

A Survey of Modern Deep Learning based Object Detection Models

Syed Sahil Abbas Zaidi^{a,*}, Mohammad Samar Ansari^b, Asra Aslam^c, Nadia Kanwal^d, Mamoona Asghar^e and Brian Lee^a

Software Research Institute, Technological University of the Shannon: Midlands Midwest, Athlone, N37 HD68, Westmeath, Ireland

Faculty of Science and Engineering, University of Chester, CH1 4BJ, United Kingdom

Insight Centre for Data Analytics, National University of Ireland, Galway, Ireland

School of Computing and Mathematics, Keele University, United Kingdom

Lahore College for Women University, Pakistan

School of Computer Science, National University of Ireland, Galway, Ireland

ARTICLE INFO

Keywords:

Object detection and recognition
convolutional neural networks (CNN)
lightweight networks
deep learning

ABSTRACT

Object Detection is the task of classification and localization of objects in an image or video. It has gained prominence in recent years due to its widespread applications. This article surveys recent developments in deep learning based object detectors. Concise overview of benchmark datasets and evaluation metrics used in detection is also provided along with some of the prominent backbone architectures used in recognition tasks. It also covers contemporary lightweight classification models used on edge devices. Lastly, we compare the performances of these architectures on multiple metrics.

1. Introduction

Object detection is a trivial task for humans. A few months old child can start recognizing common objects, however teaching it to the computer has been an uphill task until the turn of the last decade. It entails identifying and localizing all instances of an object (like cars, humans, street signs, etc.) within the field of view. Similarly, other tasks like classification, segmentation, motion estimation, scene understanding, etc. have been the fundamental problems in computer vision.

Early object detection models were built as an ensemble of hand-crafted feature extractors such as Viola-Jones detector [1], Histogram of Oriented Gradients (HOG) [2] etc. These models were slow, inaccurate and performed poorly on unfamiliar datasets. The re-introduction of convolutional neural network (CNNs) and deep learning for image classification changed the landscape of visual perception. Its use in the ImageNet Large Scale Visual Recognition Challenge (ILSVRC) 2012 challenge by AlexNet [3] inspired further research of its application in computer vision. Today, object detection finds application from self-driving cars and identity detection to security and medical uses. In recent years, it has seen exponential growth with rapid development of new tools and techniques.

This survey provides a comprehensive review of deep learning based object detectors and lightweight classification architectures. While existing reviews are quite thorough [4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15], most of them are more than 2-years old, and the two recent ones were also released an year ago, and therefore miss out on new developments in the domain. The main contributions of this paper are as follows:




Figure 1: Structure of the paper.

1. This paper provides an in-depth analysis of major object detectors in three categories – single stage, two stage and transformer based detectors. Furthermore, we take historic look at the evolution of these methods.
2. We highlight the changes in network designs and their implication on the performance on detectors. We have also added a section for detector fundamentals to assist readers.
3. We present a detailed evaluation of the landmark backbone architectures and lightweight models. While research on lightweight networks has advanced, their reviews are few and far between. This work intends to fill the void.

In this paper, we have systematically reviewed various object detection architectures and its associated technologies, as illustrated in figure 1. Rest of this paper is organized as

sahilzaidi78@gmail.com (S.S.A. Zaidi)

ORCID(s): 0000-0002-9140-6721 (S.S.A. Zaidi); 0000-0002-4368-0478 (M.S. Ansari)




Figure 2: Sample images from different datasets.

follows. In section 2, the problem of object detection and its associated challenges are discussed. Section 3 explains the need and motivation for having this survey. Various benchmark datasets and evaluation metrics are listed in Section 4. In Section 5, several milestone backbone architectures used in modern object detectors are examined. Section 6 is divided into three major sub-section, each studying a different category of object detectors. This is followed by the analysis of a special classification of object detectors, called lightweight networks in section 7 and a comparative analysis in Section 8. The future trends are mentioned in Section 9 while the paper is concluded in Section 10.

2. Background

2.1. Problem Statement

The object detection is the natural extension of object classification, which only aims at recognizing the objects in the image. The goal of the object detection is to detect all instances of the predefined classes and provide its coarse localization in the image by axis-aligned boxes. The detector should be able to identify all instances of the object classes and draw bounding box around it. It is generally seen as a supervised learning problem. Modern object detection models have access to large sets of labelled images for training and are evaluated on various canonical benchmarks.

2.2. Key challenges in Object Detection

Computer vision has come a long way in the past decade, however it still has some major challenges to overcome. Some of these key challenges faced by the networks in real life applications are:

- *Intra class variation* : Intra class variation between the instances of same object is relatively common in nature. This variation could be due to numerous reasons like occlusion, illumination, pose, viewpoint, etc. These unconstrained external factors can have

dramatic effect of the object appearance [16]. It is expected that the objects could have non-rigid deformation or be rotated, scaled or blurry. Some objects could have inconspicuous surroundings, making the extraction difficult.

- *Number of categories*: The sheer number of object classes available to classify makes it a challenging problem to solve. It also requires more high-quality annotated data, which is hard to come by. Using fewer examples for training a detector is an open research question.
- *Efficiency*: Present day models need high computation resources to generate accurate detection results. With mobile and edge devices becoming common place, efficient object detectors are crucial for further development in the field of computer vision.

3. Comparison with Existing Surveys

A number of reviews for object detectors have been published in recent years [4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15], as summarised in Table 1. The works in [4, 5, 6, 7] were early attempts to present a review of the object detection models available till the time of their respective publication (i.e. till the year 2017). The rapid advances being witnessed in the field of computer vision in general, and in object detection in particular, led to the introduction of various new and improved deep learning models since then. The review works in [8, 9, 10, 11, 12, 13] were timely additions to the literature and contained pertinent discussions on various aspects of the design and usage of the different object detection models. In particular, the survey in [12] was quite wide-ranging in the temporal sense – covering object detection approaches from 1990 to 2019. More recent surveys however have been targeted at specific approaches, with [14] considering only supervised and few-shot object detection; and [15] presenting a review of salient object detection approaches.

The present manuscript aspires to improve upon the above mentioned works in the following manner:

- We have taken a historic look at the state of developments in the field of image classification and object detectors. We have deliberately kept our explanation concise yet comprehensive to help readers get intuition on the working of prominent networks. This paper investigates the changes and their implications on the performance of the network. It also explores their evolution and relationship to each other.
- We have included some of the latest works which are missing in existing surveys. Some of the networks covered have limited explanation available in the current literature.
- This work contains a detailed discussion on lightweight object detection models, which was absent in such detail in earlier works. Such lightweight models are becoming increasingly important in the context of Edge Intelligence and TinyAI¹ where ‘compressed’ or ‘low-parameter-count’ models are desirable to facilitate their execution (both training and inference) on resource-constrained edge/IoT nodes.
- Graphical representations of the number of images in the different classes in various training datasets have been included to provide insights into the structure of the data being used for training different models. This pictorial representation can be used for readily understanding the structure of the training data. For instance, a visual comparison between Fig. 3 and Fig. 5 is expected to inform the reader that (i) the Pascal VOC dataset has lesser number of classes than the Microsoft COCO dataset, (ii) the majority of the classes in the Pascal VOC dataset have image counts ranging between 1000 and 2500, while in the Microsoft COCO dataset most classes have image counts between 1200 and 43000.

4. Datasets and Evaluation Metrics

4.1. Datasets

This section presents an overview of the popular datasets most commonly used for object detection tasks.

4.1.1. PASCAL VOC 07/12

The Pascal Visual Object Classes (VOC) challenge was a multiyear effort to accelerate the development in the field of visual perception. It started in 2005 with classification and detection tasks on four object classes [18], but two versions of this challenge are mostly used as a standard benchmark. While the VOC07 challenge had 5k training images and more than 12k labelled objects [35], the VOC12 challenge increased them to 11k training images and more

than 27k labelled objects [36]. Object classes were expanded to 20 categories and the tasks like segmentation and action detection were included as well. Pascal VOC introduced the mean Average Precision (mAP) at 0.5 IoU (Intersection over Union) to evaluate the performance of the models. Figure 3 depicts the distribution of the number of images w.r.t. the different classes in the Pascal VOC dataset.

4.1.2. ILSVRC

The ImageNet Large Scale Visual Recognition Challenge (ILSVRC) [17] was an annual challenge running from 2010 to 2017 and became a benchmark for evaluating algorithm performance. The dataset size was scaled up to more than a million images consisting of 1000 object classification classes. 200 of these classes were hand-picked for object detection task, constitute of more than 500k images. Various sources including ImageNet [37] and Flickr, were used to construct detection dataset. ILSVRC also updated the evaluation metric by loosening the IoU threshold to help include smaller object detection. Figure 4 depicts the distribution of the number of images w.r.t. the different classes in the ImageNet dataset.

4.1.3. MS-COCO

The Microsoft Common Objects in Context (MS-COCO) [22] is one of the most challenging datasets available. It has 91 common objects found in their natural context which a 4-year-old human can easily recognize. It was launched in 2015 and its popularity has only increased since then. It has more than two million instances and an average of 3.5 categories per images. Furthermore, it contains 7.7 instances per image, comfortably more than other popular datasets. MS COCO comprises of images from varied viewpoints as well. It also introduced a more stringent method to measure the performance of the detector. Unlike the Pascal VOC and ILSVRC, it calculates the IoU from 0.5 to 0.95 in steps of 0.5, then uses a combination of these 10 values as final metric, called Average Precision (AP). Apart from this, it also utilizes AP for small, medium and large objects separately to compare performance at different scales. Figure 5 depicts the distribution of the number of images w.r.t. the different classes in the MS-COCO dataset.

4.1.4. Open Image

Google’s Open Images [39] dataset is composed of 9.2 million images, annotated with image-level labels, object bounding boxes, and segmentation masks, among others. It was launched in 2017 and has since received six updates. For object detection, Open Images has 16 million bounding boxes for 600 categories on 1.9 million images, which makes it the largest dataset of object localization. Its creators took extra care to choose interesting, complex and diverse images, having 8.3 object categories per image. Several changes were made to the AP introduced in Pascal VOC like ignoring un-annotated class, detection requirement for class and its subclass, etc. Figure 6 depicts the distribution of the number of images with respect to the different classes in the Open Images dataset.

¹Tiny AI | MIT Technology Review
(URL: <https://www.technologyreview.com/technology/tiny-ai/>)

Table 1

Detailed comparison of existing surveys and reviews on object detection.

Reference	Year	Topic	Remarks
[4]	2017	Comparison of image datasets for machine learning	This work compares 7 object detection datasets, viz. PASCAL VOC, INRIA, Caltech, ImageNet, SUN, Kitti, and Microsoft COCO [17, 18, 19, 2, 20, 21, 22]
[5]	2017	CNN-based object detectors	This work contains a discussion on regional proposal and regression approaches for object detection. The discussion is restricted to CNN-based detectors only.
[6]	2017	Accuracy/Speed trade-offs for convolution-based object detectors (till 2017)	This work contained a detailed discussion on the accuracy vs. speed trade-offs in high-performance object detection models available till 2017 (Faster R-CNN[23], R-FCN[24], and SSD[25]) by varying critical hyperparameters and feature extractors used in the respective models
[7]	2017	Comparative study of object detection models	This work is very similar to [6], and considers SSD[25], Faster R-CNN[23], and R-FCN[24], for evaluating and comparing speed and accuracy
[8]	2018	Latest advances in deep convolutional neural network based object detection models	This work discussed and compared several deep neural network based object detection models (available till the time of publication: R-CNN[26], SPPNet[27], Fast-RCNN[28], Faster RCNN[23], R-FCN[24], FPN[29], Mask RCNN[30], and Deformable CNN [31]). The work mainly focused on the design differences, and variations in the dataset(s) used by the different models
[9]	2019	Comparison of the performance of several object detection models with different feature extractors	This work compares the performance of several available object detection models (SSD[25], Faster RCNN[23], and R-FCN[24]). The authors used different feature extractors (Inception V2[32], Resnet-101[33], and Mobilenet-V1[34]) for the performance comparison analyses
[10]	2019	Review on deep learning object detection models	This work contains a detailed analysis of object detection models available till 2019 by using Microsoft COCO dataset[22]. Further, the work also includes a brief overview of three common tasks most often required in modern-day computer vision applications viz. salient object detection, face detection, and pedestrian detection
[11]	2019	Survey on models for salient object detection	This work mainly discusses the models available in the technical literature for the purpose of detection and segmentation of salient objects from natural scenes. It is to be noted that this work does not discuss object proposal generation, generic scene segmentation, and fixation prediction
[12]	2019	Survey on object detection techniques from 1990 to early 2019	This work contains discussions on the history of object detection approaches, datasets, and speed up techniques besides including explanations on the existing state-of-art detection models till that time
[13]	2020	Survey on generic object detection	This work focuses on generic object detection and considers detection frameworks, context modeling, feature representation, and proposal generation
[14]	2021	Survey of Self-Supervised and Few-Shot Object Detection	This work has a narrower field of discussion, and contains a review of recent research on few-shot and self-supervised object detection only
[15]	2021	Survey on salient object detection by employing deep learning models	This survey deals with algorithm taxonomy, network architecture, level of supervision, learning paradigm, and object/instance level detection, and unresolved issues in that area of research
This Work	2022	A Survey of Modern Deep Learning based Object Detection Models	This work provides a comprehensive review of various state-of-the-art backbone architectures and object detectors present in the technical literature up to August 2021. We also take a detailed look at the contemporary lightweight networks suitable for computer vision tasks in resource-constrained environments

Table 2

Comparison of various object detection datasets.

