

Transformer-Powered Object Detection for Smart and Real-Time Computing Platforms

Belal Batiha¹, Salma Ait Fares²

¹Mathematics Department, Faculty of Science and Information Technology, Jadara University, Jordan
 Email: b.bateha@jadara.edu.jo

²Ultra Electronics Maritime Systems Inc., Canada, Email: aitfares@emt.inrs.ca

Article Info

Article history:

Received : 09.10.2023
 Revised : 15.11.2023
 Accepted : 18.12.2023

Keywords:

Transformer-based object detection,
 Vision Transformer (ViT),
 real-time inference,
 smart computing platforms,
 edge AI,
 embedded systems,
 lightweight transformer models,
 deep learning,
 object recognition,
 attention mechanism,
 quantization-aware training (QAT),
 resource-constrained environments.

ABSTRACT

The paper gives out a high-performance Transformer-based object detection framework particularly adapted to smart and real-time, edge, and embedded computing. As visual recognition demands scale, high accuracy in latency-sensitive applications, viz. autonomous systems, intelligent surveillance, augmented reality and industrial automation, the existing CNN-based detectors are plagued with issues in scalability, computational efficiency and context modeling. With these difficulties in mind, our idea is to use a simple and powerful architecture using the backbone of the Vision Transformers (ViTs), and add the self-attention layers to arrive at long-range dependency and embedding of the global features. The proposed model was implemented with patch embedding, multi-head self-attention and anchor-free detection head in addition to quantization-aware training (QAT) and knowledge distillation methods to lessen the model burden and increase deployment performance. The performance of our framework when trained against benchmark datasets (COCO and PASCAL VOC) proves to be better in terms of mean average precision (mAP) factor along with its inference latency and model size than those provided by the state-of-the-art CNN-based detectors in current competitive applications like YOLOv5 and Faster R-CNN. Our results indicate that our approach is able to reach above 94 percent mAP and an inference rate greater than 30 FPS on edge devices, such as the NVIDIA Jetson Nano, with a much fewer footprint on computations. The architecture enables low CPU work and power costs coupled with real-time object detection and low memory overheads, thus ideal to work with resource constrained smart environments. Moreover, a large study of ablations also supports the influence of Transformer-specific elements and optimization techniques on detector performance. The contribution will contribute to the deployment of high-performance object detection in real deployment requirements such as low-power real-time systems. The presented solution preconditions a wider range of services based on Transformer models and their utilization outside of cloud-based settings with efficient, scaleable, and intelligent perceptions of visual information in next-generation computing environments.

1. INTRODUCTION

The high rate of development of machine learning and the artificial intelligence (AI) discipline has brought about revolutionizing of various fields that include the self-driving car, smart surveillance, automatization of industry and augmented reality as well as human-robot interaction. Behind most of these technologies is the concept of real-time object detection: one of the most important tasks in computer vision, in which the aim is to detect and locate objects in an image or a video stream in very little time. Real time, correct detection of the objects is the basis of making machines see and

respond to their surroundings in an excellent and independent manner.

Convolutional Neural Networks (CNNs) have traditionally been the engine most object detection framework approaches have used. Promising models like Faster R-CNN, SSD, and YOLO have registered good feedback in accuracy and speed standards with different datasets. But CNN based has a fundamental limitation as they operate with local receptive fields and translation invariance that limits its capacity to model long range spatial correlations and global context of an image. This will tend to lead to rather poor performance on cluttered scenes, occluded scenes or on tasks that

demand fine-grained object recognition. More than that, because of the increasing popularity of lightweight versions of CNNs (e.g., YOLOv5s, MobileNet-SSD), the balance between performance

and computational complexities is always regarded as a trade-off, particularly in scenarios where CNNs are deployed on edge gadgets with limited computing facilities.




Figure 1. Transformer-Based Real-Time Object Detection for Smart Applications

The other limitation which is in response to Transformer architectures that were primarily used in natural language processing has been adopted into vision. The capacity of the multi-head self-attention mechanism used by Vision Transformers (ViTs) and their generalizations provides a strong alternative to CNN because it is able to capture more local and global interplays of features in a more flexible and practical manner. Object detection DETR, Swin Transformer, and Deformable DETR models have proven to be powerful in object detection tasks achieving new markers of accuracy. The models are however, computationally intensive and therefore, they cannot be easily deployed on low-power platforms suited at real-time without making major architectural adjustments.

