

Rethinking Boundary Discontinuity Problem for Oriented Object Detection

Hang Xu^{1,2*}, Xinyuan Liu^{2,3*}, Haonan Xu^{2,3}, Yike Ma², Zunjie Zhu^{1,4}, Chenggang Yan¹, Feng Dai^{2†}

¹Hangzhou Dianzi University, Hangzhou, China

²Institute of Computing Technology, Chinese Academy of Sciences, Beijing, China

³University of Chinese Academy of Sciences, Beijing, China

⁴Lishui Institute of Hangzhou Dianzi University, Lishui, China

{h xu, zunjie zhu, cgyan}@hdu.edu.cn {liuxinyuan21s, xuhaonan23s, ykma, fdai}@ict.ac.cn

Abstract

Oriented object detection has been developed rapidly in the past few years, where rotation equivariance is crucial for detectors to predict rotated boxes. It is expected that the prediction can maintain the corresponding rotation when objects rotate, but severe mutation in angular prediction is sometimes observed when objects rotate near the boundary angle, which is well-known boundary discontinuity problem. The problem has been long believed to be caused by the sharp loss increase at the angular boundary, and widely used joint-optim IoU-like methods deal with this problem by loss-smoothing. However, we experimentally find that even state-of-the-art IoU-like methods actually fail to solve the problem. On further analysis, we find that the key to solution lies in encoding mode of the smoothing function rather than in joint or independent optimization. In existing IoU-like methods, the model essentially attempts to fit the angular relationship between box and object, where the break point at angular boundary makes the predictions highly unstable. To deal with this issue, we propose a dual-optimization paradigm for angles. We decouple reversibility and joint-optim from single smoothing function into two distinct entities, which for the first time achieves the objectives of both correcting angular boundary and blending angle with other parameters. Extensive experiments on multiple datasets show that boundary discontinuity problem is well-addressed. Moreover, typical IoU-like methods are improved to the same level without obvious performance gap. The code is available at <https://github.com/hangxu-cv/cvpr24acm>.

1. Introduction

As an expansion of horizontal object detection [11, 12, 20], oriented object detection has a wider applications in many

*Equal contribution

†Corresponding author




Figure 1. Two optimization paradigms for angle in oriented object detection: (a) in joint-optim methods [35, 36, 38, 45], smoothing function is *explicitly* applied for detector’s output θ_p during loss calculation; (b) while in independent-optim methods [31, 41], smoothing function is *implicitly* embedded in the model, and θ_p is decoded from detector’s output f_p . According to our analysis, only the latter can really solve boundary discontinuity problem.

scenes, such as aerial images [3, 30], panoramic images [13, 28, 29], scene text [9], 3D objects [43], etc, since it can achieve a good balance between fine localization and low labeling cost. In oriented object detection, a detector needs to predict the minimal rotated bounding boxes for objects, so it has a high requirement for rotation equivariance. However, researchers have observed mutation in angular prediction when objects rotate near the boundary angle, which is commonly known as boundary discontinuity problem [31, 35].

In previous works, the boundary discontinuity problem has been long believed to be caused by the sharp loss increase at the angular boundary during training. To address this problem, researchers designed a series of smooth loss functions to prevent the sharp loss increase, and these methods can be divided into two categories, i.e., independent-optim loss [31, 33, 41] and joint-optim loss (dominated by IoU-like loss) [35, 36, 38, 45]. Due to the negative impact of the low consistency between loss and IoU-metric, the detectors trained through the independent-optim loss are usually worse than IoU-like loss. It has long been a consensus in object detection [21, 40, 44], so increasing IoU-like loss methods become mainstream choices for oriented object detectors.

However, we experimentally find that even state-of-the-art IoU-like methods do not actually solve the boundary

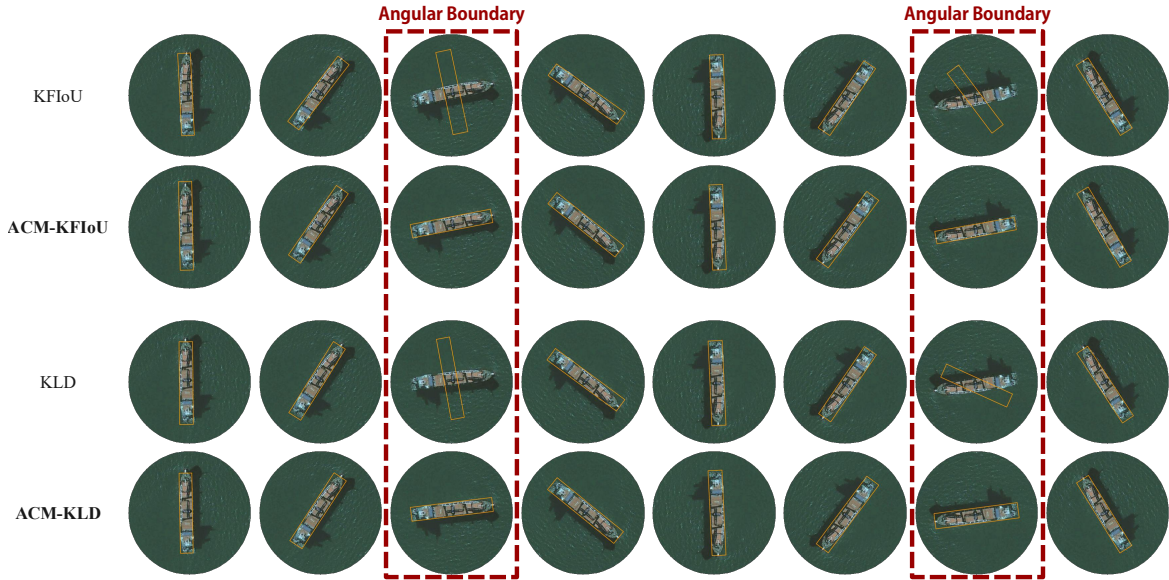


Figure 2. When objects rotate near the boundary angle, state-of-the-art IoU-like methods (e.g., KFIoU [38], KLD [36]) actually suffer from severe mutation in angular prediction. With the correction for angle by our ACM, the prediction achieves rotation equivariance.

discontinuity problem. Specifically, we select an image containing only a single object, and rotate it 360° at 1° intervals to obtain a series of images. These images are sequentially fed into a well-trained detector (with state-of-the-art IoU-like methods) for inference. As is shown in Fig. 2, visualized results show that the predicted boxes can tightly enclose object in most cases, but collapse with a seriously deviated angle in some cases near the angular boundary.

