will show later, adaptive tracker configuration is beneficial to track such many objects in real-time. We analyze the threshold of small objects by running the local object detector and remote object detector on the dataset. We use the detection results of the remote object detector as the ground truth and the areas of bounding boxes as the sizes of objects to calculate the recall-size curves of each class. Fig. 3 shows the results of the car class and pedestrian class. We set the threshold of the small car as 5150 pixels and the threshold of the small person as 1698 pixels. From Fig. 10b and Fig. 10c, 75.1% of cars and 71.8% of pedestrians are small objects. We will show the performance gain of offloading small object detection over local detection shortly.

8.1.2 Compared Methods

We compare our EdgeDuet with the following object detection schemes.

- `Glimpse` [11]: a continuous, real-time object detection system that first proposes the “detect + track” framework on mobile devices. It offloads frames to the cloud and uses optical flow based tracker for real-time tracking.
- `EAAR` [5]: a state-of-the-art real-time object detection system with offloading. It exploits parallel streaming and inference as well as motion vector based object tracking.
- `LaT`: a variant of EdgeDuet (Local object detector + adaptive Tracking configuration) that only performs local object detection and tracks with our adaptive tracker configuration.




Fig. 10. Characteristics of VisDrone dataset. There are numbers of objects in each frame, and most objects are small in size.

TABLE 2
Dataset Statistics

scenario	#videos	#avg duration	#frames	#categories	#avg objects of each frame	resolution	year
drone	6	10.5s	1886	2	132	3840x2160	2018

To demonstrate the effectiveness of RoI frame encoding module and Content-prioritized tile offloading module in our EdgeDuet, we also compare the performance of the following variants of EdgeDuet.

- EdgeT The variant only utilizes tile-level parallelism for small object detection.
- EdgeTR The variant is EdgeT enhanced with RoI frame encoding.

8.1.3 Implementation and Settings of Compared Methods

The implementation of EdgeDuet and LaT can be found in Sec. 7. We briefly explain the implementation of Glimpse, EAAR, and parameter settings for all methods below.

- *Frames Encoding.* To ignore the difference of JPEG encoders and video encoders, we use Kvazaar to get compressed JPEGs for Glimpse and video frames for EAAR and EdgeDuet. This setting ensures the same frame quality for a fair comparison. For Glimpse, we encode each frame to I frame by setting Group of pictures (GOP) as 1. For EAAR and EdgeDuet, the GOP is set to 10. We use the ultrafast preset for real-time video encoding. We tune high-quality QP as 22 and low-quality QP as 44 for RoI encoding in EAAR and EdgeDuet. EAAR uses Kvazaar to encode one frame to 1x4 tiles and pack each tile into one slice, as the default setting. EdgeDuet splits the frame into 5x3 tiles for offloading, as in Fig. 5c.
- *Remote Object Detector.* We use the same remote object detection model, i.e., YOLOv3-spp, for Glimpse, EAAR, and EdgeDuet. Glimpse operates on the input size of 2560x2560. EAAR operates on 1280x1280 but returns the detection results in 2560x2560. Specifically, we first cache the object detection results of each 2560x2560 frame in the memory and implement the *Dependency Aware Inference* by detecting on only one tile and overlooking other tiles for convenience. Since we return the detection results of the whole frame, our setting has the same accuracy but allows

a lower latency of EAAR compared with its origin dependency aware inference. EdgeDuet operates on 960x960 for overlap-tiling inference. The overlap size is set to 128 (2 macroblocks) as in Sec. 4.2.1.

- *Real-time Object Tracking.* We implement the optical flow based tracker with `calcOpticalFlowPyrLK` for Glimpse. For EAAR, the motion vector based tracker is implemented as an offline process. When received from the server, each track frame is associated with a refer frame ID. We use `ffmpeg` to compress the refer and track frames and extract the motion vector based on the refer frame to simulate RPS Control in EAAR. LaT uses the same tracking module of EdgeDuet. They both split the speed range into two groups and set the weight as 2.0 for the fast speed range and 1.0 for the slow speed range, as explained in Sec. 6.0.2. The fast speed range is set as $[30 \text{ pixels/s}, +\infty)$, and the slow speed range is set as $[0, 30 \text{ pixels/s})$. We query the tracking results of the previous frame when the current frame is fed for all methods.