Dataset	Classes	Train			Validation			Test
		Images	Objects	Objects/Image	Images	Objects	Objects/Image	
PASCAL VOC 12	20	5,717	13,609	2.38	5,823	13,841	2.37	10,991
MS-COCO	80	118,287	860,001	7.27	5,000	36,781	7.35	40,670
ILSVRC	200	456,567	478,807	1.05	20,121	55,501	2.76	40,152
OpenImage	600	1,743,042	14,610,229	8.38	41,620	204,621	4.92	125,436




Figure 3: (This image is best viewed in PDF form with magnification) Number of images for different classes annotated in the PascalVOC dataset [38]

Issues of Data Skew/Bias

While observing Fig. 3 through Fig. 6, an alert reader would certainly notice that the number of images for difference classes vary significantly in all the datasets [38]. Three (Pascal VOC, MS-COCO, and Open Images Dataset) of the four datasets discussed above have a very significant drop in the number of images beyond the top-5 most frequent classes. As can be readily observed for Fig. 3, there are 13775 images which contain a ‘person’ and then 2829 images which contain a ‘car’. The number of images for the remaining 18 classes in this dataset almost fall linearly to the 55 images of ‘sheep’. Similarly, for the MS-COCO dataset, the class ‘person’ has 262465 images, and the next most-frequent class ‘car’ has 43867 images. The downward

trend continues till there are only 198 images for the class ‘hair drier’. A similar phenomenon is also observed in the Open Images Dataset, wherein the class ‘Man’ is the most frequent with 378077 images, and the class ‘Paper Cutter’ has only 3 images. This clearly represents a skew in the datasets and is bound to create a bias in the training process of any object detection model. Therefore, an object detection model trained on these skewed datasets will in all probability show better detection performance for the classes with more number of images in the training data. Although still present, this issue is slightly less pronounced in the ImageNet dataset, as can be observed from Fig. 4 from where it can be seen that the most frequent class i.e. ‘koala bear’ has 2469 images, and the least frequent class i.e. ‘cart’ has 624 images. However, this




Figure 4: (This image is best viewed in PDF form with magnification) Number of images for different classes annotated in the ImageNet dataset [38]




Figure 5: (This image is best viewed in PDF form with magnification) Number of images for different classes annotated in the MS-COCO dataset [38]



Figure 6: (This image is best viewed in PDF form with magnification) Number of images for different classes annotated in the Open Images dataset [38]

leads to another point of concern in the ImageNet dataset: the most frequent class is for ‘koala’ and the next most-appearing class is ‘computer keyboard’, which are clearly not the most sought after objects in a real-world object detection scenario (where person, cars, traffic signs, etc. are of higher concern).

4.2. Metrics

Object detectors use multiple criteria to measure the performance of the detectors viz., frames per second (FPS), precision and recall. However, mean Average Precision (*mAP*) is the most common evaluation metric. Precision is derived from Intersection over Union (IoU), which is the ratio of the

area of overlap and the area of union between the ground truth and the predicted bounding box. A threshold is set to determine if the detection is correct. If the IoU is more than the threshold, it is classified as True Positive while an IoU below it is classified as False Positive. If the model fails to detect an object present in the ground truth, it is termed as False Negative. Precision measures the percentage of correct predictions while the recall measure the correct predictions with respect to the ground truth.

$$\begin{aligned} \text{Precision} &= \frac{\text{True Positive}}{\text{True Positive} + \text{False Positive}} \\ &= \frac{\text{True Positive}}{\text{All Observations}} \end{aligned} \quad (1)$$

$$\begin{aligned} \text{Recall} &= \frac{\text{True Positive}}{\text{True Positive} + \text{False Negative}} \\ &= \frac{\text{True Positive}}{\text{All Ground Truth}} \end{aligned} \quad (2)$$

Based on the above equation, average precision is computed separately for each class. To compare performance between the detectors, the mean of average precision of all classes, called mean average precision (*mAP*) is used, which acts as a single metric for final evaluation.

5. Backbone architectures

Backbone architectures are one of the most important component of the object detector. These networks extract feature from the input image used by the model. Here, we have discussed some milestone backbone architectures used in modern detectors:

5.1. AlexNet

Krizhevsky et al. proposed AlexNet [3], a convolutional neural network based architecture for image classification, and won the ImageNet Large-Scale Visual Recognition Challenge (ILSVRC) 2012 challenge. It achieved a considerably higher accuracy (more than 26%) than the contemporary models. AlexNet is composed of eight learnable layers - five convolutional and three fully connected layers. The last layer of the fully connected layer is connected to an *N*-way (*N*: number of classes) softmax classifier. It uses multiple convolutional kernels throughout the network to obtain features from the image. It also uses dropout and ReLU for regularization and faster training convergence respectively. The convolutional neural networks were given a new life by its reintroduction in AlexNet and it soon became the go-to technique in processing imaging data.

5.2. VGG

While AlexNet [3] and its successors like [40] focused on smaller receptive window size to improve accuracy, Simonyan and Zisserman investigated the effects of network depth on it. They proposed VGG [41], which used small convolution filters to construct networks of varying depths.

While a larger receptive field can be captured by a set of smaller convolutional filters, it drastically reduces network parameters and converges sooner. The paper demonstrated how deep network architecture (16-19 layers) can be used to perform classification and localization with superior accuracy. VGG was created by adding a stack of convolutional layers with three fully connected layers, followed by a softmax layer. The number of convolutional layers in VGG can vary from 8 to 16. VGG is trained in multiple iterations; first, the smallest 11-layer architecture is trained with random initialization whose weights are then used to train larger networks to prevent gradient instability. VGG outperformed ILSVRC 2014 winner GoogLeNet [42] in the single network performance category. It soon became one of the most used network backbones for object classification and detection models.

5.3. GoogLeNet/Inception

Even though classification networks were making inroads towards faster and more accurate networks, deploying them in real-world applications was still a distant possibility due to their resource-intensive nature. As networks are scaled for better performance, the computation cost increases exponentially. Szegedy et al. in [42] postulated the wastage of computations in the network as a major reason for it. Bigger models also have a large number of parameters and tend to overfit the data. They proposed using locally sparse connected architecture instead of a fully connected one to solve these issues. GoogLeNet is thus a 22 layer deep network, made up by stacking multiple Inception modules on top of each other. Inception modules are networks that have multiple sized filters at the same level. Input feature maps pass through these filters and are concatenated at the end before being forwarded to the next layer. The network also has auxiliary classifiers in the intermediate layers to help regularize and propagate gradient. GoogLeNet showed how efficient use of computation blocks can perform at par with other parameter-heavy networks. It achieved 93.3% top-5 accuracy on ImageNet [17] dataset without using external data, while being faster than other contemporary models. Updated versions of Inception like [32],[43] were also published in the following years which further improved its performance and presented evidence of the applications of refined sparsely connected architectures.

5.4. ResNets

As convolutional neural networks become deeper and deeper, Kaiming He et al. in [33] showed how their accuracy first saturates and then degrades rapidly. They proposed the use of skip connections to the stacked convolution layers to mitigate the performance decay. It is realized by addition of a skip connection between the layers. This connection is an element-wise addition between input and output of the block and does not add extra parameter or computational complexity to the network. A typical 34 layer ResNet [33] is basically a large (7x7) convolution filter followed by 16 bottleneck modules (pair of small 3x3 filters with identity shortcut across them) and ultimately a fully connected layer. The

Table 3

Comparison of Backbone architectures.

Model	Year	Layers	Parameters (Million)	Top-1 acc%	FLOPs (Billion)
AlexNet	2012	7	62.4	63.3	1.5
VGG-16	2014	16	138.4	73	15.5
GoogLeNet	2014	22	6.7	-	1.6
ResNet-50	2015	50	25.6	76	3.8
ResNeXt-50	2016	50	25	77.8	4.2
CSPResNeXt-50	2019	59	20.5	78.2	7.9
EfficientNet-B4	2019	160	19	83	4.2

bottleneck architecture can be adapted for deeper networks by stacking 3 convolutional layers ($1 \times 1, 3 \times 3, 1 \times 3$) instead of 2. Kaiming He et al. also demonstrated how the 16-layer VGG net had higher complexity than their considerably deeper 101 and 152 layer ResNet architectures while having lower accuracy. In subsequent paper, the authors proposed Resnetv2 [44] which used batch normalization and ReLU layer in the blocks. It is more generalized and easier to train. ResNets are widely used in classification and detection backbones, and its core principles have inspired many networks ([45, 46, 43]).

5.5. ResNeXt

The existing conventional methods of improving the accuracy of a model were by either increasing the depth or the width of the model. However, increasing any of these leads to higher model complexity and number of parameters while the gain margins diminish rapidly. Xie et al. introduced ResNeXt [46] architecture which is simpler and more efficient than other existing models. ResNeXt was inspired by the stacking of similar blocks in VGG/ResNet[33, 3] and “split-transform-merge” behavior of Inception module[42]. It is essentially a ResNet where each ResNet block is replaced by an inception-like ResNeXt module. The complicated, tailored transformation modules from the Inception is replaced by topologically same modules in the ResNeXt blocks, making the network easier to scale and generalize. Xie et al. also emphasize that the cardinality (topological paths in the ResNeXt block) can be considered as a third dimension, along with depth and width, to improve model accuracy. ResNeXt is elegant and concise. It achieved higher accuracy while having considerably fewer hyperparameters than a similar depth ResNet architecture. It was also the first runner up to the ILSVRC 2016 challenge.

5.6. CSPNet

Existing neural networks have shown incredible results in achieving high accuracy in computer vision tasks; however, they rely on excessive computational resources. Wang et al. believe that heavy inference computations can be reduced by cutting down the duplicate gradient information in the network. They proposed CSPNet [47] which creates different paths for the gradient flow within the network. CSPNet separates feature maps at the base layer into two parts. One part is passed through the partial convolution

network block (e.g., Dense and Transition block in DenseNet [45] or Res(X) block in ResNeXt [46]) while the other part is combined with its outputs at a later stage. This reduces the number of parameters, increases the utilization of computation units and eases memory footprint. It is easy to implement and general enough to be applicable on other architectures like ResNet [33], ResNeXt [46], DenseNet [45], Scaled-YOLOv4 [48] etc. Applying CSPNet on these networks reduced computations from 10% to 20%, while the accuracy remained constant or improved. Memory cost and computational bottleneck is also reduced significantly with this method. It is leveraged in many state of the art detector models, while also being used for mobile and edge devices.

5.7. EfficientNet

Tan et al. systematically studied network scaling and its effects on the model performance in [49]. They summarized how altering network parameters like depth, width and resolution influence its accuracy. They showed how scaling any parameter individually comes with an associated cost. Increasing depth of a network can help in capturing richer and more complex features, but they are difficult to train due to vanishing gradient problem. Similarly, scaling network width will make it easier to capture fine grained features but have difficulty in obtaining high level features. Gains from increasing the image resolution, like depth and width, saturate as model scales. In the paper [49], Tan et al. proposed the use of a compound coefficient that can uniformly scale all three dimensions. Each model parameter has an associated constant, which is found by fixing the coefficient as 1 and performing a grid search on a baseline network. The baseline architecture, inspired by their previous work [50], is developed by neural architecture search on a search target while optimizing accuracy and computations. EfficientNet is a simple and efficient architecture. It outperformed existing models in accuracy and speed while being considerably smaller. By providing a monumental increase in efficiency, it could potentially open a new era in the field of efficient networks.

6. Object Detectors

We have divided this review based on the three types of detectors — two-stage, single-stage and transformer-based detectors. However, we also discussed the pioneer work,




Figure 7: Visualization of CNN Architectures³. Left to Right: AlexNet, VGG-16, GoogLeNet, ResNet-50, CSPResNeXt-50, EfficientNet-B4.

where we briefly examine a few traditional object detectors. A network which has a separate module to generate region proposals is termed as a two-stage detector. These models try to find an arbitrary number of objects proposals in an image during the first stage and then classify and localize them in the second. As these systems have two separate steps, they generally take longer to generate proposals, have complicated architecture and lacks global context. Single-stage detectors classify and localize semantic objects in a single shot using dense sampling. They use predefined boxes/keypoints of various scale and aspect ratio to localize objects. It edges two-stage detectors in real-time performance and simpler design. While Transformers were initially introduced for NLP, their general-purpose nature helped in migrating to vision tasks. Although transformer based detectors have achieved SOTA in many vision tasks, they require comparatively more data to train.

