

# REACT: Streaming Video Analytics On The Edge With Asynchronous Cloud Support

Anurag Ghosh\*  
Carnegie Mellon University  
Pittsburgh, PA, USA  
anuraggh@andrew.cmu.edu

Srinivasan Iyengar  
Microsoft  
Bangalore, India  
sriyengar@microsoft.com

Stephen Lee  
University of Pittsburgh  
Pittsburgh, PA, USA  
stephen.lee@pitt.edu

Anuj Rathore\*  
Clutterbot  
Bangalore, India  
anuj@clutterbot.com

Venkata N Padmanabhan  
Microsoft Research  
Bangalore, India  
padmanab@microsoft.com

## ABSTRACT

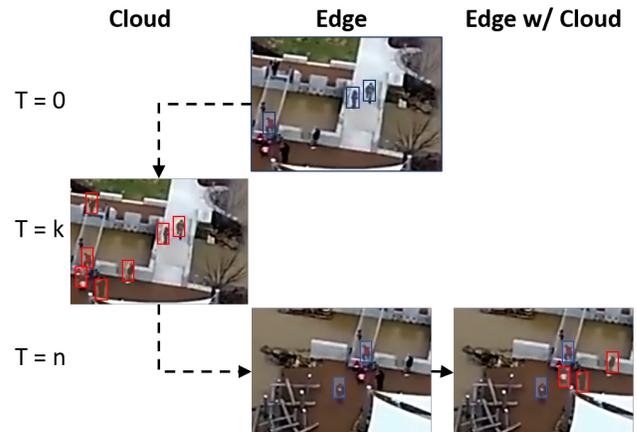
Emerging Internet of Things (IoT) and mobile computing applications are expected to support latency-sensitive deep neural network (DNN) workloads. To realize this vision, the Internet is evolving towards an edge-computing architecture, where computing infrastructure is located closer to the end device to help achieve low latency. However, edge computing may have limited resources compared to cloud environments and thus, cannot run large DNN models that often have high accuracy. In this work, we develop REACT, a framework that leverages cloud resources to execute large DNN models with higher accuracy to improve the accuracy of models running on edge devices. To do so, we propose a novel edge-cloud fusion algorithm that fuses edge and cloud predictions, achieving low latency and high accuracy. We extensively evaluate our approach and show that our approach can significantly improve the accuracy compared to baseline approaches. We focus specifically on object detection in videos (applicable in many video analytics scenarios) and show that the fused edge-cloud predictions can outperform the accuracy of edge-only and cloud-only scenarios by as much as 50%. REACT shows that for Edge AI, the choice between offloading and on-device inference is not binary – redundant execution at cloud and edge locations complement each other when carefully employed.

### ACM Reference Format:

Anurag Ghosh, Srinivasan Iyengar, Stephen Lee, Anuj Rathore, and Venkata N Padmanabhan. 2023. REACT: Streaming Video Analytics On The Edge With Asynchronous Cloud Support. In *International Conference on Internet-of-Things Design and Implementation (IoTDI '23)*, May 09–12, 2023, San Antonio, TX, USA. ACM, New York, NY, USA, 14 pages. <https://doi.org/10.1145/3576842.3582385>

## 1 INTRODUCTION

Many emerging smart video analytics applications, such as traffic state detection, health monitoring, surveillance and assistive technology require fast processing and real-time response to work effectively. Such applications rely on deep learning-based object detection models as a core part of their processing and decision-making pipeline. Unfortunately, the models are compute-intensive and tend to have large memory requirements, which limits their application in resource-constrained local environments.



**Figure 1: Illustrates the efficacy of asynchronous cloud response to improve edge performance. Note that objects are undetected on edge but detected in the cloud. Thus, cloud responses can be cascaded to improve system performance.**

Prior works have looked at *offloading* object detection to the cloud [11, 30]. By transferring data, the inference is either entirely or partially offloaded to use the computing available in the cloud. However, processing data in the cloud increases latency, making it unsuitable for near real-time analysis. For example, low latency objection detection that achieves high accuracy is highly advantageous for intelligent drones [21] or smartphone-based driver assistance [5] applications. Thus, designing architectures that achieve low latency and high accuracy would benefit all these applications.

Edge computing has emerged as an approach to address latency, where information is processed and analyzed closer to the data source. In many cases, small form-factor hardware that is low-cost and consumes lower power is often used as edge devices. However, these devices often fall short of the heavy computing needs of deep learning models. As such, there has been a significant focus on special-purpose devices – e.g., Nvidia Jetson, Google Coral – optimized to run specific DNN workloads. While edge accelerators

\*Work done while at Microsoft Research

provide improved performance over a general-purpose edge computing platform, they are still limited in their support<sup>1</sup> compared to cloud-based GPUs. Further, due to system constraints, these approaches run smaller and quantized models at the edge, with lower accuracy, compared to the larger models, with significantly higher accuracy, run on the cloud [18].

In this paper, we seek to answer the following research question: *Can cloud resources complement edge computing to achieve low latency and high accuracy?* In other words, can we achieve low latency of the edge models and high accuracy of the cloud models? In contrast to cloud-only and edge-only approaches, our key idea is to employ edge-based and cloud-based models in tandem with the cloud resources accessible over a wide-area network that may have high latency. By having *redundant computation* of object detections, we can use cloud-based inferences *asynchronously* to *course correct* edge-based inferences, thereby improving accuracy without sacrificing latency. Table 1 distinguishes our work from the prior work involving cloud-only and edge-only approaches.

Past works [18, 43] in the computer vision community have proposed using model ensemble approaches. However, they combine detections from different models with comparable performance and do so on the same frame without latency considerations. REACT’s novel *fusion algorithm* in contrast combines higher accuracy cloud-based detections on recent frames with current inference on the less-accurate edge detections while removing irrelevant stale results from the cloud.

Figure 1 illustrates how redundant computation helps improve overall accuracy for object detection. The models detect people on a flood-affected riverbank area collected from an intelligent drone at two different points in time. As shown, a cloud-based detection model achieves higher accuracy but comes with significant latency, wherein the results of a frame sent at  $t = 0$  are obtained at  $t = k$ . On the other hand, the edge-based detection model has *lower accuracy*, as several humans are not detected. Note that at  $t = n$ , even though the scene has changed, some people are still common across the current and previous frames. However, the edge model still does not detect these people. Moreover, edge results may be false positives. Thus, we use cloud-based models to improve the overall accuracy by considering detections from the accurate cloud model at time  $k < n$  and *merging* these with the frame at  $t = n$  on the edge. We note that this merge operation is not trivial. We need to consider cases where both detectors don’t agree with each other. Moreover, combining results will not work if the edge receives a cloud response after all the objects of interest within the frame change. It is necessary to ensure that approaches must work in highly dynamic environments, where objects of interest change frequently.

In this paper, we describe REACT — our system that builds on these intuitions to exploit cloud’s accuracy with the low latency of the edge. Below are our contributions.

**REACT System Design:** We designed an edge-cloud video pipeline system capable of exploiting the performance gap of object detection models between the cloud and the edge. Our approach is designed to scale to multiple edge devices and is resilient to network

variability. Finally, we develop APIs that edge-based systems can use to leverage cloud-based models and improve overall accuracy. **Edge-Cloud Fusion Algorithm:** We develop a novel fusion algorithm that combines predictions from edge and cloud object detection models to achieve higher accuracy than edge-only and cloud-only scenarios. To the best of our knowledge, we are the first to leverage redundant computations to improve the accuracy of on-edge object detection.

**Real-world Evaluation:** We evaluate REACT on two challenging real-world datasets — data collected from car dashcams [9] and drones [42]. These datasets span different cities and exhibit high variations in scene characteristics and dynamics. Our results show REACT can significantly improve accuracy by 50% over baseline methods. Further, REACT can tradeoff edge and cloud computation while maintaining the same level of accuracy. For instance, by reducing the edge detection frequency by a fourth (from every 5th frame to every 20th frame) and increasing cloud frequency (from every 100th frame to 30th frame), REACT can achieve similar accuracy.

**Scalability and Resilience Analysis:** We analyze the scalability of our approach and show REACT can support 60+ concurrent edge devices on a single machine with a server-class GPU. We also show that REACT is resilient to network variability. That is, it can function on varying network conditions and leverages cloud models when feasible. We evaluate REACT over different network types (WiFi and LTE) with varying latency using a network emulator. Our results show that even with varying response latency from the cloud, REACT performs better than the edge-only scenario.

## 2 BACKGROUND

In this section, we provide background on video-based applications and challenges in cloud or edge-based video analytics applications.

Video-analytics systems collect rich visual information that offers insights into the environment. These systems can be broadly categorized as: (i) devices that send all video to the cloud for processing, and (ii) devices that have limited processing capabilities constrained by its small form-factor, cost, or energy. In this case, the video processing can be split between the device and the cloud. That is, the device can perform either some or possibly all the processing before it sends the video to the cloud. Deep learning inference for object detection forms the core aspect of such systems.

