

The MIS Check-Dam Dataset for Object Detection and Instance Segmentation Tasks

Chintan Tundia^{a,*}, Rajiv Kumar^{b,*}, Om Damani^c and G. Sivakumar^d

Indian Institute of Technology Bombay, Mumbai, India

Keywords: Object Detection, Instance Segmentation, Remote Sensing, Image Transformers.

Abstract: Deep learning has led to many recent advances in object detection and instance segmentation, among other computer vision tasks. These advancements have led to wide application of deep learning based methods and related methodologies in object detection tasks for satellite imagery. In this paper, we introduce MIS Check-Dam, a new dataset of check-dams from satellite imagery for building an automated system for the detection and mapping of check-dams, focusing on the importance of irrigation structures used for agriculture. We review some of the most recent object detection and instance segmentation methods and assess their performance on our new dataset. We evaluate several single stage, two-stage and attention based methods under various network configurations and backbone architectures. The dataset and the pre-trained models are available at <https://www.cse.iitb.ac.in/gramdrishti/>.

1 INTRODUCTION

Machine learning and deep learning have evolved exponentially over the past decade contributing a lot to the fields of image and video processing (Jiao and Zhao, 2019), computer vision (Kumar et al., 2021), natural language processing (Fedus et al., 2021), reinforcement learning (Henderson et al., 2018) and more. Understanding a scene involves detecting and identifying objects (object detection) or keypoints and different regions in images (instance and semantic segmentation). With the improvements in imaging technologies of the satellite sensors, large volumes of hyper-spectral and spatial resolution data are available, resulting in the application of object detection in remote sensing for identifying objects from satellite, aerial and SAR imagery. The availability of very large remote sensing datasets (Li et al., 2020a) have led to applications like land monitoring, target identification, change detection, building detection, road detection, vehicle detection, and so on.

With the increased spatial resolution, more objects are visible in satellite images, while been subjected to viewpoint variation, occlusion, background clutter,

illumination, shadow, etc. This leads to various challenges in addition to the common object detection challenges like small spatial extent of foreground target objects, large scale search space, variety of perspectives and viewpoints, etc. The objects in remote sensing images are also subject to rotation, intra-class variations, similarity to surroundings, etc. making object detection on remote sensing images a challenging task.



Figure 1: Photograph (left) and satellite image (right) of a wall based check-dam.



Figure 2: Photograph (left) and satellite image (right) of a gate based check-dam.

In general, images from satellites have small, densely clustered objects of interest with a small spa-

^a <https://orcid.org/0000-0003-3169-1775>

^b <https://orcid.org/0000-0003-4174-8587>

^c <https://orcid.org/0000-0002-4043-9806>

^d <https://orcid.org/0000-0003-2890-6421>

* Authors *a, b* made equal contribution

tial extent, unlike that of COCO (Lin et al., 2014) or Imagenet (Russakovsky et al., 2015) with large and less clustered objects. For e.g. In satellite images, objects like cars will be a few tens of pixels in extent even at the highest resolution. In addition, the object images captured from above can have different orientations. i.e. objects classes like planes, vehicles, etc. can be oriented between 0 and 360 degrees, whereas object classes in COCO or Imagenet data are mostly vertical.

Recent works (Tundia et al., 2020) have focused and attended to identifying man-made agricultural and irrigation structures used for farming. Check-dams (See Fig.1 and Fig.2) are artificial irrigation structures built across water bodies to provide irrigation to nearby and surrounding places. They have a cultivable command area of up to 2000 hectares. Minor irrigation census is conducted for purposes of rural planning, formulating various policies and in reserving and allocating funds for agricultural development. However, maintaining a record of the minor irrigation structures makes irrigation census a time-consuming, expensive and laborious task, which is prone to errors. In the light of this, we introduce a new dataset of check-dams from satellite images for building an automated system for detection and mapping of check-dams to ease the process of census. Our work will complement existing object detection techniques on agricultural structures (Tundia et al., 2020) like farm ponds and wells. In this paper, we primarily focus on deep learning based object detection and instance segmentation methods and benchmark the latest methods in computer vision. Our contribution in this paper is two-fold :

1. **Minor Irrigation Structures Check-Dam Dataset:** a public dataset annotated by domain experts using images from Google static map APIs (Google Static Maps, 2021) for instance segmentation and object detection tasks.
2. A benchmark and assessment of the performance of various object detection and instance segmentation methods on the check-dam dataset.

