This CVPR workshop paper is the Open Access version, provided by the Computer Vision Foundation. Except for this watermark, it is identical to the accepted version; the final published version of the proceedings is available on IEEE Xplore.



Augmentation of Atmospheric Turbulence Effects on Thermal Adapted Object Detection Models

Engin Uzun^{1,2*} Ahmet Anıl Dursun^{1*} Erdem Akagündüz² ¹Dept. of Image Proc. & Computer Vis. Technologies, ASELSAN Inc., Turkey ²Graduate School of Informatics, Middle East Technical University, Turkey {enginuzun, aadursun}@aselsan.com.tr, akaerdem@metu.edu.tr

Abstract

Atmospheric turbulence has a degrading effect on the image quality of long-range observation systems. As a result of various elements such as temperature, wind velocity, humidity, etc., turbulence is characterized by random fluctuations in the refractive index of the atmosphere. It is a phenomenon that may occur in various imaging spectra such as the visible or the infrared bands. In this paper, we analyze the effects of atmospheric turbulence on object detection performance in thermal imagery. We use a geometric turbulence model to simulate turbulence effects on a medium-scale thermal image set, namely "FLIR ADAS v2". We apply thermal domain adaptation to state-of-theart object detectors and propose a data augmentation strategy to increase the performance of object detectors which utilizes turbulent images in different severity levels as training data. Our results show that the proposed data augmentation strategy yields an increase in performance for both turbulent and non-turbulent thermal test images.

1. Introduction

Caused by various atmospheric elements such as temperature, humidity or wind, atmospheric turbulence usually refers to the three-dimensional chaotic flow of air. Turbulence has a distorting effect on the propagation of light through air. It changes the refractive index of air in a spatiotemporal manner. Especially in long-range imaging systems, such distortions become apparent in various electromagnetic spectra. In astrophotography, the turbulence effect is assumed to be isoplanatic, meaning that the distortions are spatially non-varying. Adaptive optics is a common solution for mitigating the isoplanatic turbulence effects [1]. However, for ground-to-ground imaging systems, turbulence is anisoplanatic, which results in non-rigid distortions and blurring on the image. Such non-rigid distor-



Figure 1. The turbulence generation processes on an IR and a chessboard images are illustrated. t{} is the turbulence operation and the vertical $(d_n^u(x, y) \text{ in red})$ and the horizontal $(d_n^v(x, y) \text{ in green})$ elements of the 2D distortion field are depicted in false color.

tions possibly affect the performance of any computer vision algorithm, which is a fact that is not widely investigated in the literature.

Theoretical studies on the effects of atmospheric turbulence in the infrared (IR) spectra can be found in the literature [17, 20, 37, 38]. For example [17] proposes the Fried coherence diameter as a measure of the optical transmission quality through the turbulent air, which reflects the relation of wavelength and imaging quality. According to this measure, as the wavelength increases, the resultant effect of turbulence tends to decrease. Therefore, the visible spectrum suffers more than the IR spectra from turbulence. However, the effect of turbulence on images also becomes significant in the IR spectra, when the refractive index structure parameter, C_n^2 [5, 17], and the path length are sufficiently high.

In this work, the effects of anisoplanatic turbulence on object detection are analyzed for IR imagery. Object detection is one of the trending challenges of computer vision and deep learning, which is being widely applied in several

^{*}indicates equal contribution

problem domains [19, 24]. The literature for object detection can be chronologically scrutinized in two mainstream approaches, namely the two-stage and the one-stage object detectors. One-stage approaches such as "YOLO" [29] and its variants [3, 30, 31, 36], are the successors of two-stage approaches and have lower complexities; hence, are popular for real-time applications. Studies on object detection in IR imagery are relatively new [8, 21, 22]. Similar architectures to visible domain are being utilized. However, due to relatively low-scales of IR image sets, detection success on IR imagery is lower when compared to visible band. In order to tackle the data problem in IR imagery, recent studies focus on visible-to-IR domain adaptation and achieve satisfactory results using thermal-RGB pairs and supervised learning [16, 27]. Moreover, adversarial-based unsupervised methods are also utilized to thermal domain adaptation [2]. Hence, thermal adaption is currently a defacto practice in IR computer vision problems.