8.1.4 Network Setting

Since autonomous mobile vision applications are often deployed outdoors, the network connections vary. Accordingly, we compare the methods in different network settings. We connect the mobile device and the edge with WiFi 5GHz and emulate different types of networks with `Network Link Conditioner`, a developer tool provided by Apple. We use it to simulate different network conditions (LTE, WiFi 2.4GHz, WiFi 5GHz), and network bandwidths, Table 3 summarizes the parameters of the three network conditions. Network conditions with high network propagation delay is overlooked since our system requires offloading high-resolution frame.

8.1.5 Metrics

We evaluate the performance of different methods with the following metrics.

TABLE 3
Parameters of three network conditions.

Parameter	LTE	WiFi 2.4GHz	WiFi 5GHz
In Bandwidth (Kbps)	50000	40000	250000
Out Bandwidth (Kbps)	10000	33000	100000
In Delay (ms)	50	1	1
Out Delay (ms)	65	1	1

- *Latency*: The evaluation metric measures the delay of the detected objects. Lower latency will benefit the accuracy of object tracking. Since our detection results are composed of medium- to large-sized objects in frames from the local object detector and small objects in tiles from the remote object detector. We average latency for all objects, which is compatible with the definition of latency in EAAR. Specifically, supposing that there are M frames F_m uploaded to the edge server, the M -th frame is captured at time $t(F_m)$ and contains N_m objects. The bounding box and class label of each object O_m^n are sent to the client at time $t(O_m^n)$. Therefore, the latency can be obtained by:

$$Latency = \frac{1}{\sum_{m=1}^M N_m} \sum_{m=1}^M \sum_{n=1}^{N_m} t(O_m^n) - t(F_m) \quad (2)$$

- *Accuracy*: We use the average IoU [68] to measure the real-time object detection accuracy as in *Glimpse* and EAAR. The IoU of object o is defined as the Intersection over Union overlap with its ground-truth box.

$$IoU_o = \frac{|BBox(o) \cap BBox(g)|}{|BBox(o) \cup BBox(g)|} \quad (3)$$

where $|\cdot|$ is the area of the geometry, $BBox(\cdot)$ is the bounding box of the geometry. g is the ground truth of the object o . We take the object accurately detected if $IoU \geq 0.8$. The IoU is averaged over all objects in all frames.

8.2 End-to-End Performance

Fig. 11 summarizes the accuracy and latency of different methods under LTE, WiFi 2.4GHz, WiFi 5GHz network conditions. Fig. 12 highlights the accuracy of small objects. We present our observations and explain the results below. Fig. 13 shows the sample of detection results.

8.2.1 Overall Comparison

EdgeDuet notably outperforms the two offloading schemes, *Glimpse* and EAAR, in both accuracy and latency under all three network conditions. *LaT* is the fastest because it only performs local detection. However, pure local detection has the worst accuracy, especially for small object detection. EdgeDuet achieves 160.2%, 242.9%, 287.8% improvement in IoU metric for small object detection accuracy under the three network conditions, respectively. Under slow network connection, *e.g.*, LTE, *LaT* achieves similar accuracy with *Glimpse* and EAAR. This indicates the necessity of a local object detector when network conditions vary, which is

common outdoors. Since *LaT* performs badly for small objects, we exclude it for the subsequent evaluations.

Among the variants of EdgeDuet, EdgeRT achieves a higher IoU (1.47%, 2.80%, 1.43%) and lower latency (5.42%, 2.16%, 8.28%) than EdgeT. This demonstrates the effectiveness of RoI Frame Encoding. EdgeDuet achieves a higher IoU (4.06%, 5.20%, 8.90%) and lower latency (1.47%, 18.79%, 11.41%) than EdgeTR. This shows the effectiveness of Content-prioritized tile offloading.

8.2.2 Comparison with Offloading Schemes on Latency

EdgeDuet achieves 48.8%, 39.6%, 38.6% latency improvement than *Glimpse* and 35.4%, 35.2%, 32.2% latency improvement than EAAR under the three network conditions. The improvement in latency is more notable under slower networks, *e.g.*, LTE. EAAR achieves shorter latency than *Glimpse* since it transmits encoded videos instead of raw JPEGs. EdgeDuet is faster than EAAR for the following reasons.