The paper is structured as follows: section 2 covers the related work with an overview of the object detection and instance segmentation methods, section 3 covers the details of the proposed dataset, section 4 covers the evaluation criteria, architecture and training details, section 5 covers the experiments, results, observation and analysis, and finally conclusion in section 6.

2 RELATED WORK

Deep learning with convolutional layers has led to learning of high-level feature representations of images, followed by the trend of overlapping detectors and proposal generators leading to powerful object detectors. On the other hand, instance segmentation gives a finer inference for every pixel of the object in the input image, labelled using the segmentation labels.

2.1 Object Detection Methods

In general, object detection based methods are either based on region proposals or regression. Object proposals refer to the candidate boxes that possibly contain objects and are used to detect objects of multiple aspect ratios while avoiding sliding window searches. In two stage models, the first stage generates object proposals, followed by the refinement of these bounding boxes in the subsequent stages. In single stage models, the whole detection pipeline is performed in a single step, formulated as dense classification, generally optimized by a focal loss and localization of the object as bounding box regression. These regression based methods do not produce candidate region proposals or use any feature re-sampling methods, which lead to more efficient models. In bounding box regression, the location of a predicted bounding box is refined based on the anchor box and has been integrated into the detector while training. Deep regression applies deep learning to directly regress the bounding box coordinates based on deep learning features, but it is prone to difficulties in the localization of small objects. In this section, we briefly mention some of the object detection methods and summarize the important changes that each method introduced. Many of these object detection methods also double as instance segmentation methods.

RCNN (Girshick et al., 2014) extracts a set of object proposals by a selective search on a given image, and then the region proposals are re-scaled and fed to a backbone to extract features. These features are then used by SVM classifiers to predict and recognize the object categories within a region. Fast RCNN (Girshick, 2015) simultaneously trains a detector and performs bounding box regression, leading to detection speeds over 200 times faster than the RCNN. However, most of its computational time and resources get expended in computing the fully connected layers. Later, faster region based CNN (Faster RCNN) (Ren et al., 2015) introduces Region Proposal Method (RPN) and a detector that runs detection only on the network's top layer to further improve the in-

ference speed.

YOLO (Redmon et al., 2016) introduces the single stage detection paradigm without using proposal detection and verification. Though YOLO is extremely fast, it lacks the localization accuracy in comparison to two-stage object detection methods as well as the performance on small objects. SSD (Liu et al., 2016) focuses on YOLO’s issues by focusing on different layers of the network rather than just the topmost layer. YOLOv2 (Redmon and Farhadi, 2017) improves the detection of small objects over YOLO by using larger input sizes. YOLOv3 (Redmon and Farhadi, 2018) improves on the various metrics over YOLOv2 (Redmon and Farhadi, 2017) and improves on the real-time inference and execution time. RetinaNet (Lin et al., 2017) solves the issue of foreground-background class imbalance during training of dense one-stage detectors, by reshaping the cross entropy loss to focus on hard examples and down-weighting the loss assigned to well classified examples.

A few of the single-stage methods alleviate the anchor box imbalance between positive and negative anchor boxes by treating bounding boxes as keypoints pair. Cornernet (Law and Deng, 2018), a keypoint based one-stage detector eliminates the need for designing anchor boxes by using corner pooling to improve the corner localization. CascadeRPN (Vu et al., 2019) improves the quality of region-proposals by relying on single anchor per location and refines it in multiple stages, where each stage progressively defines positive samples by starting with anchor-free metric followed by anchor-based metrics. Free Anchor (Zhang et al., 2019) focuses on the idea of allowing objects to match the anchors flexibly and optimizes the detection customized likelihood. Fully Convolutional One Stage (FCOS) object detector (Tian et al., 2019), an anchor-free, proposal-free method approaches object detection in a per-pixel fashion, with only NMS post-processing step. Empirical Attention (Zhu et al., 2019) compares spatial attention elements from Transformer attention, dynamic convolution and deformable convolution in a generalized attention formulation. CentripetalNet (Dong et al., 2020) improves the corner point matching by using centripetal shift and by designing cross-star deformable convolution. DetectoRS (Qiao et al., 2020) combines Recursive Feature Pyramid (RFP) and Switchable Atrous Convolution (SAC) to achieve SoTA object detection performance. RFP incorporates feedback connections into the bottom-up backbone layers from the feature pyramid networks (FPN), while SAC uses the different atrous gates to convolve the features and uses switch

functions to gather the results.