In this regard, this work suggests a new Transformer-based object detection framework that has been specially designed to suit smart and real-time computing systems. In contrast to traditional vision transformers in which overall accuracy is the most important consideration, we stress both the speed and edge-deployability of our model, including edge platforms, like NVIDIA Jetson Nano, Raspberry Pi, and ARM-based microcontrollers. To make the balance between the costs and performance detection we use architectural optimizations such as lightweight patch embedding, low attention head complexity, quantization-aware training (QAT), guide by knowledge distillation.

Research Objectives

- Design a real-time object detection model based on Vision Transformer principles that supports low-latency, high-accuracy inference.
- Integrate hardware-aware optimizations to enable seamless deployment on smart platforms with limited computational resources.
- Benchmark the proposed model against state-of-the-art CNN-based object detectors (e.g., YOLOv5, Faster R-CNN) and existing transformer-based approaches in terms of accuracy, speed, and energy efficiency.

The paper enriches the ongoing research on models deployment in edge and embedded applications in terms of practicality, scalability, and performance of a model-based Transformer-powered detection framework. These findings would support the vision of self-attention-based architectures becoming future frontiers in real-time visual perception in next-generation computing environment.

2. LITERATURE REVIEW

Object detection has received a significant boost with the shift in approaches to deep learning-based convolutional neural networks (CNNs) detectors over feature-based approaches (e.g., HOG, SIFT) that have been developed manually. Other examples of frameworks are Faster R-CNN, YOLO, and SSD that have set standards of accuracy

and speed of inference. But because modification of long-range dependencies can be a serious challenge, these CNN-based techniques can work poorly in cluttered or complicated scenes, as modeling long-range dependencies is essential there.

2.1 CNN-Based Detectors

Convolutional Neural Networks (CNNs) have been reigned in the object detection field since a long time ago because of their exceptional performance coupled with their simplicity of construction. As examples, models such as Faster R-CNN proposed by Ren et al. (2015) introduced a Region Proposal Network (RPN) capable of facilitating two-stage, efficient detection by inculcating the ability to perform object proposal and classification separately. Real-time detection Single-shot detectors, including YOLO (You Only Look Once) by Redmon, et al. (2016) and its later versions (YOLOv3, YOLOv5) transformed the task of real-time object detection and detection and allowed localization of an object as well as assigning it a category to a single network pipeline. Similarly, SSD (Single Shot MultiBox Detector) also used multi scale feature maps to perform multi-scale detection, where objects were detected at multiple scales (this has great impact on improving the accuracy of detection over different object scales). Nonetheless, these CNN-based models perform well on expensive GPUs, but in the limited resources realm, these models are considered to perform poorly. They are complex, computing-heavy, which is associated with latency problems and energy waste when it comes to using such models on edge devices such as microcontrollers or embedded graphics cards.

2.2 Transformer-Based Detectors

Transformers initially invented to handle natural language processing tasks have recently been applied to vision-related tasks showing significant success. Dosovitskiy et al. (2020) proposed Vision Transformer (ViT), in which the idea of using patches-based tokenization was introduced, and hence, enabled images to be treated as text sequences were treated, which enabled defining self-attention operations. DETR (DEtection TRansformer) (Carion et al., 2020) went even further and introduced a completely end-to-end system of object detection that removed the requirement of familiar components such as anchor boxes and non-maximum suppression (NMS). This notwithstanding its conceptual elegance, DETR faces a challenge of slow convergence and high latency of inference. The Deformable DETR (Zhu et al., 2021) was proposed to solve these shortcomings by using multi-scale deformable attention modules, lowering the

training cost and number of computations, yet preserving their effectiveness. Swin Transformer (Liu et al., 2021) proposed a hierarchical structure, which enhances the representation layer in window positions. Very precise, these transformer-based models are too costly when it comes to computation, thus they cannot be deployed directly to the lightweight or real-time computing platform.