Although several IR image and video sets are publicly available [9], to the best of our knowledge, there is no public IR image set that includes turbulence effects, for object detection or any other computer vision tasks. This is mainly because creating a set consisting of both turbulent and nonturbulent images of a given scene is a very expensive and difficult task. A feasible option is to create synthetic turbulence on existing image sets, using mathematical models of turbulence. The existing models in the literature [13, 18, 32]generally represent the dynamics of atmospheric turbulence in 3D coordinates and project synthetic 3D scenes to 2D images. However, creating a large corpus of images with different 3D scene settings is also a very difficult task to achieve. Furthermore, in the thermal domain, creating 3D models with realistic material properties is problematic as well. Another approach can be working on existing thermal object detection datasets that additionally include 3D objects models. However, this case would require a reconstruction of the 3D scene with depth information, which is not available for almost any image set.

Previous work on generating synthetic atmospheric turbulence on images mainly focuses on the visible spectrum. In [28], generative adversarial networks (GANs) are utilized for turbulent image generation in the visible spectrum, where sufficiently large-scale data can easily be collected. [7] uses neural networks to simulate a geometric model on visible band images. In [14, 34], physics-based models are proposed. There are recent studies that focus on the effects of atmospheric turbulence on long-range observation systems for limited applications fields [10]. Nevertheless, none of the aforementioned literature propose solutions to turbulence related issues of infrared computer vision problems.

In this paper, we use a geometric model to create turbulence distortions on IR images. The model is operationalized on a set of IR images selected from the "FLIR ADAS v2" dataset [12]. Both the generated (turbulent) and original (non-turbulent) images are used so as to improve the performance of deep learning-based object detectors. In our experiments, we benchmark 3 different one-stage deep object detectors, which we transfer as pretrained and attempt to adapt to thermal domain via fine-tuning. Moreover, we propose a data augmentation strategy that yields an increase in detection performance for both turbulent and non-turbulent thermal test images. The details of the geometric turbulence model, the benchmarked architectures, the experiments and our results are provided in the following sections.

2. The Turbulence Model

Previous studies that model the turbulence effect geometrically usually utilize a composition of blurring and random distortions [4, 26, 41]. Blurring can be applied using basic low-pass filtering. Random distortions can be emulated using image warping and can be efficiently performed on GPUs [33]. We utilize the following model for turbulent image generation using non-turbulent images:

$$F_{n}(x,y) = D((G_{B}(x,y) \circledast I_{n}(x,y)), ... \\ d_{n}^{u}(x,y), d_{n}^{v}(x,y))$$
(1)

where $F_n(x, y)$ is the source image, D is the warping function, \circledast is the convolution operation and $G_B(x, y)$ is a Gaussian kernel with variance σ_B^2 , responsible for the blurring operation. Note that the warping operation is applied on both horizontal and vertical directions using the random distortion fields, $d_n^u(x, y)$ and $d_n^v(x, y)$, respectively, which are defined as:

$$d_n^u(x,y) = \gamma * (G_D(x,y) \circledast v_n^u(x,y))$$
(2)

$$d_n^v(x,y) = \gamma * (G_D(x,y) \circledast v_n^v(x,y))$$
(3)

where γ is the amplitude of the random distortion and $G_D(x, y)$ is the Gaussian kernel with variance σ_D^2 . $v_n^u(x, y)$ and $v_n^v(x, y)$ are random vectors with zero-mean, unitvariance normal distributions. Convolution operation with $G_D(x, y)$ provides spatial correlation of the random distortions over the image. σ_D^2 is used to adjust the strength of the spatial correlation while γ is the amplitude of the distortions in the model. σ_B^2 is used to adjust the amount of blurring. In Fig. 1, the proposed geometric turbulence model is depicted with a sample IR image and a chessboard image.

2.1. Physical Approximation

In order to relate the proposed geometric turbulence model parameters with real world turbulence conditions, in this section we formulate the relation between the magnitude of the pixel shift (distortion vectors) and the model parameters, namely γ and σ_D^2 . Note that, we do not utilize σ_B^2 for this approximation since it is not related with the distortion process. Let z(x, y) be the random distortion vector over an arbitrary image region,

$$\boldsymbol{z}(\boldsymbol{x}, \boldsymbol{y}) = [d_n^u(\boldsymbol{x}, \boldsymbol{y}), d_n^v(\boldsymbol{x}, \boldsymbol{y})]^T$$
(4)