Since deep learning is compute-intensive, existing systems typically send data to the cloud for processing. However, cloud analysis may incur significant delays and may be unsuitable for live applications. *Edge computing* has emerged as an alternative to complement the cloud, where data processing is done close to the devices to avoid these delays. A variety of edge computing architectures exist, depending on where the edge servers are located relative to the end-devices [38]. Our work assumes the edge device is of low latency, and limited computing capabilities, such as hubs in smart homes, routers, and mobile phones and IoT devices such as intelligent drones and wearable VR headsets. We assume that some form of resource constrained AI-based workloads can be run on these edge devices. Modern devices like Raspberry Pi or Jetson are devices are capable of running lightweight models [37] with a smaller memory footprint. Pairing specialized accelerators (such as Google Coral or

<sup>1</sup>Google Coral only supports integer (INT8) operations. Support for some specialized DNN layers/operations is not available in Jetson devices for FLOAT16 and INT8 operations.

**Table 1: A comparison of our approach with existing video analytics techniques. Earlier methods treat execution as a binary choice and do not consider variations in models employed on the edge and cloud. Considering such variations in REACT leads to improved performance along with resilience to network variability.**

| Features           | Our Approach | Glimpse [11] | Marlin [2] | Edge-Ast.[30] |
|--------------------|--------------|--------------|------------|---------------|
| detection at edge  | ✓            | ✗            | ✓          | ✗             |
| detection at cloud | ✓            | ✓            | ✗          | ✓             |
| n/w resilience     | ✓            | ✓            | ✓          | ✗             |

Intel Movidius) speeds up the inference time of small models without affecting accuracy for a class of model. Unfortunately, larger deep learning models (having higher accuracy than smaller models) are still not within the latency and memory budget of these devices. Larger models require cloud GPU resources, but this comes at the cost of network delays. This is unacceptable for live and streaming applications. In summary, edge processing provides a latency advantage but there remains a significant accuracy gap between real-time prediction on an edge device and offline prediction in a resource-rich setting [25]. Our goal in REACT is to leverage cloud processing in tandem with edge processing to bridge the accuracy gap while preserving the latency advantage of edge processing.

### 3 REACT DESIGN

For real-time edge inference, we propose a system that uses an edge-cloud architecture while retaining the low latency of edge devices but achieving higher accuracy than an edge-only approach. In this section, we discuss how we leverage the cloud models to influence and improve edge results.

**Basic Approach:** It is known that video frames are spatiotemporally correlated. Typically, it is sufficient to invoke edge object detection once every few frames. As illustrated in Figure 2(a), edge detection runs every 5th frame. As shown in the Figure, to interpolate the intermediate frames, a comparatively lightweight operation of object tracking can be employed. Additionally, to improve the accuracy of inference, select frames are *asynchronously* transmitted to the cloud for inference. Depending on network conditions (RTT, bandwidth, etc.) and the cloud server configuration (GPU type, memory, etc.), cloud detections are available to the edge device only after a few frames. The newer cloud detections, which were previously undetected, can be brought to the current frame using another instance of an object tracker running on the past buffered images. Video frames retain the spatial and temporal context depending on scene and camera dynamics. Our key insight is that these asynchronous detections from the cloud can help improve overall system performance as the scene usually does not change abruptly. See Figure 2(b) for a visual result of the approach.

**Challenges:** Nevertheless, designing a system that utilizes the above approach would require addressing several challenges. First, combining the detections from two sources, i.e., local edge detections and the delayed cloud detections is not straightforward. Each of these two detections contain separate list of objects represented by a  $\langle class\_label, bounding\_box, confidence\_score \rangle$  tuple. A fusion algorithm must consider several cases – such as class label mismatch, misaligned bounding boxes, etc. – to consolidate the edge and cloud detections into a single list. Second, some or all of the cloud objects

may be “stale”, outside the current edge frame. The longer it takes to perform fusion, the greater the risk of such staleness, especially if the scene changes rapidly. Thus, to minimize this risk, once the old cloud annotations are received, they must be quickly processed at the edge to help with the current frame.

Another challenge when running detection models on live videos at the edge is minimizing resource utilization while maintaining detection accuracy. Previous studies with edge-only detection systems have shown that running a deep neural network (DNN) for every frame in a video can drain system resources (e.g., battery) quickly [2]. In our case, with a distributed edge-cloud architecture, several resource constraints need to be simultaneously considered. For example, cloud detections are more accurate as one can run computationally expensive models with access to server-class GPU resources. However, bandwidth constraints or a limited cloud budget might restrict their use to once every few frames. Moreover, if the scene change is insignificant, it would be prudent not to invoke object detections at the edge and the cloud. On the contrary, for more dynamic scenes, increasing the frequency of edge detection might result in excessive heat generation from the modest GPUs used on edge devices leading to throttling.

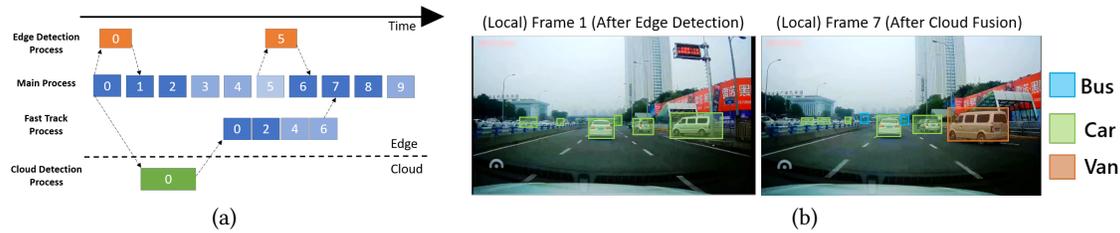
Next, we present our system called REACT, which overcomes the above challenges. Primarily, REACT consists of three components – i) REACT Edge Manager, ii) Cloud-Edge Fusion Unit, iii) REACT Model Server. Below, we describe them in more detail.

#### 3.1 REACT Edge Manager

The REACT Edge Manager (REM) consists of different modules, and put together, enables fast and accurate object detection at the edge.

**Change detector:** Previous studies have shown that running a object detection on every frame in a video can drain system resources (e.g., battery) quickly [2]. REM provides two parameters, i.e., the detection frequency at the edge ( $k$ ) and the cloud ( $m$ ) – to modulate the number of frames between object detection. Intuitively, if there is little object displacement across frames, running detection models frequently will lead to wastage of resources. REM employs a change detector that computes the *optical flow* on successive frames. This represents the relative motion of the scene consisting of objects and the camera, similar to [2, 11, 22]. Thus, the object detection invocations will only occur at a detection frequency of every  $k^{th}$  and  $m^{th}$  frame at the edge and the cloud, respectively, if this motion is greater than a pre-decided threshold.

**Edge Object Detector:** Every  $k^{th}$  frame, REM triggers the edge object detector module, which in turn outputs a list of  $\langle l, p, c \rangle$  tuples. Here,  $l$  and  $c$  are class labels (e.g., cars, person) and confidence scores (between 0 and 1) associated with the detected objects, respectively.



**Figure 2: (a) REACT System Process Flow. Orange and Green boxes indicate edge and cloud detections. Tracking performance degrades with age, which is indicated by the lighter shades of the color blue. It should be noted that Cloud Detection and Fast Track are both asynchronous events. Edge detections ensure lower tracking age, which improves localization. Cloud detections supplement edge detections, detecting missed objects and correcting class labels. (b) REACT uses *asynchronous* cloud detections to correct the box labels and detect more objects.**

$p = (x, y, w, h)$  represents the bounding box for each of the detected objects, where  $x, y$  is the center coordinate of the object;  $w, h$  is the width and height of the bounding box. To avoid multiple bounding boxes for the same object, we use Non-max suppression, which removes locally repeated detections.

**Main Object tracker:** REM employs an CPU-based object tracker, a computationally cheaper technique, between frames for which the object detections are available. For example, a CSRT [31] tracker can process images at  $>40$  fps (on Nvidia Jetson Xavier). However, as the quantum of associated displacement of objects increases, the tracker accuracy also reduces. The tracker module accounts for this degradation by multiplying every tracked object’s confidence scores by a decay rate  $\delta \in [0, 1]$ . As the confidence scores reduce with every passing frame with this multiplier, the module sweeps over the list of objects to discard the ones with lower confidence scores (i.e.,  $c < 0.5$ ).

**Cloud communicator:** The REM consists of a communication module responsible for sending every  $m^{th}$  frame (cloud detection frequency) to the cloud and receive the associated output annotations. Similar to edge detections, the cloud annotations consist of a list of  $(l, p, c)$  tuples. Since the cloud can execute larger object detection models, it provides better accuracy over lightweight models running at the edge. The communication module transmits frames asynchronously to the cloud. Again, the cloud detection frequency is based on objects’ motion and leverages the change detector module. If the change is below threshold, we do not transmit frames to the cloud for object detection. As the cloud always processes an older frame due to network latency, the predictions might become stale (i.e., fall outside the frame) by the time it reaches the edge.