6.1. Pioneer Work

6.1.1. Viola-Jones

Primarily designed for face detection, Viola-Jones object detector [1], proposed in 2001, was an accurate and powerful detector. It combined multiple techniques like Haar-like features, integral image, Adaboost and cascading classifier. First step is to search for Haar-like features by sliding a window on the input image and uses integral image to calculate. It then uses a trained Adaboost to find the classifier of each haar feature and cascades them. Viola Jones algorithm is still used in small devices as it is efficient and fast.

6.1.2. HOG Detector

In 2005, Dalal and Triggs proposed the Histogram of Oriented Gradients (HOG) [2] feature descriptor to extract image features for object detection. It was an improvement over other detectors like [51, 52, 53, 54]. HOG extracts gradient and its orientation of the edges to create a feature table. The image is divided into grids and the feature table is then used to create histogram for each cell in the grid. HOG features are generated for the region of interest and fed into a linear SVM classifier for detection. The detector was proposed for pedestrian detection; however, it could be trained to detect various classes.

6.1.3. DPM

Deformable Parts Model (DPM) [55] was introduced by Felzenszwalb et al. and was the winner Pascal VOC challenge in 2009. It used individual “part” of the object for detection and achieved higher accuracy than HOG. It follows the philosophy of *divide and rule*; parts of the object are individually detected during inference time and a probable arrangement of them is marked as detection. For example, a human body can be considered as a collection of parts like head, arms, legs and torso. One model will be assigned to capture one of the parts in the whole image and the process is repeated for all such parts. A model then removes improbable configurations of the combination of these parts

³Tool Used: <https://netron.app/>

to produce detection. DPM based models [56, 57] were one of the most successful algorithms before the era of deep learning.

6.2. Two-Stage Detectors

6.2.1. R-CNN

The Region-based Convolutional Neural Network (R-CNN) [26] was the first paper in the R-CNN family, and demonstrated how CNNs can be used to immensely improve the detection performance. R-CNN use a class agnostic region proposal module with CNNs to convert detection into classification and localization problem. A mean-subtracted input image is first passed through the region proposal module, which produces 2000 object candidates. This module find parts of the image which has a higher probability of finding an object using Selective Search [58]. These candidates are then warped and propagated through a CNN network, which extracts a 4096-dimension feature vector for each proposal. Girshick et al. used AlexNet [3] as the backbone architecture of the detector. The feature vectors are then passed to the trained, class-specific Support Vector Machines (SVMs) to obtain confidence scores. Non-maximum suppression (NMS) is later applied to the scored regions, based on its IoU and class. Once the class has been identified, the algorithm predicts its bounding box using a trained bounding-box regressor, which predicts four parameters i.e., center coordinates of box along with its width and height.

R-CNN has a complicated multistage training process. The first stage is pre-training the CNN with a large classification dataset. It is then fine-tuned for detection using domain-specific images (mean-subtracted, warped proposals) by replacing of the classification layer with a randomly initialized $N+1$ -way classifier, N being the number of classes, using stochastic gradient descent (SGD) [59]. One liner SVM and bounding box regressor is trained for each class.

R-CNN ushered a new wave in the field of object detection, but it was slow (47 sec per image) and expensive in time and space [28]. It had complex training process and took days to train on small datasets even when some of the computations were shared.

6.2.2. SPP-Net

He et al. proposed the use of Spatial Pyramid Pooling (SPP) layer [60] to process image of arbitrary size or aspect ratio. They realized that only the fully connected part of the CNN required a fixed input. SPP-net [27] merely shifted the convolution layers of CNN before the region proposal module and added a pooling layer, thereby making the network independent of size/aspect ratio and reducing the computations. The selective search [58] algorithm is used to generate candidate windows. Feature maps are obtained by passing the input image through the convolution layers of a ZF-5 [40] network. The candidate windows are then mapped on to the feature maps, which are subsequently converted into fixed length representations by spatial bins of a pyramidal pooling layer. This vector is passed to the fully connected layer and ultimately, to SVM classifiers to predict

class and score. Similar to R-CNN [26], SPP-net has a post processing layer to improve localization by bounding box regression. It also uses the same multistage training process, except that the fine tuning is done only on the fully connected layers.

SPP-Net is considerably faster than the R-CNN model with comparable accuracy. It can process images of any shape/aspect ratio and thus, avoid object deformation due to input warping. However, as its architecture is analogous to R-CNN, it shared R-CNN's disadvantages too like multistage training, computationally expensive and longer training time.

6.2.3. Fast R-CNN

One of the major issues with R-CNN/SPP-Net was the need to train multiple systems separately. Fast R-CNN [28] solved this by creating a single end-to-end trainable system. The network takes as input an image and its object proposals. The image is passed through a set of convolution layers to obtain feature maps and the object proposals are mapped to it. Girshick et al. replaced pyramidal structure of pooling layers from SPP-net [27] with a single spatial bin, called RoI pooling layer. This layer is connected to 2 fully connected layer and then branches out into a $N+1$ -class SoftMax layer and a bounding box regressor layer, which has a fully connected layer as well. The model also changed the loss function of bounding box regressor from L2 to smooth L1 to improve performance and introduced a multi-task loss to better train the network.

Girshick et al. used modified version of existing state-of-art pre-trained models like [3], [41] and [61] as backbone. The network was trained in a single step by stochastic gradient descent (SGD) and a mini-batch of 2 images. This helped the network converge faster as the back-propagation shared computations among the RoIs from the two images.

Fast R-CNN was introduced as an improvement in speed (146x) on R-CNN while the increase in accuracy was supplementary. It simplified training procedure, removed pyramidal pooling and presented a new loss function. The object detector, without the region proposal network, reported near real time speed with considerable accuracy.

6.2.4. Faster R-CNN

Even though Fast R-CNN inched closer to real time object detection, its region proposal generation was still an order of magnitude slower (2 sec per image compared to 0.2 sec per image). Ren et al. suggested a fully convoluted network [62] as a region proposal network (RPN) in [23] that takes an arbitrary input image and outputs a set of candidate windows. Each such window has an associated *objectness score* which determines likelihood of an object. Unlike its predecessors ([56, 28, 33]) which used image pyramids to solve size variance of objects, RPN introduces Anchor boxes. It used multiple bounding boxes of different aspect ratios and regressed over them to localize object. The input image is first passed through the CNN to obtain a set of feature maps. These are forwarded to the RPN, which




Figure 8: Illustration of the internal architecture of different two stage object detectors⁵.

produces bounding boxes and their classification. Selected proposals are then mapped back to the feature maps obtained from the previous CNN layer in the RoI pooling layer, and ultimately fed to the fully connected layer. The result is then sent to classifier and bounding box regressor. Faster R-CNN is essentially Fast R-CNN with RPN as region proposal module.

Training of Faster R-CNN is more convoluted, due to the presence of shared layers between two models which perform very different tasks. Firstly, RPN is pre-trained on ImageNet dataset [37] and fine-tuned on PASCAL VOC dataset [18]. A Fast R-CNN is trained from the region proposals of RPN from the first step. Till this point, the networks do not have shared convolution layer. Now, we fix the convolution layers of the detector and fine-tune the unique layers in RPN. And finally, Fast R-CNN is fine-tuned from the updated RPN.

Faster R-CNN improved the detection accuracy over the previous state-of-art [28] by more than 3% and decreased inference time by an order of magnitude. It fixed the bottleneck of slow region proposal and ran in near real time at 5 frames per second. Another advantage of having a CNN in region proposal was that it could learn to produce better proposals and thereby increase accuracy.

6.2.5. FPN

Use of image pyramid to obtain feature pyramid (or *featurized image pyramids*) at multiple levels is a common

method to increase detection of small objects. Even though it increases Average Precision of the detector, the increase in the inference time is substantial. Lin et al. proposed the Feature Pyramid Network (FPN) [63], which has a top-down architecture with lateral connections to build high-level semantic features at different scales. The FPN has two pathways, a bottom-up pathway which is a ConvNet computing feature hierarchy at several scales and a top-down pathway which upsamples coarse feature maps from higher level into high-resolution features. These pathways are connected by lateral connection by a 1×1 convolution operation to enhance the semantic information in the features. FPN is used as a region proposal network (RPN) of a ResNet-101 [33] based Faster R-CNN here.

FPN could provide high-level semantics at all scales, which reduced the error rate in detection. It became a standard building block in future detections models and improved their accuracy across the table. It also lead to development of other improved networks like PANet [64], NAS-FPN [65] and EfficientNet [49].

6.2.6. R-FCN

Dai et al. proposed Region-based Fully Convolutional Network (R-FCN) [24] that shared almost all computations within the network, unlike previous two stage detectors which applied resource intensive techniques on each proposal. They argued against the use of fully connected layers and instead used convolutional layers. However, deeper

layers in the convolutional network are translation-invariant, making them ineffective for localization tasks. The authors proposed the use of position-sensitive score maps to remedy it. These sensitive score maps encode relative spatial information of the subject and are later pooled to identify exact localization. R-FCN does it by dividing the region of interest into $k \times k$ grid and scoring the likeliness of each cell with the detection class feature map. These scores are later averaged and used to predict the object class. R-FCN detector is a combination of four convolutional networks. The input image is first passed through the ResNet-101[33] to get feature maps. An intermediate output (Conv4 layer) is passed to a Region Proposal Network (RPN) to identify the RoI proposals while the final output is further processed through a convolutional layer and is input to classifier and regressor. The classification layer combines the created position-sensitive map with the RoI proposals to generate predictions while the regression network outputs the bounding box attributes. R-FCN is trained in a similar 4 step fashion as Faster-RCNN [23] whilst using a combined cross-entropy and box regression loss. It also adopts online hard example mining (OHEM) [66] during the training.

Dai et al. offered a novel method to solve the problem of translation invariance in convolutional neural networks. R-FCN combines Faster R-CNN and FCN to create a fast, accurate detector. Even though it did not improve accuracy by much, it was 2.5-20 times faster than its counterpart.

6.2.7. Mask R-CNN

Mask R-CNN [30] extends on the Faster R-CNN by adding another branch in parallel for pixel-level object instance segmentation. The branch is a fully connected network applied on the RoIs to classify each pixel into segments with little overall computation cost. It uses similar basic Faster R-CNN architecture for object proposal, but adds a mask head parallel to classification and bounding box regressor head. One major difference is the use of RoIAlign layer, instead of RoIPool layer, to avoid pixel level misalignment due to spatial quantization. The authors chose the ResNeXt-101 [46] as its backbone along with the feature Pyramid Network (FPN) for better accuracy and speed. The loss function of Faster R-CNN is updated with the mask loss and uses 5 anchor boxes with 3 aspect ratio, similar to FPN. Overall training of Mask R-CNN is similar to faster R-CNN.

Mask R-CNN performed better than the existing state of the art single-model architectures, added an extra functionality of instance segmentation with little overhead computations. It is simple to train, flexible and generalizes well in applications like keypoint detection, human pose estimation, etc. However, it is still below the real time performance (>30 fps).

6.2.8. DetectoRS

Many contemporary two stage detectors like [23, 67, 68] use the mechanism of looking and thinking twice i.e. calculating object proposals first and using them to extract features to detect objects. DetectoRS [69] applies this mechanism at

both macro and micro level of the network. At macro level, they propose Recursive Feature Pyramid (RFP), formed by stacking multiple feature pyramid network (FPN) with extra feedback connection from the top-down level path in FPN to the bottom-up layer. The output of the FPN is processed by the Atrous Spatial Pyramid Pooling layer (ASPP) [70] before passing it to the next FPN layer. A Fusion module is used to combine FPN outputs from different modules by creating an attention map. At micro level, Qiao et al. presented the Switchable Atrous Convolution (SAC) to regulate the dilation rate of convolution. An average pooling layer with 5×5 filter and a 1×1 convolution is used as a switch function to decide the rate of atrous convolution [71], helping the backbone detect objects at various scale on the fly. They also packed the SAC in between two global context modules [72] as it helps in making more stable switching. The combination of these two techniques, Recursive Feature Pyramid and Switchable Atrous Convolution results in DetectoRS. The authors incorporated the above techniques in the Hybrid Task Cascade (HTC) [67] as the baseline model and a ResNext-101 backbone.

DetectoRS combined multiple systems to improve performance of the detector and sets the state-of-the-art for the two stage detectors. Its RFP and SAC modules generalized well and can be used in other detection models. However, it is not suitable for real time detections as it can only process about 4 frames per second.

6.3. Single Stage Detectors

6.3.1. YOLO

Two stage detectors solve the object detection as a classification problem, a module presents candidates which the network classifies as either an object or background. However, YOLO or You Only Look Once [73] reframed it as a regression problem, directly predicting the image pixels as objects and its bounding box attributes. In YOLO, the input image is divided into a $S \times S$ grid and the cell where the object's center falls is responsible for detecting it. A grid cell predicts multiple bounding boxes, and each prediction array consists of 5 elements: center of bounding box – x and y, dimensions of the box – w and h, and the confidence score.