- Detection of medium- to large-sized objects of EdgeDuet is from a local object detector. The latency of the local object detector is lower than offloading.
- Only small object detection is offloaded in EdgeDuet. Therefore fewer data need to be transmitted.
- EdgeDuet accelerates the offloading pipeline with tile-level parallelism. EAAR only implements task-level parallelism, so that the detection results have to wait for processing the entire frame.

8.2.3 Comparison with Offloading Schemes on Accuracy

EdgeDuet achieves 52.8%, 45.5%, 47.6% IoU improvement gain over *Glimpse*, 50.2%, 43.6%, 41.3% IoU improvement over EAAR under the three network conditions. EAAR achieves slightly better accuracy than *Glimpse* under LTE and WiFi 2.4GHz. The reason might be that motion vector based tracker behaves badly when latency increases. EdgeDuet yields the highest accuracy because it trades off between the tracker's accuracy and efficiency and employs adaptive tracking to update fast-moving objects.

8.2.4 Comparison with Offloading Schemes on Small Object Detection.

EdgeDuet achieves 34.9%, 34.9%, 44.5% IoU improvement for small objects over *Glimpse*, 72.3%, 68.0%, 61.7% IoU improvement for small objects over EAAR. EAAR is worse than *Glimpse* for small object detection, although it has a higher overall detection accuracy. This is because small objects contain very few macroblocks to extract motion vectors, making it inaccurate to represent the object.

Among the variants of EdgeDuet, EdgeRT achieves a little higher IoU (0.9%, 0.7%, 3.4%) than EdgeT. EdgeDuet achieves higher IoU (11.4%, 13.5%, 14.5%) than EdgeTR. This demonstrates Content-prioritized offloading module contributes a lot for small object detection.

8.3 Impacting Factors on Overall Performance

8.3.1 Impact of Bandwidth

Fig. 14 shows the accuracy of different methods under different bandwidths. Thanks to the local object detector




Fig. 11. End-to-end object detection accuracy (bars in red) and latency (bars in blue) of different methods under three network connections.




Fig. 12. Small object detection accuracy of different methods under three network connections. Pure local object detection (LaT) is excluded from subsequent evaluations for its low detection accuracy on small objects.




Fig. 13. Output Sample of EdgeDuet.

and optimized offloading, EdgeDuet consistently achieves better accuracy than Glimpse and EAAR. Particularly, when the bandwidth is limited (below 10Mbps), the accuracy of Glimpse and EAAR drops dramatically.

8.3.2 Impact of Frame Rate

Fig. 16 shows the accuracy of different methods when feeding videos of different frame rates. EdgeDuet consistently achieves higher accuracy than Glimpse and EAAR, even at 120fps. With the increase of frame rate, the accuracy of Glimpse drops. This is because the real-time tracker of Glimpse only works in an fps lower than 30fps. As for EAAR, increasing the frame rate motion does not impact the motion vector based tracker and thus the accuracy. An interesting finding is the accuracy of EdgeDuet increases with the frame rate. The reason may be that we use adaptive tracker configuration to update trackers of high speed objects frequently to reduce the influence of the skipped frames. Our tile-level parallelism may also help since once each tile's results are received, they do not wait for the new frame fed with high fps video input.

8.4 Benefits of Individual Modules in EdgeDuet

8.4.1 Benefits of RoI Frame Encoding

We evaluate the benefits of RoI frame encoding by comparing the offloading file size with EAAR and Glimpse. We average the bits count of frames with the same index in GOP. Since Glimpse only contains I frame, we only average the corresponding frames with the same frame index. Fig. 15 shows the average frame size of GOP. Since Glimpse does not apply inter-frame prediction and RoI frame encoding, its frame size is the largest, especially when the frame is encoded to P frame in EAAR and EdgeDuet. Since EdgeDuet does not offload medium- to large-sized objects, its frame size is smaller than EAAR.

8.4.2 Benefits of Content-prioritized Tile Offloading

We show the benefits of content-prioritized tile offloading by comparing EdgeDuet with two variants. The variant Frame-Level encodes frames without splitting into tiles. The variant Tile-Level splits frames into tiles, but does not change their priority. Fig. 17 shows the accuracy and latency of EdgeDuet and the two variants. EdgeDuet achieves




Fig. 14. Impact of Network Bandwidth




Fig. 15. Benefits of RoI Frame Encoding




Fig. 16. Impact of Frame Rate




Fig. 17. Benefits of Content-Prioritized Tile Offloading




Fig. 18. Benefits of Adaptive Tracker

8.1% and 5.0% accuracy improvement over Frame-Level and Tile-Level. EdgeDuet achieves 12.1% and 5.1% latency improvement over Frame-Level and Tile-Level.