z(x, y) is a bivariate Gaussian distribution, where u and v components are independent of each other and are zeromean. Variance of u and v components depends on the values of both γ^2 and σ_D^2 . In Equations 2 and 3, the vertical and horizontal distortions are derived from zero-mean, unitvariance Gaussian distributions. In the proposed model, these random variables, $v_n^u(x, y)$ and $v_n^v(x, y)$, are firstly filtered spatially by a 2D Gaussian kernel with σ_D^2 variance, then multiplied with the turbulence severity level, γ . Spatial filtering operation on $d_n^u(x, y)$ and $d_n^v(x, y)$ corresponds to the weighted linear combination of zero-mean unit-variance uncorrelated Gaussian random variables, where the weights are the parameters in $G_D(x, y)$. The equivalent variance of the spatially filtered random distortions are the sum of the squared kernel parameters. For a sufficiently large kernel, this summation is equivalent to the integration of the corresponding squared continuous bivariate Gaussian function in 2D space. Then, γ is simply a multiplier over the variance. The resultant variance of $d_n^u(x, y)$ and $d_n^v(x, y)$, σ_z^2 , can be expressed as,

$$\sigma_z^2 = \gamma^2 \int \int G(p, \mu_p, \Sigma_p)^2 dp \tag{5}$$

where μ_p is the mean and Σ_p is the covariance matrix of the bivariate Gaussian kernel. For our case, the Gaussian kernel $G_D(x, y)$ is assumed to be zero-mean and have constant diagonal covariance matrix Σ_p . Given these conditions, in explicit form 5 reduces to,

$$\sigma_z^2 = \gamma^2 \int \int (\frac{1}{2\pi\sigma_D^2} e^{-\frac{p_0^2 + p_1^2}{2\sigma_D^2}})^2 dp_0 dp_1 \qquad (6)$$

The the resultant expression for σ_z^2 becomes,

$$\sigma_z^2 = \frac{\gamma^2}{4\sigma_D^2 \pi} \tag{7}$$

Let l(x, y) be the Euclidean norm of the random distortion vector z(x, y), which is given as,

$$l(x,y) = \sqrt{d_n^u(x,y)^2 + d_n^v(x,y)^2}$$
(8)

l(x, y) is also a random variable and because it is the Euclidean norm of a vector sampled from a normal distribution with zero-mean and constant diagonal covariance matrix, l(x, y) is expected to have a Rayleigh distribution, which can be written as,

$$f_{l(x,y)}(l(x,y)) = \frac{l(x,y)}{\sigma_z^2} e^{-l(x,y)^2 (2\sigma_z^2)}$$
(9)

where $f_{l(x,y)}(l(x,y))$ is the probability density function of random variable l(x,y). Mean of the l(x,y), namely μ_l , is the measure of how much the pixels are distorted on an image for the given parameters γ and σ_D^2 , and given as,

$$\mu_l = \sigma_z \sqrt{\frac{\pi}{2}} = \frac{\gamma}{2\sqrt{2}\sigma_D} \tag{10}$$

Note that, μ_l is a measure in pixel dimensions. In order to relate the distortions in the pixel dimensions to the real world measurements, a pin-hole camera model can be utilized using internal parameters of the thermal camera used in collecting the FLIR ADAS v2 dataset. Relation between pixel distortions and the corresponding real world distortion for both horizontal and vertical axes, t_s^h and t_s^v respectively, can be written as,

$$t_s^h = \frac{2\mu_l d_s tan(\alpha_h/2)}{W} = \frac{\gamma d_s tan(\alpha_h/2)}{\sqrt{2}\sigma_D W}$$
(11)

$$t_s^v = \frac{2\mu_l d_s tan(\alpha_v/2)}{H} = \frac{\gamma d_s tan(\alpha_v/2)}{\sqrt{2}\sigma_D H}$$
(12)

where d_s is the depth of the scene region, α_h and α_v are the horizontal and the vertical field of views respectively, W is the horizontal pixel count and H is the vertical pixel count of the camera. In [12], horizontal and vertical fields of views of the thermal camera are given as 45° and 37° and the image resolution is given as 640x512. Then, 11 and 12 reduces to,

$$t_s^h = 0.0004576 \frac{\gamma d_s}{\sigma_D} \tag{13}$$

$$t_s^v = 0.0004621 \frac{\gamma d_s}{\sigma_D} \tag{14}$$

In the following sections, we will use the approximations given in 13 and 14 to determine the γ levels to be used in our experimental setup.