Through theoretical analysis, we find that the key to addressing the problem lies in the encoding mode of the smoothing function rather than in joint or independent optimization. Although both optimization paradigms insist on loss-smoothing, the joint-optim methods have a subtle technical detail differing with independent-optim methods. As is shown as Fig. 1, in joint-optim methods [35, 36, 38, 45], smoothing function is *explicitly* applied for detector’s output θ_p during loss calculation; while in independent-optim method [31, 41], smoothing function is *implicitly* embedded in the model, and θ_p is decoded from detector’s output f_p . For example, in typical joint-optim method KLD [36], Gaussian distribution is transformed from predicted angle and other parameters, not directly output from the model. It is this detail that makes those IoU-like methods not really solve boundary discontinuity problem as they expect, even though they indeed improve the overall detection performance with the benefit of joint optimization. Specifically, the model still attempts to fit the angular relationship between box and object. The relationship is actually a piecewise function with a break point at the angular boundary as Fig. 3b, which is difficult to fit for intrinsically continuous neural networks [2, 15, 46]. It makes angles highly unstable near breakpoints,

and results in the boundary discontinuity problem. Such being the case, an intuitive idea occurs that lets the model output a Gaussian distribution. However, it is challenging to recover the original rotation angles of bounding boxes from Gaussian distributions. If we want to have one’s cake and eat it too, we must find a coding function that simultaneously satisfies the smooth, joint, and reversible characteristics.

To deal with this issue, we propose a dual-optimization paradigm for angles as Fig. 4. We decouple reversibility and joint-optim from single smoothing function into two distinct entities f and g . The former corrects angular boundary, while the latter blends angle with other parameters. In this paradigm, the model outputs angular encoding f_p , subject to explicit supervision. On this basis, another joint-optim g is applied into decoded angle $\theta_p = f_p^{-1}$. Obviously, the role of g can be played by existing joint-optim methods. However, given that f^{-1} is involved in loss calculation, it is necessary to ensure that f^{-1} is differentiable, which is not satisfied for lots of existing encoding. Inspired by the continuous encoding of PSC [41], we propose a coding function based on the complex-exponential function, achieving the goal of differentiability of the inverse function. Finally, boundary discontinuity problem is well-addressed as Fig. 2. **Overall, our contribution can be summarized as following:**

- We extract and induce the optimization logic of existing methods from mathematical perspective, for the first time clarifying the long-standing misunderstanding that IoU-like methods can solve boundary problem.
- We propose a novel dual-optimization paradigm for angles, which for the first time achieves the objectives of both correcting angular boundary and blending param-

ters, achieving rotational equivariance for detection.

- Extensive experiments on multiple datasets show that boundary discontinuity problem is well-addressed. Moreover, typical IoU-like methods are improved to the same level without obvious performance gap.

2. Related Works

2.1. Rotated Object Detection

In oriented detection, the minimal enclosing rotated bounding box (x, y, w, h, θ) is adopted widely to represent an oriented object, where (x, y) is center *position*, (w, h) is *scale* (i.e., width & height) and θ is rotated *angle* of box. There are many algorithms inherited from classic horizontal detection [5, 11, 12, 20] to predict the rotated boxes, where ROI-Transformer [3], SCRDet [32], ReDet [7] are two-stage mainstreamed methods, while DRN [19], R³Det [34], S²A-Net [6] are single-stage methods. However, these detectors suffer from boundary discontinuity problems in varying degrees, as the issue itself is unrelated to the detectors.

2.2. Boundary Discontinuity Problem

The boundary discontinuity problem has been a persistent challenge, requiring a comprehensive understanding of the antecedents and consequences of each milestone to grasp the essence of this paper. In horizontal detection, bbox-regression loss typically employs joint-optimum IoU-Loss, which has reached a consensus without controversy. Due to the complexity and non-differentiability of IoU calculation for rotated box, it was initially considered that IoU-Loss can not be available for oriented detection. Therefore, early methods in oriented detection usually used L1-Loss for each parameters (x, y, w, h, θ) .

CSL [31] pointed out that using L1-Loss would lead to sharp increases in angle-regression loss at angle boundaries, termed "boundary discontinuity problem". By using angle classification instead of angle regression, CSL avoids the intractable problem. Subsequently, a series of methods (e.g., DCL [33] / GF-CSL [25] / MGAR [22]) based on angle classification have sprung up.

GWD [35] argued that while CSL solved the "boundary discontinuity problem" caused by sharp loss increases, independently optimizing parameters was unreasonable. This is because IoU-Loss was already established as the best choice in horizontal detection. However, since rotated IoU is non-differentiable, GWD proposed a Gaussian-based joint-optimum loss to approximately replace it. Hence, GWD claimed that it can address the "boundary discontinuity problem" and achieve joint optimization. KLD [36] and KFIOU [38] inherit the advantages of GWD's Gaussian encoding, and improve it from distribution measurement. Due to the remarkable effect of these methods, more and more Gaussian methods have emerged, which indicates that joint-optimum methods

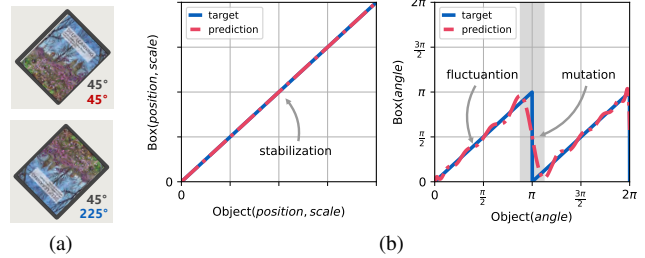


Figure 3. $\text{Box} \neq \text{Object}$: (a) objects rotated with 45° and 225° [colorful mark] share the same box rotated with 45° [black mark], which causes (b) the relationship [blue line] between *angle* of box and object to become a piecewise function with a breakpoint [gray region], differing from the $(\text{position}, \text{scale})$. Not only is the prediction [red line] of the breakpoint region mutational, but the prediction of other regions also becomes fluctuant.

have become mainstream. **Notably, the perception of the "boundary discontinuity problem" remained limited to sharp loss increases up to this point.**

Recently, PSC [41] borrows phase-shift-coding from the field of communications to improve the performance of angle prediction. It uses continuous coding to avoid quantization errors in classification methods, but it still belongs to independent optimization. Notably, **PSC focuses on coding design without new insight about boundary discontinuity problem** (e.g., it explicitly mentioned that GWD/KLD solved the boundary problem).

