

ClusterNet: Detecting Small Objects in Large Scenes by Exploiting Spatio-Temporal Information

Rodney LaLonde, Dong Zhang, Mubarak Shah

Center for Research in Computer Vision (CRCV), University of Central Florida (UCF)

lalonde@knights.ucf.edu, dromstons@gmail.com, shah@crcv.ucf.edu

Abstract

Object detection in wide area motion imagery (WAMI) has drawn the attention of the computer vision research community for a number of years. WAMI proposes a number of unique challenges including extremely small object sizes, both sparse and densely-packed objects, and extremely large search spaces (large video frames). Nearly all state-of-the-art methods in WAMI object detection report that appearance-based classifiers fail in this challenging data and instead rely almost entirely on motion information in the form of background subtraction or frame-differencing. In this work, we experimentally verify the failure of appearance-based classifiers in WAMI, such as Faster R-CNN and a heatmap-based fully convolutional neural network (CNN), and propose a novel two-stage spatio-temporal CNN which effectively and efficiently combines both appearance and motion information to significantly surpass the state-of-the-art in WAMI object detection. To reduce the large search space, the first stage (ClusterNet) takes in a set of extremely large video frames, combines the motion and appearance information within the convolutional architecture, and proposes regions of objects of interest (ROOBI). These ROOBI can contain from one to clusters of several hundred objects due to the large video frame size and varying object density in WAMI. The second stage (FoveaNet) then estimates the centroid location of all objects in that given ROOBI simultaneously via heatmap estimation. The proposed method exceeds state-of-the-art results on the WPAFB 2009 dataset by 5-16% for moving objects and nearly 50% for stopped objects, as well as being the first proposed method in wide area motion imagery to detect completely stationary objects.

1. Introduction

Object detection is a large and active area of research in computer vision. In wide area motion imagery (WAMI), performing object detection has drawn the attention of

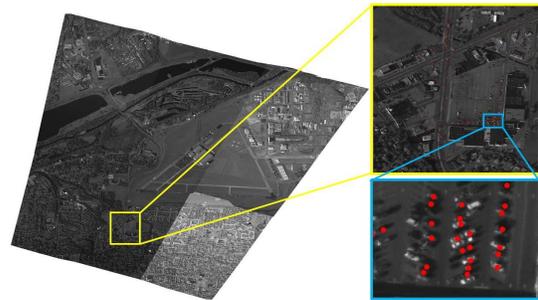


Figure 1: Example WAMI video frame. The yellow-boxed region is enlarged, then a blue-boxed region is further enlarged. Ground truth annotations are marked with red dots.

the computer vision community for a number of years [14, 19, 23, 27, 30]. Numerous applications exist in both the civilian and military domains. In urban planning, applications include automatic traffic monitoring, driver behavior analysis, and road verification for assisting both scene understanding and land use classification. Civilian and military security is another area to benefit with applications including military reconnaissance, detection of abnormal or dangerous behavior, border protection, and surveillance of restricted areas. With increases in the use and affordability of drones and other unmanned aerial platforms, the desire for building a robust system to detect objects in wide-area and low-resolution aerial videos has developed considerably in recent years.

1.1. Object Detection in WAMI

The goal of object detection in images or videos is to place a bounding box (*i.e.* the tightest fitting rectangle which contains all elements of an object while containing as few extraneous elements as possible) around all objects in the scene. Object detection in WAMI differs from the typical object detection problem in three major ways: 1) Ground-truth (*i.e.* human-generated) annotations are single (x, y) coordinates placed at the objects' centers, rather than a bounding box. Therefore, scale and orientation invari-

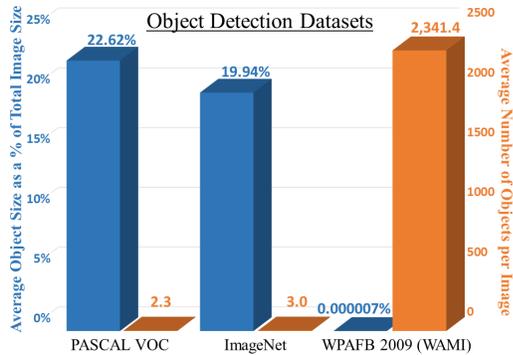


Figure 2: Benchmark datasets in object detection. Two quantities are measured for each dataset: average object size (blue plots, left y-axis) and average number of objects (orange plots, right y-axis).

ance must be learned in order to locate objects’ centers, but this information cannot be provided during supervised training. 2) In typical object detection datasets, images or video frames most often contain only one to three objects, with no more than 15 objects, while these objects take up a large percentage of the image or video frame [13]. In WAMI, video frames can contain thousands of small objects, each object accounting for less than 0.000007% of the total pixels in a given frame. Quantitative analysis of this is shown in Fig. 2. 3) Majority of object detection frameworks deal with images at 256 x 256 up to 500 x 500 pixel resolutions. Video frames in WAMI are significantly larger, typically on the order of several, to hundreds of, megapixels. This creates an extremely large search space, especially given the extremely small typical object size in WAMI being on the order of 9×18 pixels. An example WAMI video frame with ground-truth annotations is shown in Fig. 1.