SparseRCNN (Sun et al., 2020) reduces the number of hand-designed object candidates from thousands to a few hundred learnable proposals and predicts the final output directly without NMS post-processing. Dynamic RCNN (Zhang et al., 2020) proposes a dynamic design to alleviate the inconsistency problem between the fixed network and dynamic training procedures by adjusting the IoU threshold and the shape of regression loss function, based on the training statistics of the object proposals. GFL (Li et al., 2020b) designs a joint representation of localization quality and classification to eliminate the inconsistency risk and depicts the flexible distribution in real data, by merging the quality estimation into class prediction vector. The resulting labels are continuous, which are optimized using a generalized focal loss. Deformable DETR (Zhu et al., 2020) mitigates the issues of DETR by attending to key sampling points around a reference to overcome slow convergence and the limitations of feature spatial resolution.

2.2 Instance Segmentation Methods

In this section, we briefly mention some of the instance segmentation methods and summarize the important changes in each method. All these instance segmentation methods can also perform object detection under the same pipeline.

Cascade RCNN (Cai and Vasconcelos, 2019) uses a sequence of detectors that are trained sequentially using the output of a detector as the training set for the next detector. In cascaded detection, a coarse to fine technique is used, improving localization accuracy for small objects. Instaboost (Fang et al., 2019), a data augmentation technique, explores feasible locations where objects could be placed based on the similarity of local appearances and proposes a location probability map. Hybrid Task Cascade (Chen et al., 2019a) introduces a framework for joint multi-stage processing of both detection and segmentation pipelines by progressively learning discriminative features by integrating complementary features.

YOLACT (Bolya et al., 2019) performs real-time instance segmentation by having two parallel pipelines, where one generates a set of prototype masks, while the other predicts the coefficients of masks per-instance. General ROI Extraction (Rossi et al., 2020) introduces non-local building blocks and attention mechanisms to attend to multiple layers of FPN, to extract a coherent subset of features for integrating in two stage methods. Prime Sample Attention (PISA) in object detection (Cao et al., 2020) assesses how different samples from the dataset con-

tribute to the overall performance in mean AP. Swin Transformer (Liu et al., 2021) presents a hierarchical architecture that uses shifted windows for computing representations by limiting self-attention computation to non-overlapping local windows, while also permitting cross-window connections.

3 PROPOSED DATASET

Here, we introduce and give the details of our dataset of satellite images on check-dams, an irrigation structure constructed for agricultural needs. Our dataset has images comparable in image dimensions to that of Imagenet and COCO, instead of commonly used single-view remote sensing image with dimensions in the range of tens of thousands of pixels. Our dataset falls in the category of small-scale optical satellite image dataset, whose ground truth instance level annotations are not easily available. Our dataset complements minor irrigation structures dataset (Tundia et al., 2020) for wells and farm ponds, by adding a new category in man-made irrigation structures. While most datasets are designed keeping only one task in mind, our dataset is designed with annotations for both object detection and instance segmentation tasks. To ensure geographical diversity in the dataset, we collected images from 36 districts of Maharashtra, India, based on the availability of ground-truth data.

Related Datasets. There are several datasets in remote sensing that use images from multiple sources like airplane, drone (aerial images), satellites (optical, multispectral, hyperspectral, etc.). Some of these datasets focus on objects like buildings, vehicles, cars, ships, etc., while some have classes like pools, playgrounds, vehicles of different types, etc. While some datasets have very high resolution images, some datasets have images ranging from 10 (Maggiori et al., 2017) to a few tens of thousands. Datasets also vary in terms of the number of instances from 600 (Maggiori et al., 2017) to around 1/5 of a million (Li et al., 2020a). A comparison of the related datasets in terms of the number of classes, number of instances, number of images and the image sizes are given in the Table 1.

Scale: Contrasting the datasets (Li et al., 2020a) that have large class imbalance, we have 562 instances of wall-based check-dam and 520 instances of gate-based check-dam. While other datasets have relatively abundant and dense object instances that are easy to capture, our dataset has instances that are geographically far apart from each other and require annotations from the domain experts. In comparison to dense object datasets (e.g. cars in aerial images), our

Table 1: Optical satellite image datasets for object detection.