3. Preliminary

3.1. The Root of All Evil is "Box \neq Object"

For an oriented object detector, it accepts image of object as input, and outputs bounding box with *position*, *scale* and *angle* parameters. However, we reveal that box and object are essentially different concepts, which will produce breakpoints in the angular ground-truth. The discontinuous ground-truth cannot be fitted exactly by continuous output of the detector especially at the breakpoints, so angular prediction near the breakpoints becomes very unstable.

In the interest of brevity, we denote the object instance and bounding box as $O(x_{obj}, y_{obj}, w_{obj}, h_{obj}, \theta_{obj})$ and $B(x_{box}, y_{box}, w_{box}, h_{box}, \theta_{box})$, respectively. The difference between O and B lies in θ rather than (x, y) and (w, h) , where the range of θ_{obj} is $[0, 2\pi)$ while the range of θ_{box} is $[0, \pi)$. This is because the object holds content which needs to rotate at least one full circle to be completely overlapped, while the box is a kind of geometry without any content which just needs to rotate half of circle to be completely overlapped. For example in Fig. 3a, objects rotated with 45° and 225° can be distinguished by content, while the corresponding bounding boxes cannot as well.

In this setting, the bounding box is a truly symmetric

rectangle, whose rotations θ and $\theta \pm \pi$ are indistinguishable. As a result, the relationship between θ_{box} and θ_{obj} exhibits a piecewise function with a break point, rather than a linear relationship between $(x_{box}, y_{box}, w_{box}, h_{box})$ and $(x_{obj}, y_{obj}, w_{obj}, h_{obj})$, as is shown in Fig. 3b and Eq. (1).

$$\begin{cases} (x_{box}, y_{box}) = (x_{obj}, y_{obj}) \\ (w_{box}, h_{box}) = (w_{obj}, h_{obj}) \\ \theta_{box} = \theta_{obj} \bmod \pi \end{cases} \quad (1)$$

The detector takes the object image as input and the box as supervision, which means that the detector is actually enforced to fit Eq. (1) (blue solid lines in Fig. 3b). Obviously, θ_{box} has a step-point at $\theta_{obj} = \pi$, which makes it difficult for the detector F , a continuous function essentially, to fit it accurately. Irrespective of the quality of fit achieved by detector F , there always exists a small interval $(\pi - \epsilon_1, \pi + \epsilon_2)$ near the breakpoint (gray region in Fig. 3b), where predicted angle (red dash line) drops rapidly from π to 0, and angular prediction becomes highly unstable, resulting in a severe degradation of the AP/IoU of boxes. In addition, angular prediction tends to fluctuate even outside the interval.

3.2. The Devil is in Encoding Mode

For the problem of angle discontinuity at the boundary, the core of the mainstream solutions is to smooth loss value at the angular boundary, and these studies are usually categorized by independent or joint optimization. However, our experiments as Fig. 2 shows that even joint-optm methods do not actually solve the boundary discontinuity problem.

To understand the reason behind this finding, we make a reformulation of the existing works. For convenience, let the ground-truth and prediction of θ_{box} be denoted as θ_t and θ_p , respectively. The way to optimize angle in joint-optm methods can be reformulated as follows (also as Fig. 1):

$$\theta = \arg \min_{\theta_p} \ell(f(\theta_p); f(\theta_t)) \quad (2)$$

where model fits discontinuous θ_p , f and ℓ are the encoding function for angle and measuring function for encoded value, respectively. For example, **1**) in the case of KLD [36], $f = gaussian_{x,y,w,h}(\theta)$, $\ell = \ell_{kld}$. f encodes the angle and other parameters as a smooth Gaussian distribution, and ℓ just measures the distance of Gaussian distribution between prediction and ground-truth; **2**) in the case of SkewIoU [45], f and ℓ are implicit functions derived from $SkewIoU(\theta_p^{xywh}, \theta_t^{xywh})$. Although we cannot get explicit expression of f and ℓ , their role must be similar to $gaussian_{x,y,w,h}(\theta)$ and ℓ_{kld} .

As a contrast, the way to optimize angle in independent-optm methods can be formulated as follows (also as Fig. 1):

$$\theta = f^{-1} \left(\arg \min_{f_p} \ell(f_p; f(\theta_t)) \right) \quad (3)$$

where model fits continuous f_p , f^{-1} is the inverse function of f , and we can get angle by $\theta_p = f_p^{-1}$. For example, **1**) in the case of CSL [31], $f = onehot(\theta)$, $\ell = \ell_{focal}$. f encodes the angle into a discrete distribution, and ℓ measures quality of classification; **2**) in the case of PSC [41], $f = \cos(\theta + \varphi_i)$, $i = 1 \dots N$, $\ell = \ell_{l1}$. f encodes the angle into a continuous vector, ℓ measures the encoded vector distance.

Compared with the diverse optimization forms (independent or joint) for f , what is more noteworthy is encoding mode of f . Note that the model in Eq. (2) outputs θ , f is explicitly applied in loss calculation, while the model in Eq. (3) directly outputs the value encoded by f . For detector F , the former's fitting target is still $\theta_{box} \sim \theta_{obj}$ with a break point, while the latter's target becomes $f_{box} \sim f_{obj}$. Thanks to the periodic aggregation properties of f , the differences between box and object are eliminated, which will no longer suffer from difficulty about fitting breakpoints.