1.2. Spatial vs. Temporal Information

For the past several years, object detection has been dominated by detectors relying solely on spatial and appearance information (*e.g.* Faster R-CNN [20], ResNet [5], YOLO 9000 [18]). These methods extract low-to-high level spatial and appearance features from images to predict and classify objects. However, it has been stated in numerous recent works [21, 24, 27, 28] that these appearance- and machine-learning-based methods fail in WAMI due to several unique challenges. 1) Extremely small objects averaging 9×18 pixels in size. 2) High intra-class variation, ranging from deep black to bright white and from semi-trucks to small cars with the typical vehicle color (*i.e.* silver/gray) exactly matching the background, as well as dramatic changes in camera gain cause significant changes in objects’ appearance between consecutive frames. 3) Lacking color and with low resolution, videos are single-channel gray-scale with often blurred/unclear object boundaries. 4)

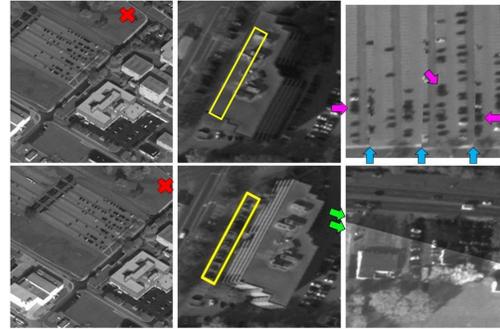


Figure 3: **Left:** Consecutive video frames showing the large object displacement (red X placed at same real-world coordinates). **Center:** Motion parallax effects: Vehicles in the yellow box (top) are occluded at various times. **Right:** Mosaic seams (blue arrows), camera gain differences (green arrows), blurred/unclear object boundaries (purple arrows).

Low frame rates of roughly 1.25 Hz make exploiting temporal information a challenge. Moving objects travel a significant distance between consecutive frames, most often with no overlap to the previous frame. Also, since the aerial recording platform is moving, background objects have significant motion causing strong parallax effects and frame-registration errors, leading to false-positive detections. Moving mosaic seams, where multiple cameras are stitched together to form a single sensor, sweep across the video, leading to even more false positives. Several of these challenges are shown in Fig. 3.

Due to the aforementioned reasons, all state-of-the-art object detection methods in WAMI are motion-based [17, 27, 28], which use background subtraction or frame differencing to find the objects in the videos. However, as with the appearance-based methods, motion-based approaches suffer from their own costly drawbacks. Frame differencing and background subtraction at their core, rely heavily on the video frame registration. Small errors in frame registration can induce large failures and attempting to remove false positives is often critical to these methods. In addition to frame registration, background subtraction requires computing median background images over a large number of frames for the entire video. This combined with the ignorance of appearance information leads to an inefficient use of information across multiple video frames. Yet, the biggest drawback is the complete inability to detect stopped vehicles. All state-of-the-art methods, due to their sole reliance on temporal information, cannot detect vehicles which are not moving relative to the background.

Some recent works [3, 6, 8, 22, 26] have attempted to begin combining spatial and temporal information in various ways for object detection and action recognition. These methods include connecting detections across frames using tracking methods or optical flow, using a sliding window

or out-of-the-box detector to perform detection then classify this result using some temporal information, as well as combining the outputs of a single frame CNN and optical flow input to a CNN. However, all of these methods rely either on a single-frame detector, which uses no temporal information, or use a sliding window to check all possible locations in a video frame for object proposals, and thus do not fully exploit temporal information for the task of object detection in video. This is discussed further in Section 2.3.

1.3. Contribution

The proposed two-stage, spatio-temporal convolutional neural network (CNN) predicts the location of multiple objects simultaneously, without using single-frame detectors or sliding-window classifiers. We show that, consistent with findings in several other works, single-frame detectors fail in this challenging WAMI data, and it is known that sliding window classifiers are terribly inefficient. The novelty of this paper is as follows: 1) Our method effectively utilizes both spatial and temporal information from a set of video frames to locate multiple objects simultaneously in WAMI. 2) This approach removes the need for computing background subtracted images, thus reducing the computational burden and the effect of errors in frame registration. 3) The two-stage network shows the potential to reduce the extremely large search space present in WAMI data with a minimal effect on accuracy. 4) The proposed method is capable of detecting completely stationary vehicles in WAMI, where no other work yet published can do so. 5) The proposed method significantly outperforms the state-of-the-art in WAMI with a **5-16% relative improvement in F₁ score** on moving object detection and a nearly **50% relative improvement** for stopping vehicles, while reducing the average error distance of true positive detections from the previous state-of-the-art 5.5 pixels to roughly 2 pixels.

2. Related Work

2.1. Frame Differencing & Background Subtraction

As stated in Section 1.2, due to the difficulties in WAMI and the reported failures of appearance- and machine-learning-based-methods, all state-of-the-art methods in WAMI are based on either frame-differencing or background subtraction. Both methods require video frames to be registered to a single coordinate system. Reilly *et al.* [19] detects Harris corners in two frames, computes the SIFT features around those corners, and matches the points using descriptors. A frame-to-frame homography is then fit and used to warp images to a common reference frame. Frame differencing is the process of computing pixel-wise differences in intensities between consecutive frames. Both two-frame and three-frame differencing methods have been proposed in literature with a number of variations [9, 16, 24, 27, 29].

Background subtraction methods focus on obtaining a background model for each frame, then subtract each video frame from its corresponding background model. These methods suffer heavily from false positives introduced by the issues discussed in Section 1.2 and cannot detect stationary vehicles. Slowing vehicles also cause a major problem as they are prone to cause split detections in frame differencing [28] while registration errors and parallax effects are increased in background subtraction models, which use more frames than frame differencing. Sudden and dramatic changes in camera gain cause illumination changes which in-turn cause problems for background modeling and frame differencing methods that require consistent global illumination [24].