YOLO was inspired from the GoogLeNet model for image classification [42], which uses cascaded modules of smaller convolution networks [74]. It is pre-trained on ImageNet data [37] till the model achieves high accuracy and consequently modified by adding randomly initialized convolution and fully connected layers. At training time, the grid cells predict only one class as the network converges better, but it is increased during the inference time. Multitask loss, i.e. combined loss of all predicted components, is used to optimize the model. Non maximum suppression (NMS) removes the class-specific multiple detections.

YOLO surpassed its contemporary single stage real time models by a huge margin in both accuracy and speed. However, it had significant shortcomings as well. Localization accuracy for small or clustered objects and limitation on the




Figure 9: Illustration of the internal architecture of different two and single stage object detectors².

number of objects per cell were its major drawbacks. These issues were fixed in later versions of YOLO [75, 76, 77].

6.3.2. SSD

Single Shot MultiBox Detector (SSD) [25] was the first single stage detector that matched accuracy of contemporary two stage detectors like Faster R-CNN [23], while maintaining real time speed. SSD was built on VGG-16 [41], with additional auxiliary structures to improve performance. These auxiliary convolution layers, added to the end of the model, decrease progressively in size. SSD detects smaller objects earlier in the network when the image features are not too crude, while the deeper layers are responsible for offset of the default boxes and aspect ratios [78].

During training, SSD match each ground truth box with the default boxes with the best jaccard overlap and train the network accordingly, similar to Multibox [78]. It also uses hard negative mining and heavy data augmentation. Similar to DPM [55], it utilized weighted sum of the localization and confidence loss to train the model. Final output is obtained by performing non maximum suppression.

Even though SSD was significantly faster and more accurate than both state-of-the-art networks like YOLO and Faster R-CNN, it had difficulty in detecting small objects. This issue was later solved by using better backbone architectures like ResNet [33] and implementing other small fixes.

6.3.3. YOLOv2 and YOLO9000

YOLOv2 [75], an improvement on the YOLO [73], offered an easy tradeoff between speed and accuracy while

the YOLO9000 model could predict 9000 object classes in real time. They replaced the backbone architecture of GoogLeNet [42] with DarkNet-19 [79]. It incorporated many impressive techniques like Batch Normalization [80] to improve convergence, joint training of classification and detection systems to increase detection classes, removing fully connected layers to increase speed and using learnt anchor boxes to improve recall and priors. Redmon et al. also combined the classification and detection datasets in hierarchical structure using WordNet [81]. This WordTree can be used to predict a higher conditional probability of hypernym, even when the hyponym is not classified correctly, thereby increasing the overall performance of the system.

YOLOv2 provided better flexibility to choose the model on speed and accuracy, and the new architecture had fewer parameters. As the title of the paper suggests, it was “*better, faster and stronger*” [75].

6.3.4. RetinaNet

Given the difference between the accuracies of single and two stage detectors, Lin et al. suggested that the reason single stage detectors lag is the “extreme foreground-background class imbalance” [29]. They proposed a reshaped cross entropy loss, called Focal loss as the means to remedy the imbalance. Focal loss parameter reduces the loss contribution from easy examples. The authors demonstrate its efficacy with the help of a simple, single stage detector, called RetinaNet [29], which predicts objects by dense sampling of the input image in location, scale and aspect ratio. It uses ResNet [33] augmented by Feature Pyramid

Network (FPN) [63] as the backbone and two similar subnets - classification and bounding box regressor. Each layer from the FPN is passed to the subnets, enabling it to detect objects at various scales. The classification subnet predicts the object score for each location while the box regression subnet regresses the offset for each anchor to the ground truth. Both subnets are small FCN and share parameters across the individual networks. Unlike most previous works, a class-agnostic bounding box regressor was employed and found to be equally effective.

RetinaNet is simple to train, converges faster and easy to implement. It achieved better performance in accuracy and run time than the two stage detectors. RetinaNet also pushed the envelope in advancing the ways object detectors are optimized by the introduction of a new loss function.

6.3.5. YOLOv3

YOLOv3 had “incremental improvements” from the previous YOLO versions [73, 75]. Redmon et al. replaced the feature extractor network with a larger Darknet-53 network [79]. They also incorporated various techniques like data augmentation, multi-scale training, batch normalization, among others. Softmax in classifier layer was replaced by a logistical classifier.

Even though YOLOv3 was faster than YOLOv2 [75], it lacked any ground breaking change from its predecessor. It even had lesser accuracy than an year old state-of-the-art detector [29].

6.3.6. CenterNet

Zhou et al. in [82] takes a very different approach of modelling objects as points, instead of the conventional bounding box representation. CenterNet predicts the object as a single point at the center of the bounding box. The input image is passed through the FCN that generates a heatmap, whose peaks correspond to center of detected object. It uses a ImageNet pretrained stacked Hourglass-101 [83] as the feature extractor network and has 3 heads – heatmap head to determine the object center, dimension head to estimate size of object and offset head to correct offset of object point. Multitask loss of all three heads is back propagated to feature extractor while training. During inference, the output from the offset head is used to determine the object point and finally a box is generated. As the predictions are points instead of bounding boxes, non-maximum suppression (NMS) is not required for post-processing.

CenterNet brings a fresh perspective and set aside years of progress in the field of object detection. It is more accurate and has lesser inference time than its predecessors. It has high precision for multiple tasks like 3D object detection, keypoint estimation, pose, instance segmentation, orientation detection and others. However, it requires different backbone architectures as general architectures that work well with other detectors give poor performance with it and vice-versa.

⁵Features created using: <https://poloclub.github.io/cnn-explainer/>

6.3.7. EfficientDet

EfficientDet [84] builds towards the idea of scalable detector with higher accuracy and efficiency. It introduces efficient multi-scale features, BiFPN and model scaling. BiFPN is bi-directional feature pyramid network with learnable weights for cross connection of input features at different scales. It improves on NAS-FPN [65], which required heavy training and had complex network, by removing one-input nodes and adding an extra lateral connection. This eliminates less efficient nodes and enhances high-level feature fusion. Unlike existing detectors which scale up with bigger, deeper backbone or stacking FPN layers, EfficientDet introduces a compounding coefficient which can be used to “jointly scale up all dimensions of backbone network, BiFPN network, class/box network and resolution” [84]. EfficientDet utilizes EfficientNet [49] as the backbone network with multiple sets of BiFPN layers stacked in series as feature extraction network. Each output from the final BiFPN layer is sent to class and box prediction network. The model is trained using SGD optimizer along with synchronized batch normalization and uses swish activation [85], instead of the standard ReLU activation, which is differentiable, more efficient and has better performance.

EfficientDet achieves better efficiency and accuracy than previous detectors while being smaller and computationally cheaper. It is easy to scale and generalizes well for other tasks.

6.3.8. YOLOv4

YOLOv4 [77] incorporated a lot of exciting ideas to design a fast and easy to train object detector that could work in existing production systems. It utilizes “bag of freebies” i.e., methods that only increase training time and do not affect the inference time. YOLOv4 utilizes data augmentation techniques, regularization methods, class label smoothing, CIoU-loss [86], Cross mini-Batch Normalization (CmBN), Self-adversarial training, Cosine annealing scheduler [87] and other tricks to improve training. Methods that only affect the inference time, called “Bag of Specials”, are also added to the network, including Mish activation [88], Cross-stage partial connections (CSP) [47], SPP-Block [27], PAN path aggregated block [64], Multi input weighted residual connections (MiWRC), etc. It also used genetic algorithm for searching hyper-parameter. It has an ImageNet pre-trained CSPNetDarknet-53 backbone, SPP and PAN block neck and YOLOv3 as detection head.

Most existing detection algorithms require multiple GPUs to train model, but YOLOv4 can be easily trained on a single GPU. It is twice as fast as EfficientDet with comparable performance [77].

6.3.9. The Curious Case of YOLOv5

YOLOv5 was released only a month after its predecessor (i.e. YOLOv4) by a company called Ultralytics [89], and claimed to have several significant improvements over existing YOLO detectors [90]. Since the YOLOv5 model was only released as a GitHub repository, and the model

was not published as a peer-reviewed research, there were doubts/apprehensions about the authenticity and the effectiveness of the proposal [90]. The model was put to detailed scrutiny by another company called Roboflow, and it was found that the only significant modification that YOLOv5 included (over YOLOv4) was integrating the anchor box selection process into the model [91]. As a result, YOLOv5 does not need to consider any of the datasets to be used as input, and possesses the capability to automatically learn the most appropriate anchor boxes for the particular dataset under consideration, and use them during training [91]. Despite the non-availability of a formal paper, the fact that the YOLOv5 model has subsequently been utilized in several applications with effective results has started to generate credibility for the model [92, 93]. Lastly, it needs to be pointed out the latest version of the YOLOv5 model is the YOLOv5-V6.0 release which is claimed to be a (further) lightweight model with an improved inference speed of 1666 fps (the original YOLOv5 claimed to have 140 fps) [94]. Another improved version of the YOLOv5 model is recently presented in [95] which integrates the salient features of the Transformer model (details of Transformers are discussed in a subsequent section) into the YOLOv5 model, and demonstrates the improvements in the context of object detection in drone-captured videos.

6.4. Transformer based Detectors

Transformers [96] have had a profound impact in the Natural Language Processing (NLP) domain since its inception. Its application in language models like BERT (Bidirectional Encoder Representation from Transformers) [97], GPT (Generative Pre-trained Transformer) [98], T5 (Text-To-Text Transfer Transformer) [99] etc. have pushed the state-of-the-art in the field. Transformers [96] uses the attention model to establish dependencies among the elements of the sequence and can attend to longer context than other sequential architectures. The success of transformers in NLP sparked interest in its application in computer vision. The work in [100] introduced the first transformer for classification named ViT. While CNNs have been the backbone on advancement in vision, transformers have inherent shortcomings like the lack of inductive bias, fixed post-training weights [101] etc.

6.4.1. DeTR

Detection Transformer or DeTR was presented by Carion et al. in [102]. It used a combination of CNN and transformer to create an end-to-end trainable detector and removed any hand-crafted modules, like non-maximum suppression or anchor generation. The input image is first passed through a CNN backbone network to obtain image features and then through a set of transformers. Final predictions, a bipartite output of class and bounding box, are obtained by passing transformer output through feed-forward networks. DeTR used ResNets [33] as backbone network and transformer similar to the original work [96]. The transformer encoder takes image features, along with the position encodings as input and directs result to the decoder. The decoder

manipulates N input embeddings, called object queries, to generate output. Object queries are position encodings that are initialized as zeros but learnt during training. Multi-head attentions in the decoder modifies these object queries with encoder embeddings to generate results, which are ultimately passed through multi-layer perceptrons to predict class and bounding boxes. As the number of outputs in DeTR is defined by object queries, a special 'no-object' or \emptyset classes is assigned to background class. DeTR uses bipartite matching loss to find the optimal one-to-one matching between detector output and padded ground truth.

Carion et al. [102] introduced a simple, general-purpose transformer-based detector which is competitive with the CNN based detectors. However, its performance on small objects leaves something to be desired.

6.4.2. Swin Transformer

Swin Transformer [103] seeks to provide a transformer-based backbone for computer vision tasks. It splits the input images in multiple, non-overlapping patches and converts them into embeddings. Numerous Swin Transformer blocks are then applied to the patches in 4 stages, with each successive stage reducing the number of patches to maintain hierarchical representation. The Swin Transformer block is composed of local multi-headed self-attention (MSA) modules, based on alternating shifted patch window in successive blocks. Computation complexity becomes linear with image size in local self-attention while shifted window enables cross-window connection. [103] also shows how shifted windows increase detection accuracy with little overhead.

Swin Transformer achieved the state-of-the-art on MS COCO dataset, but utilises comparatively higher parameters than convolutional models. Transformers present a paradigm shift from the CNN based neural networks. While its application in vision is still in a nascent stage, its potential to replace convolution from these tasks is very real.

7. Lightweight Networks

A new branch of research has shaped up in recent years, aimed at designing small and efficient networks for resource constrained environments as is common in Internet of Things (IoT) deployments [104, 105, 106, 107]. This trend has percolated to the design of potent object detectors too. It is seen that although a large number of object detectors achieve excellent accuracy and perform inference in real-time, a majority of these models require excessive computing resources and therefore cannot be deployed on edge devices.

Many different approaches have shown exciting results in the past. Utilization of efficient components and compression techniques like pruning ([108, 109]), quantization ([110, 111]), hashing [112], etc. have improved the efficiency of deep learning models. Use of trained large network to train smaller models, called distillation [113], has also shown interesting results. However in this section, we explore some prominent examples of efficient neural network design for achieving high performance on edge devices.