3. Experimental Setup

3.1. The Image Set

The experiments in this paper are implemented on the FLIR-ADAS v2 image set [12]. FLIR-ADAS v2 is a medium-scale image set annotated for object detection tasks in both thermal and visible bands. The set includes a total of 26,442 images with 15 different object classes. A total of 9,711 thermal and 9,233 RGB training/validation images are included with a 90%-10% train/validation split. In this study, we use only the thermal images of the set, which are acquired with a Teledyne FLIR Tau 2. The resolution of



Figure 2. The leftmost figure is a patch from a sample image of the FLIR ADAS v2 dataset. The middle and the rightmost figures are patches from the same image constructed with the proposed geometric turbulence model with γ values 50 and 100, respectively. Annotations for the *person* and *car* classes are depicted in colored rectangles.

the thermal images is 640×512 pixels. In our experiments, we utilize only the *car* and the *person* classes because the number of samples from the other classes is insufficient for domain adaptation.

for their IACS score function. In the original paper, VfNet is tested with various backbones, computation-wise lightest one of which is the ResNet-50 [15], as utilized in our experiments.

3.2. Object Detection Architectures

In our experiments, three state-of-the-art one-stage object detectors, namely VfNet [39], TOOD [11], and YOLOR [36], are specifically chosen for their real-time performance. The reason that we deploy these three real-time models is because they are the state-of-the-art models that show leading detection performance in various common benchmarks.

3.2.1 VfNet

Proposed by [39], VarifocalNet (VfNet) is an architecture that is designed on top of the fully convolutional one-stage object detector (FCOS) [35]. FCOS originally proposes a pixel-wise approach similar to segmentation and utilizes a one-stage anchor free object detector. In [40], the idea of adaptive training sample selection (ATSS) is developed on top of the FCOS architecture. VfNet is another improvement over this design family and has three main contributions to the baseline FCOS + ATSS approach. Firstly, VfNet proposes the so-called intersection over union-aware classification score (IACS) function, which is a scalar representation of a fusion of classification and the intersectionover-union (IoU) score. Secondly, they propose a novel loss function, namely Varifocal loss, which is designed to penalize their proposed IACS function. Varifocal loss function is based on the so-called Focal Loss [23] that is designed to handle the imbalance problem between positive and negative samples. Verifocal loss is utilized only for negative samples; whereas, for positive samples, Binary Cross-Entropy (BCE) Loss is employed with the multiplication of the ground truth class. Finally, they propose a star-shaped bounding box feature, which they claim to be more efficient

3.2.2 TOOD

Proposed by [11], the Task-aligned One-stage Object Detection (TOOD) has two important contributions over the existing one-stage object detector models. First, instead of using parallel heads for classification and localization, TOOD utilizes a single, so-called task-aligned head in order to reduce any spatial misalignment. Second, TOOD uses a novel learning scheme, namely task alignment learning, which yields optimal anchor estimates for both classification and localization tasks. Although it is not one of the backbones tested for TOOD in the original paper, we again utilized ResNet-50 on this architecture, for a fair comparison with the VfNet.

3.2.3 YOLOR

YOLOR [36] is basically a unified network architecture, which combines the explicit knowledge (i.e. shallow layer activations) and the implicit knowledge (i.e. deeper layer activations) of a network. Hence, more than an object detection specific architecture, YOLOR is designed as a multitask learning model. However, the model is operationalized for object detection and performs comparable to leading non-real-time models such as the swin transformer [25], but achieves this in real-time. Because the model includes an encoder design of its own, no external backbone such as ResNet-50 is implemented for this model. In our experiments, we utilize the version called "YOLOR-P6", which is based on the "YOLOv4-P6-light" architecture.

3.3. Training Configuration

The experiments in this paper are deployed in paralel using 4 NVIDIA GeForce RTX 2080Ti GPUs. For all models, the batch size is set for 2 images for each GPUs. Open source MMDetection [6] toolbox is utilized for training of VFNet and TOOD architectures. While training of these architectures, default hyper-parameters and original augmentation strategies are utilized. YOLOR-P6 architecture is trained using the official GitHub repository [36], which is provided by the authors. In a similar manner, default settings are utilized for training of YOLOR-P6. All experiments are run for 100 epochs.

In order to augment the training and test sets with turbulent images of varying severity levels, we utilized four different values of γ in our experiments, specifically 25, 50, 100 and 150. These different γ values correspond to different atmospheric conditions, where the effect of the atmospheric turbulence increases directly proportional to γ values. We chose σ_D^2 and σ_B^2 as 25 and 1, respectively. In Figure 2, the effects of applied turbulence on a sample image patch taken from the FLIR ADAS vs dataset are illustrated for γ values 50 and 100. For instance, using Equation 13, for a scene with average depth of 100m, the selected four γ values result in real world horizontal distortions of 0.23m, 0.46m, 0.92m and 1.37m. The vertical distortions can be calculated similarly using Equation 14 and will result in close metric values. We believe that the selected γ values span a realistic range of turbulence levels for real world applications.

