Article

Small-Object Detection in Remote Sensing Images with End-to-End Edge-Enhanced GAN and Object Detector Network

Jakaria Rabbi [1,*] iD **, Nilanjan Ray** [1] iD **, Matthias Schubert** [2] **, Subir Chowdhury** [3] iD **and Dennis Chao** [3]

[1]   Department of Computing Science, 2-32 Athabasca Hall, University of Alberta, Edmonton, AB T6G 2E8, Canada; nray1@ualberta.ca

[2]   Institute for Informatic, Ludwig-Maximilians-Universität München, Oettingenstraße 67, D-80333 Munich, Germany; schubert@dbs.ifi.lmu.de

[3]   Alberta Geological Survey, Alberta Energy Regulator, Edmonton, Alberta, Canada; subir.chowdhury@aer.ca; dennis.chao@aer.ca

*   Correspondence: jakaria@ualberta.ca

Abstract: The detection performance of small objects in remote sensing images has not been satisfactory compared to large objects, especially in low-resolution and noisy images. A generative adversarial network (GAN)-based model called enhanced super-resolution GAN (ESRGAN) showed remarkable image enhancement performance, but reconstructed images usually miss high-frequency edge information. Therefore, object detection performance showed degradation for small objects on recovered noisy and low-resolution remote sensing images. Inspired by the success of edge enhanced GAN (EEGAN) and ESRGAN, we applied a new edge-enhanced super-resolution GAN (EESRGAN) to improve the quality of remote sensing images and used different detector networks in an end-to-end manner where detector loss was backpropagated into the EESRGAN to improve the detection performance. We proposed an architecture with three components: ESRGAN, EEN, and Detection network. We used residual-in-residual dense blocks (RRDB) for both the ESRGAN and EEN, and for the detector network, we used a faster region-based convolutional network (FRCNN) (two-stage detector) and a single-shot multibox detector (SSD) (one stage detector). Extensive experiments on a public (car overhead with context) dataset and another self-assembled (oil and gas storage tank) satellite dataset showed superior performance of our method compared to the standalone state-of-the-art object detectors.

Keywords: Object detection; faster region-based convolutional neural network (FRCNN); single-shot multibox detector (SSD); super-resolution; remote sensing imagery; edge enhancement; satellites

Abbreviations

The following acronyms are used in this paper:

SRCNN	Single image Super-Resolution Convolutional Neural Network
VDSR	Very Deep Convolutional Networks
GAN	Generative Adversarial Network
SRGAN	Super-Resolution Generative Adversarial Network
ESRGAN	Enhanced Super-Resolution Generative Adversarial Network
EEGAN	Edge-Enhanced Generative Adversarial Network
EESRGAN	Edge-Enhanced Super-Resolution Generative Adversarial Network
RRDB	Residual-in-Residual Dense Blocks
EEN	Edge-Enhancement Network
SSD	Single-Shot MultiBox Detector
YOLO	You Only Look Once
CNN	Convolutional Neural Network
R-CNN	Region-based Convolutional Neural Network

FRCNN	Faster Region-based Convolutional Neural Network
VGG	Visual Geometry Group
BN	Batch Normalization
MSCOCO	Microsoft Common Objects in Context
OGST	Oil and Gas Storage Tank
COWC	Car Overhead With Context
GSD	Ground Sampling Distance
G	Generator
D	Discriminator
ISR	Intermediate Super-Resolution
SR	Super-Resolution
HR	High-Resolution
LR	Low-Resoluton
GT	Ground Truth
FPN	Feature Pyramid Network
RPN	Region Proposal Network
AER	Alberta Energy Regulator
AGS	Alberta Geological Survey
AP	Average Precision
IoU	Intersection over Union
TP	True Positive
FP	False Positive
FN	False Negative

1. Introduction

1.1. Problem Description and Motivation

Object detection on remote sensing imagery has numerous prospects in various fields, such as environmental regulation, surveillance, military [1,2], national security, traffic, forestry [3], oil and gas activity monitoring. There are many methods for detecting and locating objects from images, which are captured using satellites or drones. However, detection performance is not satisfactory for noisy and low-resolution (LR) images, especially when the objects are small [4]. Even on high-resolution (HR) images, the detection performance for small objects is lower than that for large objects [5].

Current state-of-the-art detectors have excellent accuracy on benchmark datasets, such as ImageNet [6] and Microsoft common objects in context (MSCOCO) [7]. These datasets consist of everyday natural images with distinguishable features and comparatively large objects.

On the other hand, there are various objects in satellite images like vehicles, small houses, small oil and gas storage tanks etc., only covering a small area [4]. The state-of-the-art detectors [8–11] show a significant performance gap between LR images and their HR counterparts due to a lack of input features for small objects [12]. In addition to the general object detectors, researchers have proposed specialized methods, algorithms, and network architectures to detect particular types of objects from satellite images such as vehicles [13,14], buildings [15], and storage tanks [16]. These methods are object-specific and use fixed resolution for feature extraction and detection.

To improve detection accuracy on remote sensing images, researchers have used deep convolutional neural network (CNN)-based super-resolution (SR) techniques to generate artificial images and then detect objects [5,12]. Deep CNN-based SR techniques such as single image super-resolution convolutional networks (SRCNN) [17] and accurate image super-resolution using very deep convolutional networks (VDSR) [18] showed excellent results on generating realistic HR imagery from LR input data. Generative Adversarial Network (GAN)-based [19] methods such as super-resolution GAN (SRGAN) [20] and enhanced super-resolution GAN (ESRGAN) [21] showed remarkable performance in enhancing LR images with and without noise. These models have two subnetworks: a generator and a discriminator. Both subnetworks consist of deep CNNs. Datasets containing HR and LR image pairs are used for training and testing the models. The generator generates HR images from LR input images, and the discriminator predicts whether generated image is a real HR image or an upscaled LR image. After sufficient training, the generator generates HR images

that are similar to the ground truth HR images, and the discriminator cannot correctly discriminate between real and fake images anymore.

Although the resulting images look realistic, the compensated high-frequency details such as image edges may cause inconsistency with the HR ground truth images [22]. Some works showed that this issue negatively impacts land cover classification results [23,24]. Edge information is an important feature for object detection [25], and therefore, this information needs to be preserved in the enhanced images for acceptable detection accuracy.

In order to obtain clear and distinguishable edge information, researchers proposed several methods using separate deep CNN edge extractors [26,27]. The results of these methods are sufficient for natural images, but the performance degrades on LR and noisy remote sensing images [22]. A recent method [22] used the GAN-based edge-enhancement network (EEGAN) to generate a visually pleasing result with sufficient edge information. EEGAN employs two subnetworks for the generator. One network generates intermediate HR images, and the other network generates sharp and noise-free edges from the intermediate images. The method uses a Laplacian operator [28] to extract edge information and in addition, it uses a mask branch to obtain noise-free edges. This approach preserves sufficient edge information, but sometimes the final output images are blurry compared to a current state-of-the-art GAN-based SR method [21] due to the noises introduced in the enhanced edges that might hurt object detection performance.

Another important issue with small-object detection is the huge cost of HR imagery for large areas. Many organizations are using very high-resolution satellite imagery to fulfill their purposes. When it comes to continuous monitoring of a large area for regulation or traffic purposes, it is costly to buy HR imagery frequently. Publicly available satellite imagery such as Landsat-8 [29] (30 m/pixel) and Sentinel-2 [30] (10 m/pixel) are not suitable for detecting small objects due to the high ground sampling distance (GSD). Detection of small objects (e.g., oil and gas storage tanks and buildings) is possible from commercial satellite imagery such as 1.5-m GSD SPOT-6 imagery but the detection accuracy is low compared to HR imagery, e.g., 30-cm GSD DigitalGlobe imagery in Bing map.

We have identified two main problems to detect small-objects from satellite imagery. First, the accuracy of small-object detection is lower compared to large objects, even in HR imagery due to sensor noise, atmospheric effects, and geometric distortion. Secondly, we need to have access to HR imagery, which is very costly for a vast region with frequent updates. Therefore, we need a solution to increase the accuracy of the detection of smaller objects from LR imagery. To the best of our knowledge, no work employed both SR network with edge-enhancement and object detector network in an end-to-end manner, i.e, using joint optimization to detect small remote sensing objects.

In this paper, we propose an end-to-end architecture where object detection and super-resolution is performed simultaneously. Figure 1 shows the significance of our method. State-of-the-art detectors miss objects when trained on the LR images; in comparison, our method can detect those objects. The detection performance improves when we use SR images for the detection of objects from two different datasets. Average precision (AP) versus different intersection over union (IoU) values (for both LR and SR) are plotted to visualize overall performance on test datasets. From figure 1, we observe that for both the datasets, our proposed end-to-end method yields significantly better IoU values for the same AP. In section 4.2, we discuss AP and IoU in more detail and these results are discussed in section 4.

| | (I) LR image | (II) SR image | (III) AP vs IoU curves |
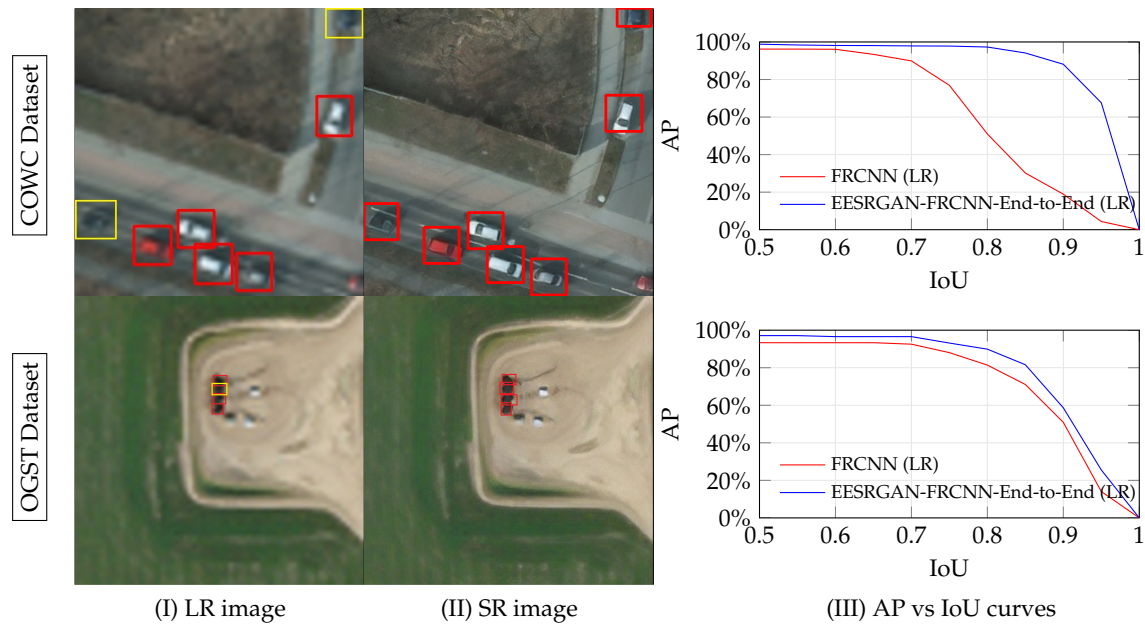
Figure 1. Detection on LR (low-resolution) images (60cm/pixel) is shown in (I); in (II), we show the detection on generated SR (super-resolution) images (15 cm/pixel). The first row of this figure represents the COWC (car overhead with context) dataset [31], and the second row represents the OGST (oil and gas storage tank) dataset [32]. AP (average precision) values versus different IoU (intersection over union) values for the LR test set and generated SR images from the LR images are shown in (III) for both the datasets. We use FRCNN (faster region-based CNN) detector on LR images for detection. Then instead of using LR images directly, we use our proposed end-to-end EESRGAN (edge-enhanced SRGAN) and FRCNN architecture (EESRGAN-FRCNN) to generate SR images and simultaneously detect objects from the SR images. The red bounding boxes represent true positives, and yellow bounding boxes represent false negatives. IoU=0.75 is used for detection.