7.1. SqueezeNet

Recent advances in the field of CNNs had mostly focused on improving the state-of-the-art accuracy on the benchmark datasets, which led to an explosion of model size and their parameters. But in 2016, Iandola et al. proposed a smaller, smarter network called SqueezeNet [114], which reduced the parameters while maintaining the performance. They achieved it by employing three main design strategies viz. using smaller filters, decreasing the number of input channels to 3×3 filters and placing downsampling layers later in the network. The first two strategies decrease the number of parameters while attempting to preserve the accuracy and the third strategy increases the accuracy of the network. The building block of SqueezeNet is called a fire module, which consist of two layers: a squeeze layer and an expand layer, each with a ReLU activation. The squeeze layer is made up of multiple 1×1 filters while the expand layer is a mix of 1×1 and 3×3 filters, thereby limiting the number of input channels. The SqueezeNet architecture is composed of a stack of 8 Fire modules squashed in between the convolution layers. Inspired by ResNet [33], SqueezeNet with residual connections was also proposed which increased the accuracy over the vanilla model. The authors also experimented with Deep Compression [110] and achieved $510\times$ reduction in model size compared to AlexNet, while maintaining the baseline accuracy. SqueezeNet presented a good candidate for improving the hardware efficiency of the neural network architectures.

7.2. MobileNets

MobileNet [34] moved away from the conventional methods of small models like shrinking, pruning, quantization or compressing, and instead used efficient network architecture. The network used depthwise separable convolution, which factorizes a standard convolution into a depthwise convolution and a 1×1 pointwise convolution. A standard convolution uses kernels on all input channels and combines them in one step while the depthwise convolution uses different kernels for each input channel and uses pointwise convolution to combine inputs. This separation of filtering and combining of features reduces the computation cost and model size. MobileNet consists of 28 separate convolutional layers, each followed by batch normalization and ReLU activation function. Howard et al. also introduced two model shrinking hyperparameters: width and resolution multiplier, in order to further improve speed and reduce size of the model. The width multiplier manipulates the width of the network uniformly by reducing the input and output channels while the resolution multiplier influences the size of the input image and its representations throughout the network. MobileNet achieves comparable accuracy to some full-fledged models while being a fraction of their size. Howard et al. also showed how it could generalize over various applications like face attribution, geolocalization and object detection. However, it was too simple and linear like VGG and therefore had fewer avenues for gradient flow. These were fixed in later iterations of this model [115, 116].

7.3. ShuffleNet

In 2017, Zhang et al. introduced ShuffleNet [117], an extremely computationally efficient neural network architecture, specifically designed for mobile devices. They recognized that many efficient networks become less effective as they scale down and purported it to be caused by expensive 1×1 convolutions. In conjunction with channel shuffle, they proposed the use of group convolution to circumvent its drawback of limited information flow. ShuffleNet consists mainly of a standard convolution followed by stacks of ShuffleNet units grouped in three stages. The ShuffleNet unit is similar to the ResNet block where they use depthwise convolution in the 3×3 layer and replace the 1×1 layer with pointwise group convolution. The depthwise convolution layer is preceded by a channel shuffle operation. The computation cost of the ShuffleNet can be administered by two hyperparameters: group number to control the connection sparsity and scaling factor to manipulate the model size. As group numbers become large, the error rate saturates as the input channels to each group decreases and therefore may reduce the representational capabilities. ShuffleNet outperformed contemporary models ([3, 42, 114, 34]) while having considerably smaller size. As the only advancement in ShuffleNet was channel shuffle, there isn't any improvement in inference speed of the model.

7.4. MobileNetv2

Improving on MobileNetv1 [34], Sandler et al. proposed MobileNetv2 [115] in 2018. It introduced the inverted residual with linear bottleneck, a novel module to reduce computations and improve accuracy. The module expands the low-dimensional representations of the input into higher dimensions, filters with a depthwise convolution and then projects it back to the lower dimensions, unlike the common residual block which performs compression, convolution and then expansion operations. The MobileNetv2 contains a convolution layer followed by 19 residual bottleneck modules and subsequently two convolutional layers. The residual bottleneck module has a shortcut connection only when the stride is 1. For higher stride, the shortcut is not used because of the difference in dimensions. They also employed the ReLU6 as the non-linearity function, instead of the simple ReLU, to limit computations. For object detection, the authors used MobileNetv2 as the feature extractor to a computationally efficient variant of the SSD [25]. This model, called SSDLite, claimed to have 8x fewer parameters than the original SSD while achieving competitive accuracy. It generalizes well over on other datasets, is easy to implement and hence, was well-received by the community.

7.5. PeleeNet

Existing lightweight deep learning models like [34, 117, 115] relied heavily on depthwise separable convolution, which lacked efficient implementation. Wang et al. proposed a novel efficient architecture based on conventional convolution, named PeleeNet [118], using an assortment of computation conserving techniques. PeleeNet was centered around the DenseNet [45] but looked at many other models for

inspiration. It introduced two-way dense layers, stem block, dynamic number of channels in a bottleneck, transition layer compression and conventional post activation to reduce computation cost and increase speed. Inspired from [42], the two-way dense layer helps in getting different scales of the receptive field, making it easier to identify larger objects. To reduce information loss, a stem block was used in the same way to [43, 119]. They also parted way with the compression factor used in [45] as it hurts the feature expression and reduces accuracy. PeleeNet consists of a stem block, four stages of modified dense and transition layers, and ultimately the classification layer. The authors also proposed a real-time object detection system, called Pelee, which was based on PeleeNet and a variant of SSD [25]. Its performance against the contemporary object detectors on mobile and edge devices was incremental but showed how simple design choices can make a huge difference in overall performance.

7.6. ShuffleNetv2

In 2018, Ningning Ma et al. present a set of comprehensive guidelines for designing efficient network architectures in ShuffleNetv2 [120]. They argued for the use of direct metrics like speed or latency to measure computational complexity, instead of indirect metrics like FLOPs. ShuffleNetv2 is built on four guiding principles – 1) equal width for input and output channels to minimize memory access cost, 2) carefully choosing group convolution based on the target platform and task, 3) multi-path structures achieve higher accuracy at the cost of efficiency and 4) element-wise operations like add and ReLU are computationally non-negligible. Following the above principles, they designed a new building block. It split the input into two parts by a channel split layer, followed by three convolutional layers which are then concatenated with the residual connection and passed through a channel shuffle layer. For the downsampling model, channel split is removed and residual connection has depthwise separable convolution layers. An ensemble of these blocks slotted in between a couple of convolutional layers results in ShuffleNetv2. Large models (50/162 layers) were also experimented to obtain superior accuracy with considerably fewer FLOPs. ShuffleNetv2 punched above its weight and outperformed other state-of-the-art models at comparable complexity.

7.7. MnasNet

With the increasing need for accurate, fast and low latency models for various edge devices, designing such a neural network is becoming more challenging than ever. In 2018, Tan et al. proposed Mnasnet [50] designed from an automated neural architecture search (NAS) approach. They formulate the search problem as multi-object optimization aimed at both high accuracy and low latency. It also factorized the search space by partitioning the CNN into unique blocks and subsequently searching for operations and connections in those blocks separately, thereby reducing the search space. This also allowed each block to have a distinctive design, unlike the earlier models [121, 122, 123]

which stacked the same blocks. Tan et al. used a RNN-based reinforcement learning agent as a controller along with a trainer to measure accuracy. Each sampled model is trained on a task to get its accuracy and run on the real devices for latency. This is used to achieve a soft reward target and the controller is updated. The process is repeated until the maximum iterations or a suitable candidate is derived. It is composed of 16 diverse blocks, some with residual connections. MnasNet was almost twice as fast as MobileNetv2 while having higher accuracy. However, like other reinforcement learning based neural architecture search models, the search time of MnasNet requires astronomical computational resources.

7.8. MobileNetv3

At the heart of MobileNetv3 [116] is the same method used to create MnasNet [50] with few modifications. A platform aware automated neural architecture search is performed in a factorized hierarchical search space and consequently optimized by NetAdapt [124], which removes the underutilized components of the network in multiple iterations. Once an architecture proposal is obtained, it trims the channels, randomly initialize the weights and then fine-tunes it to improve the target metrics. The model was further modified to remove some expensive layer in the architecture and gain additional latency improvement. Howard et al. argued that the filters in the architecture are often mirrored images of each other, and that accuracy can be maintained even after dropping half of these filters. Using this technique reduced the computations. MobileNetv3 used a blend of ReLU and hard swish as activation filters, the latter is mostly employed towards the end of the model. Hard swish has no noticeable difference from the swish function but is computationally cheaper while retaining the accuracy. For different resource use cases, [116] introduced two models – MobileNetv3-Large and MobileNetv3-Small. MobileNetv3-Large is composed of 15 bottleneck blocks while MobileNetv3-Small has 11. It also included squeeze and excitation layer [72] in its building blocks. Similar to [115], these model act as a feature detector in SSDLite and is 35% faster than earlier iterations [50, 115], whilst achieving higher *mAP*.

7.9. Once-For-All (OFA)

The use of neural architecture search (NAS) for architecture design has produced state-of-the-art models in the past few years, however, they are compute expensive because of the sampled model training. Cai et al. in [125] proposed a novel method of decoupling model training stage and the neural architecture search stage. The model is trained only once and sub-networks can be distilled from it as per the requirements. Once-for-all (OFA) network provides flexibility for such sub-networks in four important dimension of a convolutional neural network – depth, width, kernel size and dimension. As they are nested within the OFA network and interfere with the training, progressive shrinking was introduced. First, the largest network is trained with all parameters set to maximum. Subsequently, network is fine-tuned by gradually reducing the parameter dimensions like

Table 4

Performance comparison of various object detectors on MS COCO and PASCAL VOC 2012 datasets at similar input image size. Rows colored gray are real-time detectors (>30 FPS).

Model	Year	Backbone	Size	AP _[0.5:0.95]	AP _{0.5}	FPS
R-CNN*	2014	AlexNet	224	-	58.50%	~0.02
SPP-Net*	2015	ZF-5	Variable	-	59.20%	~0.23
Fast R-CNN*	2015	VGG-16	Variable	-	65.70%	~0.43
Faster R-CNN*	2016	VGG-16	600	-	67.00%	5
R-FCN	2016	ResNet-101	600	31.50%	53.20%	~3
FPN	2017	ResNet-101	800	36.20%	59.10%	5
Mask R-CNN	2018	ResNeXt-101-FPN	800	39.80%	62.30%	5
DetectoRS	2020	ResNeXt-101	1333	53.30%	71.60%	~4
YOLO*	2015	(Modified) GoogLeNet	448	-	57.90%	45
SSD	2016	VGG-16	300	23.20%	41.20%	46
YOLOv2	2016	DarkNet-19	352	21.60%	44.00%	81
RetinaNet	2018	ResNet-101-FPN	400	31.90%	49.50%	12
YOLOv3	2018	DarkNet-53	320	28.20%	51.50%	45
CenterNet	2019	Hourglass-104	512	42.10%	61.10%	7.8
EfficientDet-D2	2020	Efficient-B2	768	43.00%	62.30%	41.7
YOLOv4	2020	CSPDarkNet-53	512	43.00%	64.90%	31
DeTR	2020	ResNet-101	-	43.50%	63.80%	20
Swin-L	2021	HTC++	-	57.70%	-	-

^aModels marked with * are compared on PASCAL VOC 2012, while others on MS COCO.

Table 5

Comparison of Lightweight models.

Model	Year	Top-1 Acc%	Latency (ms)	Param. (mil.)	FLOPs (mil.)
SqueezeNet	2016	60.5	-	3.2	833
MobileNet	2017	70.6	113	4.2	569
ShuffleNet	2017	73.3	108	5.4	524
MobileNetv2	2018	74.7	143	6.9	300
PeleeNet	2018	72.6	-	2.8	508
ShuffleNetv2	2018	75.4	178	7.4	597
MnasNet	2018	76.7	103	5.2	403
MobileNetv3	2019	75.2	58	5.4	219
OFA	2020	80.0	58	7.7	595
MobileViT-S	2021	78.4	-	5.6	-

kernel size, depth and width. For elastic kernel, a center of the large kernel is used as the small kernel. As the center is shared, a kernel transformation matrix is used to maintain performance. To vary depth, the first few layers from the large network are used while the rest are skipped. Elastic width employs a channel sorting operation to reorganize channels and uses the most important ones in smaller models. OFA achieved state-of-the-art of 80% in ImageNet top-1 accuracy percentage and also won the 4th Low Power Computer Vision Challenge (LPCVC) while reducing GPU training time by many order of magnitude. It shows a new paradigm of designing lightweight models for a variety of hardware requirements.

7.10. MobileViT

Mehta and Rastegari in [126] introduced a lightweight transformer-based detector capable of running on edge devices. They combined the strengths of CNNs and ViT [100]

while keeping it mobile-friendly. While some literature exists on combining CNNs and transformers, this was the first work that achieved competitive performance to lightweight CNNs. It used a special MobileViT block to effectively find short and long-range dependencies. The MobileViT block process input in three stages – 1) Local Representations block to encode local features with a 3×3 convolution followed by a point-wise convolution, 2) Transformer as Convolution block which unfolds input into patches, which are then passed through L_x transformers and folded back, and 3) Fusion block which adds residual connection and convolutions to fuse local and global features. MobileViT also utilizes MobileNetv2 modules [115], which are spread serially along with MobileViT blocks. Unlike other transformer-based networks [100, 96], position encoding is not required for this network as they used transformer as a convolution, which implicitly includes spatial bias. MobileViT claims to be a general-purpose backbone for multiple vision tasks and performed well on detection and segmentation tasks. It performed comparatively well against full-fledged networks while on a basic data augmentation regime. Even though it achieves higher accuracy on a lower parameter budget, the latency is almost an order of magnitude lower than CNN models due to the limitations of transformers on mobile devices.