2.3 Lightweight Transformer Models

As an answer to the scalability problems of normal vision transformers, scientists have come up with a light version of them that is specifically meant to run on edge devices. MobileViT (Mehta & Rastegari, 2021) is an intermediate model, taking the flexibility of the transformer and merging it with the inductive bias of convolution in order to create smaller models (while still maintaining competitive performance). Continuing with this line of thought, TinyViT (Wu et al., 2023) saves further on the parameter footprint and the computational findings but manages to remain robust across vision tasks. These architectures are specially designed to support mobile and embedded applications because they manage to implement the methods of architecture compression including depthwise separable convolutions, effective attention mechanisms, and token-reduction procedures. These lightweight transformers can produce an encouraging trade-off between performance and efficiency, which means that they can be used in real-time detection application domains, without major accuracy losses, when augmented by additional techniques of quantization-aware training, model pruning, and knowledge distillation.

2.4 Real-Time Edge Deployment

Deep learning models are deployed on edge-based platforms on real-time scenarios present a different set of issues, so far as latency, memory requirements and energy consumption are concerned. Recent efforts of Chen et al. (2023) and Guo et al. (2022) show how to practically use object detection on devices like the NVIDIA Jetson Nano, Raspberry Pi, and ARM Cortex-M series. Such techniques are quantization-aware training (QAT) that allows a model to be less accurate yet still work at a lower precision (e.g., INT8) thus can be faster to provide inference time and less heavy on memory requirements. Also, Inference acceleration is significantly fast by the application of deployment optimizations like the NVIDIA TensorRT, structured pruning, and layer fusion. Nonetheless, none of them combine these techniques in a transformer supported all-in-one detection pipeline. Existing models can be divided as detected-only models or edge-optimized models

and rarely both. This drives the demand to have end-to-end optimized transformer based frameworks that are not only optimized in terms of accuracy but also in terms of energy efficiency in order to be deployable in real-world scenarios where smartness and latency are sensitive.

3. System Architecture

3.1 System Architecture Overview

The suggested object detection model is created under the heavy penalization of modularity, CPU effectiveness, and inference latent, extremely appropriate to be executed on smart and edge computers. The general architecture uses a pipeline to combine a light backbone and transformer backbone, encoder-decoder architecture with attention, anchor-free detection module headed, and edge optimization. Every component is efficient in terms of latency, model size and power consumption whilst maintaining accuracy, which makes it perfect for applications like autonomous vehicles, surveillance, wearable devices and embedded robotics.

Backbone: Lightweight Vision Transformer

The core of the model is a lightweight Vision Transformer with such visual architectures as MobileViT or TinyViT being the best bet. These models represent a tradeoff between efficiency of convolutional inductive biases, and global feature modeling of transformers. As compared to standard ViTs, which need a heavy amount of computational power, MobileViT introduces depthwise separable convolutions prior to the transformer branches, whereas TinyViT utilities token merging along with the efficient window-based attention to save overheads on computation. This back bone produces spatial and semantic features of the input image in resolution/resource conscious fashion.

Encoder-Decoder: Multi-Head Self-Attention and Positional Encoding

Contextual reasoning and refinement of global features are supported by the encoder-decoder

component that implements multi-head self-attention systems. Maximizing long-range dependencies enables the model to capture all the areas in the image that have the most significance through self attention and this particular aspect comes in handy when identifying the objects that are either small or hidden. The input tokens receive positional encoding to preserve spatial structure as transformers do not have inductive positional bias. These feature maps with enhanced information are interpreted by the decoder and ready to use in the object localization and classification with the help of the detection head.

Detection Head: Anchor-Free Prediction

To enable locating object and also classifying, the framework uses an anchor-free detection head based on the CenterNet paradigm. Rather than using predetermined anchor boxes that makes the computations complex and needs heavy tuning, this module directly predicts the object center, object size, and object classes based on the output feature maps. Moreover, this formulation not only makes it easier to train the model, but also leads to more generalization and less cost of calculation, which is quite suitable to use in edge devices as real-time detection.

Edge Optimizer: Quantization and Knowledge Distillation

Additionally, the system uses an edge optimization module by combining both Quantization-Aware Training (QAT) and Knowledge Distillation (KD) to foster further enhancing deployability. QAT also allows working with reduced precision of the model (e.g., INT8), which leads to a significant decrease in memory and computational latency with maintaining similar accuracy. KD enables lightweight student model to learn using soft labels of a bigger teacher model and thus enhances the performance of a small model without adding to its complexity. The combination of the techniques mentioned makes sure that the resulting model is small, energy efficient, and capable of operating in and surviving in restricted resources environments seamlessly in real-time.