**Asynchronous “Cloud”-Objects tracker:** Proposed edge-cloud fusion presents another challenge. While the detections from the edge are available for immediate use, the detections received from the cloud are delayed and may not align with objects in the current frame. This is because objects may have moved in the current frame. Using cloud detections in this scenario may lead to localization errors. To use these detections and improve the current frame’s detection, we *fast track* cloud object predictions asynchronously. Here, we start a new instance of the tracker on a new process separate from the main tracker. Specifically, we initialize this instance with cloud predictions and track the objects on every alternate frame, until it reaches the currently processed frame. We use a

stride parameter that skips frames to speed up the localization of objects detected in the cloud. Note that the stride can be increased at the cost of decreased localization accuracy. In practice, we observe tracking on every alternate frame using our algorithm performed the best.

### 3.2 REACT Model Server

The REACT Model Server’s primary goal is to respond to edge inference requests by executing the object detection models on the cloud and sending annotations of the detected objects back to the edge device. The server may be shared across numerous edge devices to handle multiple requests at any given time. A request queue is maintained with multiple worker threads (parameterized by *num\_workers*) to maximize throughput while adhering to a latency constraint. Server class GPU architectures can efficiently operate in parallel on a batch of images (say, *batch\_size* image tensors) that are dispatched together for inference. Requests are preprocessed and batched by the worker threads, and a batch is sent for inference to the GPU(s) either when a batch has *batch\_size* images for inference or when a *max\_delay* wait threshold is reached. Optimal parameter choices depend on the GPU hardware characteristics and the distribution of incoming requests. For simplicity, we do not consider dynamic batching scenarios.

### 3.3 Edge-Cloud Fusion Unit

The primary function of this component is to enhance overall detection accuracy by combining predictions from both edge and cloud models. As shown in Table 2, the cloud model (CenterNet) outperforms the edge model (TinyYolo) by detecting more objects. Our analysis indicates that the edge model struggles to detect smaller objects, resulting in lower overall object detections. Additionally, the cloud-based model performs better in terms of overall detection rate, demonstrating higher bounding box overlap (which affects localization error), confidence scores, and predicted labels (which affect classification rate) than the edge model. Even if both the models detect an object, the edge models has larger classification error compared to the cloud model, as we can observe. Also, this variation is present across classes, TinyYolo only correctly classifies 74.3% of non-car objects (biased towards classifying the majority of vehicles as cars), while CenterNet correctly classifies 82.0% of these objects.

**Algorithm 1:** Edge-Cloud Fusion Algorithm

---

```

1  $M = [][]$ 
2  $det\_source = \text{GetDetectionSource}(objects_{new})$ 
3  $objects_{current} =$ 
   RemoveOldDetections( $objects_{current}, det\_source$ )
4 for  $o_c \in objects_{current}$  do
5   for  $o_n \in objects_{new}$  do
6      $iou = \text{ComputeIOU}(o_c.bbox, o_n.bbox)$ 
7     if  $iou \geq threshold$  then
8        $M[o_c][o_n] = iou$ 
9     else
10       $M[o_c][o_n] = 0$ 
11    end
12  end
13 end
14  $curr\_objs, new\_objs = \text{LinearSumAssignment}(M)$ 
15  $updated\_curr\_objs = []$ 
16 for  $o_c, o_n \in \text{zip}(curr\_objs, new\_objs)$  do
17   if  $M[o_c][o_n] \neq 0$  then
18      $o = \{\}$ 
19     if  $det\_source == \text{"cloud"}$  then
20        $o.label = o_n.label$ 
21        $o.bbox = o_c.bbox$ 
22     end
23     if  $det\_source == \text{"edge"}$  then
24        $o.label = o_c.label$ 
25        $o.bbox = o_n.bbox$ 
26     end
27      $o.score = o_n.score$ 
28      $o.score = \text{decay}(o.score)$ 
29      $o.last\_det\_source = det\_source$ 
30      $updated\_curr\_objs += o$ 
31   else
32      $o_n.last\_det\_source = det\_source$ 
33      $updated\_curr\_objs += o_n$ 
34   end
35 end
36 return  $updated\_curr\_objs$ 

```

---

It is worth noting that box fusion techniques, such as non-maximum weighted (NMW) [40] or NMS, combine predictions based on class labels and only consider a match if the overlap of the bounding box is high for the same class. If the labels are different, these techniques will consider them as two separate objects. Furthermore, our analysis reveals that edge detection models could correctly localize objects but often had false positives, i.e., they assigned class labels incorrectly. If we used the techniques above, the same object would be counted twice, thereby decreasing the overall accuracy. We use these insights to develop a novel bounding box fusion algorithm to combine cloud-edge predictions that prioritizes cloud-based detections and draws inspiration from earlier works on IoU based tracking [6] and association strategies [8, 24] and detection ensembling [18, 40] approaches.

Our box fusion technique works as follows. In the edge, we maintain a current list of objects (in the form of tuples described earlier) for the present frame. Whenever any new detections, either

from the cloud or the edge, are available, we first delete the old objects from the current list that were last submitted by the same detection source. For example, we delete old objects detected by the cloud (or edge) when newer cloud (or edge) detections are available.

Next, we create an Intersection over Union (IoU) matrix that indicates the overlap between current objects and the detections received. IoU is the ratio of overlapped area with the union of the area between the two sets of objects. Any value smaller than a threshold ( $\geq 0.5$ ) is set to 0. We then perform a linear sum assignment [8], which matches two objects with the maximum overlap. This matrix provides a list of objects that were already present in the current object list. We modify the confidence values, bounding box, and class label based on the new detections' source. For example, objects from the cloud obtained from running bigger models will be more accurate in predicting the class correctly. We present the pseudo-code to determine the merging of the boxes in Algorithm 1.

## 4 IMPLEMENTATION

**REACT Edge Manager:** Our implementation uses OpenCV to receive a stream of video images. Further, the tracking module is built upon OpenCV's object tracker API. We train two object detection models (MobileNetV2-SSD [37] and TinyYOLO [35]) for the edge scenario. We also deployed these models on an Nvidia Jetson Xavier device for inference.

**REACT Model Server:** Our object detection models (Faster R-CNN [36] and RetinaNet [28]) are trained using the mmdetection [10] library, and for training CenterNet [41] models we utilize the official implementation written in Pytorch. Pytorch's default object\_detector handler only supports torchvision models, so we implemented custom handlers for generic mmdetection models and the CenterNet model to serve them on the cloud server using Torch Serve. We created two handlers as the mmdetection library and the CenterNet library expose and utilize very different model initialization, preprocessing and postprocessing programming paradigms in their implementations. Thus, REACT can be used with newer object detection algorithms in the future by modifying the sample handlers for reflecting model specific changes. We serve these models as HTTPS/JSON endpoints over an API.

**REACT fusion API:** We expose two classes CloudServerInference and ReactEdgeInference. CloudServerInference class can be instantiated by providing the address of the HTTP endpoint, the image resolution, and the model along with optional frequency parameter, tracker type and number of tracker threads (if used in Server only inference mode). To instantiate ReactEdgeInference we specify parameters such as the model to run, image resolution, tracker type, number of tracker threads, the frequency parameters, and CloudServerInference object to use. Both the classes expose a get\_annotations method and use the image as input, returning the annotation output as a JSON object. We believe this API design facilitates adoption by application developers due to its simplicity and ease of use.

## 5 EVALUATION METHODOLOGY

In this section, we give a detailed description of the datasets used and the evaluation setup.

**Table 2: Edge models perform worse on detecting and classifying objects, specially for small objects. mAP (capturing overall performance) for small objects is far lower. Percentage of objects correctly detected (overlap greater than 0.5 IoU; label ignored) is far higher for the representative Cloud Model compared to the Edge Model. Among the correctly detected objects, we observe that a far lower percentage are correctly classified.**

| Model             | $mAP_{Small}$ | % Detected (All) | % Detected (Small) | Classification Accuracy (Small) |
|-------------------|---------------|------------------|--------------------|---------------------------------|
| CenterNet (Cloud) | 12.9          | 57.8%            | 43.6%              | 72.3%                           |
| TinyYOLO (Edge)   | 1.2           | 41.5%            | 23.0%              | 52.0%                           |

## 5.1 Dataset Description

We extensively evaluate the proposed system’s efficacy on two datasets in built environment monitoring domain highlighting its potential in different use cases (drone-based surveillance and dashcam-based driver assist). Both these datasets are popular and are among the largest available dataset for edge-based object detection. These datasets are quite challenging as they exhibit significant scene change and have a varied number and size of objects. Table 3 provides a summary of the two datasets.

**$D^2$ -City [9]:** The video dataset is created from front-facing car dashcams and captures the dynamic complexity of real-world traffic conditions. The dataset is crowdsourced from passenger vehicles registered on DiDi’s platform and intended for improving vision technologies, driving intelligence, and similar use cases. It has 1000 driving videos taken in five different cities under various scene conditions and video resolution. Objects in each video frame are annotated and include their bounding boxes and class ids.

**VisDrone [42]:** The videos in the dataset are captured using drones flown over different cities under various weather and lighting conditions. It contains 79 video clips with around 1.5 million manually annotated objects. We use the evaluation protocol followed in VisDrone-VDT 2018 challenge for video object detection task [43], which focuses on detecting specific objects (e.g., pedestrian, car, van) taken from drones. The object detection task in these videos is considered challenging due to the density of really small objects, dynamic scene conditions, and drones’ movement.