To summary up, it is a better choice to make model directly fit the smooth value rather than utilize it just in loss calculation. As for the reason why joint-optm methods do not adopt such design, it is most likely because it is difficult to recover the angle from the joint-encoding of the model output. Dramatically, the advantages of joint optimization outweigh the disadvantages of loss-smoothing, which eventually misleads researchers to believe that the boundary problem can be solved by joint optimization.

4. Method

4.1. Dual-Optimization for Angle

Considering that joint optimization has become the mainstream scheme at present, a convenient improvement strategy is to correct angle by angle-smoothing independent optimization with reversible coding, as well as to blend angle with other parameters by joint optimization based on corrected angle, which can be formulated as follows:

$$\begin{aligned} \theta &= f^{-1} \left(\arg \min_{z=f_p} [\ell_f + \ell_g] \right) \\ s.t. \ell_f &= \ell(z; f(\theta_t)) \\ \ell_g &= \ell(g(f^{-1}(z)); g(\theta_t)) \end{aligned} \quad (4)$$

where model still fits continuous f_p , and f, g are encoding function in independent/joint optimization.

Since f^{-1} participates in loss calculation in joint-optimization, f not only needs to be continuous, differentiable, and reversible (GWD/KLD/FKIoU/SkewIoU fail to satisfy), but its inverse function f^{-1} also needs to satisfy these properties. Discrete encodings like CSL rely on argmin to make f^{-1} nondifferentiable, so only continuous encodings like PSC remains a chance to be differentiable.

To this end, we propose a **Angle Correct Module (ACM)** based on complex-exponential function. The module implements the above f and f^{-1} , which can be easily plugged into

Table 1. Ablation study of different encoding length.

Method	Encoding Length	HRSC2016	
		AP ₅₀	AP ₇₅
Direct	1	88.26	62.95
CSL	3	19.50	2.56
	60	48.96	13.58
	90	90.49	61.43
	180	90.53	77.76
PSC	3	89.91	79.20
	20	90.55	79.54
	60	90.62	79.86
	180	90.56	79.51
ACM	2	90.57	86.33

on independently-optim encoding methods, especially PSC [41], which is a continuous encoding as ACM. ACM and PSC are distinct under any circumstances. Basically, the encoding formula of PSC is $\{\cos(\omega\theta + \varphi_i) \mid i = 1 \dots N, N \geq 3\}$, while ACM is $e^{j\omega\theta} = \cos(\omega\theta) + j \sin(\omega\theta)$. Obviously, ACM has a different theory and form. Moreover, they cannot be interconverted, even with PSC’s N at 4. Specifically, when PSC’s N equals 4, the encoding vector becomes $\mathbf{f} = (-c, -s, c, s)$, where c & s are abbreviation of \cos & \sin . Note that \mathbf{f} is just ground-truth, its prediction $\hat{\mathbf{f}}$ has some deviation e compared with \mathbf{f} , i.e. $\hat{\mathbf{f}} = \mathbf{f} + e$. But corresponding prediction-heads are independent modules without pairwise-constraints ($e_1 + e_3 = 0, e_2 + e_4 = 0$), so $\hat{\mathbf{f}}$ is no longer ensured as two duplicate values (c, s), i.e., $\hat{f}_1 \neq -\hat{f}_3, \hat{f}_2 \neq -\hat{f}_4$. In contrast, complex encoding $e^{j\omega\theta}$ of ACM is equivalent to vector (c, s) , where pairwise-constraints have been implied. Consequently, the decoded angular prediction will naturally differ from ACM. **The stronger constraints/priors implied in ACM reduce optimization difficulty of entire model, making ACM better.**

5. Experiment

5.1. Datasets

DOTA[27] is one of the largest datasets for oriented object detection in aerial images, which contains 2,806 images with fifteen categories of 188,282 instances in total. The training, validation and testing set include 1411, 458 and 937 images, respectively. The categories are defined as: Plane (PL), Baseball Diamond (BD), Bridge (BR), Ground Field Track (GTF), Small Vehicle (SV), Large Vehicle (LV), Ship (SH), Tennis Court (TC), Basketball Court (BC), Storage Tank (ST), Soccer-Ball Field (SBF), Roundabout (RA), Harbor (HA), Swimming Pool (SP), and Helicopter (HC). We crop training images into the patches of size 1024×1024 pixels with an overlap of 256 pixels. When testing, we crop the testing set images into 4000×4000 patches with an overlap of 1024 pixels, to mitigate the negative impact of the cutting.

Table 2. Ablation study of different encoding mode.

Method	Encoding Mode	HRSC2016		DOTA-v1.0	
		AP ₅₀	AP ₇₅	AP ₅₀	AP ₇₅
Direct	n/a	88.26	62.95	71.97	26.11
CSL	implicit	90.53	77.76	70.83	38.71
	explicit	6.06	1.05	33.29	10.90
PSC	implicit	89.91	79.20	71.41	39.35
	explicit	53.43	33.65	50.02	23.08
ACM	implicit	90.57	86.33	74.99	41.44
	explicit	54.66	31.45	50.67	19.91

Table 3. Ablation study of different supervision.

Model Output	\mathcal{L}_{acm}	\mathcal{L}_{box}	HRSC2016		DOTA	
			AP ₅₀	AP ₇₅	AP ₅₀	AP ₇₅
θ	n/a	n/a	88.26	62.95	71.97	26.11
f_p	✓		90.57	86.33	74.99	41.44
		✓	37.37	13.98	54.67	19.67
	✓	✓	90.47	88.33	74.21	42.83

HRSC2016[14] contains images from two scenarios with ships on sea and close inshore. The training, validation and testing set include 436, 181 and 444 images, with the image size ranging from 300×300 to 1500×900 . We adjust the long side of each image to a fixed size (640 pixels) and keep the original aspect ratio for training and testing.

UCAS-AOD[49] contains two categories: Car and Plane, which includes 1,510 aerial images of about $659 \times 1,280$ pixels, with two categories of 14,596 instances in total. We randomly select 1,110 for training and 400 for testing. Finally, we adopt the same data processing strategy as HRSC2016.