8. Comparative Results

We compare the performance of various object detectors on PASCAL VOC 2012 [36] and Microsoft COCO [22] datasets. Performance of object detectors is influenced by a number of factors like input image size and scale, feature extractor, GPU architecture, number of proposals, training methodology, loss function etc., which makes it difficult




Figure 10: Performance of Object Detectors on MS COCO dataset.

to compare various models without a common benchmark environment. Here in table 4, we evaluate performance of models based on the results from their papers. Models are compared on average precision (AP) and processed frames per second (FPS) at inference time. $AP_{0.5}$ is the average precision of all classes when predicted bounding box has an IoU > 0.5 with ground truth. COCO dataset introduced another performance metric $AP_{[0.5:0.95]}$, or simply AP, which is the average AP for IoU from 0.5 to 0.95 in step size of 0.5. We intentionally compare the performances of detectors on similarly size input image, where possible, to provide a reasonable account. This is done as the authors often introduce an array of models to provide flexibility between accuracy and inference time. In fig. 10, we use only the state-of-the-art model from the possible array of object detector family of models. Lightweight models are compared in table 5 where we compare them on ImageNet Top-1 classification accuracy, latency, number of parameters and complexity in MFLOPs. Models with MFLOPs lesser than 600 are expected to perform adequately on mobile devices.

9. Future Trends

Object detection has seen tremendous progress in the last decade. The algorithm have almost reached human level accuracy in some narrow domains, however it still has many exciting challenges to tackle. In this section, we discuss some of the open problem in the field of object detection.

AutoML: The use of automatic neural architecture search (NAS) for determining the characteristics of object detector is already an actively growing area. We have shown some detectors designed by NAS in earlier sections, however, it is still in its nascency. Searching for an algorithm is complex and resource intensive.

Transformers: While the transformers were introduced in computer vision fairly recently, they have achieved state-of-the-art performance on multiple benchmarks. Their application in combination with convolutional neural networks has shown promising results [127, 128, 129].

Lightweight detectors: While lightweight networks have shown great promise by matching classification errors with the full-fledged models, they still lack in detection accuracy by more than 50%. As more and more on-device machine learning applications are added to the market, need for small, efficient and equally accurate models will rise.

Weakly supervised/few shot detection: Most of the state-of-the-art object detection models are trained on millions of bounding box annotated data, which is unscalable as annotating data requires time and resources. Ability to train on weakly supervised data, i.e. image level labelled data, could result in considerable reduction in these costs.

Domain transfer: Domain transfer refers to use of a model trained on labeled image of a particular source task on a separate, but related target task. It encourages reuse of trained model and reduces reliance on the availability of a large dataset to achieve high accuracy.

3D object detection: 3D object detection is a particularly critical problem for autonomous driving. Even though models have achieved high accuracy, deployment of anything below human level performance will raise safety concerns.

Object detection in video: Object detectors are designed to perform on individual image which lack correlation between themselves. Using spatial and temporal relationship between the frames for object recognition is an open problem.

10. Conclusion

Even though object detection has come a long way in the past decade, the best detectors are still far from saturation in performance. As its applications increase in real world, the need for lightweight models that can be deployed on mobile and embedded systems is going to increase exponentially. There has been a rising interest in this domain, but it is still an open challenge. In this paper, we have shown how various object detectors developed over their predecessors. While the two stage detectors are generally more accurate, they are slow and cannot be used for real-time applications like self-driving cars or security. However, this has changed in the last few year where one stage detectors are equally accurate and much faster than the former. As evident in Figure 10, a transformer based detector i.e. Swin Transformer is the most accurate detector till date. With the current positive trend in the accuracy of detectors, we have high hopes for more accurate and faster detectors.

References

- [1] P. Viola, M. Jones, Rapid object detection using a boosted cascade of simple features, in: Proceedings of the 2001 IEEE Computer Society Conference on

- Computer Vision and Pattern Recognition. CVPR 2001, Vol. 1, IEEE Comput. Soc, 2001, pp. I-511–I-518. doi:10.1109/CVPR.2001.990517.
URL <http://ieeexplore.ieee.org/document/990517/>
- [2] N. Dalal, B. Triggs, Histograms of oriented gradients for human detection, in: 2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05), Vol. 1, 2005-06, pp. 886–893 vol. 1, ISSN: 1063-6919. doi:10.1109/CVPR.2005.177.
- [3] A. Krizhevsky, I. Sutskever, G. E. Hinton, Imagenet classification with deep convolutional neural networks, in: F. Pereira, C. J. C. Burges, L. Bottou, K. Q. Weinberger (Eds.), Advances in Neural Information Processing Systems, Curran Associates, Inc., p. 9.
- [4] K. Gauen, R. Dailey, J. Laiman, Y. Zi, N. Asokan, Y.-H. Lu, G. K. Thiruvathukal, M.-L. Shyu, S.-C. Chen, Comparison of visual datasets for machine learning, in: 2017 IEEE International Conference on Information Reuse and Integration (IRI), 2017, pp. 346–355. doi:10.1109/IRI.2017.59.
- [5] W. Zhiqiang, L. Jun, A review of object detection based on convolutional neural network, in: 2017 36th Chinese Control Conference (CCC), 2017, pp. 11104–11109. doi:10.23919/ChiCC.2017.8029130.
- [6] J. Huang, V. Rathod, C. Sun, M. Zhu, A. Korattikara, A. Fathi, I. Fischer, Z. Wojna, Y. Song, S. Guadarrama, K. Murphy, Speed/accuracy trade-offs for modern convolutional object detectors (2017). arXiv:1611.10012.
- [7] N. Yadav, U. Binay, Comparative study of object detection algorithms, International Research Journal of Engineering and Technology (IRJET) 4 (11) (2017) 586–591.
- [8] S. Agarwal, J. O. D. Terrail, F. Jurie, Recent advances in object detection in the age of deep convolutional neural networks, arXiv preprint arXiv:1809.03193 (2018).
- [9] A. Gupta, R. Puri, M. Verma, S. Gunjyal, A. Kumar, Performance comparison of object detection algorithms with different feature extractors, in: 2019 6th International Conference on Signal Processing and Integrated Networks (SPIN), IEEE, 2019, pp. 472–477.
- [10] Z.-Q. Zhao, P. Zheng, S.-t. Xu, X. Wu, Object detection with deep learning: A review, IEEE transactions on neural networks and learning systems (2019).
- [11] A. Borji, M.-M. Cheng, Q. Hou, H. Jiang, J. Li, Salient object detection: A survey, Computational visual media 5 (2) (2019) 117–150.
- [12] Z. Zou, Z. Shi, Y. Guo, J. Ye, Object detection in 20 years: A survey, CoRR abs/1905.05055 (2019). arXiv:1905.05055.
URL <http://arxiv.org/abs/1905.05055>
- [13] L. Liu, W. Ouyang, X. Wang, P. Fieguth, J. Chen, X. Liu, M. Pietikäinen, Deep learning for generic object detection: A survey, International journal of computer vision 128 (2) (2020) 261–318.
- [14] G. Huang, I. H. Laradji, D. Vázquez, S. Lacoste-Julien, P. Rodríguez, A survey of self-supervised and few-shot object detection, CoRR abs/2110.14711 (2021). arXiv:2110.14711.
URL <https://arxiv.org/abs/2110.14711>
- [15] W. Wang, Q. Lai, H. Fu, J. Shen, H. Ling, R. Yang, Salient object detection in the deep learning era: An in-depth survey, IEEE Transactions on Pattern Analysis and Machine Intelligence (2021) 1–1doi:10.1109/TPAMI.2021.3051099.
- [16] L. Liu, W. Ouyang, X. Wang, P. Fieguth, J. Chen, X. Liu, M. Pietikäinen, Deep learning for generic object detection: A surveyVersion: 1 (2018). arXiv:1809.02165.
URL <http://arxiv.org/abs/1809.02165>
- [17] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, A. C. Berg, L. Fei-Fei, ImageNet large scale visual recognition challenge 115 (3) 211–252. doi:10.1007/s11263-015-0816-y.
URL <https://doi.org/10.1007/s11263-015-0816-y>
- [18] M. Everingham, L. Van Gool, C. K. I. Williams, J. Winn, A. Zisserman, The pascal visual object classes (VOC) challenge 88 (2) 303–338. doi:10.1007/s11263-009-0275-4.
URL <http://link.springer.com/10.1007/s11263-009-0275-4>
- [19] J. Xiao, J. Hays, K. A. Ehinger, A. Oliva, A. Torralba, Sun database: Large-scale scene recognition from abbey to zoo, in: 2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition, IEEE, 2010, pp. 3485–3492.
- [20] A. Geiger, P. Lenz, C. Stiller, R. Urtasun, Vision meets robotics: The kitti dataset, International Journal of Robotics Research (IJRR) (2013).
- [21] P. Dollár, C. Wojek, B. Schiele, P. Perona, Pedestrian detection: A benchmark (2009).
- [22] T.-Y. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, C. L. Zitnick, Microsoft coco: Common objects in context, in: D. Fleet, T. Pajdla, B. Schiele, T. Tuytelaars (Eds.), Computer Vision – ECCV 2014, Springer International Publishing, Cham, 2014, pp. 740–755.

- [23] S. Ren, K. He, R. Girshick, J. Sun, Faster r-cnn: Towards real-time object detection with region proposal networks (2016). [arXiv:1506.01497](https://arxiv.org/abs/1506.01497).
- [24] J. Dai, Y. Li, K. He, J. Sun, R-fcn: Object detection via region-based fully convolutional networks (2016). [arXiv:1605.06409](https://arxiv.org/abs/1605.06409).
- [25] W. Liu, D. Anguelov, D. Erhan, C. Szegedy, S. Reed, C.-Y. Fu, A. C. Berg, SSD: Single shot MultiBox detector, in: B. Leibe, J. Matas, N. Sebe, M. Welling (Eds.), *Computer Vision – ECCV 2016*, Springer International Publishing, pp. 21–37.
- [26] R. Girshick, J. Donahue, T. Darrell, J. Malik, Rich feature hierarchies for accurate object detection and semantic segmentation, in: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2014.
- [27] K. He, X. Zhang, S. Ren, J. Sun, Spatial pyramid pooling in deep convolutional networks for visual recognition 37 (9) 1904–1916, conference Name: *IEEE Transactions on Pattern Analysis and Machine Intelligence*. doi:10.1109/TPAMI.2015.2389824.
- [28] R. Girshick, Fast r-cnn, in: *2015 IEEE International Conference on Computer Vision (ICCV)*, 2015, pp. 1440–1448. doi:10.1109/ICCV.2015.169.
- [29] T. Lin, P. Goyal, R. Girshick, K. He, P. Dollár, Focal loss for dense object detection, *IEEE Transactions on Pattern Analysis and Machine Intelligence* 42 (2) (2020) 318–327. doi:10.1109/TPAMI.2018.2858826.
- [30] K. He, G. Gkioxari, P. Dollár, R. Girshick, Mask r-cnn (2018). [arXiv:1703.06870](https://arxiv.org/abs/1703.06870).
- [31] J. Dai, H. Qi, Y. Xiong, Y. Li, G. Zhang, H. Hu, Y. Wei, Deformable convolutional networks (2017). [arXiv:1703.06211](https://arxiv.org/abs/1703.06211).
- [32] C. Szegedy, V. Vanhoucke, S. Ioffe, J. Shlens, Z. Wojna, Rethinking the inception architecture for computer vision, in: *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, IEEE, pp. 2818–2826. doi:10.1109/CVPR.2016.308.
URL <http://ieeexplore.ieee.org/document/7780677/>
- [33] K. He, X. Zhang, S. Ren, J. Sun, Deep residual learning for image recognition, in: *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016, pp. 770–778. doi:10.1109/CVPR.2016.90.
- [34] A. G. Howard, M. Zhu, B. Chen, D. Kalenichenko, W. Wang, T. Weyand, M. Andreetto, H. Adam, MobileNets: Efficient convolutional neural networks for mobile vision applications [arXiv:1704.04861](https://arxiv.org/abs/1704.04861).
URL <http://arxiv.org/abs/1704.04861>
- [35] M. Everingham, L. Van Gool, C. K. I. Williams, J. Winn, A. Zisserman, The PASCAL Visual Object Classes Challenge 2007 (VOC2007) Results, <http://www.pascal-network.org/challenges/VOC/voc2007/workshop/index.html>.
- [36] M. Everingham, J. Winn, The PASCAL visual object classes challenge 2012 (VOC2012) development kit 32.
- [37] J. Deng, W. Dong, R. Socher, L. Li, Kai Li, Li Fei-Fei, ImageNet: A large-scale hierarchical image database, in: *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pp. 248–255, ISSN: 1063-6919. doi:10.1109/CVPR.2009.5206848.
- [38] A. Aslam, E. Curry, A survey on object detection for the internet of multimedia things (iomt) using deep learning and event-based middleware: Approaches, challenges, and future directions, *Image and Vision Computing* 106 (2021) 104095.
- [39] A. Kuznetsova, H. Rom, N. Alldrin, J. Uijlings, I. Krasin, J. Pont-Tuset, S. Kamali, S. Popov, M. Mallocci, A. Kolesnikov, T. Duerig, V. Ferrari, The open images dataset v4 128 (7) 1956–1981. doi:10.1007/s11263-020-01316-z.
URL <https://doi.org/10.1007/s11263-020-01316-z>
- [40] M. D. Zeiler, R. Fergus, Visualizing and understanding convolutional networks, in: D. Fleet, T. Pajdla, B. Schiele, T. Tuytelaars (Eds.), *Computer Vision – ECCV 2014*, Lecture Notes in Computer Science, Springer International Publishing, pp. 818–833. doi:10.1007/978-3-319-10590-1_53.
- [41] K. Simonyan, A. Zisserman, Very deep convolutional networks for large-scale image recognition [arXiv:1409.1556](https://arxiv.org/abs/1409.1556).
URL <http://arxiv.org/abs/1409.1556>
- [42] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, A. Rabinovich, Going deeper with convolutions [arXiv:1409.4842](https://arxiv.org/abs/1409.4842).
URL <http://arxiv.org/abs/1409.4842>
- [43] C. Szegedy, S. Ioffe, V. Vanhoucke, A. Alemi, Inception-v4, inception-ResNet and the impact of residual connections on learning [arXiv:1602.07261](https://arxiv.org/abs/1602.07261).
URL <http://arxiv.org/abs/1602.07261>
- [44] K. He, X. Zhang, S. Ren, J. Sun, Identity mappings in deep residual networks [arXiv:1603.05027](https://arxiv.org/abs/1603.05027).
URL <http://arxiv.org/abs/1603.05027>
- [45] G. Huang, Z. Liu, L. van der Maaten, K. Q. Weinberger, Densely connected convolutional networks [arXiv:1608.06993](https://arxiv.org/abs/1608.06993).
URL <http://arxiv.org/abs/1608.06993>