1.2. Contributions of Our Method

Our proposed architecture consists of two parts: EESRGAN network and a detector network. Our approach is inspired by EEGAN and ESRGAN networks and showed a remarkable improvement over EEGAN to generate visually pleasing SR satellite images with enough edge information. We employed a generator subnetwork, a discriminator subnetwork, and an edge-enhancement subnetwork [22] for the SR network. For the generator and edge-enhancement network, we used residual-in-residual dense blocks (RRDB) [21]. These blocks contain multi-level residual networks with dense connections that showed good performance on image enhancement.

We used a relativistic discriminator [33] instead of a normal discriminator. Besides GAN loss and discriminator loss, we employed Charbonnier loss [34] for the edge-enhancement network. Finally, we used different detectors [8,10] to detect small objects from the SR images. The detectors acted like the discriminator as we backpropagated the detection loss into the SR network and, therefore, it improved the quality of the SR images.

We created the oil and gas storage tank (OGST) dataset [32] from satellite imagery (Bing map), which has 30 cm and 1.2 m GSD. The dataset contains labeled oil and gas storage tanks from the Canadian province of Alberta, and we detected the tanks on SR images. Detection and counting of the tanks are essential for the Alberta Energy Regulator (AER) [35] to ensure safe, efficient, orderly, and environmentally responsible development of energy resources. Therefore, there is a potential use of our method for detecting small objects from LR satellite imagery. The OGST dataset is available on Mendeley [32].

In addition to the OGST dataset, we applied our method on the publicly available car overhead with context (COWC) [31] dataset to compare the performance of detection for varying use-cases.

During training, we used HR and LR image pairs but only required LR images for testing. Our method outperformed standalone state-of-the-art detectors for both datasets.

The remainder of this paper is structured as follows. We discuss related work in section 2. In section 3, we introduce our proposed method and describe every part of the method. The description of datasets and experimental results are shown in section 4, final discussion is stated in section 5 and section 6 concludes our paper with a summary.

2. Related Works

Our work consists of an end-to-end edge enhanced image SR network with an object detector network. In this section, we discuss existing methods related to our work.

2.1. Image Super-Resolution

Many methods were proposed on SR using deep CNNs. Dong et al. proposed super-resolution CNN (SRCNN) [17] to enhance LR images in an end-to-end training outperforming previous SR techniques. The deep CNNs for SR evolved rapidly, and researchers introduced residual blocks [20], densely connected networks [36], and residual dense block [37] for improving SR results. He et al. [38] and Lim et al. [39] used deep CNNs without the batch normalization (BN) layer and observed significant performance improvement and stable training with a deeper network. These works were done on natural images.

Liebel et al. [40] proposed deep CNN-based SR network for multi-spectral remote sensing imagery. Jiang et al. [22] proposed a new SR architecture for satellite imagery that was based on GAN. They introduced an edge-enhancement subnetwork to acquire smooth edge details in the final SR images.

2.2. Object Detection

Deep learning-based object detectors can be categorized into two subgroups, region-based CNN (R-CNN) models that employ two-stage detection and uniform models using single stage detection [41]. Two-stage detectors comprise of R-CNN [42], Fast R-CNN [43], Faster R-CNN [8] and the most used single stage detectors are SSD [10], You only look once (YOLO) [11] and RetinaNet [9]. In the first stage of a two-stage detector, regions of interest are determined by selective search or a region proposal network. Then, in the second stage, the selected regions are checked for particular types of objects and minimal bounding boxes for the detected objects are predicted. In contrast, single-stage detectors omit the region proposal network and run detection on a dense sampling of all possible locations. Therefore, single-stage detectors are faster but, usually less accurate. RetinaNet [9] uses a focal loss function to deal with the data imbalance problem caused by many background objects and often showed similar performance as the two-stage approaches.

Many deep CNN-based object detectors were proposed on remote sensing imagery to detect and count small objects, such as vehicles [13,44,45]. Tayara et al. [13] introduced a convolutional regression neural network to detect vehicles from satellite imagery. Furthermore, a deep CNN-based detector was proposed [44] to detect multi oriented vehicles from remote sensing imagery. A method combining a deep CNN for feature extraction and a support vector machine (SVM) for object classification was proposed [45]. Ren et al. [46] modified the faster R-CNN detector to detect small objects in remote sensing images. They changed the region proposal network and incorporated context information into the detector. Another modified faster R-CNN detector was proposed by Tang et al. [47]. They used a hyper region proposal network to improve recall and used a cascade boosted classifier to verify candidate regions. This classifier can reduce false detection by mining hard negative examples.

An SSD-based end-to-end airplane detector with transfer learning was proposed, where, the authors used a limited number of airplane images for training [48]. They also proposed a method to solve the input size restrictions by dividing a large image into smaller tiles. Then they detected objects on smaller tiles and finally, mapped each image tile to the original image. They showed that their method performed better than the SSD model. In [49], the authors showed that finding a suitable

parameter setting helped to boost the object detection performance of convolutional neural networks on remote sensing imagery. They used YOLO [11] as object detector to optimize the parameters and infer the results.

In [3], the authors detected conifer seedlings along recovering seismic lines from drone imagery. They used a dataset from different seasons and used faster R-CNN to infer the detection accuracy. There is another work [50] related to plant detection, where authors detected palm trees from satellite imagery using sliding window techniques and an optimized convolutional neural network.

Some works produced excellent results in detecting small objects. Lin et al. [51] proposed feature pyramid networks, which is a top-down architecture with lateral connections. The architecture could build high-level semantic feature maps at all scales. These feature maps boosted the object detection performance, especially for small object detection, when used as a feature extractor for faster R-CNN. Inspired by the receptive fields in human visual systems, Liu et al. [52] proposed a receptive field block (RFB) module that used the relationship between the size and eccentricity of receptive fields to enhance the feature discrimination and robustness. Hence, the module increased the detection performance of objects with various sizes when used as the replacement of the top convolutional layers of SSD.

A one-stage detector called single-shot refinement neural network (RefineDet) [53] was proposed to increase the detection accuracy and also enhance the inference speed. The detector worked well for small object detection. RefineDet used two modules in its architecture: an anchor refinement module to remove negative anchors and an object detection module that took refined anchors as the input. The refinement helped to detect small objects more efficiently than previous methods. In [54], feature fusion SSD (FSSD) was proposed where features from different layers with different scales were concatenated together, and then some downsampling blocks were used to generate new feature pyramids. Finally, the features were fed to multibox detector for prediction. The feature fusion in FSSD increased the detection performance for both large and small objects. Zhu et al. [55] trained single-shot object detectors from scratch and obtained state-of-the-art performance on various benchmark datasets. They removed the first downsampling layer of SSD and introduced root block (with modified convolutional filters) to exploit more local information from an image. Therefore, the detector was able to extract powerful features for small object detection.

All of the aforementioned works were proposed for natural images. A method related to small object detection on remote sensing imagery was proposed by Yang et al. [56]. They used modified faster R-CNN to detect both large and small objects. They proposed rotation dense feature pyramid networks (R-DFPN), and the use of this network helped to improve the detection performance of small objects.

There is an excellent review paper by Zhao et al. [57], where the authors showed a thorough review of object detectors and also showed the advantages and disadvantages of different object detectors. The effect of object size was also discussed in the paper. Another survey paper about object detection in remote sensing images by Li et al. [58] showed review and comparison of different methods.

### 2.3. Super-resolution along with Object Detection

The positive effects of SR on object detection tasks was discussed in [5] where the authors used remote sensing datasets for their experiments. Simultaneous CNN-based image enhancement with object detection using single-shot multibox detector (SSD) [10] was done in [59]. Haris et al. [60] proposed a GAN-based generator to generate a HR image from a LR image and then used a multi-task network as a discriminator and also for localization and classification of objects. These works were done on natural images, and LR and HR image pairs were required. In another work [12], a method using simultaneous super-resolution with object detection on satellite imagery was proposed. The SR network in this approach was inspired by the cycle-consistent adversarial network [61]. A modified faster R-CNN architecture was used to detect vehicles from enhanced images produced by the SR network.

3. Method

In this paper, we aim to improve the detection performance of small objects on remote sensing imagery. Towards this goal, we propose an end-to-end network architecture that consists of two modules: A GAN based SR network and a detector network. The whole network is trained in an end-to-end manner and HR and LR image pairs are needed for training.

The SR network has three components: generator (G), discriminator (D_{Ra}), and edge-enhancement network (EEN). Our method uses end-to-end training as the gradient of the detection loss from the dectector is backpropagated into the generator. Therefore, the detector also works like a discriminator and encourages the generator G to generate realistic images similar to the ground truth. Our entire network structure can also be divided into two parts: A generator consisting of the EEN and a discriminator, which includes the D_{Ra} and the detector network. In figure 2, we show the role of the detector as a discriminator.