2.2. Region Proposal Networks

Region proposal networks (RPN), such as Faster R-CNN [20] which has in some ways become the standard in object detection, have shown the ability to generate object proposals with high accuracy and efficiency. Unfortunately, Faster R-CNN fails in WAMI due to four main reasons. 1) Faster R-CNN acts only on single frames, thus does not exploit the available temporal information, which proves to be extremely important. 2) WAMI video frames are extremely large, thus cannot be sent in their entirety to a Faster R-CNN network on any reasonable number of GPUs. This requires spatially-chipping videos into smaller sections and checking these sections individually, dramatically hurting the computational efficiency benefit supposed to be provided by a RPN. 3) If one changed the RPN stage of Faster R-CNN to extremely downsample the images in the earliest layers in order to fit the large WAMI video frames within GPU memory, object proposals would become impossible. Due to the extremely small object size combined with the areas of high object density means any significant amount of downsampling in the network immediately makes object locations indistinguishable, as they are often separated by only a few pixels or even less. 4) WAMI data is ill-suited for Faster R-CNN as the ground-truth locations are single points, not bounding boxes. We experimentally verify that Faster R-CNN fails in WAMI, even when given the benefit of spatially-chipping the video frames to manageable sizes.

2.3. Spatio-Temporal CNNs

In the past few years, partially due to the enormous success of deep learning methods in a vast array of problems, several works have been proposed for combining spatial and temporal information in various ways within deep learning frameworks. Baccouche *et al.* [3] and Ji *et al.* [6] both propose using 3D CNNs for action recognition. Simonyan and Zisserman [26] propose a "two-stream" CNN, one branch receiving individual video frames as input and the other receiving optical flow image stacks where the output of the

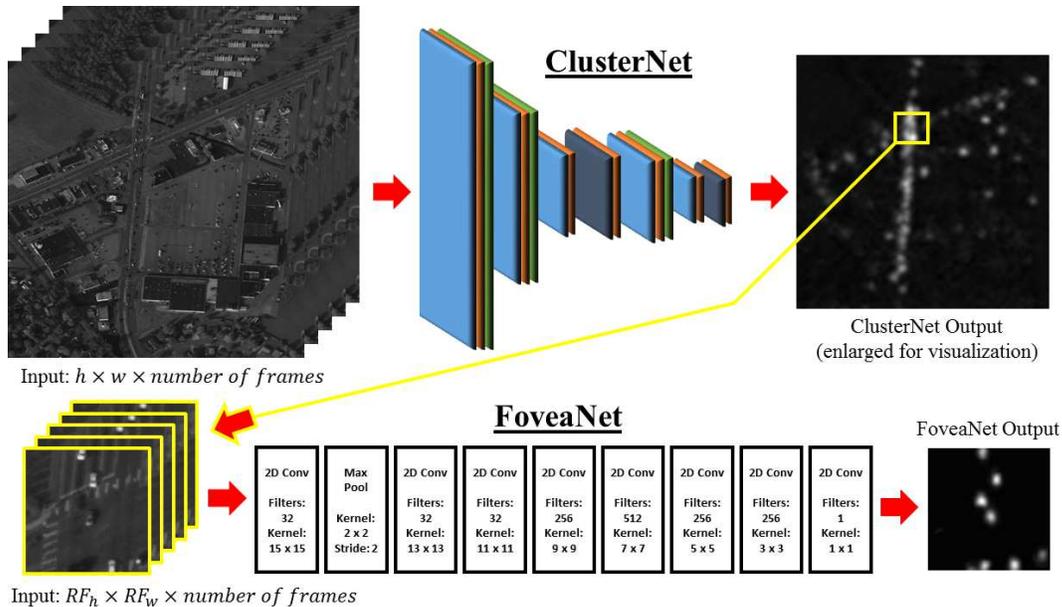


Figure 4: RF_h and RF_w are the height and width of the receptive field of a given set of output neurons. In ClusterNet: light blue and dark blue are 3×3 and 1×1 convolutional layers respectively with PReLU activation functions, orange is batch normalization and green is 2×2 MaxPooling. The 1st and 2nd convolutional layers have stride 2. All FoveaNet convolutional layers have ReLU activation functions, and the 6th and 7th have 50% dropout.

two streams are combined at the end of the network. Kang *et al.* [8] proposes several methods to connect object detections in individual frames across time, including using tracking algorithms, optical-flow-guided propagation, and a long short-term memory (LSTM) sub-network. Rozantsev *et al.* [22] detects flying drones using sliding-window proposals, input to two CNNs multiple times to align each frame, then performs binary classification of the object or non-object in the sliding window.

Our proposed work differs from all of the above in several key ways. Baccouche *et al.* and Ji *et al.* both use stacks of frames as input to a 3D CNN. However, these works do not perform object detection. Both first assume an object of interest is already detected and perfectly centered in each input video frame. To accomplish this, these works use out-of-the-box single-frame human detector algorithms to find the objects of interest in their videos. Our method proposes to solve this object detection problem where single-frame detectors fail, in the challenging WAMI domain. Simonyan and Zisserman keep spatial and temporal information separate during feature extraction, simply combining the extracted features at the end of the network. As stated, single-frame detectors fail in WAMI. Also, due to the extremely large object displacements between consecutive frames, the optical flow stream would likely struggle significantly. The work by Kang *et al.* also relies on first acquiring single-frame object detections before applying their tracking or LSTM methods. The work by Rozantsev *et al.*

is the only one of these methods which does not rely on single-frame detections, instead opting for a sliding window to first generate its object proposals before using a 3D CNN for classification. However, sliding-window-based methods are extremely inefficient. Our work proposes to generate all object proposals simultaneously using a multi-frame, two-stage CNN for videos in WAMI in a more computationally efficient manner than background subtraction or sliding-windows, effectively combining both spatial and temporal information in a deep-learning-based algorithm.