5.2. Implementation Details

Evaluation Metric. Methods are evaluated using the standard COCO style Average Precision (AP) [10], which is the convention throughout the field of object detection. It is worth noting that AP₇₅ is gradually replacing AP₅₀ as the most reliable metric for oriented object detection due to AP₇₅’s higher sensitivity to angle deviation than AP₅₀. Following mainstream works [36, 42], we adopt AP₇₅ as main metric, while AP₅₀ is auxiliary metric.

Training Details. All approaches are implemented in PyTorch, and training is done on NVIDIA RTX 3090 GPUs. We choose the anchor-free method CenterNet [48] to build the rotated detector and ImageNet pretrained ResNet-50 [8] as the backbone. The network is optimized by Adam for 140 epochs with the learning rate dropped by $10 \times$ at 100 and 130 epochs. As the DOTA dataset takes a large image resolution as an input, the batch size is set as 12 with an initial learning rate 1.25×10^{-4} . For the HRSC2016 and UCAS-AOD datasets, the batch size is set as 32, and the initial learning rates are set as 2×10^{-4} and 1×10^{-4} , respectively.

Table 4. Ablation study on various datasets with different joint-optim loss. Base detector is CenterNet.

Method	HRSC2016 (Ship)		UCAS-AOD (Car)		UCAS-AOD (Plane)		DOTA-v1.0	
	AP ₅₀	AP ₇₅	AP ₅₀	AP ₇₅	AP ₅₀	AP ₇₅	AP ₅₀	AP ₇₅
GWD [35]	84.94	61.87	87.25	28.46	90.34	38.22	73.12	34.98
ACM-GWD	90.63 (+5.69)	86.71 (+24.84)	88.69 (+1.44)	29.15 (+0.69)	90.35 (+0.01)	76.00 (+37.78)	73.71 (+0.59)	41.97 (+6.99)
KLD [36]	90.01	79.29	87.54	29.99	90.33	29.19	73.41	35.25
ACM-KLD	90.55 (+0.54)	87.45 (+8.16)	88.76 (+1.22)	30.40 (+0.41)	90.39 (+0.06)	75.65 (+46.46)	73.95 (+0.54)	42.97 (+7.72)
KFIoU [38]	88.26	62.95	85.74	24.44	90.34	16.81	71.97	26.11
ACM-KFIoU	90.55 (+2.29)	87.77 (+24.82)	88.31 (+2.57)	34.81 (+10.37)	90.40 (+0.06)	74.48 (+57.67)	74.51 (+2.54)	40.49 (+14.38)
SkewIoU [45]	89.39	76.43	87.73	27.59	90.34	63.64	73.62	38.01
ACM-SkewIoU	90.47 (+1.08)	88.33 (+11.09)	88.27 (+0.54)	29.13 (+1.74)	90.37 (+0.03)	75.13 (+11.49)	74.21 (+0.59)	42.83 (+4.37)




Figure 5. Visualized comparison of detection results between KFIoU [38] and enhanced ACM-KFIoU. The images are arranged from left to right in order of increasing aspect-ratio of objects, and the first-col and bottom-col are the results of KFIoU and ACM-KFIoU, respectively. Our ACM greatly eliminates the angular prediction errors in the original KFIoU.

5.3. Ablation Studies

Different encoding length. To comprehensively compare the costs and performance of various encoding methods, we conducted experiments as Tab. 1. The results indicate that CSL relies on relatively long encoding, while the choice of coding length for PSC is challenging. Compared with these methods, ACM achieves the best performance without the need to choose the encoding length, which is fixed to 2.

Different encoding mode. To validate our motivation, that is, implicit encoding is superior to explicit encoding, we conducted comparative experiments based on various encoding methods as Tab. 2. The results show that, even for the identical encoding methods, performance is notably weaker when predicting the angle itself (*explicit*) rather than angular encoding (*implicit*). This is an observation that has long been overlooked, yet it points towards a viable direction for addressing boundary discontinuity problem in the future.

Different supervision. To validate the necessity of each supervision of proposed dual-optimization, we conducted experiments as Tab. 3. The results show that, compared with full dual-optimization, performance decreased (yet still surpass baseline) after the removal of refine supervision, but there was a severe decline in performance after the removal of correct supervision. We think it is difficult to ensure that predicted frequencies align with the preset decoded frequencies if without explicit supervision, and such mismatch will

make decoded angles fall into a wrong range, so performance degenerates. By the way, since the size of DOTA dataset is much larger than HRSC2016 dataset, the mismatch can get some mitigation, but just like a drop in the ocean.

Different joint-optim loss on various datasets. To eliminate the influence of the classification branch on the detection results, we conducted experiments on the datasets comprising simple scenes with only single-category objects per image. The HRSC2016 dataset contains large aspect ratio ships, and the UCAS-AOD dataset contains rectangle cars and square-like planes. As is shown in Tab. 4, both AP₅₀ and AP₇₅ get significant improvement from ACM on HRSC2016(Ship) and UCAS-AOD(Car).

It is worth noting that the improvement of AP₅₀ are negligible on the UCAS-AOD(Plane) dataset, while the improvement of AP₇₅ are tremendous. It is never an accident, and the reasons include: **1)** For square-like objects, the IoU is always over 0.5 regardless of the angle of the predicted box, making AP₅₀ insensitive to square-like objects. **2)** When a square-like box is converted to a 2D-Gaussian distribution, the 2D-Gaussian distribution is completely symmetric like a circle, which makes it impossible for these methods (GWD, KLD, KFIoU) based on 2D-Gaussian distribution to accurately predict the angle of square-like objects. Since our ACM is friendly to square-like objects, it greatly improves these baseline methods based on 2D-Gaussian distribution

Table 5. Performance of different detectors on DOTA-v1.0 dataset. MS indicates that multi-scale training/testing is used.