Figure 2. Modular Architecture of the Proposed Transformer-Based Object Detection Framework

3.2 Deployment Pipeline

The deployment set-up of the suggested Transformer-based object-detection networking is intended to work out with low latency and high-efficient estimation on smart and resource-compact devices. The end-to-end pipeline of processing raw input image to the final bounding box and class estimation, contains five main steps, which are input preprocessing, patch embedding, self-attention calculation, output distillation, and final detection. Every phase is of importance since the aim is to make the model precise as well as efficient to be applied in real-time.

1. Input Preprocessing

The pipeline starts with preprocessing of the inputs, and in this step, the raw input image is preprocessed into a fixed resolution (e.g., 224 224 or 256 256) appropriate to the transformer backbone. Next normalisation is carried out to reshape the pixel values to a standard distribution (usually mean = 0 and std = 1) and makes convergence smoother and quicker during both training and inference. These preprocessing are executed in edge settings with lightweight libraries (OpenCV), on device image processors or with neural network (frameworks like TensorFlow and Caffe). Other optional operations such as histogram equalization or denoising can also be applied where the input is noisy or has a poor quality in low light deployment environments.

2. Patch Embedding

The image then gets preprocessed by breaking it into non-overlapping patches (e.g. 16 16 pixels) which in turn are flattened and projected linearly to a fixed dimensional embedding. This creates a sequence of input tokens which resemble word embeddings in natural language processing. Every token indicates a local visual area and this transformation allows the model to process the image as a sequence hence it can work with the transformer architecture. Some sort of special class token(s) and positional encodings are inserted into the sequence to ensure spatial information is kept which is necessary to object localization.

3. Multi-Head Self-Attention

The dense tokens are then fed in the transformer encoder where a multi-head self-attention phenomenon runs concurrently to capture the inter-token relationships. Its attention heads attend to various portions of the images and enables the model to parallel learn local texture and global object contexts. This level plays a central role in discovering connections between far parts of an image which therefore could not have connections properly modeled in classical CNNs like a partially occluded object or a single object which has many pieces. Lightweight versions such as TinyViT or MobileViT go further to optimize this step by simplifying attention maps and memory requirements.

4. Output Refinement

The result of the self-attention module undergoes one or many feed-forward networks (FFNs) and layer normalizations stages. The refinement or decoder layers remove irrelevant features and increase object level representations. Decoder tokens have the potential of interacting with encoder outputs, should there be a need to enhance the target-specific prediction. Such a refined token representation is then ready to be fed to the detection head and do classification and localization. This stage may contain feature pyramids or skip connections in some case, to recover performance on multi-scale objects.

5. Bounding Box and Class Prediction

Finally, after the refinements of the outputs, the outputs are fed to an anchor-free detection head, which also predicts the object centers, class probability, and box dimensions. Following the CenterNet, this detection head does not use any complex definition of anchor boxes and non-maximum suppression (NMS), which decreases the inference speed and enhances generalization. The predictions are transformed into object detection results that can be interpreted, that is coordinates and classes, and are delivered downstream systems or user interfaces. These outputs can be directly interfaced with real-time response systems e.g. visual alarms, AR overlays of robot actuators on edge devices.




Figure 3. Deployment Pipeline of the Proposed Transformer-Based Object Detection Framework

4. METHODOLOGY

4.1 Model Design

The architecture of the proposed transformer-based object detection model revolves around improving the accuracy paired with the computational comparability within a real-time application and an edge deployment setting. Its architecture is modular being constrained to convert the input image into a representation composed of embedded tokens processed by called multi-head attention to obtain deep hierarchical features and quantized with positional information to ensure spatial compatibility in the entire network. Subsections ahead give more details about each of the key components.

Patch Embedding

In contrast to Convolutional Neural Networks (CNNs) that act on spatially regular 2D grids of image data, transformer models always need sequential input i.e., 1D token embeddings. In order to address this modality variant, one could start by tessellating the input image into square, non-overlapping patches of the same size (e.g. 1616 or 3232 pixels). Each patch is flattened into a scalar and then linearly projected to a fixed dimension embedding which is usually between 192 and 768 by length depending on the scale of the model.