4. Results

The results of all the experiments are presented as average precision (AP) values in Table 1. In this table, each row represents a separate experiment. The first column, namely Aug, indicates if turbulent data is augmented during training (w/) or not (w/o). The following two columns state the Model and the Backbone used. The AP column is the mean average precision calculated over the two categories; whereas AP_{50} and AP_{75} represent the cases when IoU is selected as 0.5 and 0.75, respectively. The rightmost three columns set forth the AP values for the test subsets, where the test samples are distributed according to their pixel size. AP_S is the performance score for the subset that includes test objects smaller than 32×32 pixels; whereas when calculating AP_L , test objects larger than 96×96 pixels are used. AP_M represents the performance of the test objects that have sizes in between.

A rigorous analysis of Table 1 provides several important implications. To begin with, for all models and for all turbulence levels of the test sets, augmenting turbulent images of different γ values during training always increase detection performance. This is one very important result of our

Aug	Model	Backbone	AP	AP_{50}	AP_{75}	AP_S	AP_M	AP_L
test set without turbulence								
w/o	VfNet	ResNet -50	52.4	81.9	54.8	41.1	73.4	76.8
w/	VfNet	ResNet -50	54.7	83.8	57.1	43.7	75.2	79.9
w/o	TOOD	ResNet -50	46.1	74.6	47.3	33.1	71.1	76.3
w/	TOOD	ResNet -50	46.8	75.4	47.9	33.7	71.3	77.8
w/o	YOLOR	YOLOR-P6	53.1	80.5	56.2	40.6	77.3	82.9
w/	YOLOR	YOLOR-P6	53.8	81.0	57.2	41.4	77.5	83.3
test set with turbulence $\gamma = 25$								
w/o	VfNet	ResNet -50	51.0	80.4	53.1	39.2	73	77.0
w/	VfNet	ResNet -50	53.9	83.1	56.1	42.7	74.7	79.4
w/o	TOOD	ResNet -50	45.3	74.0	45.7	31.9	70.9	76.8
w/	TOOD	ResNet -50	46.2	75.0	46.7	33.2	71.1	77.0
w/o	YOLOR	YOLOR-P6	52.3	79.9	54.7	39.5	77.1	83.3
w/	YOLOR	YOLOR-P6	53.1	80.5	56.0	40.6	77.2	83.8
test set with turbulence $\gamma = 50$								
w/o	VfNet	ResNet -50	47.5	77.0	48.3	35.1	70.8	76.2
w/	VfNet	ResNet -50	51.8	81.4	53.5	39.8	73.9	79.8
w/o	TOOD	ResNet -50	43.4	71.9	44.0	29.7	69.7	77.0
w/	TOOD	ResNet -50	44.9	73.9	45.1	31.5	69.8	77.2
w/o	YOLOR	YOLOR-P6	50.0	77.8	51.8	36.5	75.7	82.8
w/	YOLOR	YOLOR-P6	51.4	79.4	53.7	38.5	76.2	83.7
test set with turbulence $\gamma = 100$								
w/o	VfNet	ResNet -50	36.7	64.4	36.1	23.4	61.3	72.5
w/	VfNet	ResNet -50	45.6	76.0	45.1	32.3	69.9	78.1
w/o	TOOD	ResNet -50	35.7	62.8	34.4	21.3	62.8	73.3
w/	TOOD	ResNet -50	40.4	69.7	39.0	26.4	66.2	75.3
w/o	YOLOR	YOLOR-P6	41.4	69.1	40.9	26.4	69.7	82.4
w/	YOLOR	YOLOR-P6	45.9	74.5	46.3	31.8	72.5	82.6
test set with turbulence $\gamma{=}150$								
w/o	VfNet	ResNet -50	23.0	43.8	21.0	12.1	43.7	61.1
w/	VfNet	ResNet -50	38.9	68.6	36.5	25.0	64.2	76.2
w/o	TOOD	ResNet -50	26.0	49.3	24.1	13.0	49.4	68.1
w/	TOOD	ResNet -50	35.4	63.7	33.2	20.9	61.6	74.1
w/o	YOLOR	YOLOR-P6	31.1	55.8	29.9	16.2	58.2	78.6
w/	YOLOR	YOLOR-P6	39.7	67.9	38.4	24.5	67.4	81.3

Table 1. Mean Average Precision results obtained for different experiments with different models, with or without turbulent image augmentation for varying levels of turbulence gain γ .

experiments showing that the proposed augmentation strategy can be generalized for any thermal adaptation or model training experiment. Another important observation is that as the turbulence level of the test set increase, the positive effect of the proposed augmentation strategy also inevitably increase. We believe that this is an indicator of the robustness of the proposed augmentation strategy.