Dataset	#Classes	#Instances	#Images	Image Size (pixels)
RSOD (Long et al., 2017)	4	6,950	976	~1000
CARPPK (Hsieh et al., 2017)	1	89,777	1448	1280
ITCVD (Yang, 2018)	1	228	23,543	5616
LEVIR (Chen and Shi, 2020)	3	11,000	22,000	600-800
SpaceNet MVOI (Weir et al., 2019)	1	126,747	60,000	900
DIOR (Li et al., 2020a)	20	90,228	23,463	800
RarePlanes (Shermeyer et al., 2020)	1	644,258	50,253	1080
Well (Tundia et al., 2020)	1	1614	1,011	640
Farmpond (Tundia et al., 2020)	4	715	1,018	640
MIS Check-Dam	2	1082	1,037	640

dataset has images consisting of sparsely located objects. The object instances in our dataset come from two categories. Our dataset has 1082 instances of check-dams across 1037 images on images of 640 x 640 pixel dimensions. Our dataset also has images with visually similar objects to check-dams like buildings, huts, etc. that makes the object detection and instance segmentation task difficult. In terms of spatial resolution, Google maps API uses zoom levels ranging from 1 to 22 that correspond to the image spatial resolution. However, most objects of interest are identifiable to humans at zoom levels of 17, 18 and 19 with resolutions of 1.262, 0.315 and 0.078 respectively. While a heavy traffic image can have up to hundreds of instances, only 8-10 farm ponds and wells are visible at zoom level 18, and upto four check-dams are visible at zoom level 18.

Size Variations: Check-dams are subject to size variations due to their surrounding environments. Unlike other dataset classes, our dataset has instances with large intra-class diversity and high inter-class similarity. Since a check-dam is built across a river or water body, it is easy for a human with domain knowledge to identify the structure and distinguish between a check-dam or a bridge. However, the width of a water body varies along its trajectory giving rise to large variations in check-dam sizes in comparison to objects like cars or planes in other datasets, which may have a standard shape and size.

Image Variations: The satellite images collected from Google maps span seasons and collection times. The imaging conditions also vary due to the presence of air pollution, clouds, weather conditions and image sensor types, which leads to blurry images or image tinting. Moreover, there are wide variations in the surrounding vegetation even under a small geographical region. Unlike nadir (overhead) satellite imagery, off-nadir imagery captures different perspectives of the

same object. An object captured from one angle may cast the shadow and appear differently based on the time of image capture. Also, the presence of trees and surrounding vegetation may occlude the check-dam. Moreover, a check-dam and its surrounding appearances vary depending upon the water level, seasonal rains and the irrigation purposes and uses.

4 IMPLEMENTATION DETAILS

4.1 Evaluation Criteria

The two main metrics used for object detection and instance segmentation tasks are PascalVOC metric (Everingham et al., 2012) and COCO metric (Lin et al., 2014). These metrics rely on Intersection Over Union (IOU) given by the overlapping area between the predicted bounding box and the ground truth bounding box, divided by the area common between them.

PascalVOC metric takes the average precision at a fixed IoU of 0.5, whereas COCO metric measures the average over multiple IoU thresholds. For multi-class detectors, final mAP (AP@[0.5:0.95]) is computed by taking the mean over 10 IoU levels (starting from 0.5 to 0.95 with a step size of 0.05) on each class. We base all our experiments on COCO metric, since COCO metric provides metrics at different IOU thresholds.

4.2 Architecture and Training Details

For training our dataset, we made a train-test split of 80-20, which consists of 450 wall-based check-dam instances and 418 gate-based check-dam instances in the training set and 110 wall-based check-dam instances and 103 gate-based check-dam instances in the test set. The details of architecture and backbone used in some of the methods are given in Table 3. The specific details of the method, the model, and the backbone used for training are given in Table 2 and Table 4. We used MMDetection (Chen et al., 2019b) framework for implementing the various methods in our benchmark. We used ResNet 50, ResNet 101, ResNext 101, HourglassNet, DetectoRS-R50 pretrained backbones in object detection experiments and Resnet-50, Resnet-101, and Swin Transformer pretrained backbones in instance segmentation experiments. We used pretrained backbones primarily trained on Imagenet, COCO, and cityscapes for training different models.