3. ClusterNet & FoveaNet: Two-Stage CNN

We propose a new region proposal network which combines spatial and temporal information within a deep CNN to propose object locations. Where in Faster R-CNN, each 3×3 region of the output map of the RPN proposes nine possible objects, our network generalizes this to propose regions of objects of interest (ROOBI) containing varying amounts of objects, from a single object to potentially over 300 objects, for each 4×4 region of the output map of the RPN. We then focus the second stage of the network on each proposed ROOBI to predict the location of all object(s) simultaneously, again combining spatial and temporal information in this network. This two-stage approach is loosely inspired by biological vision where a large field of vision takes in information, then cues, one of the strongest being motion-detection, determine where to focus the much smaller *fovea centralis*.

3.1. Region Proposal: Exploiting Motion

To reduce the extremely large search space in WAMI, several works proposed using road-overlay maps. This dramatically reduces the search area but severely limits to applicability of the method. Road maps must be known in advance and must be fit perfectly to each video frame, in addition to removing the possibility for detecting objects which do not fall on the road. Instead, we propose a method to learn this search space reduction, without any prior knowledge of road maps. We create a fully-convolutional neural network shown in Fig. 4 which dramatically downsamples the very large WAMI video frames using convolutional strides and max pooling. To exploit temporal information, rather than sending an individual frame to the CNN, we input consecutive adjoining frames with the frame we want to generate proposals for. These adjoining and central frames are input to a 2D convolutional network. The advantage of using a 2D CNN over a 3D CNN as in [3, 6] is the preservation of the temporal relationship between frames. Each frame learns its own convolutional filter set, then these are combined to produce feature maps which maximize information related to the frame we care most about (in our case we chose to train the network to maximize the central frame). Instead of a sliding temporal convolution, our method uses the following equation,

$$f_{x,y}^m = \sum_{n=1}^N \left[\sum_{i=1}^{k_h} \sum_{j=1}^{k_w} V_n(i,j) \times K_n(k_h - i, k_w - j) \right] + b_m \quad (1)$$

where V_n is the n^{th} video frame temporally in the stack and K_n is the convolutional kernel for frame n of size (k_h, k_w) , to produce our feature map values $f^m \in \mathbb{R}^M$, where M is the set of feature maps, $n \in N$ is a frame in the set of temporal frames input to the network, and b_m is a learned bias for the feature map m . This formulation differs from both the standard 2D single-frame CNN and 3D CNNs by allowing us to choose which frame n we want to maximize via the backpropagation of the Euclidean or cross-entropy loss between the output scoremap and the ground truth heatmap for that desired frame.

We formulate the problem in two different ways. In one, we estimate ROOBIs, or object locations in the second stage, via a heatmap-based formulation using the Euclidean loss between the network output and a heatmap created in the following manner,

$$H = \sum_{n=1}^N \frac{1}{2\pi\sigma^2} e^{-\frac{(x/2^d)^2 + (y/2^d)^2}{2\sigma^2}} \quad (2)$$

where $n \in N$ are single (x, y) ground-truth coordinates, d is the amount of downsampling in the network, and σ is the variance of the Gaussian blur fit to each transformed objects

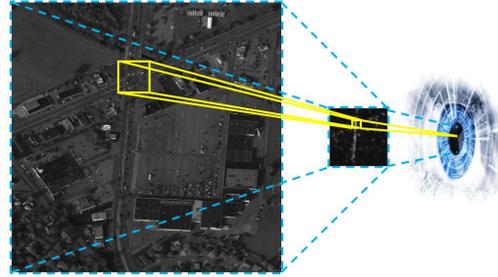


Figure 5: FoveaNet passing over the proposed objects and object clusters from ClusterNet, following up the effective receptive field of high-voting neurons to the initial input.

location. This gives the loss a smooth gradient to follow for estimating the object/region locations rather than single points in space. H is then clipped at 1 in order to equally weight regions with single objects and clusters of hundreds of objects. Segmentation maps are created by thresholding the Gaussian heatmaps for our two classes. The segmentation formulation, using a softmax-cross-entropy loss, is useful when object locations are mutually exclusive. If object locations do not overlap, one could predict the location of a high number of classes of objects using a single output. If locations are not mutually exclusive, the Gaussian heatmap formulation can be employed where each class of object has a corresponding heatmap and the network produces this number of outputs. The results of these experiments show very similar results, demonstrating either formulation can be used, given the specific problem, and thus allows our method more flexibility and a wider range of possible applications.

