ProbIoU+

Enhanced Probabilistic IoU Loss for Oriented Object Detection

Yasmine Sakas

Pascal Houssam Salmane • Josué Rivera

Patrick Danès • Guillaume Saint Pierre

Université de Toulouse

LAAS CNRS

Cerema
CLIMAT & TERRITOIRES DE DEMAIN

Presentation Overview

1 BACKGROUND

2 MOTIVATION

3 PROPOSED APPROACH

4 EVALUATION RESULTS

Object Detection

Object Detection

Classification

Localization

Horizontal Bounding Boxes
HBB (x,y,w,h)

w h

Oriented Bounding Boxes
OBB (x,y,w,h,θ)

w h θ

Horizontal Object Detection

Oriented Object Detection

Aerial image from the DOTA[ref.1] dataset

ref.1: *Xia et al., CVPR 2018 (DOTA)*

Localization Losses for Horizontal Object Detection

Smooth L1 Loss[ref.2]

The standard regression loss used in Fast-RCNN[ref.2].

Provides a balanced trade-off between the L1 and L2 losses.

Assumes parameter independence and lacks scale invariance.

$$L(z_{pred}, z_{gt}) = \begin{cases} 0.5(z_{pred} - z_{gt})^2 & if\ |z_{pred} - z_{gt}| < 1 \\ |z_{pred} - z_{gt}| - 0.5 & otherwise \end{cases}$$

$$Loss_{smooth\ L1} = \sum_{z \in \{x,y,w,h\}} L(z_{pred}, z_{gt})$$

ref.2: *Girshick, R., "Fast R-CNN," ICCV 2015*

Localization Losses for Horizontal Object Detection

Smooth L1 Loss[ref.2]

The standard regression loss used in Fast-RCNN[ref.2].

Provides a balanced trade-off between the L1 and L2 losses.

Assumes parameter independence and lacks scale invariance.

$$L(z_{pred}, z_{gt}) = \begin{cases} 0.5(z_{pred} - z_{gt})^2 & if\ |z_{pred} - z_{gt}| < 1 \\ |z_{pred} - z_{gt}| - 0.5 & otherwise \end{cases}$$

$$Loss_{smooth\ L1} = \sum_{z \in \{x,y,w,h\}} L(z_{pred}, z_{gt})$$

Intersection over Union (IoU) Loss[ref.3]

Directly optimizes the geometric overlap between boxes.

Scale-invariant.

Zero gradient for non-overlapping boxes (IoU = 0).

$$IoU = \frac{Area\ of\ Overlap}{Area\ of\ Union} = \frac{|B_{pred} \cap B_{gt}|}{|B_{pred} \cup B_{gt}|}$$

$$Loss_{IoU} = 1 - IoU$$

ref.2: Girshick, R., "Fast R-CNN," ICCV 2015 **ref.3:** Yu et al., ACM MM 2016 (UnitBox)

Localization Losses for Horizontal Object Detection

Smooth L1 Loss[ref.2]

The standard regression loss used in Fast-RCNN[ref.2].

Provides a balanced trade-off between the L1 and L2 losses.

Assumes parameter independence and lacks scale invariance.

$$L(z_{pred}, z_{gt}) = \begin{cases} 0.5(z_{pred} - z_{gt})^2 & if\ |z_{pred} - z_{gt}| < 1 \\ |z_{pred} - z_{gt}| - 0.5 & otherwise \end{cases}$$

$$Loss_{smooth\ L1} = \sum_{z \in \{x,y,w,h\}} L(z_{pred}, z_{gt})$$

Intersection over Union (IoU) Loss[ref.3]

Directly optimizes the geometric overlap between boxes.

Scale-invariant.

Zero gradient for non-overlapping boxes (IoU = 0).

$$IoU = \frac{Area\ of\ Overlap}{Area\ of\ Union} = \frac{|B_{pred} \cap B_{gt}|}{|B_{pred} \cup B_{gt}|}$$

$$Loss_{IoU} = 1 - IoU$$

Generalized IoU (GIoU) Loss[ref.4]

Addresses the vanishing gradients of IoU Loss.

HOW? Introduces a penalty based on the smallest enclosing box covering both boxes.

Distance IoU (DIoU) Loss[ref.5]

Addresses the slow convergence of GIoU Loss.

HOW? Introduces a penalty based on the distance between box centers.

Complete IoU (CIoU) Loss[ref.5]

Improves over DIoU Loss.