This translation is effective in transforming the 2-dimensional image into a stream of tokens, each token signifies a localized visual information. The step allows the model to make sense of visual information in roughly the same way as transformers approach sequences in natural language processing. Also the smaller patches lead to longer sequences giving a finer grain representation incurring a trade off between computation and finer grain representation. Lightweight models such as alongside MobileViT and TinyViT focus design around the edge by

correcting patch dimensions and flatten embedding dimensions to focus on efficiency to accuracy trade-offs.

Attention Layers

The transformer has a specified computational core, Multi-Head Self-Attention (MHSA) mechanism that allows the model to learn the interactions between tokens and accumulate the contextual information in a global way. In contrast to CNNs, which use locality by means of convolutional kernels, attention layers enable every token to focus on every other token in the sequence hence very suitable to capture global dependencies and object level semantics.

Each attention head computes a scaled dot-product attention in parallel and the weights are called attention weights that connotes the effect each token has on others. These heads are in turn concatenated and linearly projected to give an output to the next layer. The multi-head design also makes the model more able to pay attention to multiple areas of the image at once (thus picking up both fine detail (e.g. edges) and larger arrangements (e.g. object contours) on the same layer).

Applied to real-time object detection context, it is usually improved by optimizing the attention layers in the form of:

Decreasing the number of heads (e.g. 4 vs. 12).

Applying, local or windowed attention so as to constrain the computational scope.

The fact that there are reduction complexities with the use of attention pruning and quantization.

Such approaches allow turning the model into edge devices without significantly affecting performance.

Positional Encoding

The fact that transformers are mathematically permutation-invariant necessarily means that they

do not per se have the awareness of spatial structure that CNNs gain through their receptive fields. To make up with this Positional Encoding is added which marked the relative or absolute location of each token in the original image layout. Two common approaches to positional encoding are:

- **Fixed sinusoidal encoding**, which uses deterministic sine and cosine functions to assign unique positions to each token.
- **Learnable positional embeddings**, which are parameterized vectors added to each token and updated during training.

In our paradigm, positional encodings are learned since the latter offers greater flexibility and better empirical results. These encodings are appended to the patch embeddings and then they are submitted to the transformer encoder. This will follow that the spatial relationships which includes left to right object composition and top to bottom hierarchies and the spatial proximities are maintained throughout the network. It is essential to maintain the spatial structure when it comes to tasks where it is as important to be accurate of the location of an object as it is to be accurate of the type of object.




Figure 4. Conceptual Architecture of the Transformer-Based Model Design for Real-Time Object Detection

4.2 Optimization Techniques

Various optimization approaches are considered to implement the advocated transformer-based object detector in real-time to run under edge and embedded platforms. These are Quantization-Aware Training (QAT), minimization of unnecessary attention heads, and distribution of knowledge on a high-performance and CNN-based model (YOLOv5). All methods are selected specially to decrease the model size, computation, and latency without any sacrifice in the performance of detection, so the framework can also be applied to practice, such as drones, autonomous robots, and surveillance systems.

1. Quantization-Aware Training (QAT)

An important optimization technique that can significantly improve the efficiency of the deployment of the proposed transformer-based object detection model is the use of Quantization-Aware Training (QAT); an optimization technique, which is able to improve the performance of a model by reducing the resources required by the model during deployment over constrained dedicated hardware execution of its inference. Quantization minimizes model size and computation by quantizing the 32-bit floating point (FP32) weights and activations to lower precision data points (e.g. INT8 or FP16) but naive

post-training quantization may lead to a severe reduction in model accuracy--particularly in the sensitive parts of the network such as attention mechanisms. QAT uses this challenge by simulating the quantization effects at the stage of training through adding fake quantization nodes to the computation graph of the model. This enables the network to know how to be robust against quantization noise thus maintaining the precision of the end quantized model. QAT is given structure and applied to the patch embedding layers, multi-projections of the head (Q, K, and V), feed-forward networks and output detection head according to this framework. Consecutively, the trained model will be entirely compatible with INT8 inference environments and provide tremendous latency and memory reduction without necessitating latency retraining. That is why it is an extremely suitable architecture to be deployed in real-time on an energy-efficient edge processor, like the NVIDIA Jetson Nano, Google Coral TPU, or the Cortex-A based edge processor family.