3.2. FoveaNet: Predicting Object Locations

The FoveaNet stage of our two-stage CNN works on the principle of the effective receptive field of neurons in ClusterNet. Each output neuron in the final 1×1 convolutional layer essentially gives a vote, whether there is a vehicle or cluster of vehicles within that given region or whether there are none. These neurons vote based on the information of the neurons they are connected to in the previous layer which in turn are connected back to the layer before them and so on until the initial input. FoveaNet calculates the region of input information each neuron in the final layer is using to make its final vote. For any ClusterNet output values above a set threshold, this input region, across all input frames, is sent through FoveaNet for high-resolution analysis, as FoveaNet has only a single down-sample in its network. The effect is ClusterNet allows us to ignore large regions of the search space while focusing a small high-resolution *fovea centralis* over regions which contain at least one to several hundred vehicles, illustrated in Fig. 5. FoveaNet then predicts the location all of vehicles

within that region to a high degree of accuracy for the given temporal frame of interest. A full example is provided in the supplementary materials.

Since our FoveaNet input can be much smaller thanks to ClusterNet reducing the search space, we opted to use large kernels within the convolutional layers of FoveaNet, decreasing in size to the final 1×1 convolutional layer, see Fig 4. This was inspired by the recent work by Peng *et al.* [15] as well as a large amount of experimentation. For the options of large kernels ascending in size, descending in size, or fixed in size, as well as small kernels, we found the proposed network to consistently perform the best.

4. Experimental Setup

Experiments were performed on the WPAFB 2009 dataset [1]. This dataset is the benchmark by which all methods in WAMI compare as it is one of the most varied and challenging, as well as one of the only publicly available with human-annotated vehicle locations. The video is taken from a single sensor, comprised of six slightly-overlapping cameras, covering an area of over 19 sq. km., at a frame rate of roughly 1.25 Hz. The average vehicle in these single-channel images make up only approximately 9×18 out of the over 315 million pixels per frame, with each pixel corresponding to roughly 1/4 meter. With almost 2.4 million vehicle detections spread across only 1,025 frames of video, it averages out to be well over 2,000 vehicles to detect in every frame.

Frames are registered to compensate for camera motion following the method by Reilly *et al.* [19] as discussed in Section 2.1. After registration, eight areas of interest (AOI) were cropped out in accordance to those used in testing other state-of-the-art methods [4, 17, 27, 28]. AOIs 01 – 04 are 2278×2278 pixels, covering different types of surroundings and varying levels of traffic. AOI 34 is 4260×2604 . AOI 40 is 3265×2542 . AOI 41 is 3207×2892 . AOI 42 is simply a sub-region of AOI 41 but was included to test our method against the one proposed by Prokaj *et al.* [17] on persistent detections where slowing and stopped vehicles were not removed from the ground truth, even though Prokaj *et al.* uses tracking methods to maintain detections. All other AOIs have any vehicle which moved fewer than 15 pixels ($2/3$ a car length) over the course of 5 frames removed as to be consistent in testing against other methods for moving object detection. All cropped AOIs are shown with their ground-truth and our results in the supplementary materials.

Data was split into training and testing splits in the following way. For training, only tiles which contain vehicles were included. The splits were as follows: AOIs 02, 03, and 34 were trained on AOIs 40, 41, and 42; AOIs 01 and 40 were trained on AOIs 34, 41, and 42; and AOIs 04, 41, and 42 were trained on 34 and 40. Both ClusterNet and

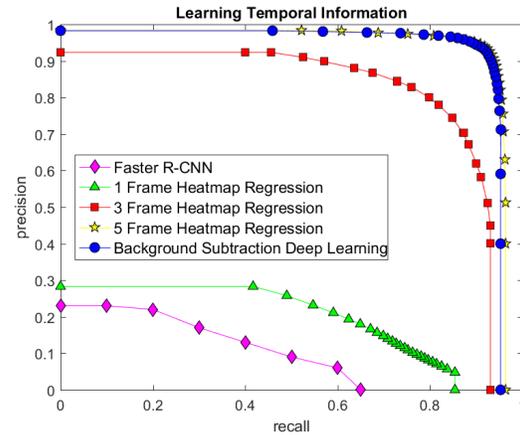


Figure 6: Results on AOI 41 testing the ability of the deep CNN to learn explicitly or implicitly given temporal information, and its necessity.

FoveaNet were trained separately from scratch using Caffe [7]. ClusterNet uses stochastic gradient descent with Nesterov momentum, a base learning rate of 0.01, a batch size of 8, and decreased the learning rate by a factor of 0.1 upon validation loss plateaus. FoveaNet used Adam [11] with a base learning rate of 0.00001 and a batch size of 32. Training and testing was performed on a single Titan X GPU.

To turn the final network output back to single (x, y) coordinates for comparison against the ground-truth, the output is thresholded (either by set levels for creating precision-recall curves, or by Otsu thresholding to find the best threshold level during deployment). Connected components are obtained, weak responses (*i.e.* < 100 pixels) are removed, and large responses (*i.e.* > 900 pixels; assumed to be merged detections) are split into multiple detections by finding circular centers in a bounding box surrounding that connected component. The centroid of each connected component is considered as a positive detection. It should be noted merged detections are quite rare; completely removing this component results in a F_1 score decrease of less than 0.01 across all AOIs. Completely removing small detection removal results in a decrease in F_1 score of 0.01 to 0.05 depending on the AOI tested; however, this parameter is quite robust. Values in the range of 60 to 180 pixels show a change of less than 0.01 in F_1 score across all AOIs.