HOW? Adds a penalty to the DIoU based on the box aspect ratios.

ref.2: *Girshick, R., "Fast R-CNN," ICCV 2015* **ref.3:** *Yu et al., ACM MM 2016 (UnitBox)* **ref.4:** *Rezatofighi et al., CVPR 2019 (GIoU)* **ref.5:** *Zheng et al., AAAI 2020 (DIoU)*

Localization Losses for Oriented Object Detection

Rotated Intersection over Union (RIoU) Loss[ref.6]

$$RIoU = \frac{\left|B_{pred}^{rot} \cap B_{gt}^{rot}\right|}{\left|B_{pred}^{rot} \cup B_{gt}^{rot}\right|} \qquad Loss_{RIoU} = 1 - RIoU$$

Oriented Bounding Box OBB

Rotated Box $B(x, y, w, h, \theta)$

ref.7

Often combined with Smooth L1 loss for precise parameter refinement (e.g., AO2-DETR[ref.8]).

Angle periodicity issue
Gradient instability

ref.6: *Zhou et al., 3DV 2019 (IoU)* ***ref.7:*** *Yang et al., NeurIPS 2021 (KLD)* ***ref.8:*** *Dai et al., TCSVT 2023 (AO2-DETR)*

Localization Losses for Oriented Object Detection

Rotated Intersection over Union (RIoU) Loss[ref.6]

$$RIoU = \frac{\left|B^{rot}_{pred} \cap B^{rot}_{gt}\right|}{\left|B^{rot}_{pred} \cup B^{rot}_{gt}\right|} \qquad Loss_{RIoU} = 1 - RIoU$$

Oriented Bounding Box OBB

Rotated Box $B(x, y, w, h, \theta)$

ref.7

Often combined with Smooth L1 loss for precise parameter refinement (e.g., AO2-DETR[ref.8]).

Angle periodicity issue
Gradient instability

Gaussian-based distance losses

- Based on the distance between 2D-Gaussian distributions.
- Continuous & Differentiable with respect to the parameters

→ **Convergence stability improvement during optimization.**

Gaussian Bounding Box GBB

2-D Gaussian Distribution $N(\mu, \Sigma)$

ref.7

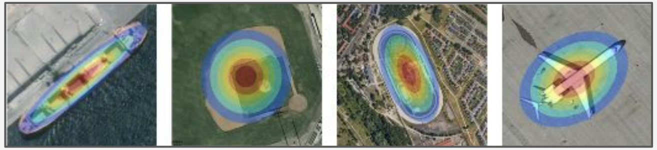

- **G**aussian **W**asserstein **D**istance (**GWD**) Loss[ref.9]
- **K**ullback-**L**eibler **D**ivergence (**KLD**) Loss[ref.7]

Necessitate nonlinear mapping function
Involve additional hyperparameter tuning
Lack intuitive geometric interpretation

***ref.6**: Zhou et al., 3DV 2019 (IoU) **ref.7:** Yang et al., NeurIPS 2021 (KLD) **ref.8:** Dai et al., TCSVT 2023 (AO2-DETR) **ref.9:** Yang et al., ICML 2021 (GWD)*

Localization Losses for Oriented Object Detection

Rotated Intersection over Union (RIoU) Loss[ref.6]

$$RIoU = \frac{\left|B_{pred}^{rot} \cap B_{gt}^{rot}\right|}{\left|B_{pred}^{rot} \cup B_{gt}^{rot}\right|} \qquad Loss_{RIoU} = 1 - RIoU$$

Oriented Bounding Box OBB

Rotated Box $B(x, y, w, h, \theta)$

ref.7

Often combined with Smooth L1 loss for precise parameter refinement (e.g., AO2-DETR[ref.8]).

Angle periodicity issue
Gradient instability

Gaussian-based distance losses

- Based on the distance between 2D-Gaussian distributions.
- Continuous & Differentiable with respect to the parameters

→ Convergence stability improvement during optimization.