## 5.2 Performance Metrics

We use *mean average precision at intersection over union (IoU) = 0.5* ( $mAP@0.5$ ) – a popular metric used for object detection tasks [13, 29]. Note that the IoU measures the ratio of the intersection area and the area of union of the predicted bounding box and ground truth bounding box. Thus, a prediction is considered a true positive if the predicted label matches the ground truth, and the IoU is greater than or equal to the threshold ( $\geq 0.5$ ).

## 5.3 Evaluation Setup

In this section, we discuss the training process, baseline techniques and environment.

**5.3.1 Model selection and training.** We use a combination of deep learning models to evaluate our approach, where we execute different models on the edge and cloud. For our cloud-based models, we use Faster-RCNN [36], RetinaNet [28]. For edge models, we use TinyYOLO [35] and MobileNetV2-SSD [37]. Table 4 provides

a summary of the different models. To train our models, we follow the protocols described in the  $D^2$ -City and VisDrone datasets. As these datasets are released as part of ongoing challenges, the test set annotations are not publicly available. Hence, we evaluate our models on the released validation data set. For our validation dataset during training, we use 15% from the train data set to tune the hyper-parameters and select the final model.

**5.3.2 Baseline Techniques.** We use the following baseline techniques to compare with our proposed approach. Since the code for existing work was not available, the baselines we use is comparable to prior work.

**Edge-only Inference:** Here, we run the object detection only at the edge and do not offload detection tasks to cloud resources. This baseline is similar to *Marlin* [2] that employs the edge-only strategy (see Table 1). However, unlike Marlin, which minimizes edge inference frequency to conserve energy, our edge-only baseline invokes edge inference at regular intervals. As a result, this baseline achieves better performance than Marlin. In the edge-only baseline, we use lightweight detection models – TinyYOLO and MobilNetV2-SSD – as they consume less memory and computation and are well-suited for resource-constrained edge devices. From hereon, we refer to this baseline as *edge-only*.

**Cloud-only Inference:** For this baseline, we run the object detection task on the cloud. The edge is a thin client that offloads the detection tasks to the cloud while using a tracker to compensate for intermediate frames. The performance of this baseline setup is comparable to existing systems such as *Edge-Assisted* [30] and *Glimpse* [11] (See Table 1) that offloads trigger frames to the cloud and uses an optical flow based object tracking method to update the object bounding boxes on the other frames. Note that cloud-only inference suffers from higher network delays compared to the edge-only scenario [30]. Such high network latency may be undesirable for latency-sensitive applications as dynamic changes in scenes may render responses from the cloud unusable. We use computationally expensive detection models on cloud, namely RetineNet, Faster RCNN, and CenterNet, due to their good performance. From hereon, we refer to this baseline as *cloud-only*.

**Every Frame Edge Inference:** In this scenario, we compare REACT with the case where one can run detectors using edge models (TinyYOLO and SSD-MobileNetv2) on every frame. Unlike the edge-only baseline, we do not interpolate predictions with any tracker. In practice, this baseline is infeasible as edge devices cannot run detections on all frames due to latency and energy constraints. We call these baselines *ef-edge-det (tinyyolo)* and *ef-edge-det (ssdmv2)*. We do not compare with Every Frame Cloud Inference as it neither meets the computational budget nor the latency budget.

**Table 3: We evaluate our approach on two large-scale real world datasets.  $D^2$ -City dashcam videos simulate driver assistance scenarios, and exhibit heavy occlusions and large variations in object sizes (far-away small objects). Visdrone videos act as drone sensing scenario, exhibiting variations in altitude (affecting average size of objects), camera viewpoint, and environmental clutter with large number of objects.**

| Name        | Type    | Size (#videos, #frames) | # of Classes | Remarks              |
|-------------|---------|-------------------------|--------------|----------------------|
| VisDrone    | Drone   | (79, 33.3K)             | 12           | Altitude, View Angle |
| $D^2$ -City | DashCam | (1000, 700K)            | 8            | Varied object sizes  |

**Table 4: We utilize the following models in our experiments. The choices of hardware are dictated by the memory use and latency of the model [18]. TinyYOLO and MobileNetV2-SSD are the archetypal model choices for mobile hardware in the vision community. Faster R-CNN, RetinaNet and CenterNet are popular detectors that are expected to be employed on server class GPUs.**

| Detector     | Backbone     | Where | #params |
|--------------|--------------|-------|---------|
| Faster R-CNN | ResNet50-FPN | Cloud | 41.5M   |
| RetinaNet    | ResNet50-FPN | Cloud | 36.1M   |
| CenterNet    | DLA34        | Cloud | 20.1M   |
| TinyYOLOv3   | DN19         | Edge  | 8.7M    |
| SSD          | MobileNetV2  | Edge  | 3.4M    |

5.3.3 *Network Emulation.* We use Mahimahi [32] and traffic control (tc) Linux utility to emulate different network traffic, in particular, LTE and WiFi. For LTE, we use the Verizon LTE uplink and downlink traces in MahiMahi to emulate LTE link between the edge and cloud [32] (hereon, we refer it as LTE). For WiFi, we throttle the traffic to 24Mbps and also introduce delay of 30ms and 50ms using tc. Hereon, we refer them as WiFi (30 ms) and WiFi (50 ms). Thus, we emulate three different network conditions between the client and the server. Unless stated otherwise, we report our results using the WiFi (30 ms) network.

## 6 EXPERIMENTAL RESULTS

In this section, we compare REACT with other baseline techniques. We also study the impact of network conditions and the tradeoff opportunities from adjusting the detection frequency at both the cloud and the edge. Further, we evaluate the scalability of our approach and its performance on an edge accelerator device.

### 6.1 Performance Comparison

We first evaluate how REACT’s use of redundant detections running asynchronously on the cloud help achieve low latency and improves accuracy. In this experiment, we set the edge and cloud object detection frequency to 5 and 30, respectively. We compare REACT to our three baseline approaches and report our results for both **D2-City** and **Visdrone** datasets. For a fair comparison, the baseline methods also use the same cloud/edge object detection frequency.

Figure 3(a) compares baseline algorithms with REACT (i.e. cloud-edge) with respect to the object detection accuracy (mAP@0.5)

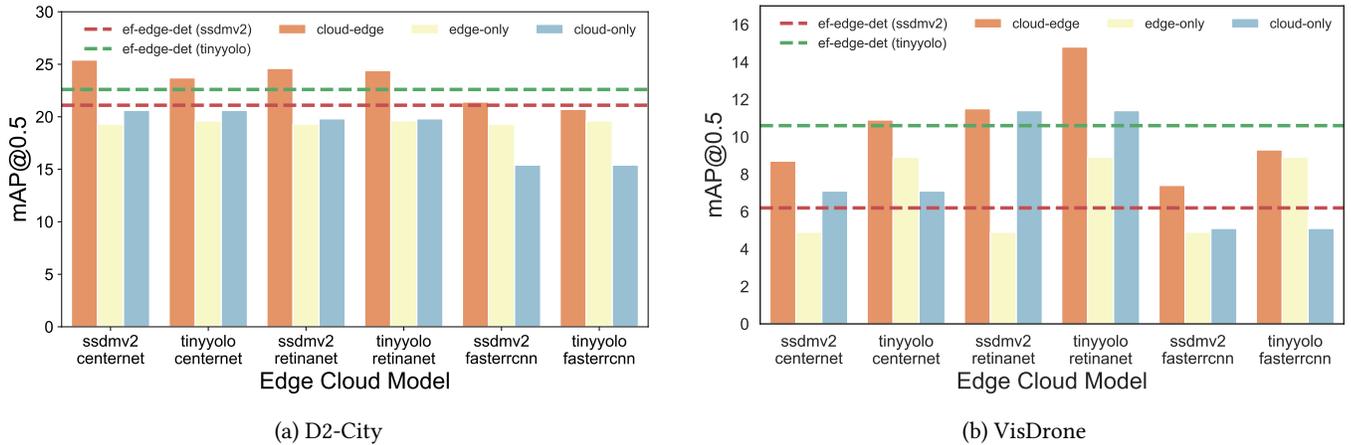
for D2City dataset. We create distinct pairs of object model combinations – one running at the edge and the other on the cloud. Specifically, we evaluate using two edge models and three cloud models, a total of six combination pairs. Our results show that REACT outperforms the edge-only and cloud-only baselines by 20-40% for all combination pairs. Different object detection models exhibit different kind of errors, due to their DNN architectural design decisions, and REACT is able to combine these detections to reduce overall error and improve performance. This is akin to using an ensemble of cascading detection models in tandem to reduce error. We also observe that our approach’s mAP is marginally better than the scenario where edge models are executed on every frame (i.e., *ef-edge-det*), where no latency constraints on edge device is assumed. This implies that edge models exhibit certain kinds of errors that can be rectified by cloud models (See discussion in Section 6.4.2). This impractical scenario shows us that redundant computation on the cloud is complementary to on-device execution on the edge. In particular, the cloud-edge pair of CenterNet and SSD MobileNetv2 are very complementary and together achieve the best performance. Qualitative results can be seen in Section 6.6, we observe that cloud models are able to detect small sized and heavily occluded objects that edge models tend to miss.