8.4.3 Benefits of Adaptive Tracker Configuration

We evaluate the benefits of adaptive tracker configuration by comparing EdgeDuet with a variant SeqTracking which sequentially updates each tracker. Fig. 18 shows the accuracy of EdgeDuet and SeqTracking. Our adaptive tracker configuration improves the overall accuracy by 4.7%.

9 LIMITATIONS AND DISCUSSION

- **Commercial HEVC Codec Support.** EdgeDuet utilizes Kvazaar as the video codec. Since Kvazaar is a CPU-based video encoder, it requires lots of computing resources for high-resolution video compression, which burdens a lot for most mobile devices. Hardware acceleration will help it for efficient video compression, however, the commercial codecs have not supported the tile feature yet. We hope some AI-driven solutions like EdgeDuet will give more attention to the feature.
- **Adaptive Local Object Detector.** The performance of EdgeDuet is influenced by the local object detector. An efficient and accurate local object detector will reduce the size of the uploaded video frame and improve the accuracy of the trackers. However, the local object detector will compete for the computing resources with the trackers, which is not analyzed in EdgeDuet. We don't search the optimal local object detector choice according to the computing resources and the video input.

10 CONCLUSION

This paper presents EdgeDuet, the first splits object detection between the mobile device and the edge for accurate, real-time object detection on resource-constrained devices. Specifically, EdgeDuet offloads small object detection to the edge while detecting medium- to large-sized objects locally on the mobile device. EdgeDuet exploits RoI frame encoding and priority-based tile offloading to reduce the network traffic and accelerate the offloading pipeline. It also optimizes the cache detection results and tracker configurations for real-time object tracking. Evaluations on VisDrone, a video dataset from drone-mounted cameras, show that EdgeDuet outperforms local object detection in small object detection accuracy by 233.0%. It also improves the overall accuracy by 44.7% and end-to-end latency by 34.2% over the state-of-the-art offloading schemes, especially in low bandwidth and high frame-rate input.

ACKNOWLEDGEMENT

This work is supported in part by the NSFC under grant 61832010, 61632008, 61872081, 61972131, as well as the Shuimu Tsinghua Scholar Program. Zimu Zhou's research is supported by the Singapore Ministry of Education (MOE) Academic Research Fund (AcRF) Tier 1 grant.