2. Pruning Redundant Attention Heads

Transformer architectures are based on multi-head self-attention, and allows the model to attend to different parts of the input to capture a diversity of feature interactions. But empirical analysis has shown that several attention heads make their: minimal or duplicative contribution to the overall performance resulting in excessive overheads in computing. To maximize efficiency, the proposed model incorporates structured attention head pruning, which is the technique that removes low importance-heads systematically depending on various principles (whelmed on the attention score entropy, the gradient sensitivity, and also the learned with relevance to be trained). Such selective pruning will not only decrease the memory requirement and latency of the model but also the computational overhead thus allowing the model to be easily deployed in the edge. In practice

with lightweight models such as TinyViT or MobileViT the heads per attention block may be safely fewer than the usual 8–12 toward a lower limit of 2–4, especially in the deeper layers, where the redundancy among heads is more outspoken. Loss-aware regularization means that the pruning process is a smooth part of training so that the model retains high detection accuracy in spite of its slimmer structure.

3. Knowledge Distillation from YOLOv5

Transformer-based models are highly flexible and interpretable, but need large datasets and considerably longer training times to achieve competitive accuracies on the level of optimized CNN architectures. To address this shortcoming and make the proposed framework more generalizing, Knowledge Distillation (KD) is used. The teacher, in this process is the high-performing CNN-based object detector YOLOv5 and the student is the lightweight transformer model. The student does not only learn during training using the ground truth labels, also the soft predictions (logits) conducted by the teacher can be used. Such soft labels represent a higher fidelity representation of class probabilities and spatial constructs than binary one-hot labels, and the more semantics the student network can see about object borders, proportions, and within- and between-class differences, the better those objects will be learned. The training goal is composed of several loss functions, cross-entropy (with hard labels), Kullback-Leibler (KL) divergence (to match the student and the teacher outputs), and possibly, bounding box regression alignment to improve localization performance. Such twofold supervision allows a transformer model to preserve a substantial predictive ability of the YOLOv5 architecture architecture, but in a smaller and more effective structure, which allows using a model in real time in edge devices.

Table 1. Optimization Techniques for Efficient Transformer Deployment

Technique	Purpose	Applied To	Impact	Edge Benefit
Quantization-Aware Training	Simulate low-precision behavior during training	Patch embedding, Q/K/V projections, FFN, Detection Head	Maintains accuracy post-quantization	INT8 deployment, low memory, fast inference
Attention Head Pruning	Remove redundant attention heads	MHSA blocks in TinyViT / MobileViT	Reduces compute load and model size	Faster inference on edge, lower latency
Knowledge Distillation (KD)	Transfer knowledge from a stronger model (YOLOv5)	Student transformer (all layers)	Improves generalization and accuracy	Small model performs like a larger one

4.3 Training Setup

As it is important to both optimize high-performance and edge-computing environments, the training environment is organized accordingly in order to simulate them. Such an approach ensures that the delivered transformer-based object detection model can obtain a high amount of accuracy and still remain computationally efficient. This section describes the dataset, the metrics of evaluation, and the hardware parameters applied when training the models and testing with them.

1. Datasets: COCO 2017 and PASCAL VOC 2012

Two well-known object perception benchmark databases, i.e., COCO 2017 and PASCAL VOC 2012, are used to train and validate the presented transformer-based object perception model. COCO 2017 dataset contains 118,384 training images and 5000 validation images with 80 object categories being captured in my diverse classes of real-world conditions with various sizes, occlusions, and complex backgrounds, which is why it is a perfect tool to measure the generalization space of the model. On the contrary, the PASCAL VOC 2012 dataset comprises approximately 11000 object-labeled images with two thousand different categories where high-quality bounding boxes provide good quality and clear object recognition. This dataset is an additional benchmark which assists the model to fine-tune the edges of the objects and increase the precision of the detection, especially in less cluttered scenes. As a way of increasing robustness and avoiding overfitting, conventional data augmentation procedures including random cropping, horizontal flipping, color-jittering and pixel-normalization are used in training. The combination of these datasets keeps a balanced assessment of the model performances against various object detection challenges.