When we analyse the individual behaviour of the models, different characteristics can be observed. Although VfNet exhibits slightly better performance for the test sets with lower levels of turbulence, we see that as the turbulence level of the test sets increases, YOLOR becomes the leading model in general. However, the difference in performance between the VfNet and YOLOR models can be considered negligible. On the other hand, TOOD performs poorer than the other two models in general, which is consistent with its relative performance on RGB, i.e. before thermal adaptation. The only consistent behaviour for all models is that regardless of their contrasting performance when trained without turbulent data augmentation, our proposed augmentation strategy brings the models to comparable performance. This can be observed for the experiments with the highest turbulence levels of test sets, when VfNet



Figure 3. "IoU vs Recall" curves for all models, with and without turbulent data augmentation, for selected γ levels.

performance becomes comparable to YOLOR as a result of our augmentation strategy, following a dramatic increase of mean *AP* value of 23.0 to 38.9.

The AP_S values in Table 1 show that for small objects, detection performance falls dramatically at high levels of turbulence. This is not a surprising fact, because as the object is farther from the camera, not only its pixel size decreases, but the effect of turbulence increases as well. Nonetheless, Table 1 clearly shows that the impact of the proposed augmentation strategy is the strongest when the object size is small and the turbulence levels are higher. The reader should note that, before their adaptation, the benchmarked models are already trained with various augmentation strategies, which do not help them to overcome highlevel turbulence effects. For medium and large objects, the positive effect of the augmentation strategy is minimal if the turbulence levels are low. However, for none of the cases, it is diminishing, i.e. AP_M and AP_L values are always higher for experiments where turbulent image augmentation is utilized.

In order to examine the localization success of each

experiment, AP_{50} and AP_{75} columns are provided in Table 1. Furthermore, "IoU vs Recall" graphs are given in Figure 3. In this figure, for different levels of γ , recall curves with respect to varying IoU are depicted, for all models, with or without turbulent data augmentation. Both the AP_{75} values in Table 1 and recall curves in Figure 3 show that YOLOR model provides finer localization. In addition, we see from Figures 3a, 3b, 3c and 3d that as the turbulence levels increase, the impact of the proposed augmentation strategy becomes dramatically effective for any level of localization (i.e. IoU).

5. Conclusions and Future Directions

We propose a data augmentation strategy to increase the performance of thermal-adapted, real-time, one-stage object detectors under varying atmospheric turbulence conditions. In order to create turbulent images, we use a computation-friendly geometric turbulence model that can easily be implemented for online learning systems. Our results show that for all turbulence levels of the test sets, including the experiment, where there is no turbulent images in the test set, augmenting turbulent images of different severity levels always increases the detection performance for all models. Especially when the turbulence levels are high, the impact of the proposed augmentation scheme becomes much clear, such that it brings different model performances to a similar satisfactory level.

The geometric distortions created by atmospheric turbulence usually demonstrate an unusual character in images. Conventional augmentation methods, such as affine transformations, do not have the capacity to simulate the effects of atmospheric turbulence, especially for long-range vision problems and small objects. We believe that turbulent image augmentation has the capacity to become a de facto practice for visible and thermal vision systems that utilize deep learning models.

In this paper, atmospheric turbulence effects are augmented using a geometric model for only the object detection problem. We believe that the approach can be expanded to other deep learning-based solutions of several IR vision problems. Furthermore, although using a computationfriendly geometric turbulence model such as ours has its advantages in practice, by using physics-based turbulence models and collecting data from calibrated scenes, more realistic turbulent images can be obtained. Such an effort would necessarily increase the impact of the proposed augmentation strategy.