5 EXPERIMENTS AND RESULTS

Localization Distillation: The network size and the number of parameters are crucial for network inference speeds. Network distillation compresses the knowledge from a teacher network to a smaller student network, by overcoming the limitation of distilling only localization information, for arbitrary teacher and student architectures. While network pruning reduces the network sizes by removing the unimportant weights, this cannot be applied directly to object detection methods, as it may result in sparse connectivity patterns in a CNN. We experiment to see whether knowledge distillation is possible with localization distillation (Zheng et al., 2021). We used three different settings based on GFL (Li et al., 2020b), training a teacher with a pretrained backbone of ResNet-101 and a student with different backbone settings of ResNet-18, ResNet-34 and ResNet-50. The details of the teacher-student backbone and the results of the knowledge distillation on object detection task are given in Table 5. We can observe from Table 5 that the bbox mAP (0.50) of 0.966 is close to 0.967 of DetectoRS (Qiao et al., 2020) given in Table 4. Also, performing the localization distillation on the same model as the teacher can further improve the performance of the model.

5.1 Results

The quantitative results of comparison of various object detection methods are given in Table 4 and the results of comparison of various instance segmentation methods are given in Table 2. The results of training a Faster-RCNN model from scratch is given in the topmost row of Table 4 while the remaining table rows compare the performance of methods trained using pre-trained backbone. We observe that the model trained from scratch has performed poorly even after training for more than a hundred epochs, while the metric scores have improved only gradually. We observe that using pre-trained backbone not only improves the convergence speed, but also helps to achieve the best possible performance from a method. We observe that DetectoRS has the highest values of 0.625 in mAP (0.50:0.95), 0.967 in mAP (0.5) and 0.749 in mAP (0.75) for the object detection task. We attribute the performance of DetectoRS to hybrid task cascade along with recursive feature pyramid and switchable atrous convolution. Though RFP and SAC can be independently applied, the best performance is observed when both are applied. For the instance segmentation task, both Swin transformer and HTC records the highest values of 0.478 in mAP

Table 2: Comparison of the performance of various instance segmentation methods on checkdam dataset.

Method	Model Name	Backbone	bbox mAP (0.50:0.95)	bbox mAP (0.50)	bbox mAP (0.75)	segm mAP (0.50:0.95)	segm mAP (0.50)	segm mAP (0.75)
Hybrid Task Cascade (Chen et al., 2019a)	HTC	R-50	0.605	0.940	0.705	0.478	0.918	0.472
	HTC	R-101	0.599	0.936	0.705	0.474	0.905	0.446
YOLOACT (Bolya et al., 2019)	YOLOACT	R-50	0.539	0.966	0.542	0.458	0.939	0.409
	YOLOACT	R-101	0.545	0.940	0.574	0.435	0.908	0.377
Instaboost (Fang et al., 2019)	Cascade MRCNN	R-50	0.613	0.933	0.722	0.477	0.915	0.451
	Cascade MRCNN	R-101	0.607	0.942	0.707	0.457	0.898	0.444
Cascade RCNN (Cai and Vasconcelos, 2019)	Cascade MRCNN	R-50	0.601	0.932	0.692	0.466	0.902	0.449
	Cascade MRCNN	R-101	0.567	0.933	0.638	0.460	0.893	0.423
General ROI Extraction (Rossi et al., 2020)	MRCNN	R-50	0.587	0.950	0.697	0.472	0.908	0.441
Prime Sample Attention (Cao et al., 2020)	MRCNN	R-50	0.603	0.941	0.703	0.477	0.914	0.477
Swin Transformer (Liu et al., 2021)	MRCNN (FP-16)	Swin Transformer	0.603	0.971	0.699	0.465	0.921	0.451
	MRCNN	Swin Transformer	0.617	0.978	0.697	0.478	0.920	0.464



Figure 3: Object Detection results on the test data. (From left to right) a, b: Wall based check-dams detected among vegetation and dry surroundings; c, d: Detection on large size varieties of gate based check-dams.



Figure 4: Instance Segmentation results on the test data. (From left to right) a, b: Wall based check-dams segmented in dry stream and wet stream surroundings; c, d: Gate based check-dams segmented in dry stream and wet stream.

Table 3: Summary of the different architectures used in the comparison study.

Method	Architecture Details
YOLOv3	Darknet
CornerNet	Stacked Hourglass, Corner pooling
Centripetal Net	Stacked Hourglass, Corner pooling
Empirical Attention	Deformable Convolution(DCN), Spatial Attention(1111, 0010)
FCOS	Group Normalization(GN), DCN
Generalized Focal Loss	Generalized Focal Loss, DCNv2
General ROI Extraction	Non-local building block + Attention
Prime Sample Attention	PISA, ROI Pool
Hybrid Task Cascade	DCN, Multiscale Training
DetectoRS	RFP, SAC

(0.50:0.95), while YOLOACT records the highest of 0.939 in mAP (0.5) and PISA records the highest of 0.477 in mAP (0.75).