Method	MS	PL	BD	BR	GTF	SV	LV	SH	TC	BC	ST	SBF	RA	HA	SP	HC	AP ₅₀
PloU [1]		80.90	69.70	24.10	60.20	38.30	64.40	64.80	90.90	77.20	70.40	46.50	37.10	57.10	61.90	64.00	60.50
RoI-Trans. [4]	✓	88.64	78.52	43.44	75.92	68.81	73.68	83.59	90.74	77.27	81.46	58.39	53.54	62.83	58.93	47.67	69.56
O ² -DNet [26]	✓	89.31	82.14	47.33	61.21	71.32	74.03	78.62	90.76	82.23	81.36	60.93	60.17	58.21	66.98	61.03	71.04
DAL [18]	✓	88.61	79.69	46.27	70.37	65.89	76.10	78.53	90.84	79.98	78.41	58.71	62.02	69.23	71.32	60.65	71.78
P-RSDet [47]	✓	88.58	77.83	50.44	69.29	71.10	75.79	78.66	90.88	80.10	81.71	57.92	63.03	66.30	69.77	63.13	72.30
BBAVectors [39]	✓	88.35	79.96	50.69	62.18	78.43	78.98	87.94	90.85	83.58	84.35	54.13	60.24	65.22	64.28	55.70	72.32
DRN [19]	✓	89.71	82.34	47.22	64.10	76.22	74.43	85.84	90.57	86.18	84.89	57.65	61.93	69.30	69.63	58.48	73.23
CFC-Net [16]	✓	89.08	80.41	52.41	70.02	76.28	78.11	87.21	90.89	84.47	85.64	60.51	61.52	67.82	68.02	50.09	73.50
Gliding Vertex [30]		89.64	85.00	52.26	77.34	73.01	73.14	86.82	90.74	79.02	86.81	59.55	70.91	72.94	70.86	57.32	75.02
Mask OBB [23]	✓	89.56	85.95	54.21	72.90	76.52	74.16	85.63	89.85	83.81	86.48	54.89	69.64	73.94	69.06	63.32	75.33
CenterMap [24]	✓	89.83	84.41	54.60	70.25	77.66	78.32	87.19	90.66	84.89	85.27	56.46	69.23	74.13	71.56	66.06	76.03
CSL [31]	✓	90.25	85.53	54.64	75.31	70.44	73.51	77.62	90.84	86.15	86.69	69.60	68.04	73.83	71.10	68.93	76.17
R ³ Det [34]	✓	89.80	83.77	48.11	66.77	78.76	83.27	87.84	90.82	85.38	85.51	65.67	62.68	67.53	78.56	72.62	76.47
GWD [35]	✓	86.96	83.88	54.36	77.53	74.41	68.48	80.34	86.62	83.41	85.55	73.47	67.77	72.57	75.76	73.40	76.30
SCRDet++ [37]	✓	90.05	84.39	55.44	73.99	77.54	71.11	86.05	90.67	87.32	87.08	69.62	68.90	73.74	71.29	65.08	76.81
KFIoU [38]	✓	89.46	85.72	54.94	80.37	77.16	69.23	80.90	90.79	87.79	86.13	73.32	68.11	75.23	71.61	69.49	77.35
DCL [33]	✓	89.26	83.60	53.54	72.76	79.04	82.56	87.31	90.67	86.59	86.98	67.49	66.88	73.29	70.56	69.99	77.37
RIDet [17]	✓	89.31	80.77	54.07	76.38	79.81	81.99	89.13	90.72	83.58	87.22	64.42	67.56	78.08	79.17	62.07	77.62
PSC [41]	✓	89.86	86.02	54.94	62.02	81.90	85.48	88.39	90.73	86.90	88.82	63.94	69.19	76.84	82.75	63.24	78.07
KLD [36]	✓	88.91	85.23	53.64	81.23	78.20	76.99	84.58	89.50	86.84	86.38	71.69	68.06	75.95	72.23	75.42	78.32
CenterNet-ACM	✓	89.84	85.50	53.84	74.78	80.77	82.81	88.92	90.82	87.18	86.53	64.09	66.27	77.51	79.62	69.57	78.53
RoI-Trans.-ACM	✓	85.55	80.53	61.21	75.40	80.35	85.60	88.32	89.88	87.13	87.10	68.15	67.94	78.75	79.82	75.96	79.45

by **37.78%** (GWD), **46.46%** (KLD) and **57.56%** (KFIoU) on AP₇₅ on UCAS-AOD(Plane) dataset.

To explore performance in a more general cases, we conducted experiments on a dataset comprising complex scenes with only single-category objects in each image. DOTA dataset contains a considerable number of categories and diverse environments. Experimental results at Tab. 4 show that the performance of all IoU-like methods are improved by **6.99%** (GWD), **7.72%** (KLD), **14.34%** (KFIoU) and **14.34%** (SkewIoU) on AP₇₅ after the ACM module is used. We also unexpectedly find that after the ACM module enhancement, both Gaussian-based loss and SkewIoU loss become very close in terms of AP₅₀ (**73.71%**, **73.95%**, **74.51%**, **74.21%**) and AP₇₅ (**41.97%**, **42.97%**, **40.49%**, **42.83%**), indicating that the primary distinction between them lies in their optimization capabilities for the angle.

Visualized results. We provide some visualization results on the DOTA v1.0 dataset as Fig. 5. From detection results obtained by the KFIoU-based detector, we select some cases of poor angular prediction. Note that there exists slight angular deviations for boxes in KFIoU results sometimes, and significant angular errors in other times. Fortunately, most angular errors are corrected in the results of ACM-KFIoU. It is also worth noting that ACM addresses the square-like-object case where KFIoU based on 2D Gaussian distribution fails for angular prediction. It is consistent with the quantitative results in Tab. 4, and further verifies the effectiveness of our methods.

5.4. Comparison with the State-of-the-Art

Tab. 5 presents a comprehensive comparison of recent detectors on DOTA-v1.0 dataset. It is important to note that the

performance of different methods may vary due to several factors, including image resolution, network architecture, detection framework, training strategies, and various optimization techniques employed. In light of these variations, it becomes challenging to establish completely fair comparisons among the different approaches. However, despite these challenges, our method has managed to achieve competitive results, at around **78.53%** / **79.45%** on AP₅₀.