References

- Adaptive Optics in Astronomy. Cambridge University Press, 1999. 1
- [2] Ibrahim Batuhan Akkaya, Fazil Altinel, and Ugur Halici. Self-training guided adversarial domain adaptation for thermal imagery. In *Proceedings of the IEEE/CVF Conference* on Computer Vision and Pattern Recognition (CVPR) Workshops, pages 4322–4331, June 2021. 2
- [3] Tolga Aksoy and Ugur Halici. Analysis of visual reasoning on one-stage object detection. arXiv preprint arXiv:2202.13115, 2022. 2
- [4] Nantheera Anantrasirichai, Alin Achim, Nick G. Kingsbury, and David R. Bull. Atmospheric turbulence mitigation using complex wavelet-based fusion. *IEEE Transactions on Image Processing*, 22(6):2398–2408, 2013. 2
- [5] Larry C. Andrews and Ronald L. Phillips. Laser beam propagation through random media. 1998. 1
- [6] Kai Chen, Jiaqi Wang, Jiangmiao Pang, Yuhang Cao, Yu Xiong, Xiaoxiao Li, Shuyang Sun, Wansen Feng, Ziwei Liu, Jiarui Xu, Zheng Zhang, Dazhi Cheng, Chenchen Zhu, Tianheng Cheng, Qijie Zhao, Buyu Li, Xin Lu, Rui Zhu, Yue Wu, Jifeng Dai, Jingdong Wang, Jianping Shi, Wanli Ouyang, Chen Change Loy, and Dahua Lin. MMDetection: Open mmlab detection toolbox and benchmark. *arXiv preprint arXiv:1906.07155*, 2019. 5
- [7] Nicholas Chimitt and Stanley H. Chan. Simulating anisoplanatic turbulence by sampling intermodal and spatially cor-

related Zernike coefficients. *Optical Engineering*, 59(8):1 – 26, 2020. 2

- [8] Xuerui Dai, Xue Yuan, and Xueye Wei. Tirnet: Object detection in thermal infrared images for autonomous driving. *Applied Intelligence*, 51:1–18, 03 2021. 2
- [9] Kevser Irem Danaci and Erdem Akagunduz. A survey on infrared image and video sets. *arXiv*, abs/2203.08581, 2022.
 2
- [10] Richard L. Espinola, Kevin R. Leonard, Kenneth A. Byrd, and Guy Potvin. Atmospheric turbulence and sensor system effects on biometric algorithm performance. In Gerald C. Holst and Keith A. Krapels, editors, *Infrared Imaging Systems: Design, Analysis, Modeling, and Testing XXVI*, volume 9452, pages 120 – 129. International Society for Optics and Photonics, SPIE, 2015. 2
- [11] Chengjian Feng, Yujie Zhong, Yu Gao, Matthew R Scott, and Weilin Huang. Tood: Task-aligned one-stage object detection. In 2021 IEEE/CVF International Conference on Computer Vision (ICCV), pages 3490–3499. IEEE Computer Society, 2021. 4
- [12] FLIR. Free flir thermal dataset for algorithm training. https://www.flir.com/oem/adas/adas-dataset-form/. 2, 3
- [13] Diego Gutierrez, Francisco J. Serón, Adolfo Muñoz, and Oscar Anson. Simulation of atmospheric phenomena. *Comput. Graph.*, 30:994–1010, 2006. 2
- [14] Russell C. Hardie and Daniel A. LeMaster. On the simulation and mitigation of anisoplanatic optical turbulence for long range imaging. In Eric J. Kelmelis, editor, *Long-Range Imaging II*, volume 10204, pages 40 – 59. International Society for Optics and Photonics, SPIE, 2017. 2
- [15] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pages 770–778, 2016. 4
- [16] Christian Herrmann, Miriam Ruf, and Jürgen Beyerer. CNNbased thermal infrared person detection by domain adaptation. In Michael C. Dudzik and Jennifer C. Ricklin, editors, *Autonomous Systems: Sensors, Vehicles, Security, and the Internet of Everything*, volume 10643, pages 38 – 43. International Society for Optics and Photonics, SPIE, 2018. 2
- [17] Gerald Holst. Electro-optical imaging system performance. 01 2000. 1
- [18] I. Ihrke, G. Ziegler, A. Tevs, C. Theobalt, M. Magnor, and H.-P. Seidel. Eikonal rendering: Efficient light transport in refractive objects. *ACM Trans. on Graphics (Siggraph'07)*, page to appear, Aug. 2007. 2
- [19] Licheng Jiao, Fan Zhang, Fang Liu, Shuyuan Yang, Lingling Li, Zhixi Feng, and Rong Qu. A survey of deep learningbased object detection. *IEEE Access*, PP:1–1, 09 2019. 2
- [20] Keith Krapels, Ronald G. Driggers, Richard H. Vollmerhausen, Norman S. Kopeika, and Carl E. Halford. Atmospheric turbulence modulation transfer function for infrared target acquisition modeling. *Optical Engineering*, 40(9):1906–1913, Jan. 2001. 1
- [21] Mate Krišto, Marina Ivasic-Kos, and Miran Pobar. Thermal object detection in difficult weather conditions using yolo. *IEEE Access*, 8:125459–125476, 2020. 2