The generator G generates intermediate super-resolution (ISR) images, and then final SR images are generated after applying the EEN network. The discriminator (D_{Ra}) discriminates between ground truth (GT) HR images and ISR. The inverted gradients of D_{Ra} are backpropagated into the generator G in order to create SR images allowing for accurate object detection. Edge information is extracted from ISR, and the EEN network enhances these edges. Afterwards, the enhanced edges are again added to the ISR after subtracting the original edges extracted by the Laplacian operator and we get the output SR images with enhanced edges. Finally, we detect objects from the SR images using the detector network.
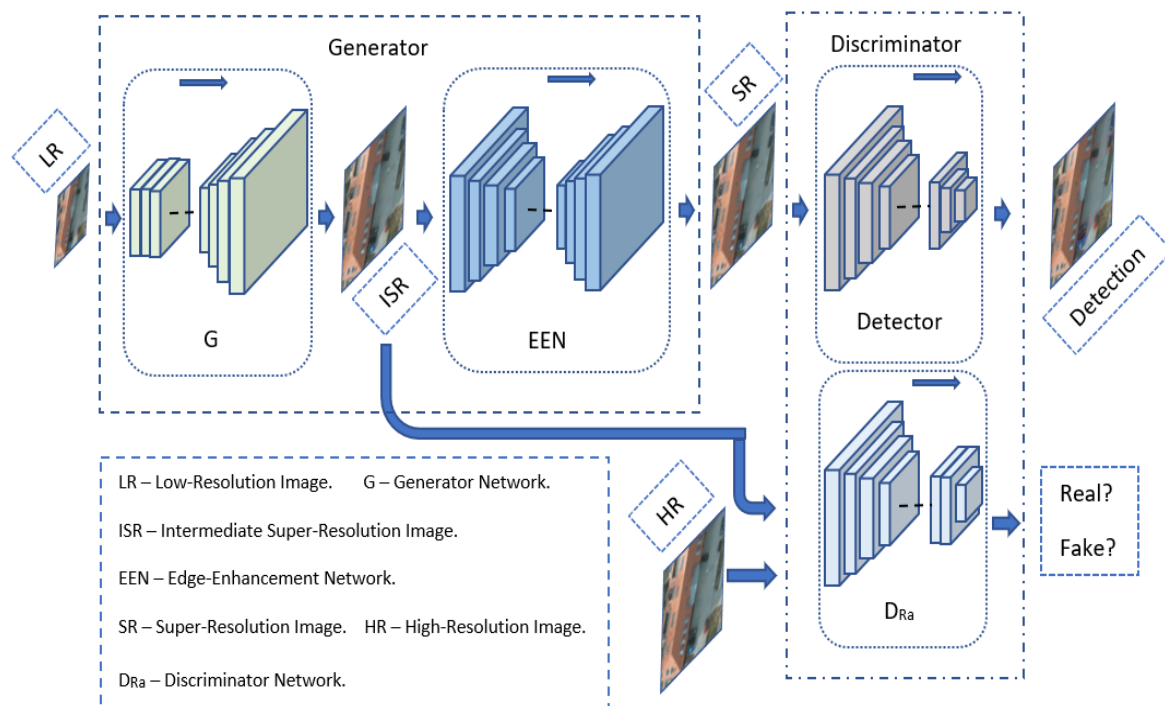

Figure 2. Overall network architecture with a generator and a discriminator module.

We use two different loss functions for EEN: one compares the difference between SR and ground truth images, and the other compares the difference between the extracted edge from ISR and ground truth. We also use the VGG19 [62] network for feature extraction that is used for perceptual loss [21]. Hence, it generates more realistic images with more accurate edge information. We divide the whole pipeline as a generator, and a discriminator, and these two components are elaborated in the following.
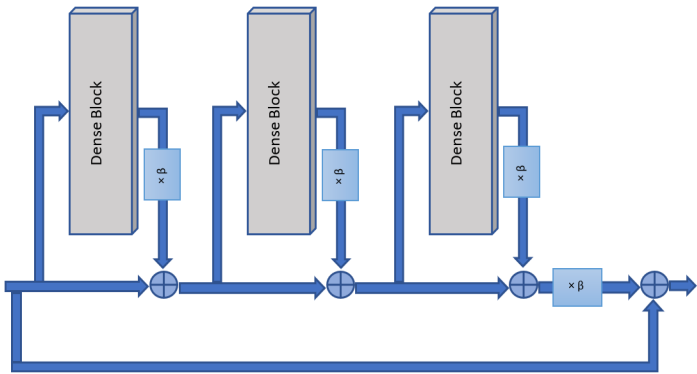
3.1. Generator

Our generator consists of a generator network G and an edge-enhancement network EEN. In this section, we describe the architectures of both networks and the corresponding loss function.

3.1.1. Generator Network *G*
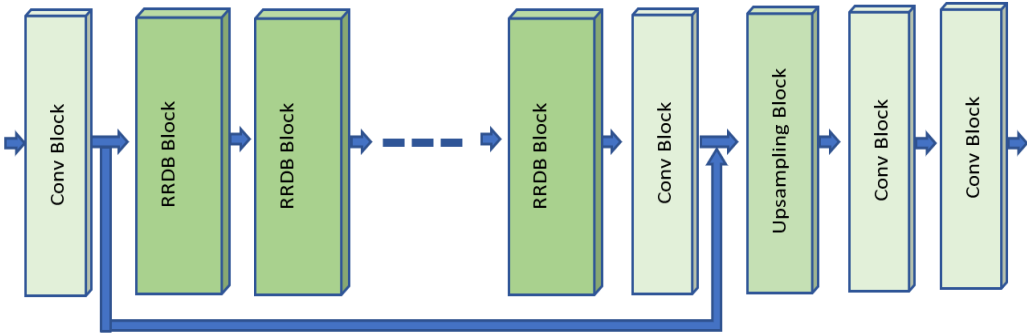
Figure 3. Generator *G* with RRDB (residual-in-residual dense blocks), convolutional and upsampling blocks.

We use the generator architecture from ESRGAN [21], where all batch normalization (BN) layers are removed, and RRDB is used. The overall architecture of generator *G* is shown in figure 3, and the RRDB is depicted in figure 4.

(a) RRDB from generator.
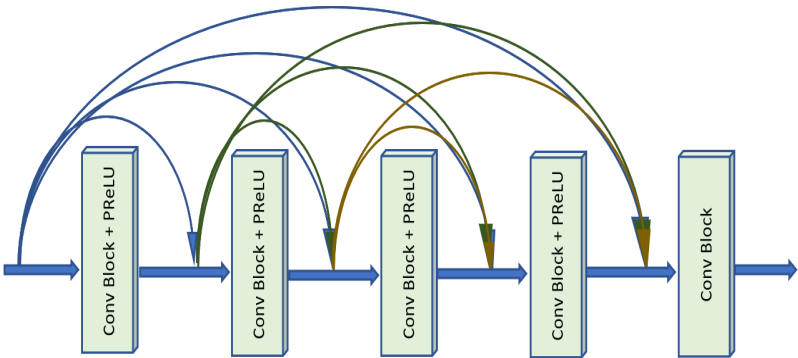
(b) Dense block from RRDB.

Figure 4. Internal diagram of RRDB (residual-in-residual dense blocks).

Inspired by the architecture of ESRGAN, we remove BN layers to increase the performance of the generator G and to reduce the computational complexity. The authors of ESRGAN also state that the BN layers tend to introduce unpleasant artifacts and limit the generalization ability of the generator when the statistics of training and testing datasets differ significantly.

We use RRDB as the basic blocks of the generator network G that uses a multi-level residual network with dense connections. Those dense connections increase network capacity, and we also use residual scaling to prevent unstable conditions during the training phase [21]. We use the parametric rectified linear unit (PReLU) [63] for the dense blocks to learn the parameter with the other neural network parameters. As discriminator (D_{Ra}), we employ a relativistic average discriminator similar to the work represented in [21].

In equation 1 and 2, the relativistic average discriminator is formulated for our architecture. Our generator G depends on the discriminator D_{Ra}, and hence we briefly discuss the discriminator D_{Ra} here and then, describe all details in section 3.2. The discriminator predicts the probability that a real image (I_{HR}) is relatively more realistic than a generated intermediate image (I_{ISR}).

$$D_{Ra}(I_{HR}, I_{ISR}) = \sigma(C(I_{HR}) - \mathbb{E}_{I_{ISR}}[C(I_{ISR})]) \rightarrow 1 \quad \text{More Realistic than fake data?} \tag{1}$$

$$D_{Ra}(I_{ISR}, I_{HR}) = \sigma(C(I_{ISR}) - \mathbb{E}_{I_{HR}}[C(I_{HR})]) \rightarrow 0 \quad \text{Less realistic than real data?} \tag{2}$$

In equation 1 and 2, σ, $C(\cdot)$ and $\mathbb{E}_{I_{ISR}}$ represents the sigmoid function, discriminator output and operation of calculating mean for all generated intermediate images in a mini-batch. The generated intermediate images are created by the generator where $I_{ISR} = G(I_{LR})$. It is evident from equation 3 that the adversarial loss of the generator contains both I_{HR} and I_{ISR} and hence, it benefits from the gradients of generated and ground truth images during the training process. The discriminator loss is depicted in equation 4.

$$L_G^{Ra} = -\mathbb{E}_{I_{HR}}[\log(1 - D_{Ra}(I_{HR}, I_{ISR}))] - \mathbb{E}_{I_{ISR}}[\log(D_{Ra}(I_{ISR}, I_{HR}))] \tag{3}$$

$$L_D^{Ra} = -\mathbb{E}_{I_{HR}}[\log(D_{Ra}(I_{HR}, I_{ISR}))] - \mathbb{E}_{I_{ISR}}[\log(1 - D_{Ra}(I_{ISR}, I_{HR}))] \tag{4}$$

We use two more losses for generator G: one is perceptual loss (L_{percep}), and another is content loss (L_1) [21]. The perceptual loss is calculated using the feature map ($vgg_{fea}(\cdot)$) before the activation layers of a fine-tuned VGG19 [62] network, and the content loss calculates the 1-norm distance between I_{ISR} and I_{HR}. Perceptual loss and content loss is shown in equation 5 and equation 6.

$$L_{percep} = \mathbb{E}_{I_{LR}} ||vgg_{fea}(G(I_{LR}) - vgg_{fea}(I_{HR})||_1 \tag{5}$$

$$L_1 = \mathbb{E}_{I_{LR}} ||G(I_{LR}) - I_{HR}||_1 \tag{6}$$

3.1.2. Edge-Enhancement Network EEN

The EEN network removes noise and enhances the extracted edges from an image. An overview of the network is depicted in figure 5. In the beginning, Laplacian operator [28] is used to extract edges from the input image. After the edge information is extracted, it is passed through convolutional, RRDB, and upsampling blocks. There is a mask branch with sigmoid activation to remove edge noise as described in [22]. Finally, the enhanced edges are added to the input images where the edges extracted by the Laplacian operator were subtracted.

The EEN network is similar to the edge-enhancement subnetwork proposed in [22] with two improvements. First, we replace the dense blocks with RRDB. The RRDB shows improved performance according to ESRGAN [21]. Hence, we replace the dense block for improved performance of the EEN network. Secondly, we introduce a new loss term to improve the reconstruction of the edge information.
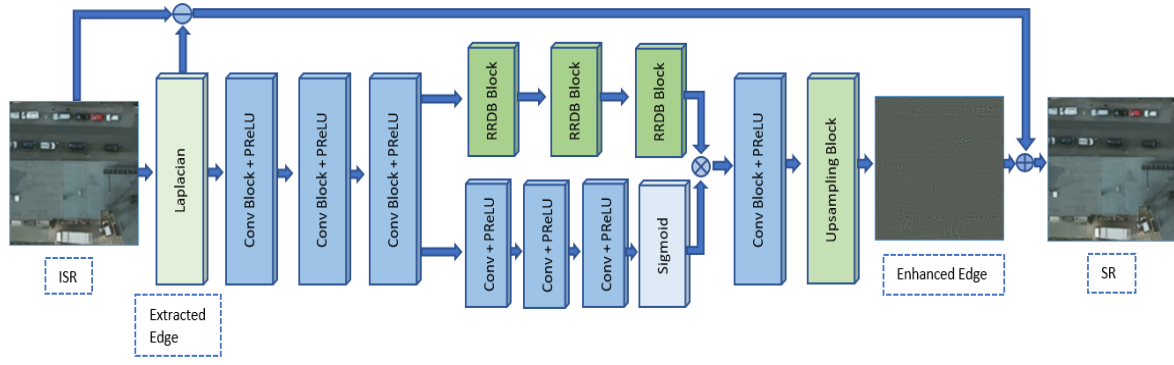
Figure 5. Edge-enhancement network where input is an ISR (intermediate super-resolution) image and output is a SR (super-resolution) image.