- [46] S. Xie, R. Girshick, P. Dollár, Z. Tu, K. He, Aggregated residual transformations for deep neural networks *arXiv:1611.05431*.
URL <http://arxiv.org/abs/1611.05431>
- [47] C.-Y. Wang, H.-Y. M. Liao, I.-H. Yeh, Y.-H. Wu, P.-Y. Chen, J.-W. Hsieh, CSPNet: A new backbone that can enhance learning capability of CNN *Version: 1*. *arXiv:1911.11929*.
URL <http://arxiv.org/abs/1911.11929>
- [48] C.-Y. Wang, A. Bochkovskiy, H.-Y. M. Liao, Scaled-YOLOv4: Scaling cross stage partial network *arXiv:2011.08036*.
URL <http://arxiv.org/abs/2011.08036>
- [49] M. Tan, Q. V. Le, EfficientNet: Rethinking model scaling for convolutional neural networks *arXiv:1905.11946*.
URL <http://arxiv.org/abs/1905.11946>
- [50] M. Tan, B. Chen, R. Pang, V. Vasudevan, M. Sandler, A. Howard, Q. V. Le, MnasNet: Platform-aware neural architecture search for mobile.
URL <https://arxiv.org/abs/1807.11626v3>
- [51] D. G. Lowe, Distinctive image features from scale-invariant keypoints 60 (2) 91–110. doi:10.1023/B:VISI.0000029664.99615.94.
URL <https://doi.org/10.1023/B:VISI.0000029664.99615.94>
- [52] D. G. Lowe, Object recognition from local scale-invariant features, in: *Proceedings of the Seventh IEEE International Conference on Computer Vision*, Vol. 2, pp. 1150–1157 vol.2. doi:10.1109/ICCV.1999.790410.
- [53] A. Mohan, C. Papageorgiou, T. Poggio, Example-based object detection in images by components 23 (4) 349–361. doi:10.1109/34.917571.
URL <http://ieeexplore.ieee.org/document/917571/>
- [54] Yan Ke, R. Sukthankar, PCA-SIFT: a more distinctive representation for local image descriptors, in: *Proceedings of the 2004 IEEE Computer Society Conference on Computer Vision and Pattern Recognition, 2004. CVPR 2004.*, Vol. 2, pp. II–II, ISSN: 1063-6919. doi:10.1109/CVPR.2004.1315206.
- [55] P. Felzenszwalb, D. McAllester, D. Ramanan, A discriminatively trained, multiscale, deformable part model, in: *2008 IEEE Conference on Computer Vision and Pattern Recognition*, pp. 1–8, ISSN: 1063-6919. doi:10.1109/CVPR.2008.4587597.
- [56] P. F. Felzenszwalb, R. B. Girshick, D. McAllester, D. Ramanan, Object detection with discriminatively trained part-based models 32 (9) 1627–1645. doi:10.1109/TPAMI.2009.167.
- [57] P. F. Felzenszwalb, R. B. Girshick, D. McAllester, Cascade object detection with deformable part models, in: *2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition, IEEE*, pp. 2241–2248. doi:10.1109/CVPR.2010.5539906.
URL <http://ieeexplore.ieee.org/document/5539906/>
- [58] J. R. R. Uijlings, T. Gevers, A. W. M. Smeulders, Selective search for object recognition 18.
- [59] Y. LeCun, B. Boser, J. S. Denker, D. Henderson, R. E. Howard, W. Hubbard, L. D. Jackel, Backpropagation applied to handwritten zip code recognition 1 (4) 541–551, publisher: MIT Press. doi:10.1162/neco.1989.1.4.541.
URL <https://doi.org/10.1162/neco.1989.1.4.541>
- [60] K. Grauman, T. Darrell, The pyramid match kernel: discriminative classification with sets of image features, in: *Tenth IEEE International Conference on Computer Vision (ICCV'05) Volume 1, Vol. 2*, pp. 1458–1465 Vol. 2, ISSN: 2380-7504. doi:10.1109/ICCV.2005.239.
- [61] Y. Jia, E. Shelhamer, J. Donahue, S. Karayev, J. Long, R. Girshick, S. Guadarrama, T. Darrell, Caffe: Convolutional architecture for fast feature embedding, in: *Proceedings of the 22nd ACM international conference on Multimedia, MM '14*, Association for Computing Machinery, pp. 675–678. doi:10.1145/2647868.2654889.
URL <https://doi.org/10.1145/2647868.2654889>
- [62] J. Long, E. Shelhamer, T. Darrell, Fully convolutional networks for semantic segmentation 10.
- [63] T.-Y. Lin, P. Dollár, R. Girshick, K. He, B. Hariharan, S. Belongie, Feature pyramid networks for object detection (2017). *arXiv:1612.03144*.
- [64] S. Liu, L. Qi, H. Qin, J. Shi, J. Jia, Path aggregation network for instance segmentation *arXiv:1803.01534*.
URL <http://arxiv.org/abs/1803.01534>
- [65] G. Ghiasi, T.-Y. Lin, Q. V. Le, NAS-FPN: Learning scalable feature pyramid architecture for object detection, in: *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), IEEE*, pp. 7029–7038. doi:10.1109/CVPR.2019.00720.
URL <https://ieeexplore.ieee.org/document/8954436/>
- [66] A. Shrivastava, A. Gupta, R. Girshick, Training region-based object detectors with online hard example mining (2016). *arXiv:1604.03540*.
- [67] K. Chen, J. Pang, J. Wang, Y. Xiong, X. Li, S. Sun, W. Feng, Z. Liu, J. Shi, W. Ouyang, C. C. Loy, D. Lin, Hybrid task cascade for instance segmentation.
URL <https://arxiv.org/abs/1901.07518v2>

- [68] Z. Cai, N. Vasconcelos, Cascade r-CNN: Delving into high quality object detection. URL <https://arxiv.org/abs/1712.00726v1>
- [69] S. Qiao, L.-C. Chen, A. Yuille, DetectoRS: Detecting objects with recursive feature pyramid and switchable atrous convolutionarXiv:2006.02334. URL <http://arxiv.org/abs/2006.02334>
- [70] L.-C. Chen, G. Papandreou, I. Kokkinos, K. Murphy, A. L. Yuille, DeepLab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected CRFsarXiv:1606.00915. URL <http://arxiv.org/abs/1606.00915>
- [71] M. Holschneider, R. Kronland-Martinet, J. Morlet, P. Tchamitchian, A real-time algorithm for signal analysis with the help of the wavelet transform, in: J.-M. Combes, A. Grossmann, P. Tchamitchian (Eds.), *Wavelets, inverse problems and theoretical imaging*, Springer, pp. 286–297. doi:10.1007/978-3-642-75988-8_28.
- [72] J. Hu, L. Shen, S. Albanie, G. Sun, E. Wu, Squeeze-and-excitation networksarXiv:1709.01507. URL <http://arxiv.org/abs/1709.01507>
- [73] J. Redmon, S. Divvala, R. Girshick, A. Farhadi, You only look once: Unified, real-time object detection, in: 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2016, pp. 779–788. doi:10.1109/CVPR.2016.91.
- [74] M. Lin, Q. Chen, S. Yan, Network in networkarXiv:1312.4400. URL <http://arxiv.org/abs/1312.4400>
- [75] J. Redmon, A. Farhadi, YOLO9000: Better, faster, stronger. URL <https://arxiv.org/abs/1612.08242v1>
- [76] J. Redmon, A. Farhadi, Yolov3: An incremental improvement (2018). arXiv:1804.02767.
- [77] A. Bochkovskiy, C.-Y. Wang, H.-Y. M. Liao, YOLOv4: Optimal speed and accuracy of object detectionarXiv:2004.10934. URL <http://arxiv.org/abs/2004.10934>
- [78] D. Erhan, C. Szegedy, A. Toshev, D. Anguelov, Scalable object detection using deep neural networks (2013). arXiv:1312.2249.
- [79] J. Redmon, Darknet: Open source neural networks in c, <http://pjreddie.com/darknet/> (2013–2016).
- [80] K. He, X. Zhang, S. Ren, J. Sun, Delving deep into rectifiers: Surpassing human-level performance on ImageNet classification, in: 2015 IEEE International Conference on Computer Vision (ICCV), IEEE, pp. 1026–1034. doi:10.1109/ICCV.2015.123. URL <http://ieeexplore.ieee.org/document/7410480/>
- [81] G. Miller, R. Beckwith, C. Fellbaum, D. Gross, K. Miller, Introduction to WordNet: An on-line lexical database* 3. doi:10.1093/ijl/3.4.235.
- [82] X. Zhou, D. Wang, P. Kr̃adhenb̃ijhl, Objects as pointsarXiv:1904.07850. URL <http://arxiv.org/abs/1904.07850>
- [83] A. Newell, K. Yang, J. Deng, Stacked hourglass networks for human pose estimation, in: B. Leibe, J. Matas, N. Sebe, M. Welling (Eds.), *Computer Vision – ECCV 2016, Lecture Notes in Computer Science*, Springer International Publishing, pp. 483–499. doi:10.1007/978-3-319-46484-8_29.
- [84] M. Tan, R. Pang, Q. V. Le, EfficientDet: Scalable and efficient object detection, in: 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), IEEE, pp. 10778–10787. doi:10.1109/CVPR42600.2020.01079. URL <https://ieeexplore.ieee.org/document/9156454/>
- [85] P. Ramachandran, B. Zoph, Q. V. Le, Searching for activation functionsarXiv:1710.05941. URL <http://arxiv.org/abs/1710.05941>
- [86] Z. Zheng, P. Wang, W. Liu, J. Li, R. Ye, D. Ren, Distance-IoU loss: Faster and better learning for bounding box regressionarXiv:1911.08287. URL <http://arxiv.org/abs/1911.08287>
- [87] I. Loshchilov, F. Hutter, SGDR: Stochastic gradient descent with warm restartsarXiv:1608.03983. URL <http://arxiv.org/abs/1608.03983>
- [88] D. Misra, Mish: A self regularized non-monotonic activation functionarXiv:1908.08681. URL <http://arxiv.org/abs/1908.08681>
- [89] G. Jocher, et al., ultralytics/yolov5: v6.0 - YOLOv5n 'Nano' models, Roboflow integration, TensorFlow export, OpenCV DNN support (Oct. 2021). doi:10.5281/zenodo.5563715. URL <https://doi.org/10.5281/zenodo.5563715>
- [90] D. Thuan, Evolution of yolo algorithm and yolov5: the state-of-the-art object detection algorithm (2021).
- [91] Roboflow, Yolov5 is here: State-of-the-art object detection at 140 fps, <https://blog.roboflow.com/yolov5-is-here/>, accessed: 2022-01-22 (June 2020).
- [92] H. Wang, S. Zhang, S. Zhao, Q. Wang, D. Li, R. Zhao, Real-time detection and tracking of fish abnormal behavior based on improved yolov5 and siamrpn++, *Computers and Electronics in Agriculture* 192 (2022) 106512.
- [93] Y. Jing, Y. Ren, Y. Liu, D. Wang, L. Yu, Automatic extraction of damaged houses by earthquake based on improved yolov5: A case study in yangbi, *Remote Sensing* 14 (2) (2022) 382.