Figure 3(b) shows the same comparison using the Visdrone dataset. As noted in prior studies, object detection in this dataset is challenging, and models tend to have low mAP values [43]. Our results show that REACT achieves higher accuracy and outperforms baselines by 50%. We also observe that the pair of RetinaNet and TinyYolo outperforms all baseline techniques.

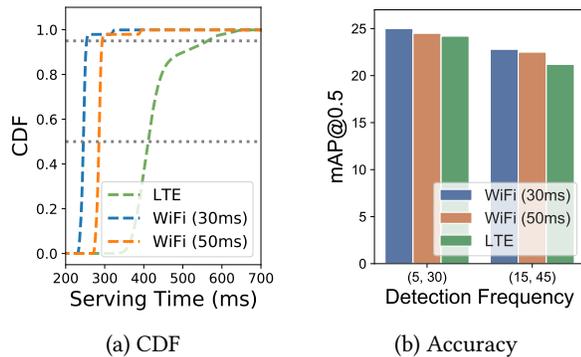
**Key Observations:** REACT *outperforms baseline algorithms by as much as 50%. Edge and Cloud models are complementary in their strengths and weaknesses, and overall performance be improved using our edge-cloud fusion algorithm.*

### 6.2 Impact of Network

As discussed earlier, REACT receives responses asynchronously from the cloud and merges its annotations with the edge detections. Clearly, stale cloud responses affect accuracy. There are three factors that affect the *servicing time* of responses from the cloud – (i) time to transmit a frame from the edge to the cloud, (ii) time to run inference on the frame at the cloud, and (iii) time to send the annotations from the cloud to the edge devices. Thus, we experiment with different networks to gauge their impact on the overall accuracy. We restrict our evaluation to the D2-City dataset with tinyYOLO and RetinaNet models running at the edge and the cloud, respectively.



**Figure 3: Comparison of REACT (cloud-edge) across two datasets compared to edge-only [2] and cloud-only [11, 30] approaches. REACT performs better than the Edge upper-bound (every frame execution), the deficiency on edge is rectified by the cloud.**



**Figure 4: Impact of Network latency on accuracy for different network conditions. Detection accuracy is marginally affected with network latency, and can be mitigated by increasing edge inference frequency.**

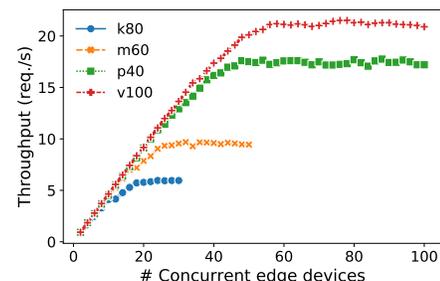
Figure 4(a) show the cumulative distribution function (CDF) of the serving times observed on the four network conditions. The two gray-colored horizontal lines represent 50<sup>th</sup> and the 95<sup>th</sup> percentiles. Using WiFi (30ms), we get the lowest serving time, i.e., around 260 ms (95<sup>th</sup> percentile). Whereas, LTE has a significantly longer serving time compared to others (420 ms for 50<sup>th</sup> and 570 ms for 95<sup>th</sup> percentile). Unlike other network types, LTE also has a much higher standard deviation. The Figure 4(b) shows the accuracy associated with the use of the four network types. Intuitively, accuracy degrades as serving times increase. This is because a change in the scene may render the stale output from the cloud useless. Thus, in the worst case, dynamic scenes where objects change frequently, such scenarios may not be able to take advantage of cloud resources. We observe this pattern in our analysis, where higher delays in serving time reduce accuracy. In particular, the model accuracy with LTE is the lowest at 21.1 — i.e., 7% lower than WiFi (30ms) in case of edge-cloud detection frequency at (15,45).

**Key Observations:** REACT performance is sensitive to different network conditions. Specifically, a 310 ms difference in 95<sup>th</sup> percentile serving time in network type results in 7% reduction in accuracy.

### 6.3 REACT’s Scalability

We discuss how the added cost of additional cloud resources be amortized over many edge devices sharing the same REACT model server.

To evaluate the scalability of REACT Model Server, we looked at four different generations of GPUs (i.e., K80, M60, P40, and V100) available on the cloud platforms. Consequently, we selected Microsoft Azure Ubuntu 18.04 VMs NC6v1 (K80), NV6v3 (M60), ND6v1 (P40), and NC6v3 (V100). As the two datasets consisted of various image resolutions, we choose a consistent image size (512 × 512) for a fair comparison. We set the inference max batch size to 4 and use the Faster-RCNN model for the results discussed here (our most expensive cloud model). We benchmark using the HTTPS/JSON endpoint and define the payload and user characteristics using the *Locust load testing library* [17]. We looked at a scenario where the edge devices send requests once every 2 seconds (once every 60 frames). The payload involved adding users at a uniformly random rate of 3 edge devices per second until we reached the maximum



**Figure 5: Throughput vs #concurrent edge devices for different GPUs.**

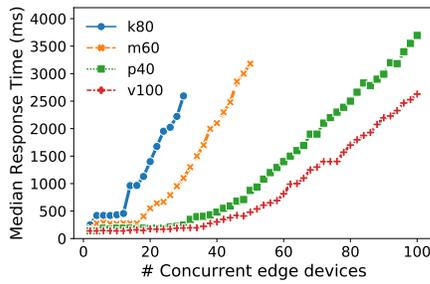


Figure 6: 50<sup>th</sup> percentile response time vs # concurrent edge devices

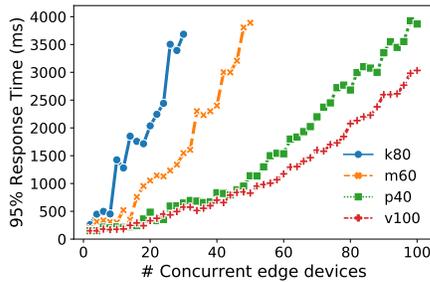


Figure 7: 95<sup>th</sup> percentile response time vs #concurrent edge devices

desired number. Specifically, we varied the concurrent number of edge devices sending requests between 2 to 100.

Figure 5 shows the throughput of the serving platform with a varying number of edge devices for the different GPU VMs. For a smaller number of devices, the GPUs are underutilized, and the throughput increases. However, each of the four GPUs will hit a maximum throughput level with the increasing number of edge devices. For newer GPU devices, such as V100 and P40, we get a maximum throughput of over 17 requests per second (req./s). Throughput can be increased by batching requests with a timeout queue at the expense of average latency. Whereas, the performance of the K80 is the worst, with throughput maxing out at slightly over 5 req./s. Thus, during lower traffic conditions, one can go with older GPUs available at a discount compared to newer ones (the pricing is dynamic and based on demand). However, the newer GPUs can provide > 3× the performance.

If an application can tolerate a median latency of 500 ms for inference on the cloud, we can support up to 60+ concurrent devices at a time using the V100 GPU (see Figure 6). If we consider a Reserved VM with a V100 GPU<sup>2</sup>, the cost is 1.63¢/hr. per concurrent device. This is a conservative analysis due to our model choice – detectors less expensive than Faster RCNN (like RetinaNet) can support greater number of concurrent devices. This number reduces to 44, 19, and 12 for P40, M60, and K80, respectively. For 95<sup>th</sup> percentile case, V100 can support 33 concurrent devices (see Figure 7). Moreover, for many video analytics applications not all edge devices are operational at all times. For example, one might

<sup>2</sup>Cost of a 3 year reserved Azure VM is 0.979\$ an hour. See <https://azure.microsoft.com/en-us/pricing/details/virtual-machines/linux/>

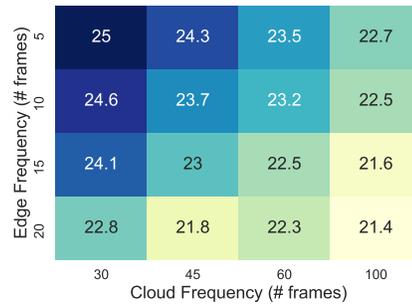


Figure 8: Higher detection frequency on the cloud and lower detection on edge may result in similar accuracy compared to lower detection frequency on the cloud and higher detection frequency on edge.

use an AR/MR app on a mobile device for just 20 minutes a day. Similarly, a dashcam-based driver-assist application will only be operated while driving (around one hour a day). The overall number of edge devices supported will be orders of magnitude greater than the current devices supported.

**Key Observations:** A single instance of the REACT Model Server can handle an excess of 60 concurrent edge devices. We can divide the cost overhead of the VMs across hundreds of edge devices as only a few devices are operated at any given time for several real-time video analytics applications.