REFERENCES

- [1] J. Xu, H. Chen, K. Qian, E. Dong, M. Sun, C. Wu, L. Zhang, and Z. Yang, "ivr: Integrated vision and radio localization with zero human effort," in *PACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, Sep 2019.
- [2] J. Xu, H. Cao, D. Li, K. Huang, C. Qian, L. Shangguan, and Z. Yang, "Edge assisted mobile semantic visual slam," in *Proceedings of the IEEE INFOCOM*, 27-30 April 2020.
- [3] K. Kanistras, G. Martins, M. J. Rutherford, and K. P. Valavanis, "A survey of unmanned aerial vehicles (uavs) for traffic monitoring," in *Proc. of IEEE ICUS*, 2013.
- [4] E. Su, "Roaming 'robodog' politely tells singapore park goers to keep apart," <https://www.reuters.com/>, 2020.
- [5] L. Liu, H. Li, and M. Gruteser, "Edge assisted real-time object detection for mobile augmented reality," in *Proc. of ACM MobiCom*, 2019.
- [6] "Augmented Reality - Apple Developer." [Online]. Available: <https://developer.apple.com/augmented-reality/>
- [7] "Build new augmented reality experiences that seamlessly blend the digital and physical worlds | ARCore | Google Developers." [Online]. Available: <https://developers.google.com/ar>
- [8] J. Xu, G. Chi, Z. Yang, D. Li, Q. Zhang, Q. Ma, and X. Miao, "Followupar: enabling follow-up effects in mobile ar applications," in *Proceedings of the 19th Annual International Conference on Mobile Systems, Applications, and Services*, 2021, pp. 1-13.
- [9] V. Sze, Y.-H. Chen, T.-J. Yang, and J. S. Emer, "Efficient processing of deep neural networks: A tutorial and survey," *Proceedings of the IEEE*, vol. 105, no. 12, pp. 2295-2329, 2017.
- [10] F. Ozge Unel, B. O. Ozkalayci, and C. Cigla, "The power of tiling for small object detection," in *Proc. of IEEE CVPR Workshops*, 2019.
- [11] T. Y.-H. Chen, L. Ravindranath, S. Deng, P. Bahl, and H. Balakrishnan, "Glimpse: Continuous, real-time object recognition on mobile devices," in *Proc. of ACM SenSys*, 2015.
- [12] M. Liu, X. Ding, and W. Du, "Continuous, real-time object detection on mobile devices without offloading," in *Proc. of IEEE ICDCS*, 2020.
- [13] A. G. Howard, M. Zhu, B. Chen, D. Kalenichenko, W. Wang, T. Weyand, M. Andreetto, and H. Adam, "Mobilenets: Efficient convolutional neural networks for mobile vision applications," *arXiv preprint arXiv:1704.04861*, 2017.
- [14] J. F. Henriques, R. Caseiro, P. Martins, and J. Batista, "High-speed tracking with kernelized correlation filters," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 37, no. 3, pp. 583-596, 2014.
- [15] P. Zhu, L. Wen, D. Du, X. Bian, Q. Hu, and H. Ling, "Vision meets drones: Past, present and future," *arXiv preprint arXiv:2001.06303*, 2020.
- [16] S. Ren, K. He, R. Girshick, and J. Sun, "Faster r-cnn: Towards real-time object detection with region proposal networks," in *Proc. of NeurIPS*, 2015.
- [17] J. Redmon, S. Divvala, R. Girshick, and A. Farhadi, "You only look once: Unified, real-time object detection," in *Proc. of IEEE CVPR*, 2016.
- [18] B. Fang, X. Zeng, and M. Zhang, "Nestdnn: resource-aware multi-tenant on-device deep learning for continuous mobile vision," in *Proc. of ACM MobiCom*, 2018.
- [19] S. Han, X. Liu, H. Mao, J. Pu, A. Pedram, M. A. Horowitz, and W. J. Dally, "Eie: efficient inference engine on compressed deep neural network," in *Proc. of ACM/IEEE ISCA*, 2016.
- [20] X. He, Z. Zhou, and L. Thiele, "Multi-task zipping via layer-wise neuron sharing," in *Proc. of NeurIPS*, 2018.
- [21] R. Xu, C.-I. Zhang, P. Wang, J. Lee, S. Mitra, S. Chatterji, Y. Li, and S. Bagchi, "Approxdet: content and contention-aware approximate object detection for mobiles," in *Proceedings of the 18th Conference on Embedded Networked Sensor Systems*, 2020, pp. 449-462.
- [22] J. Zhang, D. Zhang, X. Xu, F. Jia, Y. Liu, X. Liu, J. Ren, and Y. Zhang, "Mobipose: real-time multi-person pose estimation on mobile devices," in *Proceedings of the 18th Conference on Embedded Networked Sensor Systems*, 2020, pp. 136-149.
- [23] V. Růžička and F. Franchetti, "Fast and accurate object detection in high resolution 4k and 8k video using gpus," in *Proc. of IEEE HPEC*, 2018.
- [24] M. Gao, R. Yu, A. Li, V. I. Morariu, and L. S. Davis, "Dynamic zoom-in network for fast object detection in large images," in *Proc. of IEEE CVPR*, 2018.
- [25] S. Han, H. Shen, M. Philipose, S. Agarwal, A. Wolman, and A. Krishnamurthy, "Mcdnn: An approximation-based execution framework for deep stream processing under resource constraints," in *Proc. of ACM MobiSys*, 2016.
- [26] J. Wang, Z. Feng, Z. Chen, S. George, M. Bala, P. Pillai, S.-W. Yang, and M. Satyanarayanan, "Bandwidth-efficient live video analytics for drones via edge computing," in *Proc. of IEEE/ACM SEC*, 2018.
- [27] X. Ran, H. Chen, X. Zhu, Z. Liu, and J. Chen, "Deepdecision: A mobile deep learning framework for edge video analytics," in *Proc. of IEEE INFOCOM*, 2018.
- [28] J. Hanhiova, T. Kämäräinen, S. Seppälä, M. Siekkinen, V. Hirvisalo, and A. Ylä-Jääski, "Latency and throughput characterization of convolutional neural networks for mobile computer vision," in *Proceedings of the 9th ACM Multimedia Systems Conference*, 2018, pp. 204-215.
- [29] K. Du, A. Pervaiz, X. Yuan, A. Chowdhery, Q. Zhang, H. Hoffmann, and J. Jiang, "Server-driven video streaming for deep learning inference," in *Proc. of ACM SIGCOMM*, 2020.
- [30] J. Yi, S. Choi, and Y. Lee, "Eagleeye: wearable camera-based person identification in crowded urban spaces," in *Proc. of ACM MobiCom*, 2020.
- [31] T. Tan and G. Cao, "Fastva: Deep learning video analytics through edge processing and npu in mobile," in *Proc. of IEEE INFOCOM*, 2020.
- [32] S. P. Chinchali, E. Cidon, E. Pergament, T. Chu, and S. Katti, "Neural networks meet physical networks: Distributed inference between edge devices and the cloud," in *Proc. of ACM HotNets*, 2018.
- [33] A. Galanopoulos, V. Valls, G. Iosifidis, and D. J. Leith, "Measurement-driven analysis of an edge-assisted object recognition system," in *ICC 2020-2020 IEEE International Conference on Communications (ICC)*. IEEE, 2020, pp. 1-7.
- [34] K. Misra, A. Segall, M. Horowitz, S. Xu, A. Fuldseth, and M. Zhou, "An overview of tiles in hevc," *IEEE Journal of Selected Topics in Signal Processing*, vol. 7, no. 6, pp. 969-977, 2013.
- [35] R. Skupin, Y. Sanchez, D. Podborski, C. Hellge, and T. Schierl, "Hvc tile based streaming to head mounted displays," in *Proc. of IEEE CCNC*, 2017.
- [36] J. Son, D. Jang, and E.-S. Ryu, "Implementing motion-constrained tile and viewport extraction for vr streaming," in *Proceedings of the 28th ACM SIGMM Workshop on Network and Operating Systems Support for Digital Audio and Video*, 2018, pp. 61-66.
- [37] T. Biedert, P. Messmer, T. Fogal, and C. Garth, "Hardware-accelerated multi-tile streaming for realtime remote visualization," in *EGPGV*, 2018, pp. 33-43.
- [38] F. Qian, B. Han, Q. Xiao, and V. Gopalakrishnan, "Flare: Practical viewport-adaptive 360-degree video streaming for mobile devices," in *Proceedings of the 24th Annual International Conference on Mobile Computing and Networking*, 2018, pp. 99-114.
- [39] Y. Zhang, T. Wang, K. Liu, B. Zhang, and L. Chen, "Recent advances of single-object tracking methods: A brief survey," *Neurocomputing*, vol. 455, pp. 1-11, 2021.
- [40] B. D. Lucas, T. Kanade *et al.*, "An iterative image registration technique with an application to stereo vision." Vancouver, 1981.
- [41] B. Li, J. Yan, W. Wu, Z. Zhu, and X. Hu, "High performance visual tracking with siamese region proposal network," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 8971-8980.
- [42] H. E. V. Coding and I. Rec, "H. 265 and iso," 2013.
- [43] x265, the free H.265/HEVC encoder. [Online]. Available: <https://www.videolan.org/developers/x265.html>
- [44] NVIDIA Video Codec SDK. [Online]. Available: <https://developer.nvidia.com/nvidia-video-codec-sdk>
- [45] M. Viitanen, A. Koivula, A. Lemmetti, A. Ylä-Outinen, J. Vanne, and T. D. Härmäläinen, "Kvazaar: open-source hevc/h. 265 encoder," in *Proc. of ACM MM*, 2016.
- [46] OpenHEVC. [Online]. Available: <https://github.com/OpenHEVC/openHEVC>
- [47] "Intel Core i5-9500 vs Apple A13 Bionic." [Online]. Available: <https://gadgetversus.com/processor/intel-core-i5-9500-vs-apple-a13-bionic/>
- [48] ISO, "Iso/iec 14882: 2017 information technology—programming languages—c++." 2017.
- [49] *Introduction to Framework Programming Guide*. [Online]. Available: <https://developer.apple.com/library/archive/documentation/MacOSX/Conceptual/BPFrameworks/Frameworks.html>
- [50] <https://developer.android.com/ndk>, *Android NDK*.
- [51] "NVIDIA Embedded Systems for Next-Gen Autonomous Machines." [Online]. Available: <https://www.nvidia.com/en-us/autonomous-machines/embedded-systems/>