- [22] Mate Krišto, Marina Ivasic-Kos, and Miran Pobar. Thermal object detection in difficult weather conditions using yolo. *IEEE Access*, 8:125459–125476, 2020. 2
- [23] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollar. Focal loss for dense object detection. In *Proceedings of the IEEE international conference on computer vision*, pages 2980–2988, 2017. 4
- [24] Li Liu, Wanli Ouyang, Xiaogang Wang, Paul W. Fieguth, Jie Chen, Xinwang Liu, and Matti Pietikäinen. Deep learning for generic object detection: A survey. *International Journal* of Computer Vision, 128:261–318, 2019. 2
- [25] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In 2021 IEEE/CVF International Conference on Computer Vision (ICCV), pages 9992–10002, 2021. 4
- [26] Yifei Lou, Sung Ha Kang, Stefano Soatto, and Andrea L. Bertozzi. Video stabilization of atmospheric turbulence distortion. *Inverse Problems Imaging*, 7(3):839–861, 2013. 2
- [27] Farzeen Munir, Shoaib Azam, and Moongu Jeon. Sstn: Self-supervised domain adaptation thermal object detection for autonomous driving. In 2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pages 206–213, 2021. 2
- [28] Shyam Rai and C. Jawahar. Learning to Generate Atmospheric Turbulent Images, pages 261–271. 11 2020. 2
- [29] Joseph Redmon, Santosh Divvala, Ross Girshick, and Ali Farhadi. You only look once: Unified, real-time object detection. In 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pages 779–788, 2016. 2
- [30] Joseph Redmon and Ali Farhadi. Yolo9000: Better, faster, stronger. 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pages 6517–6525, 2017. 2
- [31] Joseph Redmon and Ali Farhadi. Yolov3: An incremental improvement. ArXiv, abs/1804.02767, 2018. 2
- [32] Endre Repasi and Robert Weiss. Computer simulation of image degradations by atmospheric turbulence for horizontal views. In Gerald C. Holst and Keith A. Krapels, editors, *Infrared Imaging Systems: Design, Analysis, Modeling, and Testing XXII*, volume 8014, pages 279 – 287. International Society for Optics and Photonics, SPIE, 2011. 2
- [33] Jakub Rosner, Hannes Fassold, Peter Schallauer, and Werner Bailer. Fast gpu-based image warping and inpainting for frame interpolation. 2012. 2
- [34] Armin Schwartzman, Marina Alterman, Rotem Zamir, and Yoav Y. Schechner. Turbulence-induced 2d correlated image distortion. In 2017 IEEE International Conference on Computational Photography (ICCP), pages 1–13, 2017. 2
- [35] Zhi Tian, Chunhua Shen, Hao Chen, and Tong He. Fcos: Fully convolutional one-stage object detection. In Proceedings of the IEEE/CVF international conference on computer vision, pages 9627–9636, 2019. 4
- [36] Chien-Yao Wang, I-Hau Yeh, and Hongpeng Liao. You only learn one representation: Unified network for multiple tasks. *ArXiv*, abs/2105.04206, 2021. 2, 4, 5
- [37] Wendell R. Watkins. Environmental bugs invade EO imaging systems. In Gerald C. Holst, editor, *Infrared Imaging*

Systems: Design, Analysis, Modeling, and Testing IV, volume 1969, pages 42 – 53. International Society for Optics and Photonics, SPIE, 1993. 1

- [38] Wendell R. Watkins, Daniel R. Billingsley, Fernando R. Palacios, Samuel B. Crow, and Jay B. Jordan. Characterization of the atmospheric modulation transfer function using the target contrast characterizer. In Wendell R. Watkins and Dieter Clement, editors, *Characterization, Propagation, and Simulation of Sources and Backgrounds*, volume 1486, pages 17 24. International Society for Optics and Photonics, SPIE, 1991. 1
- [39] Haoyang Zhang, Ying Wang, Feras Dayoub, and Niko Sünderhauf. Varifocalnet: An iou-aware dense object detector. arXiv preprint arXiv:2008.13367, 2020. 4
- [40] Shifeng Zhang, Cheng Chi, Yongqiang Yao, Zhen Lei, and Stan Z Li. Bridging the gap between anchor-based and anchor-free detection via adaptive training sample selection. In Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, pages 9759–9768, 2020. 4
- [41] Xiang Zhu and Peyman Milanfar. Removing atmospheric turbulence via space-invariant deconvolution. *IEEE Trans*actions on Pattern Analysis and Machine Intelligence, 35(1):157–170, 2013. 2