In [22], authors extracted the edge information from I_{ISR} and enhanced the edges using an edge-enhancement subnetwork which is afterwards added to the edge-subtracted I_{ISR}. To train the network, [22] proposed to use Charbonnier loss [34] between the I_{ISR} and I_{HR}. This function is called consistency loss for images (L_{img_cst}) and helps to get visually pleasant outputs with good edge information. However, sometimes the edges of some objects are distorted and produce some noises and consequently, do not give good edge information. Therefore, we introduce a consistency loss for the edges (L_{edge_cst}) as well. To compute L_{edge_cst} we evaluate the Charbonnier loss between the extracted edges (I_{edge_SR}) from I_{SR} and the extracted edges (I_{edge_HR}) from I_{HR}. The two consistency losses are depicted in equation 7 and equation 8 where $\rho(\cdot)$ is the Charbonnier penalty function [64]. The total consistency loss is finally calculated for both images and edges by summing up the individual loss. The loss of our EEN is shown in equation 9.

$$L_{img\_cst} = \mathbb{E}_{I_{SR}}[\rho(I_{HR} - I_{SR})] \tag{7}$$

$$L_{edge\_cst} = \mathbb{E}_{I_{edge\_SR}}[\rho(I_{edge\_HR} - I_{edge\_SR})] \tag{8}$$

$$L_{een} = L_{img\_cst} + L_{edge\_cst} \tag{9}$$

Finally, we get the overall loss for the generator module by adding the losses of the generator G and the EEN network. The overall loss for the generator module is shown in equation 10 where λ_1, λ_2, λ_3, and λ_4 are the weight parameters to balance different loss components. We empirically set the values as $\lambda_1 = 1$, $\lambda_2 = .001$, $\lambda_3 = .01$, and $\lambda_4 = 5$.

$$L_{G\_een} = \lambda_1 L_{percep} + \lambda_2 L_G^{Ra} + \lambda_3 L_1 + \lambda_4 L_{een} \tag{10}$$

3.2. Discriminator

As described in the previous section, we use the relativistic discriminator D_{Ra} for training the generator G. The architecture of the discriminator is taken from ESRGAN [21] which employs the VGG-19 [62] architecture. We use Faster R-CNN [8] and SSD [10] for our detector networks. The discriminator (D_{Ra}) and the detector network jointly act as discriminator for the generator module. We briefly describe these two detectors in the next two sections.

3.2.1. Faster R-CNN

The Faster R-CNN [8] is a two-stage object detector and contains two networks: a region proposal network (RPN) to generate region proposals from an image and another network to detect objects from these proposals. In addition, the second network also tries to fit the bounding boxes around the detected objects.

The task of the RPN is to return image regions that have a high probability of containing an object. The RPN network uses a backbone network such as VGG [62], ResNet, or ResNet with feature pyramid network [51]. These networks are used as feature extractors, and different types of feature extractors can be chosen based on their performance on public datasets. We use ResNet-50-FPN [51] as a backbone network for our faster R-CNN. We use this network because it displayed a higher precision than VGG-19 and ResNet-50 without FPN (especially for small object detection) [51]. Even though the use of a larger network might lead to a further performance improvement, we chose ResNet-50-FPN due to its comparably moderate hardware requirements and more efficient convergence times.

After the RPN, there are two branches for detection: a classifier and a regressor. The classification branch is responsible for classifying a proposal to a specific object, and the regression branch finds the accurate bounding box of the object. In our case, both datasets contain objects with only one class, and therefore, our classifier infers only two classes: the background class and the object class.

3.2.2. SSD

The SSD [10] is a single-shot multibox detector that detects objects in a single stage. Here, single-stage means that classification and localization are done in a single forward pass through the network. Like Faster R-CNN, SSD also has a feature extractor network, and different types of networks can be used. To serve the primary purpose of SSD, which is speed, we use VGG-16 [62] as a feature extractor network. After this network, SSD has several convolutional feature layers of decreasing sizes. This representation can seem like a pyramid representation of images at different scales. Therefore, the detection of objects happens in every layer, and finally, we get the object detection output as class values and coordinates of bounding boxes.

3.2.3. Loss of the discriminator

The relativistic discriminator loss (L_D^{Ra}) is already described in the previous section and depicted in equation 4. This loss is added to the detector loss to get the final discriminator loss.

Both Faster R-CNN and SSD have similar regression/localization losses but different classification losses. For regression/localization, both use smooth L_1 [8] loss between detected and ground truth bounding box coordinates (t_*). Classification (L_{cls_frcnn}) and regression loss (L_{reg_frcnn}) and overall loss (L_{det_frcnn}) of Faster R-CNN are given in the following:

$$L_{cls\_frcnn} = \mathbb{E}_{I_{LR}}[-\log(Det_{cls\_frcnn}(G_{G\_een}(I_{LR})))] \tag{11}$$

$$L_{reg\_frcnn} = \mathbb{E}_{I_{LR}}[smooth_{L1}(Det_{reg\_frcnn}(G_{G\_een}(I_{LR})), t_*)] \tag{12}$$

$$L_{det\_frcnn} = L_{cls\_frcnn} + \lambda L_{reg\_frcnn} \tag{13}$$

Here, λ is used to balance the losses, and it is set to 1 empirically. Det_{cls_frcnn} and Det_{reg_frcnn} are the classifier and regressor for the Faster R-CNN. Classification (L_{cls_ssd}), regression loss (L_{reg_ssd}) and overall loss (L_{det_ssd}) of SSD are as following:

$$L_{cls\_ssd} = \mathbb{E}_{I_{LR}}[-\log(softmax(Det_{cls\_ssd}(G_{G\_een}(I_{LR}))))] \tag{14}$$

$$L_{reg\_ssd} = \mathbb{E}_{I_{LR}}[smooth_{L1}(Det_{reg\_ssd}(G_{G\_een}(I_{LR})), t_*)] \tag{15}$$

$$L_{det\_ssd} = L_{cls\_ssd} + \alpha L_{reg\_ssd} \tag{16}$$

Here, α is used to balance the losses, and it is set to 1 empirically. Det_{cls_ssd} and Det_{reg_ssd} are the classifier and regressor for the SSD.

3.3. Training

Our architecture can be trained in separate steps or jointly in an end-to-end way. We discuss the details of these two types of training in the next two sections.

3.3.1. Separate Training

In separate training, we train the SR network (generator module and discriminator D_{Ra}) and the detector separately. Detector loss is not backpropagated to the generator module. Therefore, the generator is not aware of the detector and thus, it only gets feedback from the discriminator D_{Ra}. For example, in equation 11, no error is backpropagated to the G_{G_een} network (the network is detached during the calculation of the detector loss) while calculating the loss L_{cls_frcnn}.

3.3.2. End-to-End Training

In end-to-end training, we train the whole architecture end-to-end that means the detector loss is backpropagated to the generator module. Therefore, the generator module revceives gradients from both detector and discriminator D_{Ra}. We get the final discriminator loss (L_{D_det}) as following:

$$L_{D\_det} = L_D^{Ra} + \eta L_{det} \tag{17}$$

Here, η is the parameter to balance the contribution of the detector loss and we empirically set it to 1. Finally, we get an overall loss ($L_{overall}$) for our architecture as follows.

$$L_{overall} = L_{G\_een} + L_{D\_det} \tag{18}$$

4. Experiments

As mentioned above, we trained our architecture separately and in an end-to-end manner. For separate training, we first trained the SR network until convergence and then trained the detector networks based on the SR images. For end-to-end training, we also employed separate training as pre-training step for weight initialization. Afterwards SR and object detection networks were jointly trained, i.e., the gradients from the the object detector were propagated into the generator network.

In the training process, the learning rate was set to 0.0001 and halved after every 50k iterations. The batch size was set to 5. We used Adam [65] as optimizer with $\beta_1 = 0.9$, $\beta_2 = 0.999$ and updated the whole architecture weights until convergence. We used 23 RRDB blocks for the generator G and 5 RRDB blocks for the EEN network. We implemented our architecture with the PyTorch framework [66] and trained/tested using two NVIDIA Titan X GPUs. The end-to-end training with COWC took 96 hours for 200 epochs. The average inference speed using faster R-CNN was approximately 4 images/second and 7 images/second for SSD. Our implementation can be found in GitHub [67].

4.1. Datasets

4.1.1. Cars Overhead with Context Dataset

Cars overhead with context (COWC) dataset [31] contains 15 cm (one pixel cover 15 cm distance at ground level) satellite images from six different regions. The dataset contains a large number of unique cars and covers regions from Toronto in Canada, Selwyn in New Zealand, Potsdam and Vaihingen in Germany, Columbus and Utah in the United States. Out of these six regions, we used the dataset from Toronto and Potsdam. Therefore, when we refer to the COWC dataset, we refer to the dataset from these two regions. There are 12651 cars in our selected dataset. The dataset contains only RGB images, and we used these images for training and testing.

We used 256-by-256 image tiles, and every image tile contains at least one car. The average length of a car was between 24 to 48 pixels, and the width was between 10 to 20 pixels. Therefore, the area of a car was between 240 to 960 pixels, which can be considered as a small object relative to the other large satellite objects. We used bi-cubic downsampling to generate LR images from the COWC dataset. The downscale factor was 4x, and therefore, we had 64 pixels to 64 pixels size for LR images. We had a text file associated with each image tile containing the coordinates of the bounding box for each car.

Figure 6. COWC (car overhead with context) dataset: LR-HR (low-resolution and high-resolution) image pairs are shown in (a) and (b) and GT (ground truth) images with bounding boxes for cars are in (c).

Our experiments considered the dataset having only one class, car, and did not consider any other type of object. Figure 6 shows examples from the COWC dataset. We experimented with a total of 3340 tiles for training and testing. Our train/test split was 80%/20%, and the training set was further divided into a training and a validation set by an 80% to 20% ratio. We trained our end-to-end architecture with an augmented training dataset with random horizontal flips and ninety-degree rotations.