## 6.4 REACT Parameter’s Analysis

**6.4.1 Impact of Detection Frequency.** Most resource-constrained systems cannot execute deep learning-based object detections on each frame. Typically, the object detector runs only once every few frames and a lightweight object tracking is performed on intermediate frames. As we noted earlier, there are tradeoffs between executing a detector on the edge compared to the cloud. We can exploit this tradeoff between computation on edge and cloud by changing detection frequency parameters.

For our evaluation, we set the detection frequency and invoke edge and cloud models every  $X$  number of frames and use RetinaNet as our cloud model and SSD MobileNetv2 as our edge model. Figure 8 shows a heatmap indicating the accuracy of REACT using different edge and cloud detection frequencies. As expected, running more detections improves accuracy as it mitigates the degradation effects of object tracking. Moreover, if the scene changes frequently, the cloud detections may be stale, which may further contribute to degraded performance. And thus, invoking frequent detections at the edge helps in mitigating these effects.

In particular, we can reduce the frequency at the edge (or cloud) and increase at the cloud (or edge) with little impact on overall accuracy. For example, running edge detections every 5th frame and cloud detections every 100th frame results in mAP@0.5 of 22.7. However, we can instead trade-off computation and reduce the detection frequency at the edge by a fourth (e.g., run every 20th frame) and slightly more than triple the cloud frequency (e.g., every 30th frame) to achieve a similar accuracy (mAP@0.5=22.8).

Such a scenario is quite common in edge devices where excess heat generated by running detectors often might result in throttling. If cloud resources are at a premium, we can get similar accuracy ( $mAP@0.5=22.5$ ) with the edge and the cloud frequencies set to every 15th and 60th frame, respectively. Such flexibility allows application developers to perform tradeoffs to optimize for specific objectives. These changes to cloud and edge detection frequencies to maintain similar accuracy also highlight the resilience of REACT to network variability. Reducing cloud detections forced by lower bandwidth can be compensated with higher edge detections.

**6.4.2 Diagnostic Error Analysis.** REACT outperforms the baseline algorithms and also improves on the upper bound performance of using edge detections on each frame. Moreover, we can change detection frequencies and observe similar overall accuracy. However, mAP alone does not explain the effect of the various system parameters and the end-task tradeoffs they introduce.

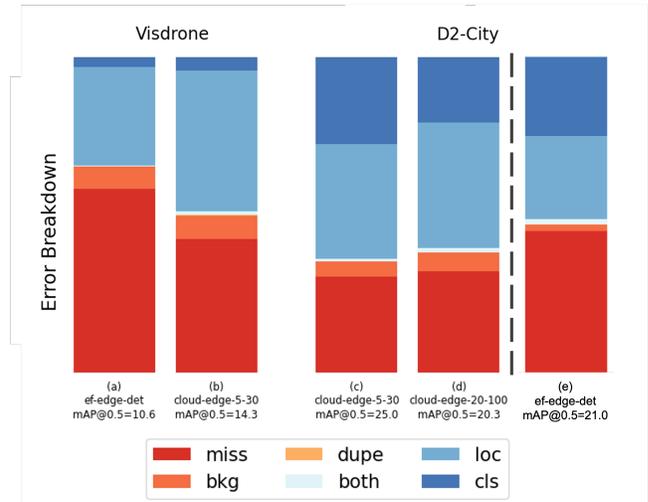
To this end, we use TIDE [7], a toolbox that helps disambiguate between six error types in object detection (Cls: classification error; Loc: localization error; Both: both cls and loc error; Dupe: duplicate predictions error; Bkg: background error; Miss: missed detections error). TIDE assigns the different error types independently for every detection error (i.e. computes change in mAP, if an error type was “fixed” by an oracle), and thus provides us relative breakdowns that can be compared. This is in contrast with earlier methods which classified all errors progressively [29] (thus summing error breakdown to  $1 - mAP$  score), but that is strongly biased toward error types fixed last [7].

We analyze the error breakdown of REACT at different detection frequencies for the tinyYOLO-RetinaNet combination (like Section 6.1). It’s clear from Figure 9 ((b) and (c)) that the kind of errors made by REACT on the two datasets are very different. On D2-City dataset, we see a substantially larger ratio of classification (class label mismatch) errors compared to VisDrone dataset, and a smaller ratio of missed detections. Thus, target domain is an important aspect in discussion of system tradeoffs.

We consider scenarios to demonstrate exactly how parameter choices effect the errors (and mitigation steps). Developers can adjust REACT’s parameters, such as changing cloud/edge detection frequency to reduce localization errors or missing detections, and how cloud detections improve the overall system performance.

**Effect of Cloud Detections:** We analyze the error breakdown between running edge model (tinyYOLO) on every frame and REACT at (5,30) on Visdrone dataset. From Figure 9 ((a) and (b)), the ratio of missed detections is substantially lower for REACT contributing to increase in mAP (See Fig 3 (b)) from 10.6  $mAP@0.5$  to 14.3  $mAP@0.5$ . This indicates is that the cloud models help in detecting objects that edge models are not able to detect (such as small and occluded objects). We can observe the same trend on the D2City dataset in Figure 9 ((c) and (e)).

**Reducing Cloud and Edge Frequencies:** Further, we analyze the results by varying edge and cloud frequency from (5,30) to (20, 100) on D2City dataset. From Figure 9 ((c) and (d)), the ratio of localization errors increases as the overall mAP decreases (See Figure 8) from 25  $mAP@0.5$  to 20.3  $mAP@0.5$ . This indicates that tracker error increases as we reduce the frequencies which could be mitigated by using a more accurate yet more expensive tracker.



**Figure 9: Error breakdown on the two datasets for different REACT configurations. Cloud detections drastically lower miss rate, and increasing edge detection frequency lowers localization errors.**

However, if localization errors are tolerable in the end use-case, for instance in person counting scenarios, then savings in both cloud cost and energy on the edge device can be made.

**Key Observations:** *The flexibility to adjust detection frequency can immensely help resource-constrained scenarios. REACT provides the flexibility to tradeoff computation at the edge and cloud, while achieving similar performance. REACT can further mitigate different types of errors by changing system parameters and iterating on specific performance bottlenecks.*

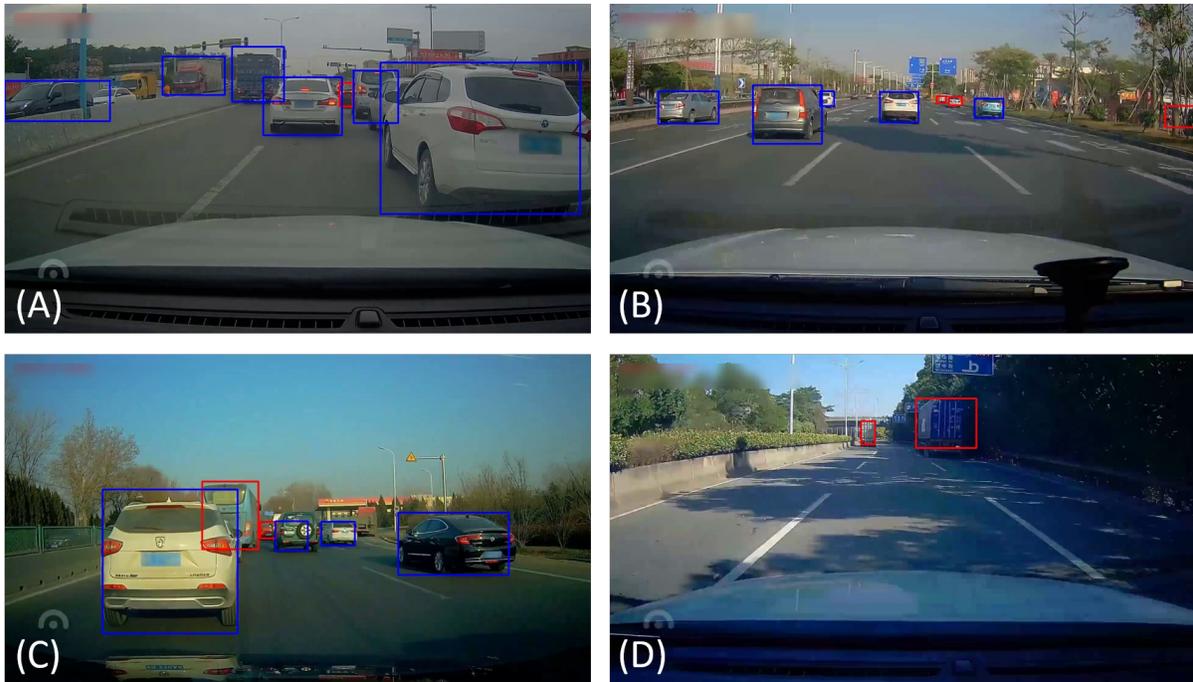
## 6.5 Performance on Edge Devices

We evaluate the feasibility of REACT on the Nvidia Jetson Xavier device with installed JetPack SDK. Specifically, we deploy REACT on the device and calculated the maximum FPS obtained for TinyYOLO edge model and the CSRT tracker employed in the REACT Edge Manager. We achieve an average detection rate of 26.1 fps for a video stream for an image resolution of  $540 \times 360$ . For streaming applications (30 fps), we cannot invoke detection very often. Additionally, our tracker algorithm achieved 36.66 fps ( $> 30$ fps). Thus, it is feasible to use REACT for many video analytics applications where object detection is a crucial block.