The qualitative results of object detection by DetectoRS on our dataset are given in Fig. 3, while the qualitative results of instance segmentation using Hy-

brid Task Cascade are given in Figure 4. From Fig. 3, we can observe that DetectoRS is able to detect bounding boxes accurately for both small and large check-dams as well as in both green and dry surroundings. From Fig. 4, we can observe that HTC is able to perform instance segmentation precisely on check-dams in both dry and wet streams.

5.2 Observations and Analysis

Recent works have integrated different techniques for improving the performance of object detection and instance segmentation tasks. Any improvement that reduces the number and size of the various network parameters without affecting the other aspects of detection or segmentation can speedup the inference stage. In sliding window based detectors, the overlap between adjacent windows can be reduced by fea-

Table 4: Comparison of the performance of various object detection methods on checkdam dataset.

Method	Model Name	Backbone	bbox mAP (0.50:0.95)	bbox mAP (0.50)	bbox mAP (0.75)
Faster RCNN-Scratch	Faster RCNN	R-50-FPN	0.145	0.381	0.077
Faster RCNN (Ren et al., 2015)	Faster RCNN	R-50-FPN	0.597	0.948	0.699
	Faster RCNN	R-101-FPN	0.576	0.94	0.625
YOLOv3 (Redmon and Farhadi, 2018)	YOLOv3	Darknet53 (Scale-416)	0.504	0.923	0.503
	YOLOv3	Darknet53 (Scale-608)	0.526	0.939	0.588
RetinaNet (Lin et al., 2017)	RetinaNet	R-50	0.580	0.940	0.670
	RetinaNet	R-101	0.586	0.946	0.681
Free Anchor (Zhang et al., 2019)	RetinaNet	R-50	0.574	0.939	0.654
Generalized Focal Loss (Li et al., 2020b)	GFL	R-101	0.519	0.918	0.519
CascadeRPN (Vu et al., 2019)	FasterRCNN	R-50	0.583	0.945	0.610
Dynamic RCNN (Zhang et al., 2020)	FasterRCNN	R-50	0.557	0.951	0.592
Empirical Attention (Zhu et al., 2019)	FasterRCNN (0010)	R-50	0.594	0.953	0.667
	FasterRCNN (1111)	R-50	0.597	0.945	0.710
FCOS Object Detector (Tian et al., 2019)	FCOS	R-50	0.262	0.612	0.161
	FCOS	R-101	0.191	0.481	0.116
SparseRCNN (Sun et al., 2020)	SparseRCNN (#Prop: 100)	R-50	0.573	0.935	0.596
	SparseRCNN (#Prop: 300)	R-50	0.580	0.944	0.647
	SparseRCNN (#Prop: 100)	R-101	0.577	0.939	0.647
	SparseRCNN (#Prop: 300)	R-101	0.554	0.937	0.609
Cascade RCNN (Cai and Vasconcelos, 2019)	CascadeRCNN	R-50	0.601	0.939	0.701
	CascadeRCNN	R-101	0.593	0.941	0.658
	CascadeRCNN	ResNeXt-101	0.600	0.943	0.685
CornerNet (Law and Deng, 2018)	CornerNet (BS: 8x6)	HourglassNet	0.570	0.894	0.672
	CornerNet (BS: 32x3)	HourglassNet	0.575	0.904	0.662
Centripetal Net (Dong et al., 2020)	CornerNet (Batch Size 16x6)	HourglassNet	0.598	0.943	0.686
DectoS (Qiao et al., 2020)	CascadeRCNN	DectoS R-50	0.615	0.967	0.711
	RFP - HTC + R-50	DectoS R-50	0.617	0.953	0.693
	SAC - HTC + R-50	DectoS R-50	0.625	0.966	0.749
	Two-stage Deformable DETR	R-50	0.540	0.930	0.584

ture map shared computation of the whole image only once before sliding windows. Similarly, grouping the feature channels into independent groups can also reduce the parameter count. Another way to reduce the complexity of a layer and filters is to approximate it with fewer filters and a non-linear activation. Factorizing convolutions is an efficient way to replace a very large filter with smaller filter sizes, which can share the same receptive fields while being efficient. The local context can help improve the object detection by referring to the visual area that surrounds the objects of interest. Similarly the global context can help integrate the information of the different scene elements by having larger receptive fields or a global pooling of the CNN features, to improve the object detection performance.