Gaussian Bounding Box GBB

2-D Gaussian Distribution $N(\mu, \Sigma)$

ref.7

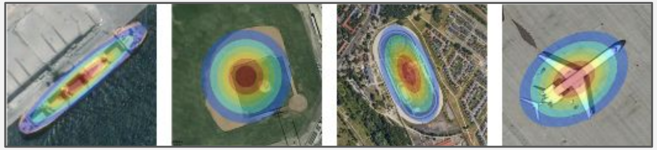

- Gaussian **W**asserstein **D**istance (**GWD**) Loss[ref.9]
- **K**ullback-**L**eibler **D**ivergence (**KLD**) Loss[ref.7]

Necessitate nonlinear mapping function
Involve additional hyperparameter tuning
Lack intuitive geometric interpretation

Probabilistic **IoU** (**ProbIoU**) **Loss**[ref.10]

- Based on **Hellinger distance**

ref.6: *Zhou et al., 3DV 2019 (IoU)* ***ref.7:*** *Yang et al., NeurIPS 2021 (KLD)* ***ref.8:*** *Dai et al., TCSVT 2023 (AO2-DETR)* ***ref.9:*** *Yang et al., ICML 2021 (GWD)* ***ref.10:*** *Murrugarra-Llerena et al., TIP 2024 (ProbIoU)*

Potentials and limitations of ProbIoU Loss

ProbIoU Potentials

- $\mathrm{ProbIoU} = 1\text{-}\mathrm{Hd}$ correlates with the IoU values between elliptical regions of the GBBs, providing a highly intuitive geometric interpretation (1 for identical distributions and 0 otherwise).
- The resulting loss (1 - $\mathrm{ProbIoU}$) is differentiable and satisfies all the properties of a distance metric.
- Does not require any nonlinear mapping or additional hyperparameters

Potentials and limitations of ProbIoU Loss

ProbIoU Potentials

- $\mathrm{ProbIoU} = 1\text{-}\mathrm{Hd}$ correlates with the IoU values between elliptical regions of the GBBs, providing a highly intuitive geometric interpretation (1 for identical distributions and 0 otherwise).
- The resulting loss $(1\ \text{-}\ \mathrm{ProbIoU})$ is differentiable and satisfies all the properties of a distance metric.
- Does not require any nonlinear mapping or additional hyperparameters

ProbIoU Limitations

- Saturation of the Hellinger distance
  - No information about the relative proximity of the boxes
  - No guidance for optimization
  - Low discriminative power of the ProbIoU loss

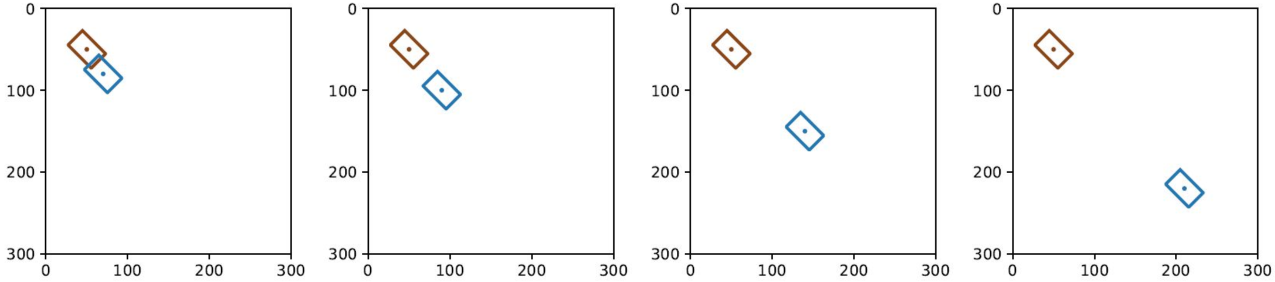

$\mathcal{L}_{ProbIoU} = 0.85$ | $\mathcal{L}_{ProbIoU} = 0.99$ | $\mathcal{L}_{ProbIoU} = 1.00$ | $\mathcal{L}_{ProbIoU} = 1.00$

Rethinking OBB-to-GBB Conversion (1/2)

OBB Representation
$B(x, y, w, h, \theta)$

GBB Representation
$N(\mu, \Sigma)$

ProbIoU

$$\mu_{ProbIoU} = (x, y)^T$$

$$\Sigma_{ProbIoU} = RD_{ProbIoU}R^T = \begin{pmatrix} \cos\theta & -\sin\theta \\ \sin\theta & \cos\theta \end{pmatrix} \begin{pmatrix} \overbrace{\frac{w^2}{12}}^{\lambda_w} & 0 \\ 0 & \underbrace{\frac{h^2}{12}}_{\lambda_h} \end{pmatrix} \begin{pmatrix} \cos\theta & \sin\theta \\ -\sin\theta & \cos\theta \end{pmatrix}$$