Quantitative results are compared in terms of precision, recall, and F_1 measure. To be consistent with literature [27] detections were considered true positives if they fell within 20 pixels (5 meters) of a ground truth coordinate. We present results in the supplementary materials reducing this radius. If multiple detections are within this radius, the closest one is taken and the rest, if they do not have any other ground truth coordinates with 20 pixels, are marked as false positives. Any detections that are not within 20 pixels of a



Figure 7: Persistent detection results for AOI 42. The video frame number is marked at the bottom of each column in yellow. **Top Row:** Yellow-boxed image region at labeled times. In frames 183 – 196, the black car (red-boxed) in the shadow of the building is nearly invisible to the naked eye. In frame 311, due to motion parallax, the white vehicle (teal-boxed) is nearly completely occluded by the building. **Middle Row:** Ground-truth heatmap. **Bottom Row:** Output heatmap without any post-processing.

ground truth coordinate are also marked as false positives. Ground truth coordinates which have no detections within 20 pixels are marked as false negatives.

5. Results

5.1. Single-Frame & Background Subtraction

To demonstrate the effect of temporal information, we ran three groups of experiments: explicit, implicit, and no temporal information. For explicit, we compute median and background-subtraction images for all frames following the method by Reilly *et al.* [19]. We then train and test FoveaNet using as input two copies of the central video frame combined with the computed background-subtracted image for that frame, each chipped into 128×128 pixel pieces. This demonstrated our deep network could outperform mere background subtraction through being given both appearance and temporal information. For implicit, we train and test our proposed method using three or five frames as input to FoveaNet to demonstrate the networks ability to learn the temporal information directly from the input images, removing the need for computing median and background-subtraction images. For none, we train and test FoveaNet using a single frame as input, and conduct experiments using Faster R-CNN. We attempted many configurations of Faster R-CNN with VGG-16 and ResNet-50, pre-trained and trained from scratch, with the proposal sizes tuned to WAMI data split into 256×256 pixel chips, where the ground-truth bounding boxes were set to 20×20 pixels centered at each objects location. Anchor box scales at 10, 20, 40 were found to perform the best and all other

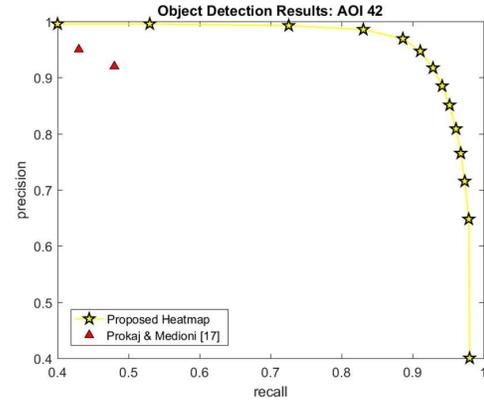


Figure 8: Precision-Recall curve for AOI 42 on persistent detection (*i.e.* no ground-truth coordinates removed).

Percentage Speed-Up From Using ClusterNet

Speed-Up	2 – 3%	5 – 6%	10 – 12%	20%	30%
F_1 Decrease	0%	< 1%	< 3%	< 5%	< 8%

Table 1: Percentage speed-up and F_1 -measure decrease from using ClusterNet at different threshold levels. Higher thresholds exclude larger portions of the input space, but can negatively impact the F_1 score if raised too high.

experimental settings were left unchanged. At test time, bounding boxes with any overlap were considered as true positives. The highest precision-recall curves for all these experiments, tested on AOI 41, are shown in Fig. 6.

5.2. ClusterNet & FoveaNet

The results of our proposed two-stage method, as compared against 13 different state-of-the-art methods, are shown in Fig. 9 and Table 2 across 7 different AOIs. On AOI 42, where stationary object are not removed, our results are shown in Fig. 7 and Fig. 8. Our final results measure the computational efficiency improvement and the effect on detection scores provided by ClusterNet. If FoveaNet had to check every single region of a given input, the time to obtain the predicted object (x, y) locations would be roughly 3 seconds per AOI. With that reference, Table 1 shows the average speed-up from ClusterNet and the associated change in F_1 measure averaged across all AOIs when varying the threshold level for ROOBIs.

6. Conclusion

We have proposed a novel two-stage convolutional neural network for detecting small objects in large scenes, validated on wide area motion imagery. Our method efficiently takes advantage of both appearance and motion cues for detecting the location of single, to hundreds of objects simultaneously. We have shown comparisons with 13 state-of-the-art methods, and the performance improvements are rel-