4.1.2. Oil and Gas Storage Tank Dataset

The oil and gas storage tank (OGST) dataset has been complied in Alberta Geological Survey (AGS) [68], a branch of the Alberta Energy Regulatory (AER) [35]. AGS provides geoscience information and support to AER's regulatory functions on energy developments to be carried out in a manner to ensure public and environmental safety. To assist AER with sustainable land management and compliance assurance [69], AGS is utilizing remote sensing imagery for identifying the number of oil and gas storage tanks inside well pad footprints in Alberta.

While the SPOT-6 satellite imagery at 1.5 m pixel resolution provided by the AGS has sufficient quality and details for many regulatory functions, it is difficult to detect small objects within well pads, e.g., oil and gas storage tanks with ordinary object detection methods. The diameter of a typical storage tank is about 3 m and their placements are usually vertical and side-by-side with less than 2 m. To train our architecture for this use-case, we needed a dataset for providing pairs of low and high-resolution images. Therefore, we have created the OGST dataset using free imagery from the Bing map [70].

| (a) LR image | (b) HR image | (c) GT image |

Figure 7. OGST (oil and gas storage tank) dataset: LR-HR (low-resolution and high-resolution) image pairs are shown in (a) and (b) and GT (ground truth) images with bounding boxes for oil and gas storage tanks are in (c).

The OGST dataset contains 30 cm resolution remote sensing images (RGB) from the Cold Lake Oil Sands region of Alberta, Canada where there is a high level of oil and gas activities and concentration of well pad footprints. The dataset contains 1671 oil and gas storage tanks from this area.

We used 512-by-512 image tiles, and there was no image without any oil and gas storage tank in our experiment. The average area covered by an individual tank was between 800 to 1600 pixels. Some industrial tanks were large, but most of the tanks covered small regions on the imagery. We downscaled the HR images using bi-cubic downsampling with the factor of 4x, and therefore, we got a LR tile of size 128-by-128 pixels. Every image tile was associated with a text file containing the coordinates of the bounding boxes for the tanks on a tile. We have showed examples from the OGST dataset in figure 7.

As with the COWC dataset, our experiments considered one unique class here, tank, and we had a total of 760 tiles for training and testing. We used a 90%/10% split for our train/test data. The training data was further divided by 90%/10% for the train/validation split. The percentage of training data was higher here compared to the previous dataset to increase the training data because of the smaller size of the dataset. The dataset is available at [67].

4.2. Evaluation Metrics for Detection

We obtained our detection output as bounding boxes with associated classes. To evaluate our results, we used average precision (AP), and calculated intersection over union (IoU), precision, and recall for obtaining AP.

We denote the set of correctly detected objects as true positives (TP) and the set of falsely detected objects of false positives (FP). The precision is now the ratio between the number of TPs relative to all predicted objects:

$$precision = \frac{|TP|}{|TP| + |FP|} \tag{19}$$

We denote the set of objects which are not detected by the detector as false negatives (FN). Then, the recall is defined as the ratio of detected objects (TP) relative to the number of all objects in the data set:

$$Recall = \frac{|TP|}{|TP| + |FN|} \tag{20}$$

To measure the localization error of predicted bounding boxes, IoU computes the overlap between two bounding boxes: the detected and the ground truth box. If we take all the boxes that have an IoU $\geq \tau$ as TP and consider all other detections as FP, then we get the precision at τ IoU. If we now vary τ from 0.5 to 0.95 IoU with a step size of 0.05, we receive ten different precision values which can be combined into the average precision (AP) at IoU=0.5:0.95 [8]. Let us note that in the case of multi-class classification, we would need to compute the AP for object each class separately. To receive a single performance measure for object detection, the mean AP (mAP) is computed which is the most common performance measure for object detection quality.

In this paper, both of our datasets only contain single class, and hence, we used AP as our evaluation metric. We mainly showed the results of AP at IoU=0.5:0.95 as our method performed increasingly better compared to other models when we increased the IoU values for AP calculation. We show this trend in section 4.3.4.

4.3. Results

4.3.1. Detection without Super-Resolution

We ran the two detectors to document the object detection performance on both LR and HR images. We used SSD with vgg16 [62] network and Faster R-CNN (FRCNN) with ResNet-50-FPN [51] detector. We trained the two models with both HR and 4x-downscaled LR images. Testing was also done with both HR and LR images.

Table 1. Detection on LR (low-resolution) and HR (high-resolution) images without using super-resolution. Detectors are trained with both LR and HR images and AP (average precision) values are calculated using 10 different IoUs (intersection over union).

Model	Training image resolution - Test image resolution	COWC dataset (Test Results) (AP at IoU=0.5:0.95) (single class - 15 cm)	OGST dataset (Test Results) (AP at IoU=0.5:0.95) (single class - 30 cm)
SSD	LR - LR	61.9%	76.5%
	HR - LR	58%	75.3%
FRCNN	LR - LR	64%	77.3%
	HR - LR	59.7%	75%
SSD-RFB	LR - LR	63.1%	76.7%
SSD	HR - HR	94.1%	82.5%
FRCNN	HR - HR	98%	84.9%

In table 1, we show the results of the detection performance of the detectors with different train/test combinations. When we only used LR images for both training and testing, we observed 64% AP for Faster R-CNN. When training on HR images and testing with LR images, the accuracy dropped for both detectors. We also added detection results (using LR images for training/testing) for both the datasets using SSD with RFB modules (SSD-RFB) [52], where accuracy slightly increased from the base SSD.

The last two rows in table 1 depict the accuracy of both detectors when training and testing on HR images. We have achieved up to 98% AP with the Faster R-CNN detector. This, shows the large impact of the resolution to the object detection quality and sets a natural upper bound on how close a SR-based method can get when working on LR images. In the next sections, we demonstrate that our approaches considerably improve the detection rate on LR imagery and get astonishingly close to the performance of directly working on HR imagery.

4.3.2. Separate Training with Super-Resolution

In this experiment, we created 4x upsampled images from the LR input images using bicubic upsampling and different SR methods. Let us note that no training was needed for applying bicubic upsampling since it is a parameter free function. We used the SR images as test data for two types of detectors. We compared three GAN architectures for generating SR images, our new EESRGAN architecture, ESRGAN [21] and EEGAN [22]. Each network was trained separately on the training set before the object detector was trained. For the evaluation, we again compared detectors being trained on the SR images from the particular architecture and detectors being directly trained on the HR images.

Table 2. Detection on SR (super-resolution) images with separately trained SR network. Detectors are trained with both SR and HR (high-resolution) images and AP (average precision) values are calculated using 10 different IoUs (intersection over union).

Model	Training image resolution - Test image resolution	COWC Dataset (Test Results) (AP at IoU=0.5:0.95) (single class - 15 cm)	OGST Dataset (Test Results) (AP at IoU=0.5:0.95) (single class - 30 cm)
Bicubic + SSD	SR - SR	72.1%	77.6%
	HR - SR	58.3%	76%
Bicubic + FRCNN	SR - SR	76.8%	78.5%
	HR - SR	61.5%	77.1%
EESRGAN + SSD	SR - SR	86%	80.2%
	HR - SR	83.1%	79.4%
EESRGAN + FRCNN	SR - SR	**93.6%**	**81.4%**
	HR - SR	92.9%	80.6%
ESRGAN + SSD	SR - SR	85.8%	80.2%
	HR - SR	82.5%	78.9%
ESRGAN + FRCNN	SR - SR	92.5%	81.1%
	HR - SR	91.8%	79.3%
EEGAN + SSD	SR - SR	86.1%	79.1%
	HR - SR	83.3%	77.5%
EEGAN + FRCNN	SR - SR	92%	79.9%
	HR - SR	91.1%	77.9%

In table 2, the detection output of the different combinations of SR methods and detectors is shown with the different combinations of train/test pairs. As can be seen, our new EESRGAN architecture displayed the best results already getting close to the detection rates which could be observed when working with HR images only. However, after training EESRGAN can be directly applied to LR imagery where no HR data is available and still achieved very good results. Furthermore, we could observe that other SR methods EEGAN and ESRGAN have already improved the AP considerably when used for preprocessing of LR images. However, for both data sets, EESRGAN have outperformed the other two methods.

4.3.3. End-to-End Training with Super-Resolution

We trained our EESRGAN network and detectors end-to-end for this experiment. The discriminator (D_{Ra}), and the detectors jointly acted as a discriminator for the entire architecture. Detector loss was backpropagated to the SR network, and therefore, the loss contributed to the enhancement of LR images. At training time, LR-HR image pairs were used to train the EEGAN part, and then the generated SR images were sent to the detector for training. At test time, only the LR images were fed to the network. Our architecture first generated a SR image of the LR input before object detection was performed.

Table 3. Detection with end-to-end SR (super-resolution) network. Detectors are trained with SR images and AP (average precision) values are calculated using 10 different IoUs (intersection over union).

| Model	Training image resolution - Test image resolution	COWC Dataset (Test Results) (AP at IoU=0.5:0.95) (single class - 15 cm)	OGST Dataset (Test Results) (AP at IoU=0.5:0.95) (single class - 30 cm)
EESRGAN + SSD	SR - SR	89.3%	81.8%
EESRGAN + FRCNN	SR - SR	**95.5%**	**83.2%**
ESRGAN + SSD	SR - SR	88.5%	81.1%
ESRGAN + FRCNN	SR - SR	93.6%	82%
EEGAN + SSD	SR - SR	88.1%	80.8%
EEGAN + FRCNN	SR - SR	93.1%	81.3%

We also compared our results with different architectures. We used ESRGAN [21] and EEGAN [22] with the detectors for comparison. Table 3 clearly shows that our method delivers superior results compared to others.

4.3.4. AP versus IoU curve

We have calculated the AP values on different IoUs. In figure 8, we plot the AP versus IoU curves for our datasets. The performance of EESRGAN-FRCNN, end-to-end EESRGAN-FRCNN, and FRCNN is shown in the figure. The end-to-end EESRGAN-FRCNN network has performed better than the separately trained network. The difference is most evident for the higher IoUs on the COWC dataset.

Our results indicate excellent performance compared to the highest possible AP values obtained from standalone FRCNN (trained and tested on HR images)

The OGST dataset has displayed less performance variation compared to the COWC dataset. The object size of the OGST dataset is larger than that of the COWC dataset. Therefore, the performance difference was not similar to the COWC dataset when we compared between standalone FRCNN and our method on the OGST dataset. To conclude, training our new architecture in an end-to-end manner has displayed an improvement for both the datasets.