Table 5: Comparing localization distillation performance using different neural network architectures between the (T)eacher and the (S)tudent.

Backbone	bbox mAP (0.50:0.95)	bbox mAP (0.50)	bbox mAP (0.75)
R-101 (T), R-18 (S)	0.532	0.934	0.572
R-101 (T), R-34 (S)	0.592	0.966	0.670
R-101 (T), R-50 (S)	0.566	0.957	0.654

6 CONCLUSIONS

We introduced MIS Check-Dam, a dataset for check-dams that belongs to the class of minor irrigation structures for agricultural use. We benchmark our dataset on some of the most recent object detection methods and instance segmentation methods. We also assess the importance of various components in the pipeline after evaluating these methods on our novel dataset. Future works can try domain adaptation methods on our dataset and other related datasets that have related classes like dams.

REFERENCES

Bolya, D., Zhou, C., Xiao, F., and Lee, Y. J. (2019). YOLACT: Real-time Instance Segmentation. In *ICCV*.

Cai, Z. and Vasconcelos, N. (2019). Cascade R-CNN: High Quality Object Detection and Instance Segmentation. *IEEE TPAMI*.

Cao, Y., Chen, K., Loy, C. C., and Lin, D. (2020). Prime Sample Attention in Object Detection. In *IEEE CVPR*.

Chen, H. and Shi, Z. (2020). A Spatial-Temporal Attention-Based Method and a New Dataset for Remote Sensing Image Change Detection. *Remote Sensing*.

Chen, K., Lin, D., and et. al (2019a). Hybrid Task Cascade for Instance Segmentation. In *IEEE CVPR*.