- [94] Roboflow, Yolov5 v6.0 is here - new nano model at 1666 fps, <https://blog.roboflow.com/yolov5-v6-0-is-here/>, accessed: 2022-01-22 (October 2021).
- [95] X. Zhu, S. Lyu, X. Wang, Q. Zhao, Tph-yolov5: Improved yolov5 based on transformer prediction head for object detection on drone-captured scenarios (2021). [arXiv:2108.11539](https://arxiv.org/abs/2108.11539).
- [96] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, I. Polosukhin, Attention is all you need (2017). [arXiv:1706.03762](https://arxiv.org/abs/1706.03762).
- [97] J. Devlin, M.-W. Chang, K. Lee, K. Toutanova, Bert: Pre-training of deep bidirectional transformers for language understanding (2019). [arXiv:1810.04805](https://arxiv.org/abs/1810.04805).
- [98] A. Radford, K. Narasimhan, T. Salimans, I. Sutskever, Improving language understanding by generative pre-training (2018).
- [99] C. Raffel, N. Shazeer, A. Roberts, K. Lee, S. Narang, M. Matena, Y. Zhou, W. Li, P. J. Liu, Exploring the limits of transfer learning with a unified text-to-text transformer, *Journal of Machine Learning Research* 21 (140) (2020) 1–67.
URL <http://jmlr.org/papers/v21/20-074.html>
- [100] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly, J. Uszkoreit, N. Houlsby, An image is worth 16x16 words: Transformers for image recognition at scale (2021). [arXiv:2010.11929](https://arxiv.org/abs/2010.11929).
- [101] S. Khan, M. Naseer, M. Hayat, S. W. Zamir, F. S. Khan, M. Shah, Transformers in vision: A survey, *arXiv preprint arXiv:2101.01169* (2021).
- [102] N. Carion, F. Massa, G. Synnaeve, N. Usunier, A. Kirillov, S. Zagoruyko, End-to-end object detection with transformers (2020). [arXiv:2005.12872](https://arxiv.org/abs/2005.12872).
- [103] Z. Liu, Y. Lin, Y. Cao, H. Hu, Y. Wei, Z. Zhang, S. Lin, B. Guo, Swin transformer: Hierarchical vision transformer using shifted windows (2021). [arXiv:2103.14030](https://arxiv.org/abs/2103.14030).
- [104] M. N. Abbas, M. S. Ansari, M. N. Asghar, N. Kanwal, T. O'Neill, B. Lee, Lightweight deep learning model for detection of copy-move image forgery with post-processed attacks, in: 2021 IEEE 19th World Symposium on Applied Machine Intelligence and Informatics (SAMI), IEEE, 2021, pp. 000125–000130.
- [105] S. Karakanis, G. Leontidis, Lightweight deep learning models for detecting covid-19 from chest x-ray images, *Computers in Biology and Medicine* 130 (2021) 104181.
- [106] A. Jadon, A. Varshney, M. S. Ansari, Low-complexity high-performance deep learning model for real-time low-cost embedded fire detection systems, *Procedia Computer Science* 171 (2020) 418–426.
- [107] A. Jadon, M. Omama, A. Varshney, M. S. Ansari, R. Sharma, Firenet: a specialized lightweight fire & smoke detection model for real-time iot applications, *arXiv preprint arXiv:1905.11922* (2019).
- [108] Y. L. Cun, J. S. Denker, S. A. Solla, *Optimal Brain Damage*, Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 1990, p. 598–605.
- [109] B. Hassibi, D. G. Stork, G. J. Wolff, Optimal brain surgeon and general network pruning, in: *IEEE International Conference on Neural Networks*, 1993, pp. 293–299 vol.1. doi:10.1109/ICNN.1993.298572.
- [110] S. Han, H. Mao, W. J. Dally, Deep compression: Compressing deep neural networks with pruning, trained quantization and huffman coding [arXiv:1510.00149](https://arxiv.org/abs/1510.00149).
URL <http://arxiv.org/abs/1510.00149>
- [111] M. Courbariaux, Y. Bengio, J.-P. David, BinaryConnect: Training deep neural networks with binary weights during propagations [arXiv:1511.00363](https://arxiv.org/abs/1511.00363).
URL <http://arxiv.org/abs/1511.00363>
- [112] W. Chen, J. T. Wilson, S. Tyree, K. Q. Weinberger, Y. Chen, Compressing neural networks with the hashing trick [arXiv:1504.04788](https://arxiv.org/abs/1504.04788).
URL <http://arxiv.org/abs/1504.04788>
- [113] G. Hinton, O. Vinyals, J. Dean, Distilling the knowledge in a neural network [arXiv:1503.02531](https://arxiv.org/abs/1503.02531).
URL <http://arxiv.org/abs/1503.02531>
- [114] F. N. Iandola, S. Han, M. W. Moskewicz, K. Ashraf, W. J. Dally, K. Keutzer, SqueezeNet: AlexNet-level accuracy with 50x fewer parameters and <0.5mb model size [arXiv:1602.07360](https://arxiv.org/abs/1602.07360).
URL <http://arxiv.org/abs/1602.07360>
- [115] M. Sandler, A. Howard, M. Zhu, A. Zhmoginov, L.-C. Chen, MobileNetV2: Inverted residuals and linear bottlenecks [arXiv:1801.04381](https://arxiv.org/abs/1801.04381).
URL <http://arxiv.org/abs/1801.04381>
- [116] A. Howard, M. Sandler, G. Chu, L.-C. Chen, B. Chen, M. Tan, W. Wang, Y. Zhu, R. Pang, V. Vasudevan, Q. V. Le, H. Adam, Searching for MobileNetV3. [arXiv:1905.02244v5](https://arxiv.org/abs/1905.02244v5)
URL <https://arxiv.org/abs/1905.02244v5>
- [117] X. Zhang, X. Zhou, M. Lin, J. Sun, ShuffleNet: An extremely efficient convolutional neural network for mobile devices, in: *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, IEEE, pp. 6848–6856. doi:10.1109/CVPR.2018.00716.
URL <https://ieeexplore.ieee.org/document/8578814/>

- [118] R. J. Wang, X. Li, C. X. Ling, Pelee: A real-time object detection system on mobile devices 10.
- [119] Z. Shen, Z. Liu, J. Li, Y.-G. Jiang, Y. Chen, X. Xue, DSOD: Learning deeply supervised object detectors from scratch *arXiv:1708.01241*.
URL <http://arxiv.org/abs/1708.01241>
- [120] N. Ma, X. Zhang, H.-T. Zheng, J. Sun, ShuffleNet v2: Practical guidelines for efficient CNN architecture design *arXiv:1807.11164*.
URL <http://arxiv.org/abs/1807.11164>
- [121] B. Zoph, Q. V. Le, Neural architecture search with reinforcement learning.
URL <https://arxiv.org/abs/1611.01578v2>
- [122] C. Liu, B. Zoph, M. Neumann, J. Shlens, W. Hua, L.-J. Li, L. Fei-Fei, A. Yuille, J. Huang, K. Murphy, Progressive neural architecture search.
URL <https://arxiv.org/abs/1712.00559v3>
- [123] E. Real, A. Aggarwal, Y. Huang, Q. V. Le, Regularized evolution for image classifier architecture search 33 (1) 4780–4789, number: 01. doi:10.1609/aaai.v33i01.33014780.
URL <https://ojs.aaai.org/index.php/AAAI/article/view/4405>
- [124] T.-J. Yang, A. Howard, B. Chen, X. Zhang, A. Go, M. Sandler, V. Sze, H. Adam, NetAdapt: Platform-aware neural network adaptation for mobile applications *arXiv:1804.03230*.
URL <http://arxiv.org/abs/1804.03230>
- [125] H. Cai, C. Gan, T. Wang, Z. Zhang, S. Han, Once-for-all: Train one network and specialize it for efficient deployment (2020). *arXiv:1908.09791*.
- [126] S. Mehta, M. Rastegari, Mobilevit: Light-weight, general-purpose, and mobile-friendly vision transformer (2021). *arXiv:2110.02178*.
- [127] T. Xiao, M. Singh, E. Mintun, T. Darrell, P. Dollár, R. Girshick, Early convolutions help transformers see better (2021). *arXiv:2106.14881*.
- [128] H. Wu, B. Xiao, N. Codella, M. Liu, X. Dai, L. Yuan, L. Zhang, Cvt: Introducing convolutions to vision transformers (2021). *arXiv:2103.15808*.
- [129] S. d'Ascoli, H. Touvron, M. Leavitt, A. Morcos, G. Biroli, L. Sagun, Convit: Improving vision transformers with soft convolutional inductive biases (2021). *arXiv:2103.10697*.



Syed Sahil Abbas Zaidi received his Bachelors in Technology degree in Electronics Engineering from Aligarh Muslim University, India in 2016. He then worked in the industry for four years before returning to academia. From 2020, he joined the Ph.D. program at Technological University of the Shannon: Midlands Midwest. His research interests include Computer Vision and Machine Learning on Edge Devices.



Mohammad Samar Ansari received B.Tech., M.Tech. and Ph.D. degrees in Electronics Engineering from Aligarh Muslim University (AMU), Aligarh, India, in 2001, 2007 and 2012, respectively. He is presently working as a Senior Lecturer in the University of Chester, UK. Prior to this, he was an Associate Professor in AMU. He has also worked as a Post-Doctoral Research Fellow in the Software Research Institute, Athlone Institute of Technology, Ireland. He has also been associated with Siemens, Defence Research and Development Organization (DRDO) and Malaviya National Institute of Technology, Jaipur. His research interests include neural networks, machine learning, and security & privacy. He has published around 110 research papers in reputed international journals & conferences and authored 2 books and contributed 6 book chapters. He is also the recipient of the prestigious Young Faculty Research Fellowship by Department of Electronics and IT, Ministry of Communications and Information Technology, Govt. of India. He is a Senior Member of IEEE.



Asra Aslam is presently working as a Machine Learning Engineer at mindtrace.ai. She completed her Ph.D. in 2021 from the INSIGHT Centre for Data Analytics, National University of Ireland (NUI), Galway. Before joining her Ph.D. she worked as a Guest Teacher in the Department of Computer Engineering, Aligarh Muslim University, India from August 2015 to March 2016. She received B.Tech. and M.Tech. degrees in Computer Engineering from Aligarh Muslim University, India, in 2013, and 2015, respectively. Her research interests include Computer Vision, Deep Neural Networks, Smart Cities, Machine Learning, and Artificial Intelligence.



Nadia Kanwal received her master's and Ph.D. degrees in Computer Sciences from the University of Essex, UK in 2009 and 2013 respectively. She has also received a prestigious Marie Skłodowska-Curie Research fellowship (for 3-years) in 2019 to do an industry led project related to secure and privacy protected CCTV video storage and retrieval as principal investigator. Her research interests are in the field of Computer Vision and Machine Learning. She is actively working on applying machine learning to develop cutting-edge solutions for health, security, and vision applications. Her publications on multimedia data security, visual privacy, low-level image features, virtual reality, EEG/ECG signal analysis, and pupillometry have been very well supported by the research community. She is an active reviewer of good-standing journals and conferences.



Mamoonah N. Asghar received a Ph.D. degree from the School of Computer Science and Electronic Engineering, University of Essex, Colchester, U.K., in 2013. In 2017, she was awarded an industry led Marie Skłodowska-Curie (MSC) Career-Fit Research Fellowship (3-years at AIT-Ireland) funded by Marie-Curie Actions and Enterprise Ireland. She has published several ISI indexed journal articles along with numerous international conference papers. She is actively involved in reviewing research articles for renowned journals/conferences and has participated as a session chair. Her research interests include security aspects of multimedia (image, audio, and video), compression, visual privacy, encryption, steganography, secure transmission in future networks, video quality metrics, key management schemes, computer vision algorithms, deep learning models, and General Data Protection Regulation (EU-GDPR).



Brian Lee received the Ph.D. degree from the Trinity College Dublin, Dublin, Ireland, in the application of programmable networking for network management. He has over 25 years R&D experience in telecommunications network monitoring, their systems and software design and development for large telecommunications products with very high impact research publications. He was the Director of research for LM Ericsson, Ireland, with responsibility for overseeing all research activities, including external collaborations and relationship management. He was an Engineering Manager with Duolog Ltd., where he was responsible for strategic and operational management of all research and development activities. He is currently the Director of the Software Research Institute, Athlone Institute of Technology, Athlone, Ireland.