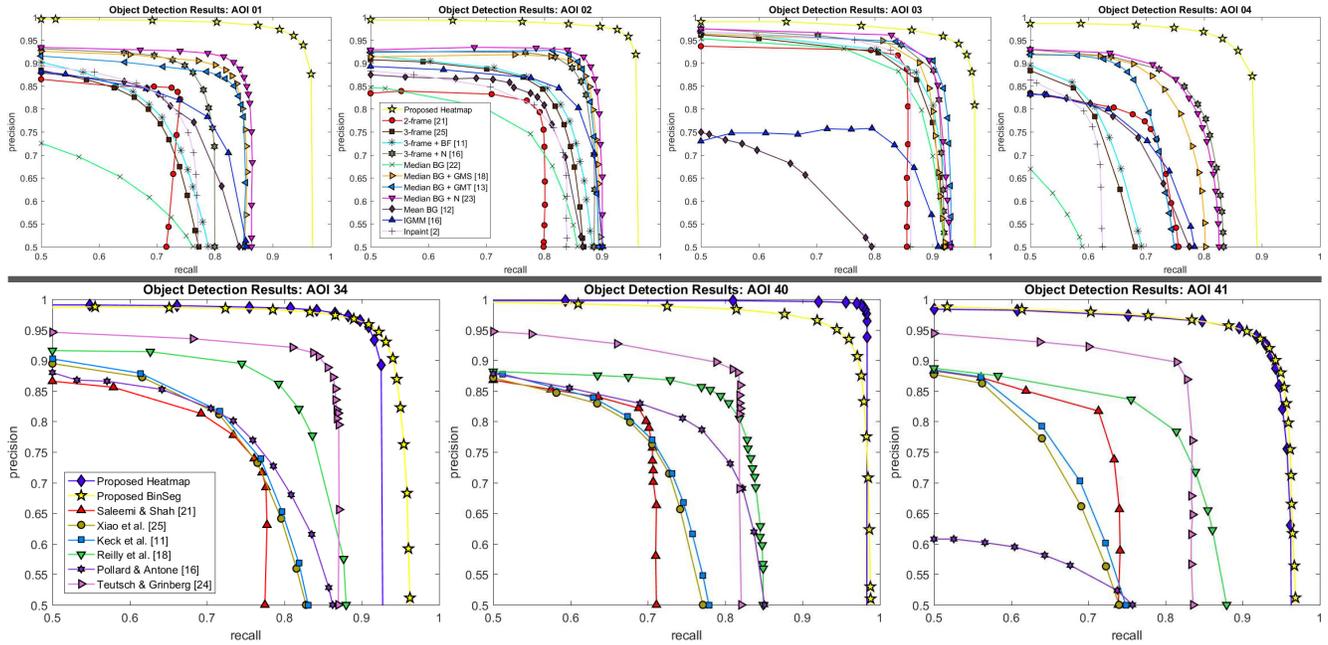


Figure 9: Moving object detection results on seven cropped AOIs with comparisons to 13 state-of-the-art approaches. If precision-recall values were not reported in the original work, the values reported in [27] and/or [28] were used.

Comparison of F_1 Scores on Eight Crop and Aligned Sections of the WPAFB 2009 Dataset

Method	01	02	03	04	34	40	41	42
Sommer <i>et al.</i> [27]	0.866	0.890	0.900	0.804	x	x	x	x
Shi [25]	0.645	0.760	0.861	0.575	x	x	x	x
Liang <i>et al.</i> [12]	0.842	0.880	0.903	0.760	x	x	x	x
Kent <i>et al.</i> [10]	0.767	0.807	0.668	0.711	x	x	x	x
Aeschliman <i>et al.</i> [2]	0.764	0.795	0.875	0.679	x	x	x	x
Pollard & Antone (3-frame + N) [16]	0.816	0.868	0.892	0.805	x	x	x	x
Saleemi & Shah [24]	0.783	0.793	0.876	0.733	0.755	0.749	0.762	x
Xiao <i>et al.</i> [29]	0.738	0.820	0.868	0.687	0.761	0.733	0.700	x
Keck <i>et al.</i> [9]	0.743	0.825	0.876	0.695	0.763	0.737	0.708	x
Reilly <i>et al.</i> [19]	0.850	0.876	0.889	0.783	0.826	0.817	0.799	x
Pollard & Antone (IGMM) [16]	0.785	0.835	0.776	0.716	0.766	0.778	0.616	x
Teutsch & Grinberg [28]	x	x	x	x	0.874	0.847	0.854	x
Prokaj & Medioni [17]	x	x	x	x	x	x	x	0.631
Proposed Method	0.947	0.951	0.942	0.887	0.933	0.983	0.928	0.927

Table 2: F_1 scores of state-of-the-art methods. If F_1 values were not reported in the original work, the values reported in [27] and/or [28] were used. Note that AOI 42 is results on persistent detection (no vehicles removed from ground truth) and is compared with one of the only other persistent detection WAMI methods currently in literature.

atively 5-16% on moving objects as measured by F_1 score and nearly 50% relative improvement on persistent detections. Additionally, the proposed method’s mean distance from ground-truth annotations to true positive detections, is roughly 2 pixels, compared to 5.5 pixels reported in [28]. We further demonstrated that the proposed method can detect stopped vehicles, which is not handled by other methods. Removing the computational burden of computing the median and background subtraction images, as well as

ClusterNet reducing the search space, are both key contributions to approaching an online method. For future work, one of the final barriers is the removal of frame-alignment computed to remove camera motion.

Acknowledgement

The authors would like to acknowledge Lockheed Martin for the funding of this research.