(a) COWC dataset   **(b)** OGST dataset
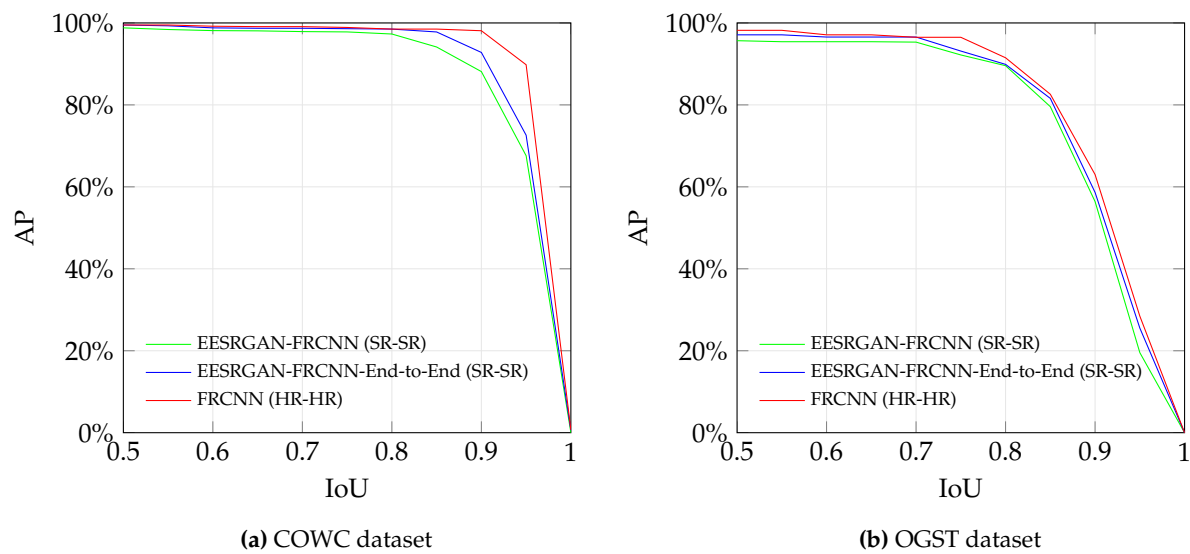
Figure 8. AP-IoU (average precision-intersection over union) curves for the datasets. Plotted results show the detection performance of standalone faster R-CNN on HR (high-resolution) images and our proposed method (with and without end-to-end training) on SR (super-resolution) images.

4.3.5. Precision versus Recall

In figure 9, precision-recall curves are shown for both of our datasets. The precision-recall curve for COWC dataset is depicted in 9a and 9b represents the curve for OGST dataset. For each dataset, we plot the curves for standalone faster R-CNN with LR training/testing images, and our method with/without end-to-end training. We used IoU=0.5 to calculate precision and recall.

**(a)** COWC dataset   **(b)** OGST dataset
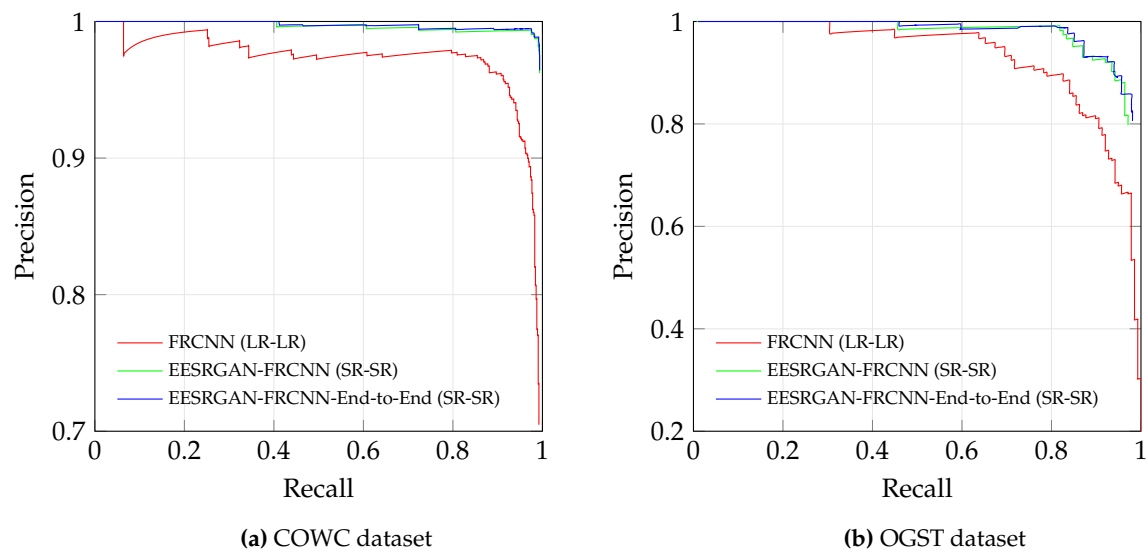
Figure 9. Precision-recall curve for the datasets. Plotted results show the detection performance of standalone faster R-CNN on LR (low-resolution) images and our proposed method (with and without end-to-end training) on SR (super-resolution) images.

The precision-recall curves for both datasets show that our method has higher precision values in higher recall values compared to the standalone faster R-CNN models. Our models with end-to-end training performed better than our models without the end-to-end training. In particular, the end-to-end models have detected more than 99% of the cars with 96% AP in the COWC dataset. For the OGST dataset, our end-to-end models have detected more than 81% of the cars with 97% AP.

4.3.6. Effects of Dataset Size

We trained our architecture with different training set sizes and tested with a fixed test set. In figure 10, we plot the AP values (IoU=0.5:0.95) against different numbers of labeled objects for both of our datasets (training data). We used five different dataset sizes: $\{500, 1000, 3000, 6000, 10000 (cars)\}$ and $\{100, 200, 400, 750, 1491 (tanks)\}$ to train our model with and without the end-to-end setting.

We got the highest AP value of 95.5% with our full COWC training dataset (10000 cars), and we used the same test dataset (1000 cars) for all combinations of the training dataset (with end-to-end setting). We also used another set of 1000 labeled cars for validation. Using 6000 cars, we got an AP value near to the highest AP, as shown with the plot of AP versus dataset size (COWC). The AP value decreased significantly when we used only 3000 labeled cars as training data. We got the lowest AP using only 500 labeled cars, and the trend of AP was further decreasing as depicted in figure 10a. Therefore, we can infer that we needed around 6000 labeled cars to get precision higher than 90% for the COWC dataset. We observed slightly lower AP values for all sizes of COWC datasets when we did not use the end-to-end setting, and we observed higher differences between the two settings (with and without end-to-end) when we used less than 6000 labeled cars.

(a) COWC dataset   **(b)** OGST dataset
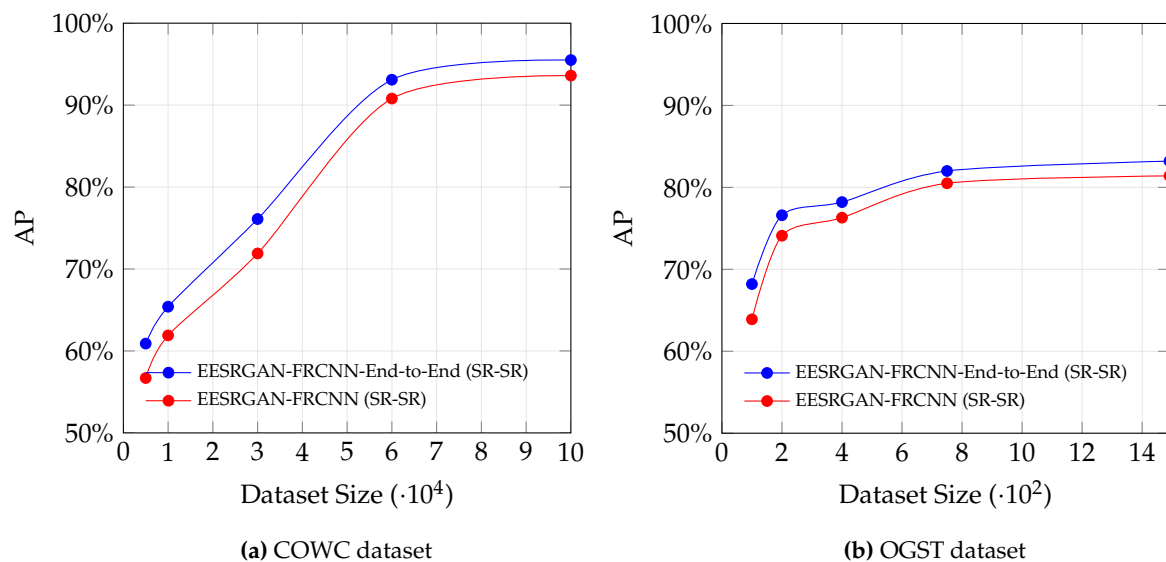
Figure 10. AP (average precision) with varying number of training sets from the datasets. Plotted results show the detection performance of our proposed method (with and without end-to-end training) on SR (super-resolution) images.

The OGST dataset gave 83.2% AP (with end-to-end setting) using the full training dataset (1491 tanks), and we used 100 labeled tanks as test and same amount as validation data for all combinations of the training dataset. We got high AP values with 50% of our full training dataset as depicted in 10b. AP values dropped below 80% when we further decreased the training data. Similar to the COWC datasets, we also got comparatively lower AP values for all sizes of OGST datasets. We observed slightly higher differences between the two settings (with and without end-to-end) when the dataset consisted of less than 400 labeled tanks, as shown in the plot of AP versus dataset size (OGST dataset).

We used 90% of the OGST dataset for training while we used the 80% of the COWC dataset for the same purpose. The accuracy of the testing data (OGST) slightly increased when we added more training data, as depicted in figure 10b. Therefore, we used a larger percentage of training data for the OGST dataset than for the COWC dataset, and it slightly helped to improve the relatively low accuracy of the OGST test data.

4.3.7. Enhancement and Detection

In figure 11, we have shown our input LR images, corresponding generated SR image, enhanced edge information and final detection. The image enhancement has helped the detectors to get high AP values and also makes the images visually good enough to identify the objects easily. It is evident from the figure that the visual quality of the generated SR images is quite good compared to the corresponding LR images, and the FRCNN detector has detected most of the objects correctly.

(a) Input LR image   (b) Generated SR image   (c) Enhanced edge   (c) Detection
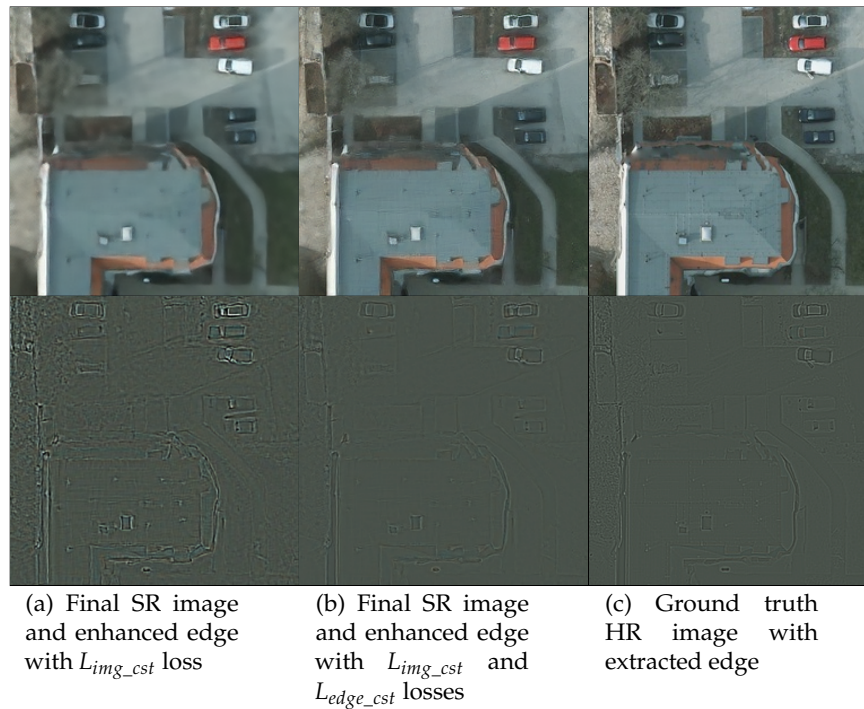
Figure 11. Examples of SR (super-resolution) images that are generated from input LR (low-resolution) images are shown in (a) and (b). The enhanced edges and detection results are shown in (c) and (d).