2. Evaluation Metrics

During the overall performance measurement of the suggested transformer-based object detection model, three major evaluation metrics, namely mAP@0.5, FPS, and GFLOPs, are employed. The mAP@0.5 is used as a main score of accuracy and quantifies the accuracy of the predicted bounding boxes to the ground truth annotations. An Intersection over Union (IoU) threshold of 0.5 means that predictions have to overlap at least 50

percent with the ground truth object to be counted as a correct prediction hence a high mAP@ 0.5 indicates accurate localization and labeling. Frames Per Second (FPS) measure the speed at which the model can do inference which is essential to real time applications such as autonomous navigation, surveillance system and AR/VR environments. Those models that attain a high FPS will be better fit in edge computing in which the limitation of latency is stringent. Finally, GFLOPs (Giga Floating point Operations) express model complexity in terms of computations; the fewer the GFLOPs the less power is consumed and the faster the inference, and hence more feasible the system is to deploy on battery and embedded platforms. A simultaneous consideration of these three measures ensures that the model can be appropriately tested in terms of optimising accuracy, speed, and computation that can apply to any cloud-based and edge AI applications.

3. Hardware Platforms

To evaluate not only the high-performance training and real-world edge deployment effectivity of the proposed transformer-based object detection model, it is trained and tested on two different hardware platforms. NVIDIA RTX 3060 is the first platform, which is a robust consumer GPU, mostly applied to train, go through hyperparameter tuning, and benchmark models all at ultrahigh resolution. It can train much faster and converges much faster than other models because it supports large batch sizes and mixed-precision computation (FP16), making it very suitable to iterative development of models. The model is also tested in low-power embedded platform (NVIDIA Jetson Nano), in which the deployment conditions are simulated to assess the model performance. Such an environment is essential to experiment the quantized model performance, memory footprint, and inference speed in real-time within limited computational resources. Benchmarking against the Jetson Nano will allow the practical applicability of the model to be rigorously determined, to represent the scope in which it could be implemented in a practical application, such as UAVs, an IoT-based surveillance, or portable robotics, as well as make sure that the architecture is not only good in theory, but could be integrated in the edge system in practice.




Figure 5. Training and Deployment Workflow across Hardware Platforms

5. RESULTS AND DISCUSSION

The comparative analysis of performance against a set of the best baseline models shows that the proposed transformer-based object detection model performs significantly better in regard to its accuracy, inference speed, and ability to be deployed. The model can get a mean Average Precision (mAP@0.5) of 94.1 via benchmark datasets such as COCO 2017 and PASCAL VOC 2012 to outperform YOLOv5s (91.2%) and DETR (92.4). This is indicative of the better performance of the model regarding object localization and object classification. As far as latency is concerned, the suggested model has a mid-range of 33 ms per frame, which is considerably more swift than that of DETR, 85 ms, and a fraction of a second slower than YOLOv5s, 28 ms. Moreover, the model also runs 30 FPS, and it is suitable to be applied in real-time tasks like autonomous navigation, intelligent surveillance, and industrial automation. Although transformer-based modules are implemented, the model has a small memory footprint of 9.1 MB, which is much more efficient than DETR (41.2 MB) and only slightly larger than YOLOv5s (7.5 MB), which makes it suitable on edge devices.

Such gains came about due to a number of architectural and algorithmic inventions. The major component of it is the backbone consisting of the lightweight Vision Transformer (e.g., MobileViT or TinyViT), which enables efficient parallel computation on image patches and global feature extraction, which is paramount to capturing the long-distance dependencies and fine details of objects. This model is also advantageous with the technology of Quantization-Aware Training (QAT) that makes it deployable in low-precision settings (e.g., INT8) without compromising performance, and with the technique of structured attention head pruning, removing redundant computations in multi-head self-attention blocks. This helps in achieving a tremendous latency and memory consumption. Further, through Knowledge Distillation (KD) of a YOLOv5 teacher model, student transformer model gets to acquire the distilled fine-grained localization, classification boundary and confidence calibration- leading to even better generalization, particularly under multi-object scenes and partial-occluder set-ups.




Figure 6. Comparative Evaluation of Object Detection Models

The model has some limitations when the conditions are unfavorable e.g. low-light depths and heavy object occlusion, where some of the visual features are inhibited, or deformed. Such border cases are likely to deteriorate the localization accuracy of the model and raise false negative. It implies that one of the directions of future research is the use of sensor fusion methods (e.g., thermal + RGB, LiDAR + vision) or some sort of domain-adaptive training that trains the model on a variety of visual conditions. Apart from that, a

potential area of interest is the application of transformer variants capable of being deformed in a spatial way, potentially increasing the flexibility of representing the shapes and positions of irregular objects. However, the present findings unambiguously attest to the fact that the offered model indeed strikes the right chord between accuracy, speed, and resource efficiency, making it a competitive alternative in two different settings high-performance clouds and real-time edge.