References

- [1] AFRL, Wright-Patterson Air Force Base (WPAFB) dataset. <http://sdms.afrl.af.mil/index.php?collection=wpafb2009>, 2009. 6
- [2] C. Aeschliman, J. Park, and A. C. Kak. Tracking vehicles through shadows and occlusions in wide-area aerial video. *IEEE Transactions on Aerospace and Electronic Systems*, 50(1):429–444, January 2014. 8
- [3] M. Baccouche, F. Mamalet, C. Wolf, C. Garcia, and A. Baskurt. Sequential deep learning for human action recognition. In *Proceedings of the Second International Conference on Human Behavior Understanding*, HBU’11, pages 29–39, Berlin, Heidelberg, 2011. Springer-Verlag. 2, 3, 5
- [4] A. Basharat, M. Turek, Y. Xu, C. Atkins, D. Stoup, K. Fieldhouse, P. Tunison, and A. Hoogs. Real-time multi-target tracking at 210 megapixels/second in wide area motion imagery. In *IEEE Winter Conference on Applications of Computer Vision*, pages 839–846, March 2014. 6
- [5] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016. 2
- [6] S. Ji, W. Xu, M. Yang, and K. Yu. 3d convolutional neural networks for human action recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(1):221–231, Jan 2013. 2, 3, 5
- [7] Y. Jia, E. Shelhamer, J. Donahue, S. Karayev, J. Long, R. Girshick, S. Guadarrama, and T. Darrell. Caffe: Convolutional architecture for fast feature embedding. *arXiv preprint arXiv:1408.5093*, 2014. 6
- [8] K. Kang, W. Ouyang, H. Li, and X. Wang. Object detection from video tubelets with convolutional neural networks. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016. 2, 4
- [9] M. Keck, L. Galup, and C. Stauffer. Real-time tracking of low-resolution vehicles for wide-area persistent surveillance. In *2013 IEEE Workshop on Applications of Computer Vision (WACV)*, pages 441–448, Jan 2013. 3, 8
- [10] P. Kent, S. Maskell, O. Payne, S. Richardson, and L. Scarff. Robust background subtraction for automated detection and tracking of targets in wide area motion imagery. volume 8546, pages 85460Q–85460Q–12, 2012. 8
- [11] D. Kingma and J. Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. 6
- [12] P. Liang, H. Ling, E. Blasch, G. Seetharaman, D. Shen, and G. Chen. Vehicle detection in wide area aerial surveillance using temporal context. In *Proceedings of the 16th International Conference on Information Fusion*, pages 181–188, July 2013. 8
- [13] T.-Y. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick. *Microsoft COCO: Common Objects in Context*, pages 740–755. Springer International Publishing, Cham, 2014. 2
- [14] T. T. Nguyen, H. Grabner, H. Bischof, and B. Gruber. Online boosting for car detection from aerial images. In *2007 IEEE International Conference on Research, Innovation and Vision for the Future*, pages 87–95, March 2007. 1
- [15] C. Peng, X. Zhang, G. Yu, G. Luo, and J. Sun. Large kernel matters – improve semantic segmentation by global convolutional network. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 6
- [16] T. Pollard and M. Antone. Detecting and tracking all moving objects in wide-area aerial video. In *2012 IEEE Computer Society Conference on Computer Vision and Pattern Recognition Workshops*, pages 15–22, June 2012. 3, 8
- [17] J. Prokaj and G. Medioni. Persistent tracking for wide area aerial surveillance. In *2014 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1186–1193, June 2014. 2, 6, 8
- [18] J. Redmon and A. Farhadi. YOLO9000: better, faster, stronger. In *Computer Vision and Pattern Recognition, 2017. CVPR 2017. IEEE Conference on*. IEEE, 2017. 2
- [19] V. Reilly, H. Idrees, and M. Shah. Detection and tracking of large number of targets in wide area surveillance. In *Proceedings of the 11th European Conference on Computer Vision Conference on Computer Vision: Part III, ECCV’10*, pages 186–199, Berlin, Heidelberg, 2010. Springer-Verlag. 1, 3, 6, 7, 8
- [20] S. Ren, K. He, R. Girshick, and J. Sun. Faster R-CNN: Towards real-time object detection with region proposal networks. In *Neural Information Processing Systems (NIPS)*, 2015. 2, 3
- [21] S. Ren, K. He, R. Girshick, and J. Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In C. Cortes, N. D. Lawrence, D. D. Lee, M. Sugiyama, and R. Garnett, editors, *Advances in Neural Information Processing Systems 28*, pages 91–99. Curran Associates, Inc., 2015. 2
- [22] A. Rozantsev, V. Lepetit, and P. Fua. Detecting flying objects using a single moving camera. *IEEE transactions on pattern analysis and machine intelligence*, 39(5):879–892, 2017. 2, 4
- [23] R. Ruskone, L. Guigues, S. Airault, and O. Jamet. Vehicle detection on aerial images: a structural approach. In *Proceedings of 13th International Conference on Pattern Recognition*, volume 3, pages 900–904 vol.3, Aug 1996. 1
- [24] I. Saleemi and M. Shah. Multiframe many—many point correspondence for vehicle tracking in high density wide area aerial videos. *Int. J. Comput. Vision*, 104(2):198–219, Sept. 2013. 2, 3, 8
- [25] X. Shi, H. Ling, E. Blasch, and W. Hu. Context-driven moving vehicle detection in wide area motion imagery. In *Proceedings of the 21st International Conference on Pattern Recognition (ICPR2012)*, pages 2512–2515, Nov 2012. 8
- [26] K. Simonyan and A. Zisserman. Two-stream convolutional networks for action recognition in videos. In *Advances in neural information processing systems*, pages 568–576, 2014. 2, 3
- [27] L. W. Sommer, M. Teutsch, T. Schuchert, and J. Beyerer. A survey on moving object detection for wide area motion imagery. In *2016 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 1–9, March 2016. 1, 2, 3, 6, 8

- [28] M. Teutsch and M. Grinberg. Robust detection of moving vehicles in wide area motion imagery. In *2016 IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pages 1434–1442, June 2016. [2](#), [3](#), [6](#), [8](#)
- [29] J. Xiao, H. Cheng, H. Sawhney, and F. Han. Vehicle detection and tracking in wide field-of-view aerial video. In *2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 679–684, June 2010. [3](#), [8](#)
- [30] T. Zhao and R. Nevatia. Car detection in low resolution aerial image. In *Proceedings Eighth IEEE International Conference on Computer Vision. ICCV 2001*, volume 1, pages 710–717 vol.1, 2001. [1](#)