4.3.8. Effects of Edge Consistency Loss (L_{edge_cst})

In EEGAN [22], only image consistency loss (L_{img_cst}) was used for enhancing the edge information. This loss generated edge information with noise, and as a result, the final SR images became blurry. The blurry output with noisy edge using only L_{img_cst} loss is shown in figure 12a. The blurry final images gave lower detection accuracy compared to sharp outputs.

Therefore, we have introduced edge consistency loss (L_{edge_cst}) in addition to L_{img_cst} loss that gives noise-free enhanced edge information similar to the edge extracted from ground truth images and the effects of the L_{edge_cst} loss is shown in figure 12b. The ground truth HR image with extracted edge is depicted in figure 12c.

(a) Final SR image and enhanced edge with L_{img_cst} loss

(b) Final SR image and enhanced edge with L_{img_cst} and L_{edge_cst} losses

(c) Ground truth HR image with extracted edge

Figure 12. Effects of edge consistency loss (L_{edge_cst}) on final SR (super-resolution) images and enhanced edges compared to the extracted edges from HR (high-resolution) images.

5. Discussion

The detection results of our method presented in the previous section have indicated that our end-to-end SR-detector network improved detection accuracy compared to several other methods. Our method outperformed the standalone state-of-the-art methods such as SSD or faster R-CNN when implemented in low-resolution remote sensing imagery. We used EESRGAN, EEGAN, and ESRGAN as the SR network with the detectors. We showed that our EESRGAN with the detectors performed better than the other methods and the edge-enhancement helped to improve the detection accuracy. The AP improvement was higher in high IoUs and not so much in the lower IoUs. We have also showed that the precision increased with higher resolution. The improvement of AP values for the OGST dataset was lower than that for the COWC dataset because the area covered by a tank was slightly bigger than that of a car, and tanks sizes and colors were less diverse than the cars.

Our experimental results indicated that AP values of the output could be improved slightly with the increase of training data. The results also demonstrated that we could use less training data for both the datasets to get a similar level of accuracy that we obtained from our total training data.

The faster R-CNN detector gave us the best result, but it took a longer time than an SSD detector. If we need detection results from a vast area, then SSD would be the right choice sacrificing some amount of accuracy.

We had large numbers of cars from different regions in the COWC dataset, and we obtained high AP values using different IoUs. On the other hand, the OGST dataset needed more data to get a general detection result because we used data from a specific area and for a specific season and this was one of the limitations of our experiment. Most likely, more data from different regions and seasons would make our method more robust for the use-case of oil and gas storage tank detection. Another limitation of our experiment was that we showed the performance of the datasets that contain only one class with less variation. We would be looking forward to exploring the performance of our method on a broader range of object types and landscapes from different satellite datasets.

We have used LR-HR image pairs to train our architecture, and the LR images were generated artificially from the HR counterparts. To our knowledge, there is no suitable public satellite dataset that contains both real HR and real LR image pairs and ground truth bounding boxes for detecting

small objects. Therefore, we have created the LR images which do not precisely correspond to true LR images. However, improvement of resolution through deep learning always improved object detection performance on remote sensing images (for both artificial and real low-resolution images), as discussed in the introduction and related works section of this paper [5]. Impressive works [61,71] exist in literature to create realistic LR images from HR images. For future work, we are looking forward to exploring the works to create more accurate LR images for training.

6. Conclusions

In this paper, we propose an end-to-end architecture that takes LR satellite imagery as input and gives object detection results as outputs. Our architecture contains a SR network and a detector network. We have used a different combination of SR systems and detectors to compare the AP values for detection using two different datasets. Our experimental results show that the proposed SR network with faster R-CNN has yielded the best results for small objects on satellite imagery. However, we need to add more diverse training data in the OGST dataset to make our model robust in detecting oil and gas storage tanks. We also need to explore diverse datasets and the techniques to create more realistic LR images. In conclusion, our method has combined different strategies to provide a better solution to the task of small-object detection on LR imagery.

Author Contributions: Conceptualization, J.R., N.R. and M.S.; methodology, J.R., N.R. and M.S.; software, J.R.; validation, J.R.; formal analysis, J.R.; investigation, J.R.; resources, N.R.; data curation, J.R., S.C. and D.C.; writing–original draft preparation, J.R.; writing–review and editing, J.R., N.R., M.S., S.C. and D.C.; visualization, J.R.; supervision, N.R. and M.S.; project administration, N.R.; funding acquisition, N.R., S.C. and D.C.

Conflicts of Interest: The authors declare no conflict of interest.

References

1. Colomina, I.; Molina, P. Unmanned aerial systems for photogrammetry and remote sensing: A review. *ISPRS Journal of photogrammetry and remote sensing* **2014**, *92*, 79–97.
2. Zhang, F.; Du, B.; Zhang, L.; Xu, M. Weakly supervised learning based on coupled convolutional neural networks for aircraft detection. *IEEE Transactions on Geoscience and Remote Sensing* **2016**, *54*, 5553–5563.
3. Fromm, M.; Schubert, M.; Castilla, G.; Linke, J.; McDermid, G. Automated Detection of Conifer Seedlings in Drone Imagery Using Convolutional Neural Networks. *Remote Sensing* **2019**, *11*, 2585.
4. Pang, J.; Li, C.; Shi, J.; Xu, Z.; Feng, H. $\mathcal{R}^2$-CNN: Fast Tiny Object Detection in Large-Scale Remote Sensing Images. *IEEE Transactions on Geoscience and Remote Sensing* **2019**, *57*, 5512–5524. doi:10.1109/TGRS.2019.2899955.
5. Shermeyer, J.; Van Etten, A. The effects of super-resolution on object detection performance in satellite imagery. Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops, 2019, pp. 0–10.
6. Russakovsky, O.; Deng, J.; Su, H.; Krause, J.; Satheesh, S.; Ma, S.; Huang, Z.; Karpathy, A.; Khosla, A.; Bernstein, M.; Berg, A.C.; Fei-Fei, L. ImageNet Large Scale Visual Recognition Challenge, 2014, [arXiv:cs.CV/1409.0575].
7. Lin, T.Y.; Maire, M.; Belongie, S.; Hays, J.; Perona, P.; Ramanan, D.; Dollár, P.; Zitnick, C.L. Microsoft coco: Common objects in context. European conference on computer vision. Springer, 2014, pp. 740–755.
8. Ren, S.; He, K.; Girshick, R.; Sun, J. Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **2017**, *39*, 1137–1149. doi:10.1109/tpami.2016.2577031.
9. Lin, T.Y.; Goyal, P.; Girshick, R.; He, K.; Dollar, P. Focal Loss for Dense Object Detection. *2017 IEEE International Conference on Computer Vision (ICCV)* **2017**. doi:10.1109/iccv.2017.324.
10. Liu, W.; Anguelov, D.; Erhan, D.; Szegedy, C.; Reed, S.; Fu, C.Y.; Berg, A.C. SSD: Single Shot MultiBox Detector. *Lecture Notes in Computer Science* **2016**, p. 21–37. doi:10.1007/978-3-319-46448-0_2.
11. Redmon, J.; Divvala, S.; Girshick, R.; Farhadi, A. You Only Look Once: Unified, Real-Time Object Detection. *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* **2016**. doi:10.1109/cvpr.2016.91.

12. Ji, H.; Gao, Z.; Mei, T.; Ramesh, B. Vehicle Detection in Remote Sensing Images Leveraging on Simultaneous Super-Resolution. *IEEE Geoscience and Remote Sensing Letters* **2019**, pp. 1–5. doi:10.1109/LGRS.2019.2930308.

13. Tayara, H.; Soo, K.G.; Chong, K.T. Vehicle detection and counting in high-resolution aerial images using convolutional regression neural network. *IEEE Access* **2017**, *6*, 2220–2230.

14. Yu, X.; Shi, Z. Vehicle detection in remote sensing imagery based on salient information and local shape feature. *Optik-International Journal for Light and Electron Optics* **2015**, *126*, 2485–2490.

15. Stankov, K.; He, D.C. Detection of buildings in multispectral very high spatial resolution images using the percentage occupancy hit-or-miss transform. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing* **2014**, *7*, 4069–4080.

16. Ok, A.O.; Başeski, E. Circular oil tank detection from panchromatic satellite images: A new automated approach. *IEEE Geoscience and Remote Sensing Letters* **2015**, *12*, 1347–1351.

17. Dong, C.; Loy, C.C.; He, K.; Tang, X. Image Super-Resolution Using Deep Convolutional Networks. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **2016**, *38*, 295–307. doi:10.1109/tpami.2015.2439281.

18. Kim, J.; Lee, J.K.; Lee, K.M. Accurate Image Super-Resolution Using Very Deep Convolutional Networks. *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* **2016**. doi:10.1109/cvpr.2016.182.

19. Goodfellow, I.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; Bengio, Y. Generative adversarial nets. Advances in neural information processing systems, 2014, pp. 2672–2680.

20. Ledig, C.; Theis, L.; Huszar, F.; Caballero, J.; Cunningham, A.; Acosta, A.; Aitken, A.; Tejani, A.; Totz, J.; Wang, Z.; et al.. Photo-Realistic Single Image Super-Resolution Using a Generative Adversarial Network. *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* **2017**. doi:10.1109/cvpr.2017.19.

21. Wang, X.; Yu, K.; Wu, S.; Gu, J.; Liu, Y.; Dong, C.; Qiao, Y.; Loy, C.C. ESRGAN: Enhanced Super-Resolution Generative Adversarial Networks. *Computer Vision – ECCV 2018 Workshops* **2019**, p. 63–79. doi:10.1007/978-3-030-11021-5_5.

22. Jiang, K.; Wang, Z.; Yi, P.; Wang, G.; Lu, T.; Jiang, J. Edge-Enhanced GAN for Remote Sensing Image Superresolution. *IEEE Transactions on Geoscience and Remote Sensing* **2019**, *57*, 5799–5812. doi:10.1109/TGRS.2019.2902431.