**Key Observations:** *It is feasible to run REACT on edge-class devices. Reducing object detection frequency at the edge (while increasing the cloud detection frequency) can offer opportunities to execute downstream tasks in the analytics pipeline.*

## 6.6 Qualitative Results

We visualize representative frames from various sequences in the D2City dataset and Visdrone dataset using REACT (TinyYOLO-RetinaNet configuration from Section 6.1). REACT’s Edge Cloud Fusion Algorithm helps in multiple scenarios. In Fig 10 (A), the



**Figure 10: Detections on D2City.** Cloud detections are colored red, whereas Edge detections are colored blue. We observe that Edge models detect larger objects, but fail when the object is small or heavily occluded, and those cases are rectified by cloud models.

edge model is able to identify and localize most of the objects, however, cloud model identifies a highly occluded car. While in Fig 10 (B), the cloud model is able to identify small objects (such as the cars far away) which the edge model could not. The cloud model is able to identify the occluded bus, which is close to the camera in Fig 10 (C). The edge model performs especially poorly in Fig 10 (D), as it’s not able to identify any of the trucks due to inconsistent lighting conditions, which our cloud model can identify and localize correctly. Similar patterns emerge in VisDrone dataset, as observed in Fig 11. Moreover, as we can see in all the sub figures (specially in Fig 11 (A)), the miss rate is significantly reduced by the detection of smaller objects by the cloud model. These results are consistent with prior observations that larger models are better at detecting small, occluded and rare kinds of objects in Section 6.4.2.

## 7 RELATED WORK

In this section, we contextualize our work with other studies.

**Edge-based hardware:** Hardware accelerators [15, 20, 33] have shown to boost the performance of deep learning inference. Such accelerators are deemed suitable for edge-AI use cases at a much lower cost and energy needs. Our work uses similar resource-constrained edge devices capable of executing lightweight deep learning-based models. Today there is a performance gap between edge and cloud computing capabilities that our work leverages. Thus, our work is applicable in current and future systems where a similar performance gap exists.

**ML Model Optimizations:** There have been a few major ways of optimizing models themselves to reduce the inference time on the resource-constrained edge devices – model pruning [16], quantization [19], distillation [4] and hardware-aware neural architecture search [39]. Unfortunately, the improvements in latency largely come at a cost of lower accuracies and generalization. Our approach is complementary to these approaches as we expect the performance arbitrage to exist and our results show that fusing the output can improve the overall accuracy. Moreover, any complementary improvement in the performance of small models reduces the dependence on the cloud for inference, increasing the concurrent clients our system can support.

**Video Analytics Optimizations:** Live video analytics is emerging as an increasingly important problem because of its applications in multiple domains [1]. However, providing efficient video inference remains a challenge due to constraints in compute, latency and bandwidth. As such, several studies have looked at optimizing several aspects within the video analytics pipeline to improve overall performance [2, 11, 30]. Several papers have considered offloading the analysis to the cloud [3, 12, 22]. Most studies that offload work to the cloud assume that there are no stringent latency requirements. Some of them focus on optimizing video queries on the cloud by selecting appropriate neural network and video configurations to save compute resources [22].

Moreover, there have been studies that look at leveraging both on-board compute and/or cloud resources to improve object detection [2, 11, 30]. RedEye [27] performs early CNN computation



**Figure 11: Detections on Visdrone. Cloud detections are colored red, whereas Edge detections are colored blue. We observe that Edge models detect larger objects, but fail when the object is small or if it’s a rare class, and those cases are rectified by cloud models.**

in the analog domain on the image sensor. Marlin [2] proposed a detection technique for mobile-based AR applications that switches between lightweight object tracking and on-device inference for object detection. Works like RedEye and Marlin’s frame selection procedure are complementary to our approach. Reducto [26] investigates on-camera filtering, and dynamically adapts filtering decisions according to the time-varying correlations. These approaches are complementary and can be used to reduce our edge detection frequency further.

Glimpse [11] presents a real-time object recognition pipeline that does object tracking locally but offloads DNN-based object detection to the cloud. DeepDecision [34] is measurement driven framework that considers running an object detector on the cloud or the edge depending on network conditions and edge hardware constraints. In contrast to these prior work, our analysis shows that redundant inference on both edge and cloud are complementary and improve accuracy compared to baseline techniques that are based on these existing works.

Separately, there have been several recent efforts to partition models across the cloud and edge (Eg. [23]). Such techniques are not suitable for live analytics because the final result is primarily computed on the cloud, which increases overall latency.

## 8 IMPLICATIONS AND DISCUSSIONS

Unlike prior works [2, 11, 30, 34] that consider the choice between offloading to the cloud and on-device execution on the edge as the only two possibilities, REACT demonstrates that for Edge AI scenarios, a nuanced approach is useful as the decision need not

be binary. Approaches leveraging these two choices have their inherent deficiencies — vision models on the edge are fast but suffer from object classification errors due to small model size, while vision models on the cloud are slow but suffer from object localization errors due to network latency. Leveraging predictions both of these choices, although *imperfect*, through our fusion algorithm of these two leads to improved overall performance.

**Flexibility:** While we evaluate network latency and analyze the impact of detection frequency on edge, network bandwidth is also important. Since our approach allows the flexibility to change cloud detection frequency, we can control the data sent across the network to conserve bandwidth. However, we can still achieve similar accuracy by increasing the detection frequency at the edge. Thus, users of REACT can achieve comparable accuracy by choosing a wide range of system parameters while satisfying use-case specific constraints, such as limited bandwidth or edge GPU cycles. Further, the modular design of REACT allows developers to swap models at the edge or the cloud as and when newer and improved DNN architectures are available. Our system also allows developers to choose a model serving system of their choice.

**Generalizability to Tasks and Applications:** Even though we evaluate our system on object detection tasks, we expect our approach to also work on human pose-estimation or instance segmentation applications. For example, human pose-estimation applications require instantaneous feedback for sports and dance activities and to understand full-body sign language — all of which require low latency analysis. In addition, instance segmentation tasks in security and surveillance applications with robots also

require low latency, making our approach useful in such scenarios. Additionally, our approach can be applied to video-based activity recognition to monitor physical activities (e.g., walking, running) in a non-intrusive way, which is useful for health monitoring purposes. Currently, these approaches rely on large machine-learning models that are compute-intensive and not suitable for edge scenarios. Our future work will involve extending our system to work for such applications.

**Adaptive parameter setting:** We note that the detection frequency was fixed for our evaluation to show trade-off opportunities. However, the detection frequency can be adaptive and change based on variations in scene dynamism. For example, if the scene changes less frequently, we can decrease the detection frequency at the edge and/or the cloud to keep up with the desired accuracy. Detection frequency can also change due to systems constraints. If there is limited cloud resource available, one can reduce the cloud detection frequency. When cloud resources are cheap, increasing the cloud detection frequency can improve detection accuracy. Likewise, if the edge device experiences thermal throttling or is constrained by power consumption, then lowering edge detection frequency is necessary (say for battery-operated drones). Concurrent work [14] has shown the feasibility of learning configurations for live streaming applications.

## 9 CONCLUSION

In this work, we introduced REACT— a novel approach that leverages both edge and cloud resources to improve live video analytics applications. Unlike prior work that optimizes cloud or edge-only inference, our approach combines edge and cloud inference results to improve overall accuracy. REACT utilizes higher accuracy object detections from the cloud to improve the past edge detections and cascade these to current predictions on edge. REACT is flexible, resilient to network latency, cost-effective, and scalable. We evaluated the efficacy of our approach on two challenging real world datasets and showed that our edge-cloud fusion approach can achieve higher accuracy than edge-only and cloud-only scenarios. In particular, by leveraging redundant computations at the cloud, our approach outperforms baseline algorithms by as much as 50%.

## REFERENCES

- [1] Ganesh Ananthanarayanan, Victor Bahl, Landon Cox, Alex Crown, Shadi Noghahi, and Yuanhao Shu. 2019. Demo: Video Analytics-Killer App for Edge Computing. In *Proc. ACM MobiSys*.
- [2] Kittipat Apicharttrisor, Xukan Ran, Jiayi Chen, Srikanth V Krishnamurthy, and Amit K Roy-Chowdhury. 2019. Frugal following: Power thrifty object detection and tracking for mobile augmented reality. In *Proc. SenSys*. 96–109.
- [3] Ashwin Ashok, Peter Steenkiste, and Fan Bai. 2015. Enabling vehicular applications using cloud services through adaptive computation offloading. In *Proceedings of the 6th International Workshop on Mobile Cloud Computing and Services*. 1–7.
- [4] Mohammad Farhadi Bajestani and Yezhou Yang. 2020. TKD: Temporal Knowledge Distillation for Active Perception. In *Proc. WACV*. 953–962.
- [5] Ravi Bhandari, Akshay Uttama Nambi, Venkata N Padmanabhan, and Bhaskaran Raman. 2018. DeepLane: camera-assisted GPS for driving lane detection. In *Proc. BuildSys*. 73–82.
- [6] Erik Bochinski, Volker Eiselein, and Thomas Sikora. 2017. High-speed tracking-by-detection without using image information. In *2017 14th IEEE international conference on advanced video and signal based surveillance (AVSS)*. IEEE, 1–6.
- [7] Daniel Bolya, Sean Foley, James Hays, and Judy Hoffman. 2020. Tide: A general toolbox for identifying object detection errors. In *Proc. ECCV*.
- [8] Rainer E Burkard and Ulrich Derigs. 1980. The linear sum assignment problem. In *Assignment and Matching Problems: Solution Methods with FORTRAN-Programs*.