- [52] *OpenCV 4.2.0*. [Online]. Available: <https://opencv.org/opencv-4-2-0>
- [53] *Kvazaar*. [Online]. Available: <https://github.com/ultravideo/kvazaar>
- [54] <https://developer.apple.com/documentation/objectivec/Objective-C Runtime>.
- [55] *CoreML*. [Online]. Available: <https://developer.apple.com/documentation/coreml>
- [56] *ThreadPool*. [Online]. Available: <https://github.com/progschj/ThreadPool>
- [57] *KCFcpp*. [Online]. Available: <https://github.com/joaofaro/KCFcpp>
- [58] P. F. Felzenszwalb, R. B. Girshick, D. McAllester, and D. Ramanan, "Object detection with discriminatively trained part-based models," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 32, no. 9, pp. 1627–1645, 2009.
- [59] *How OpenCV-Python Bindings Works?* [Online]. Available: https://docs.opencv.org/3.4.9/da/d49/tutorial_py_bindings_basics.html
- [60] *ultralytics/yolov3*. [Online]. Available: <https://github.com/ultralytics/yolov3>
- [61] A. Paszke, S. Gross, S. Chintala, G. Chanan, E. Yang, Z. DeVito, Z. Lin, A. Desmaison, L. Antiga, and A. Lerer, "Automatic differentiation in pytorch," 2017.
- [62] N. NVIDIA, "Multi-process service," 2014.
- [63] J. Postel *et al.*, "Transmission control protocol," 1981.
- [64] *sockpp*. [Online]. Available: <https://github.com/fpagliughi/sockpp>
- [65] *JSON*. [Online]. Available: <https://github.com/nlohmann/json>
- [66] *concurrentqueue*. [Online]. Available: <https://github.com/cameron314/concurrentqueue>
- [67] *spdlog*. [Online]. Available: <https://github.com/gabime/spdlog>
- [68] M. Dantone, L. Bossard, T. Quack, and L. Van Gool, "Augmented faces," in *Proc. of IEEE ICCV Workshops*, 2011.