6. Conclusion

In this paper, we experimentally find that widely used IoU-like methods do not actually solve the well-known boundary discontinuity problem. On further analysis, we find that the key to solution lies in the encoding mode of the smoothing function rather than in joint or independent optimization. Moreover, we propose a dual-optimization paradigm integrated with complex-exponential angular coding, which achieves the objectives of both correcting angular boundary and blending parameters. Finally, extensive experiments show that our methods effectively eliminate boundary problem and significantly improve detection performance for the object detector.

Acknowledgments

This work is supported by National Key R&D Program of China (2022YFD2001601), National Natural Science Foundation of China (62372433, 62072438, 61931008, U21B2024, 62071415), Zhejiang Provincial Natural Science Foundation of China(LDT23F01011F01, LDT23F01015F01, LDT23F01014F01), "Pioneer" and "Leading Goose" R&D Program of Zhejiang Province (2022C01068).

References

- [1] Zhiming Chen, Kean Chen, Weiyao Lin, John See, Hui Yu, Yan Ke, and Cong Yang. Piou loss: Towards accurate oriented object detection in complex environments. In *European Conference on Computer Vision*, pages 195–211. Springer, 2020. 8
- [2] George Cybenko. Approximation by superpositions of a sigmoidal function. *Mathematics of control, signals and systems*, 2(4):303–314, 1989. 2
- [3] Jian Ding, Nan Xue, Yang Long, Gui-Song Xia, and Qikai Lu. Learning roi transformer for oriented object detection in aerial images. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2849–2858, 2019. 1, 3
- [4] Jian Ding, Nan Xue, Yang Long, Gui-Song Xia, and Qikai Lu. Learning roi transformer for oriented object detection in aerial images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2849–2858, 2019. 8
- [5] Ross Girshick. Fast r-cnn. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1440–1448, 2015. 3
- [6] Jiaming Han, Jian Ding, Jie Li, and Gui-Song Xia. Align deep features for oriented object detection. *IEEE Transactions on Geoscience and Remote Sensing*, 2021. 3
- [7] Jiaming Han, Jian Ding, Nan Xue, and Gui-Song Xia. Redet: A rotation-equivariant detector for aerial object detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2786–2795, 2021. 3
- [8] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 770–778, 2016. 6
- [9] Yingying Jiang, Xiangyu Zhu, Xiaobing Wang, Shuli Yang, Wei Li, Hua Wang, Pei Fu, and Zhenbo Luo. R2cnn: rotational region cnn for orientation robust scene text detection. *arXiv preprint arXiv:1706.09579*, 2017. 1
- [10] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *European Conference on Computer Vision*, pages 740–755. Springer, 2014. 6
- [11] Tsung-Yi Lin, Piotr Dollár, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. Feature pyramid networks for object detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2117–2125, 2017. 1, 3
- [12] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2980–2988, 2017. 1, 3
- [13] Xinyuan Liu, Hang Xu, Bin Chen, Qiang Zhao, Yike Ma, Chenggang Yan, and Feng Dai. Sph2pob: Boosting object detection on spherical images with planar oriented boxes methods. In *International Joint Conference on Artificial Intelligence (IJCAI)*, 2023. 1
- [14] Zikun Liu, Liu Yuan, Lubin Weng, and Yiping Yang. A high resolution optical satellite image dataset for ship recognition and some new baselines. In *Proceedings of the International Conference on Pattern Recognition Applications and Methods*, pages 324–331, 2017. 6
- [15] Zhou Lu, Hongming Pu, Feicheng Wang, Zhiqiang Hu, and Liwei Wang. The expressive power of neural networks: A view from the width. *Advances in neural information processing systems*, 30, 2017. 2
- [16] Qi Ming, Lingjuan Miao, Zhiqiang Zhou, and Yunpeng Dong. Cfc-net: A critical feature capturing network for arbitrary-oriented object detection in remote sensing images. *arXiv preprint arXiv:2101.06849*, 2021. 8
- [17] Qi Ming, Lingjuan Miao, Zhiqiang Zhou, Xue Yang, and Yunpeng Dong. Optimization for arbitrary-oriented object detection via representation invariance loss. *IEEE Geoscience and Remote Sensing Letters*, 2021. 8
- [18] Qi Ming, Zhiqiang Zhou, Lingjuan Miao, Hongwei Zhang, and Linhao Li. Dynamic anchor learning for arbitrary-oriented object detection. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 2355–2363, 2021. 8
- [19] Xingjia Pan, Yuqiang Ren, Kekai Sheng, Weiming Dong, Haolei Yuan, Xiaowei Guo, Chongyang Ma, and Changsheng Xu. Dynamic refinement network for oriented and densely packed object detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 11207–11216, 2020. 3, 8
- [20] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *Advances in Neural Information Processing Systems*, pages 91–99, 2015. 1, 3
- [21] Hamid Rezaatofighi, Nathan Tsoi, JunYoung Gwak, Amir Sadeghian, Ian Reid, and Silvio Savarese. Generalized intersection over union: A metric and a loss for bounding box regression. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 658–666, 2019. 1
- [22] Hao Wang, Zhanchao Huang, Zhengchao Chen, Ying Song, and Wei Li. Multigrained angle representation for remote-sensing object detection. *IEEE Transactions on Geoscience and Remote Sensing*, 60:1–13, 2022. 3
- [23] Jinwang Wang, Jian Ding, Haowen Guo, Wensheng Cheng, Ting Pan, and Wen Yang. Mask obb: A semantic attention-based mask oriented bounding box representation for multi-category object detection in aerial images. *Remote Sensing*, 11(24):2930, 2019. 8
- [24] Jinwang Wang, Wen Yang, Heng-Chao Li, Haijian Zhang, and Gui-Song Xia. Learning center probability map for detecting objects in aerial images. *IEEE Transactions on Geoscience and Remote Sensing*, 59(5):4307–4323, 2020. 8
- [25] Jian Wang, Fan Li, and Haixia Bi. Gaussian focal loss: Learning distribution polarized angle prediction for rotated object detection in aerial images. *IEEE Transactions on Geoscience and Remote Sensing*, 60:1–13, 2022. 3
- [26] Haoran Wei, Yue Zhang, Zhonghan Chang, Hao Li, Hongqi Wang, and Xian Sun. Oriented objects as pairs of middle lines.