- [9] Zhengping Che, Guangyu Li, Tracy Li, Bo Jiang, Xuefeng Shi, Xinsheng Zhang, Ying Lu, Guobin Wu, Yan Liu, and Jieping Ye. 2019. D2-City: A Large-Scale Dashcam Video Dataset of Diverse Traffic Scenarios. *arXiv preprint arXiv:1904.01975* (2019).
- [10] Kai Chen, Jiaqi Wang, Jiangmiao Pang, Yuhang Cao, Yu Xiong, Xiaoxiao Li, Shuyang Sun, Wansen Feng, Ziwei Liu, Jiarui Xu, Zheng Zhang, Dazhi Cheng, Chenchen Zhu, Tianheng Cheng, Qijie Zhao, Buyu Li, Xin Lu, Rui Zhu, Yue Wu, Jifeng Dai, Jingdong Wang, Jianping Shi, Wanli Ouyang, Chen Change Loy, and Dahua Lin. 2019. MMDetection: Open MMLab Detection Toolbox and Benchmark. *arXiv preprint arXiv:1906.07155* (2019).
- [11] Tiffany Yu-Han Chen, Lenin Ravindranath, Shuo Deng, Paramvir Bahl, and Hari Balakrishnan. 2015. Glimpse: Continuous, real-time object recognition on mobile devices. In *Proc. SenSys*. 155–168.
- [12] Byung-Gon Chun, Sunghwan Ihm, Petros Maniatis, Mayur Naik, and Ashwin Patti. 2011. Clonecloud: elastic execution between mobile device and cloud. In *Proceedings of the sixth conference on Computer systems*. 301–314.
- [13] Mark Everingham, Luc Van Gool, Christopher KI Williams, John Winn, and Andrew Zisserman. 2010. The pascal visual object classes (voc) challenge. *IJCV* 88, 2 (2010).
- [14] Anurag Ghosh, Akshay Nambi, Aditya Singh, Harish YVS, and Tanuja Ganu. 2021. Adaptive streaming perception using deep reinforcement learning. *arXiv preprint arXiv:2106.05665* (2021).
- [15] Google. 2020. Google Coral USB Accelerator. <https://coral.ai/products/accelerator>.
- [16] Song Han, Huizi Mao, and William J Dally. 2015. Deep compression: Compressing deep neural networks with pruning, trained quantization and Huffman coding. *arXiv preprint arXiv:1510.00149* (2015).
- [17] Jonatan Heyman, Carl Byström, Joakim Hamrén, and Hugo Heyman. 2020. Locust: An Open Source Load Testing Tool. <https://locust.io/>
- [18] Jonathan Huang, Vivek Rathod, Chen Sun, Menglong Zhu, Anoop Korattikara, Alireza Fathi, Ian Fischer, Zbigniew Wojna, Yang Song, Sergio Guadarrama, et al. 2017. Speed/accuracy trade-offs for modern convolutional object detectors. In *Proc. CVPR*.
- [19] Itay Hubara, Matthieu Courbariaux, Daniel Soudry, Ran El-Yaniv, and Yoshua Bengio. 2017. Quantized neural networks: Training neural networks with low precision weights and activations. *The Journal of Machine Learning Research* 18, 1 (2017), 6869–6898.
- [20] Intel. 2020. Intel Neural Compute Stick 2. <https://software.intel.com/en-us/neural-compute-stick>.
- [21] Srinivasan Iyengar, Ravi Raj Saxena, Joydeep Pal, Bhawana Chhaglani, Anurag Ghosh, Venkata N Padmanabhan, and Prabhakar T Venkata. 2021. Holistic energy awareness for intelligent drones. In *Proc. BuildSys*.
- [22] Junchen Jiang, Ganesh Ananthanarayanan, Peter Bodik, Siddhartha Sen, and Ion Stoica. 2018. Chameleon: scalable adaptation of video analytics. In *Proc. SIGCOMM*. 253–266.
- [23] Yiping Kang, Johann Hauswald, Cao Gao, Austin Rovinski, Trevor Mudge, Jason Mars, and Lingjia Tang. 2017. Neurosurgeon: Collaborative intelligence between the cloud and mobile edge. *ACM SIGARCH Computer Architecture News* 45, 1 (2017), 615–629.
- [24] Harold W Kuhn. 1955. The Hungarian method for the assignment problem. *Naval research logistics quarterly* 2, 1-2 (1955), 83–97.
- [25] Mengtian Li, Yu-Xiong Wang, and Deva Ramanan. 2020. Towards Streaming Image Understanding. *arXiv preprint arXiv:2005.10420* (2020).
- [26] Yuanqi Li, Arthi Padmanabhan, Pengzhan Zhao, Yufei Wang, Guoqing Harry Xu, and Ravi Netravali. 2020. Reducto: On-Camera Filtering for Resource-Efficient Real-Time Video Analytics. In *Proc. SIGCOMM*.
- [27] Robert LiKamWa, Yunhui Hou, Julian Gao, Mia Polansky, and Lin Zhong. 2016. RedEye: analog ConvNet image sensor architecture for continuous mobile vision. *ACM SIGARCH Computer Architecture News* 44, 3 (2016).
- [28] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. 2017. Focal loss for dense object detection. In *Proc. ICCV*.
- [29] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. 2014. Microsoft coco: Common objects in context. In *Proc. ECCV*.
- [30] Luyang Liu, Hongyu Li, and Marco Gruteser. 2019. Edge assisted real-time object detection for mobile augmented reality. In *Proc. MobiCom*. 1–16.
- [31] Alan Lukezic, Tomas Vojir, Luka Cehovin Zajc, Jiri Matas, and Matej Kristan. 2017. Discriminative correlation filter with channel and spatial reliability. In *Proc. CVPR*.
- [32] Ravi Netravali, Anirudh Sivaraman, Somak Das, Ameesh Goyal, Keith Winstein, James Mickens, and Hari Balakrishnan. 2015. Mahimahi: Accurate record-and-replay for {HTTP}. In *USENIX ATC*. 417–429.
- [33] Nvidia. 2020. Meet Jetson, the Platform for AI at the Edge. <https://developer.nvidia.com/embedded-computing>.
- [34] Xukan Ran, Haolanz Chen, Xiaodan Zhu, Zhenming Liu, and Jiayi Chen. 2018. Deepdecision: A mobile deep learning framework for edge video analytics. In *Proc. INFOCOM*.

- [35] Joseph Redmon, Santosh Divvala, Ross Girshick, and Ali Farhadi. 2016. You only look once: Unified, real-time object detection. In *Proc. CVPR*. 779–788.
- [36] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. 2015. Faster r-cnn: Towards real-time object detection with region proposal networks. In *Proc. NeurIPS*. 91–99.
- [37] Mark Sandler, Andrew Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. 2018. Mobilenetv2: Inverted residuals and linear bottlenecks. In *Proc. CVPR*.
- [38] Mahadev Satyanarayanan, Paramvir Bahl, Ramón Caceres, and Nigel Davies. 2009. The case for vm-based cloudlets in mobile computing. *IEEE pervasive Computing* 8, 4 (2009), 14–23.
- [39] Xiaofan Zhang, Haoming Lu, Cong Hao, Jiachen Li, Bowen Cheng, Yuhong Li, Kyle Rupnow, Jinjun Xiong, Thomas Huang, Honghui Shi, et al. 2020. Skynet: a hardware-efficient method for object detection and tracking on embedded systems. In *Proc. MLSys*.
- [40] Huajun Zhou, Zechao Li, Chengcheng Ning, and Jinhui Tang. 2017. Cad: Scale invariant framework for real-time object detection. In *Proc. ICCV Workshops*.
- [41] Xingyi Zhou, Dequan Wang, and Philipp Krähenbühl. 2019. Objects as points. *arXiv preprint arXiv:1904.07850* (2019).
- [42] Pengfei Zhu, Longyin Wen, Dawei Du, Xiao Bian, Qinghua Hu, and Haibin Ling. 2020. Vision Meets Drones: Past, Present and Future. *arXiv preprint arXiv:2001.06303* (2020).
- [43] Pengfei Zhu, Longyin Wen, Dawei Du, Xiao Bian, Haibin Ling, Qinghua Hu, Haotian Wu, Qinqin Nie, Hao Cheng, Chenfeng Liu, et al. 2018. Visdrone-vdt2018: The vision meets drone video detection and tracking challenge results. In *Proc. ECCV Workshops*.