Zheng Yang received a B.E. degree in computer science from Tsinghua University in 2006 and a Ph.D. degree in computer science from Hong Kong University of Science and Technology in 2010. He is currently a professor in Tsinghua University. His main research interests include wireless ad-hoc/sensor networks and mobile computing.




Xu Wang received his Ph.D. degree in School of Software from Tsinghua University in 2020. He is currently a post-doc researcher in School of Software at Tsinghua University. He is a member of the Tsinghua National Lab for Information Science and Technology. His research interests include edge computing and machine learning.



Jiahang Wu received his B.E. degree in School of Software from Beijing Jiaotong University in 2020. He is currently an M.E. student in School of Software at Tsinghua University. He is a member of the Tsinghua National Lab for Information Science and Technology. His research interests include Internet of Things and edge computing.




Yi Zhao received his B.E. degree in School of Software from Tsinghua University in 2017. He is currently a PhD student in School of Software at Tsinghua University. He is a member of the Tsinghua National Lab for Information Science and Technology. His research interests include mobile computing and human mobility.




Li Zhang received the Ph.D. degrees in School of Computer Science from Hefei University of Technology, Hefei, China, in 2009. She is currently a Professor in School of Mathematics, Hefei University of Technology. Her current research interests include computer aided geometric design, computer graphics, and Internet of Things and mobile computing.



Qiang Ma received his BS degree in Department of Computer Science and Technology from Tsinghua University, China, in 2009, and Ph.D. degree in Department of Computer Science and Engineering at the HongKong University of Science and Technology in 2013. He is now an assistant researcher in Tsinghua University. His research interests include sensor networks, mobile computing and data privacy.



Zimu Zhou is currently a tenure-track assistant professor at the School of Information Systems (SIS), Singapore Management University (SMU). He received his Ph.D. degree from Hong Kong University of Science and Technology in 2015. His research interests include model compression, federated machine learning, applied machine learning.



Xin Miao received the B.E. degree from the Department of Computer Science and Technology, Tsinghua University, and the Ph.D. degree from the Department of Computer Science and Engineering, The Hong Kong University of Science and Technology. He is currently with the School of Software, Tsinghua University. His research interests include wireless sensor networks, RFID, and mobile computing.