- ISPRS Journal of Photogrammetry and Remote Sensing*, 169: 268–279, 2020. 8
- [27] Gui-Song Xia, Xiang Bai, Jian Ding, Zhen Zhu, Serge Belongie, Jiebo Luo, Mihai Datcu, Marcello Pelillo, and Liangpei Zhang. DOTA: A large-scale dataset for object detection in aerial images. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3974–3983, 2018. 6
- [28] Hang Xu, Qiang Zhao, Yike Ma, Xiaodong Li, Peng Yuan, Bailan Feng, Chenggang Yan, and Feng Dai. Pandora: A panoramic detection dataset for object with orientation. In *European Conference on Computer Vision (ECCV)*, pages 237–252, 2022. 1
- [29] Hang Xu, Xinyuan Liu, Qiang Zhao, Yike Ma, Chenggang Yan, and Feng Dai. Gaussian label distribution learning for spherical image object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1033–1042, 2023. 1
- [30] Yongchao Xu, Mingtao Fu, Qimeng Wang, Yukang Wang, Kai Chen, Gui-Song Xia, and Xiang Bai. Gliding vertex on the horizontal bounding box for multi-oriented object detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 43(4):1452–1459, 2020. 1, 8
- [31] Xue Yang and Junchi Yan. Arbitrary-oriented object detection with circular smooth label. In *European Conference on Computer Vision*, pages 677–694. Springer, 2020. 1, 2, 3, 4, 8
- [32] Xue Yang, Jirui Yang, Junchi Yan, Yue Zhang, Tengfei Zhang, Zhi Guo, Xian Sun, and Kun Fu. Srdet: Towards more robust detection for small, cluttered and rotated objects. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 8232–8241, 2019. 3
- [33] Xue Yang, Liping Hou, Yue Zhou, Wentao Wang, and Junchi Yan. Dense label encoding for boundary discontinuity free rotation detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 15819–15829, 2021. 1, 3, 8
- [34] Xue Yang, Junchi Yan, Ziming Feng, and Tao He. R3det: Refined single-stage detector with feature refinement for rotating object. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 3163–3171, 2021. 3, 8
- [35] Xue Yang, Junchi Yan, Qi Ming, Wentao Wang, Xiaopeng Zhang, and Qi Tian. Rethinking rotated object detection with gaussian wasserstein distance loss. In *International Conference on Machine Learning*, pages 11830–11841. PMLR, 2021. 1, 2, 3, 7, 8
- [36] Xue Yang, Xiaojiang Yang, Jirui Yang, Qi Ming, Wentao Wang, Qi Tian, and Junchi Yan. Learning high-precision bounding box for rotated object detection via kullback-leibler divergence. *Advances in Neural Information Processing Systems*, 34, 2021. 1, 2, 3, 4, 6, 7, 8
- [37] Xue Yang, Junchi Yan, Wenlong Liao, Xiaokang Yang, Jin Tang, and Tao He. Srdet++: Detecting small, cluttered and rotated objects via instance-level feature denoising and rotation loss smoothing. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2022. 8
- [38] Xue Yang, Yue Zhou, Gefan Zhang, Jirui Yang, Wentao Wang, Junchi Yan, Xiaopeng Zhang, and Qi Tian. The kfou loss for rotated object detection. In *International Conference on Learning Representations*, 2023. 1, 2, 3, 7, 8
- [39] Jingru Yi, Pengxiang Wu, Bo Liu, Qiaoying Huang, Hui Qu, and Dimitris Metaxas. Oriented object detection in aerial images with box boundary-aware vectors. In *Proceedings of the IEEE Winter Conference on Applications of Computer Vision*, pages 2150–2159, 2021. 8
- [40] Jiahui Yu, Yuning Jiang, Zhangyang Wang, Zhimin Cao, and Thomas Huang. Unitbox: An advanced object detection network. In *Proceedings of the 24th ACM international conference on Multimedia*, pages 516–520, 2016. 1
- [41] Yi Yu and Feipeng Da. Phase-shifting coder: Predicting accurate orientation in oriented object detection. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023. 1, 2, 3, 4, 6, 8
- [42] Ying Zeng, Xue Yang, Qingyun Li, Yushi Chen, and Junchi Yan. Ars-detr: Aspect ratio sensitive oriented object detection with transformer. *arXiv preprint arXiv:2303.04989*, 2023. 6
- [43] Yu Zheng, Danyang Zhang, Sinan Xie, Jiwen Lu, and Jie Zhou. Rotation-robust intersection over union for 3d object detection. In *European Conference on Computer Vision*, pages 464–480. Springer, 2020. 1
- [44] Zhaohui Zheng, Ping Wang, Wei Liu, Jinze Li, Rongguang Ye, and Dongwei Ren. Distance-iou loss: Faster and better learning for bounding box regression. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 12993–13000, 2020. 1
- [45] Dingfu Zhou, Jin Fang, Xibin Song, Chenye Guan, Junbo Yin, Yuchao Dai, and Ruigang Yang. Iou loss for 2d/3d object detection. In *2019 International Conference on 3D Vision*, pages 85–94. IEEE, 2019. 1, 2, 4, 7
- [46] Ding-Xuan Zhou. Universality of deep convolutional neural networks. *Applied and computational harmonic analysis*, 48(2):787–794, 2020. 2
- [47] Lin Zhou, Haoran Wei, Hao Li, Wenzhe Zhao, Yi Zhang, and Yue Zhang. Arbitrary-oriented object detection in remote sensing images based on polar coordinates. *IEEE Access*, 8: 223373–223384, 2020. 8
- [48] Xingyi Zhou, Dequan Wang, and Philipp Krähenbühl. Objects as points. *arXiv preprint arXiv:1904.07850*, 2019. 6
- [49] Haigang Zhu, Xiaogang Chen, Weiqun Dai, Kun Fu, Qixiang Ye, and Jianbin Jiao. Orientation robust object detection in aerial images using deep convolutional neural network. In *2015 IEEE International Conference on Image Processing*, pages 3735–3739. IEEE, 2015. 6