23. Jiang, J.; Ma, J.; Wang, Z.; Chen, C.; Liu, X. Hyperspectral Image Classification in the Presence of Noisy Labels. *IEEE Transactions on Geoscience and Remote Sensing* **2019**, *57*, 851–865. doi:10.1109/TGRS.2018.2861992.

24. Tong, F.; Tong, H.; Jiang, J.; Zhang, Y. Multiscale union regions adaptive sparse representation for hyperspectral image classification. *Remote Sensing* **2017**, *9*, 872.

25. Zhan, C.; Duan, X.; Xu, S.; Song, Z.; Luo, M. An improved moving object detection algorithm based on frame difference and edge detection. Fourth International Conference on Image and Graphics (ICIG 2007). IEEE, 2007, pp. 519–523.

26. Mao, Q.; Wang, S.; Wang, S.; Zhang, X.; Ma, S. Enhanced image decoding via edge-preserving generative adversarial networks. 2018 IEEE International Conference on Multimedia and Expo (ICME). IEEE, 2018, pp. 1–6.

27. Yang, W.; Feng, J.; Yang, J.; Zhao, F.; Liu, J.; Guo, Z.; Yan, S. Deep Edge Guided Recurrent Residual Learning for Image Super-Resolution. *IEEE Transactions on Image Processing* **2017**, *26*, 5895–5907. doi:10.1109/tip.2017.2750403.

28. Kamgar-Parsi, B.; Kamgar-Parsi, B.; Rosenfeld, A. Optimally isotropic Laplacian operator. *IEEE Transactions on Image Processing* **1999**, *8*, 1467–1472. doi:10.1109/83.791975.

29. Landsat 8. https://www.usgs.gov/land-resources/nli/landsat/landsat-8. Accessed: 2020-02-11.

30. Sentinel-2. http://www.esa.int/Applications/Observing_the_Earth/Copernicus/Sentinel-2. Accessed: 2020-02-11.

31. Mundhenk, T.N.; Konjevod, G.; Sakla, W.A.; Boakye, K. A large contextual dataset for classification, detection and counting of cars with deep learning. European Conference on Computer Vision. Springer, 2016, pp. 785–800.

32. Rabbi, J.; Chowdhury, S.; Chao, D. Oil and Gas Tank Dataset. Mendeley Data, V3, 2020. doi:10.17632/bkxj8z84m9.3.

33. Jolicoeur-Martineau, A. The relativistic discriminator: a key element missing from standard GAN, 2018, [arXiv:cs.LG/1807.00734].

34. Charbonnier, P.; Blanc-Féraud, L.; Aubert, G.; Barlaud, M. Two deterministic half-quadratic regularization algorithms for computed imaging. *Proceedings of 1st International Conference on Image Processing* **1994**, *2*, 168–172 vol.2.

35. Alberta Energy Regulator. https://www.aer.ca. Accessed: 2020-02-05.

36. Tai, Y.; Yang, J.; Liu, X.; Xu, C. MemNet: A Persistent Memory Network for Image Restoration. *2017 IEEE International Conference on Computer Vision (ICCV)* **2017**. doi:10.1109/iccv.2017.486.

37. Zhang, Y.; Tian, Y.; Kong, Y.; Zhong, B.; Fu, Y. Residual Dense Network for Image Super-Resolution. *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition* **2018**. doi:10.1109/cvpr.2018.00262.

38. He, K.; Zhang, X.; Ren, S.; Sun, J. Delving Deep into Rectifiers: Surpassing Human-Level Performance on ImageNet Classification. *2015 IEEE International Conference on Computer Vision (ICCV)* **2015**. doi:10.1109/iccv.2015.123.

39. Lim, B.; Son, S.; Kim, H.; Nah, S.; Lee, K.M. Enhanced Deep Residual Networks for Single Image Super-Resolution. *2017 IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)* **2017**. doi:10.1109/cvprw.2017.151.

40. Liebel, L.; Körner, M. Single-image super resolution for multispectral remote sensing data using convolutional neural networks. *ISPRS-International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences* **2016**, *41*, 883–890.

41. Tayara, H.; Chong, K. Object detection in very high-resolution aerial images using one-stage densely connected feature pyramid network. *Sensors* **2018**, *18*, 3341.

42. Girshick, R.; Donahue, J.; Darrell, T.; Malik, J. Rich Feature Hierarchies for Accurate Object Detection and Semantic Segmentation. *2014 IEEE Conference on Computer Vision and Pattern Recognition* **2014**. doi:10.1109/cvpr.2014.81.

43. Girshick, R. Fast R-CNN. *2015 IEEE International Conference on Computer Vision (ICCV)* **2015**. doi:10.1109/iccv.2015.169.

44. Li, Q.; Mou, L.; Xu, Q.; Zhang, Y.; Zhu, X.X. R3-Net: A Deep Network for Multioriented Vehicle Detection in Aerial Images and Videos. *IEEE Transactions on Geoscience and Remote Sensing* **2019**, *57*, 5028–5042. doi:10.1109/tgrs.2019.2895362.

45. Ammour, N.; Alhichri, H.; Bazi, Y.; Benjdira, B.; Alajlan, N.; Zuair, M. Deep learning approach for car detection in UAV imagery. *Remote Sensing* **2017**, *9*, 312.

46. Ren, Y.; Zhu, C.; Xiao, S. Small object detection in optical remote sensing images via modified faster R-CNN. *Applied Sciences* **2018**, *8*, 813.

47. Tang, T.; Zhou, S.; Deng, Z.; Zou, H.; Lei, L. Vehicle detection in aerial images based on region convolutional neural networks and hard negative example mining. *Sensors* **2017**, *17*, 336.

48. Chen, Z.; Zhang, T.; Ouyang, C. End-to-end airplane detection using transfer learning in remote sensing images. *Remote Sensing* **2018**, *10*, 139.

49. Radovic, M.; Adarkwa, O.; Wang, Q. Object recognition in aerial images using convolutional neural networks. *Journal of Imaging* **2017**, *3*, 21.

50. Li, W.; Fu, H.; Yu, L.; Cracknell, A. Deep learning based oil palm tree detection and counting for high-resolution remote sensing images. *Remote Sensing* **2017**, *9*, 22.

51. Lin, T.Y.; Dollar, P.; Girshick, R.; He, K.; Hariharan, B.; Belongie, S. Feature Pyramid Networks for Object Detection. *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* **2017**. doi:10.1109/cvpr.2017.106.

52. Liu, S.; Huang, D.; others. Receptive field block net for accurate and fast object detection. Proceedings of the European Conference on Computer Vision (ECCV), 2018, pp. 385–400.

53. Zhang, S.; Wen, L.; Bian, X.; Lei, Z.; Li, S.Z. Single-shot refinement neural network for object detection. Proceedings of the IEEE conference on computer vision and pattern recognition, 2018, pp. 4203–4212.

54. Li, Z.; Zhou, F. FSSD: feature fusion single shot multibox detector. *arXiv preprint arXiv:1712.00960* **2017**.

55. Zhu, R.; Zhang, S.; Wang, X.; Wen, L.; Shi, H.; Bo, L.; Mei, T. ScratchDet: Training single-shot object detectors from scratch. Proceedings of the IEEE conference on computer vision and pattern recognition, 2019, pp. 2268–2277.

56. Yang, X.; Sun, H.; Fu, K.; Yang, J.; Sun, X.; Yan, M.; Guo, Z. Automatic ship detection in remote sensing images from google earth of complex scenes based on multiscale rotation dense feature pyramid networks. *Remote Sensing* **2018**, *10*, 132.

57. Zhao, Z.Q.; Zheng, P.; Xu, S.t.; Wu, X. Object detection with deep learning: A review. *IEEE transactions on neural networks and learning systems* **2019**, *30*, 3212–3232.

58. Li, K.; Wan, G.; Cheng, G.; Meng, L.; Han, J. Object detection in optical remote sensing images: A survey and a new benchmark. *ISPRS Journal of Photogrammetry and Remote Sensing* **2020**, *159*, 296–307.

59. Bai, Y.; Zhang, Y.; Ding, M.; Ghanem, B. Sod-mtgan: Small object detection via multi-task generative adversarial network. Proceedings of the European Conference on Computer Vision (ECCV), 2018, pp. 206–221.

60. Haris, M.; Shakhnarovich, G.; Ukita, N. Task-driven super resolution: Object detection in low-resolution images. *arXiv preprint arXiv:1803.11316* **2018**.

61. Zhu, J.Y.; Park, T.; Isola, P.; Efros, A.A. Unpaired image-to-image translation using cycle-consistent adversarial networks. Proceedings of the IEEE international conference on computer vision, 2017, pp. 2223–2232.

62. Simonyan, K.; Zisserman, A. Very Deep Convolutional Networks for Large-Scale Image Recognition, 2014, [arXiv:cs.CV/1409.1556].

63. He, K.; Zhang, X.; Ren, S.; Sun, J. Delving Deep into Rectifiers: Surpassing Human-Level Performance on ImageNet Classification. *2015 IEEE International Conference on Computer Vision (ICCV)* **2015**. doi:10.1109/iccv.2015.123.

64. Lai, W.S.; Huang, J.B.; Ahuja, N.; Yang, M.H. Deep Laplacian Pyramid Networks for Fast and Accurate Super-Resolution. *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* **2017**. doi:10.1109/cvpr.2017.618.

65. Kingma, D.P.; Ba, J. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* **2014**.

66. Paszke, A.; Gross, S.; Massa, F.; Lerer, A.; Bradbury, J.; Chanan, G.; Killeen, T.; Lin, Z.; Gimelshein, N.; Antiga, L.; Desmaison, A.; Kopf, A.; Yang, E.; DeVito, Z.; Raison, M.; Tejani, A.; Chilamkurthy, S.; Steiner, B.; Fang, L.; Bai, J.; Chintala, S. PyTorch: An Imperative Style, High-Performance Deep Learning Library. In *Advances in Neural Information Processing Systems 32*; Wallach, H.; Larochelle, H.; Beygelzimer, A.; d'Alché Buc, F.; Fox, E.; Garnett, R., Eds.; Curran Associates, Inc., 2019; pp. 8024–8035.

67. Rabbi, J. Edge Enhanced GAN with Faster RCNN for end-to-end object detection from remote sensing imagery. https://github.com/Jakaria08/Filter_Enhance_Detect, 2020.

68. Alberta Geological Survey. https://ags.aer.ca. Accessed: 2020-02-05.

69. Chowdhury, S.; Chao, D.K.; Shipman, T.C.; Wulder, M.A. Utilization of Landsat data to quantify land-use and land-cover changes related to oil and gas activities in West-Central Alberta from 2005 to 2013. *GIScience & Remote Sensing* **2017**, *54*, 700–720.

70. Bing Map. https://www.bing.com/maps. Accessed: 2020-02-05.

71. Bulat, A.; Yang, J.; Tzimiropoulos, G. To learn image super-resolution, use a gan to learn how to do image degradation first. Proceedings of the European conference on computer vision (ECCV), 2018, pp. 185–200.