- Chen, K., Wang, J., Pang, and et. al. (2019b). MMDetection: Open MMLab Detection Toolbox and Benchmark. *arXiv:1906.07155*.
- Dong, Z., Li, G., Liao, Y., Wang, F., Ren, P., and Qian, C. (2020). CentripetalNet: Pursuing High-Quality Keypoint Pairs for Object Detection. In *IEEE CVPR*.
- Everingham, M., Van Gool, L., Williams, C. K. I., Winn, J., and Zisserman, A. (2012). The PASCAL Visual Object Classes Challenge 2012 (VOC2012) Results.
- Fang, H.-S., Sun, J., Wang, R., Gou, M., Li, Y.-L., and Lu, C. (2019). Instaboost: Boosting Instance Segmentation via Probability Map Guided Copy-Pasting. In *IEEE ICCV*.
- Fedus, W., Zoph, B., and Shazeer, N. (2021). Switch Transformers: Scaling to Trillion Parameter Models with Simple and Efficient Sparsity. *arXiv:2101.03961*.
- Girshick, R. (2015). Fast R-CNN. In *2015 IEEE ICCV*.
- Girshick, R., Donahue, J., Darrell, T., and Malik, J. (2014). Rich Feature Hierarchies for Accurate Object Detection and Semantic Segmentation. In *2014 CVPR*.
- Google Static Maps (2021). Google Maps JavaScript API Version 3 Reference.
- Henderson, P., Islam, R., Bachman, P., Pineau, J., Precup, D., and Meger, D. (2018). Deep Reinforcement Learning That Matters. In *AAAI*.
- Hsieh, M.-R., Lin, Y.-L., and Hsu, W. H. (2017). Drone-based Object Counting by Spatially Regularized Regional Proposal Networks. In *IEEE ICCV*.
- Jiao, L. and Zhao, J. (2019). A Survey on the New Generation of Deep Learning in Image Processing. *IEEE Access*.
- Kumar, R., Dabral, R., and Sivakumar, G. (2021). Learning Unsupervised Cross-domain Image-to-Image Translation using a Shared Discriminator. In *Vol. 4: VISAPP*.
- Law, H. and Deng, J. (2018). Cornernet: Detecting objects as Paired Keypoints. In *15th ECCV 2018*.
- Li, K., Wan, G., Cheng, G., Meng, L., and Han, J. (2020a). Object Detection in Optical Remote Sensing Images: A Survey and a New Benchmark. *ISPRS*.
- Li, X., Yang, J., and et. al (2020b). Generalized Focal Loss: Learning Qualified and Distributed Bounding Boxes for Dense Object Detection. *arXiv:2006.04388*.
- Lin, T., Maire, M., Zitnick, C. L., and et. al. (2014). Microsoft COCO: common objects in context. *CoRR*, abs/1405.0312.
- Lin, T.-Y., Goyal, P., Girshick, R., He, K., and Dollár, P. (2017). Focal Loss for Dense Object Detection. In *IEEE ICCV*.
- Liu, W., Anguelov, D., Erhan, D., Szegedy, C., Reed, S., Fu, C.-Y., and Berg, A. C. (2016). SSD: Single Shot MultiBox Detector. In *ECCV 2016*.
- Liu, Z., Guo, B., and et. al (2021). Swin Transformer: Hierarchical Vision Transformer using Shifted Windows. *CoRR*, abs/2103.14030.
- Long, Y., Gong, Y., Xiao, Z., and Liu, Q. (2017). Accurate Object Localization in Remote Sensing Images Based on Convolutional Neural Networks. *IEEE TGRS*.
- Maggiore, E., Tarabalka, Y., Charpiat, G., and Alliez, P. (2017). Can Semantic Labeling Methods Generalize to Any City? The Inria Aerial Image Labeling Benchmark. In *IEEE IGARSS*.
- Qiao, S., Chen, L., and Yuille, A. (2020). DetectoRS: Detecting Objects with Recursive Feature Pyramid and Switchable Atrous Convolution. *arXiv:2006.02334*.
- Redmon, J., Divvala, S., Girshick, R., and Farhadi, A. (2016). You Only Look Once: Unified, Real-Time Object Detection. In *2016 IEEE CVPR*.
- Redmon, J. and Farhadi, A. (2017). YOLO9000: Better, Faster, Stronger. In *2017 IEEE CVPR*.
- Redmon, J. and Farhadi, A. (2018). Yolov3: An incremental improvement. *CoRR*, abs/1804.02767.
- Ren, S., He, K., Girshick, R., and Sun, J. (2015). Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks. In *NeurIPS*, volume 28.
- Rossi, L., Karimi, A., and Prati, A. (2020). A novel region of interest extraction layer for instance segmentation. *CoRR*, abs/2004.13665.
- Russakovsky, O., Fei-Fei, L., and et. al (2015). Imagenet large scale visual recognition challenge. *IJCV*.
- Shermeyer, J., Hossler, T., Van Etten, A., Hogan, D., Lewis, R., and Kim, D. (2020). RarePlanes: Synthetic Data Takes Flight. *In-Q-Tel - CosmiQ Works and AI.Reverie*.
- Sun, P., Luo, P., and et. al (2020). SparseR-CNN: End-to-End Object Detection with Learnable Proposals. *arXiv:2011.12450*.
- Tian, Z., Shen, C., Chen, H., and He, T. (2019). FCOS: Fully Convolutional One-Stage Object Detection. *arXiv:1904.01355*.
- Tundia, C., Tank, P., and Damani, O. (2020). Aiding Irrigation Census in Developing Countries by Detecting Minor Irrigation Structures from Satellite Imagery. In *6th Inter. Conf. on GISTAM*.
- Vu, T., Jang, H., Pham, T. X., and Yoo, C. D. (2019). Cascade RPN: Delving into High-Quality Region Proposal Network with Adaptive Convolution. In *NeurIPS*.
- Weir, N., Lindenbaum, D., Bastidas, A., Etten, A., Kumar, V., Mcpherson, S., Shermeyer, J., and Tang, H. (2019). SpaceNet MVOI: A Multi-View Overhead Imagery Dataset. In *2019 IEEE/CVF ICCV*.
- Yang, D. M. (2018). ITCVD Dataset. DANS. *Faculty of GIS and Earth Observation (ITC)*.
- Zhang, H., Chang, H., Ma, B., Wang, N., and Chen, X. (2020). Dynamic R-CNN: Towards High Quality Object Detection via Dynamic Training. *arXiv:2004.06002*.
- Zhang, X., Wan, F., Liu, C., Ji, R., and Ye, Q. (2019). FreeAnchor: Learning to Match Anchors for Visual Object Detection. In *NeurIPS*.
- Zheng, Z., Ye, R., Wang, P., Wang, J., Ren, D., and Zuo, W. (2021). Localization Distillation for Object Detection. *arXiv:2102.12252*.
- Zhu, X., Cheng, D., Zhang, Z., Lin, S., and Dai, J. (2019). An Empirical Study of Spatial Attention Mechanisms in Deep Networks. *arXiv:1904.05873*.
- Zhu, X., Su, W., Lu, L., Li, B., Wang, X., and Dai, J. (2020). Deformable DETR: deformable transformers for end-to-end object detection. *CoRR*, abs/2010.04159.