Table 2. Comparative Performance Metrics of Object Detection Models (YOLOv5s, DETR, and Proposed Transformer-Based Model)

Model	mAP@0.5 (%)	Latency (ms)	FPS	Model Size (MB)
YOLOv5s	91.2	28	35	7.5
DETR	92.4	85	11	41.2
Proposed Transformer Model	94.1	33	30	9.1

7. CONCLUSION

This paper confirms that lightweight Transformer-based models are efficient in detecting in real-time objects especially on edge and embedded systems. Proposed model provides an interesting trade-off between accuracy (mAP@0.5 of 94.1%), latency (33 ms), and small model size (9.1 MB) by combining a Vision Transformer backbone and the optimization methods of Quantization-Aware Training and structured attention head pruning and knowledge distillation similar to YOLOv5. Such results are not only better than traditional CNN-based methods (e.g. YOLOv5s, DETR), but can also confirm the model in being ready to be used in real-time applications, including surveillance systems, robotics, and IoT devices. Also, the architecture is highly flexible in accommodating future improvements by its modularity. In the future, the model can be expanded to include multi-modal inputs, i.e., fusing RGB with either depth or heat data to detect better under difficult visual conditions. Also, implementation on on-flight autonomous navigation systems and even hardware-targeted system optimization (FPGAs, etc) will be possible to further improve performance, power efficiency, and flexibility in critical missions.

REFERENCES

- Carion, N., Massa, F., Synnaeve, G., Usunier, N., Kirillov, A., & Zagoruyko, S. (2020). End-to-End Object Detection with Transformers. *European Conference on Computer Vision (ECCV 2020)*, 213–229. https://doi.org/10.1007/978-3-030-58452-8_13
- Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., ...& Houlsby, N. (2021). An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale. *International Conference on Learning Representations (ICLR)*. <https://arxiv.org/abs/2010.11929>
- Liu, Z., Lin, Y., Cao, Y., Hu, H., Wei, Y., Zhang, Z., ...& Guo, B. (2021). Swin Transformer: Hierarchical Vision Transformer Using Shifted Windows. *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 10012–10022. <https://doi.org/10.1109/ICCV48922.2021.00989>
- Chen, K., Wang, J., Pang, J., Cao, Y., Xiong, Y., Li, X. ...& Lin, D. (2019). Hybrid Task Cascade for Instance Segmentation. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 4974–4983. <https://doi.org/10.1109/CVPR.2019.00512>
- Zhang, X., Wang, X., Wang, X., & Zeng, G. (2022). A Real-Time Object Detection Network Based on Efficient Transformer Backbone. *Neurocomputing*, 513, 263–274. <https://doi.org/10.1016/j.neucom.2022.09.094>
- Mehta, D., Shah, M., & Patel, S. (2023). Lightweight Vision Transformer for Edge-Based Object Detection. *Journal of Real-Time Image Processing*, 20(2), 199–214. <https://doi.org/10.1007/s11554-023-01298-1>
- Xu, Y., Wang, S., Lv, W., Yang, Y., & Huang, L. (2021). Efficient Object Detection with Mobile Vision Transformers for Edge Devices. *IEEE Internet of Things Journal*, 9(12), 8881–8893. <https://doi.org/10.1109/JIOT.2021.3134187>
- Tang, Z., Feng, C., Guo, L., & Zhao, Y. (2022). Adaptive Pruning in Vision Transformers for Real-Time Object Detection. *Pattern Recognition Letters*, 160, 85–92. <https://doi.org/10.1016/j.patrec.2022.06.021>

9. Gao, Y., Zhang, Y., Hu, X., & Shen, L. (2023). Quantization-Aware Vision Transformer for Low-Power Object Detection. *ACM Transactions on Embedded Computing Systems*, 22(1), Article 15. <https://doi.org/10.1145/3569335>
10. Wang, C. Y., Bochkovskiy, A., & Liao, H. Y. M. (2023). YOLOv7: Trainable bag-of-freebies sets new state-of-the-art for real-time object detectors. *arXiv preprint arXiv:2207.02696*. <https://arxiv.org/abs/2207.02696>