KLD & GWD

$$\mu_{KLD/GWD} = (x, y)^T$$

$$\Sigma_{KLD/GWD} = RD_{KLD/GWD}R^T = \begin{pmatrix} \cos\theta & -\sin\theta \\ \sin\theta & \cos\theta \end{pmatrix} \begin{pmatrix} \overbrace{\frac{w^2}{4}}^{\lambda_w} & 0 \\ 0 & \underbrace{\frac{h^2}{4}}_{\lambda_h} \end{pmatrix} \begin{pmatrix} \cos\theta & \sin\theta \\ -\sin\theta & \cos\theta \end{pmatrix}$$

Rethinking OBB-to-GBB Conversion (2/2)

Semi-axis lengths of the largest ellipse inscribed in the OBB:

$$\frac{w}{2} = \sqrt{k\lambda_w} \ ; \ \frac{h}{2} = \sqrt{k\lambda_h} \ ; \ k = \chi^2_{2,p}$$

ProbIoU

$k{=}3 \rightarrow p = 77.7\%$

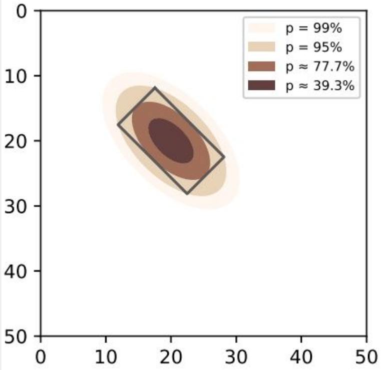

KLD & GWD

$k{=}1 \rightarrow p = 39.3\%$

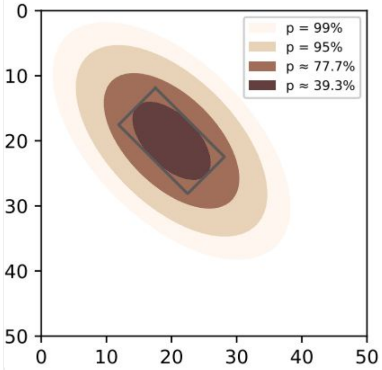

Generalized OBB-to-GBB conversion formula

$$\mu = (x, y)^T$$

$$\Sigma = RDR^T$$

$$= \begin{pmatrix} \cos\theta & -\sin\theta \\ \sin\theta & \cos\theta \end{pmatrix} \begin{pmatrix} \frac{w^2}{4k} & 0 \\ 0 & \frac{h^2}{4k} \end{pmatrix} \begin{pmatrix} \cos\theta & \sin\theta \\ -\sin\theta & \cos\theta \end{pmatrix}$$

Proposed Method: ProbIoU+ Loss (1/3)

We propose to minimize the normalized distance between the means of the GBBs.

Let $GBB_1 \sim N(\mu_1, \Sigma_1)$ and $GBB_2 \sim N(\mu_2, \Sigma_2)$

$$\mathcal{L}_{ProbIoU^+} = \mathcal{L}_{ProbIoU} + \frac{\|\mu_1 - \mu_2\|_2^2}{s^2}$$

where:

- $\|\mu_1 - \mu_2\|_2^2$: *the squared Euclidean distance between the mean vectors*
- s : *the normalizing factor*

How to compute s using a Gaussian-based approach?

Proposed Method: ProbIoU+ Loss (2/3)

STEP 1

Compute the mean and the covariance matrix of the Gaussian Mixture of GBB1 and GBB2

Mean

$$\mu_m = \pi_1\mu_1 + \pi_2\mu_2$$
$$= ½(\mu_1 + \mu_2)$$

Covariance Matrix

$$\Sigma_m = \pi_1(\Sigma_1 + (\mu_1 - \mu_m)(\mu_1 - \mu_m)^T) + \pi_2(\Sigma_2 + (\mu_2 - \mu_m)(\mu_2 - \mu_m)^T)$$
$$= \frac{1}{2}\Sigma_1 + \frac{1}{2}\Sigma_2 + \frac{1}{4}(\mu_1 - \mu_2)(\mu_1 - \mu_2)^T$$
$$=: \begin{pmatrix} \alpha & \gamma \\ \gamma & \beta \end{pmatrix}$$

Proposed Method: ProbIoU+ Loss (2/3)

STEP 1

Compute the mean and the covariance matrix of the Gaussian Mixture of GBB1 and GBB2

Mean

$$\mu_m = \pi_1\mu_1 + \pi_2\mu_2$$
$$= ½(\mu_1 + \mu_2)$$

Covariance Matrix

$$\Sigma_m = \pi_1(\Sigma_1 + (\mu_1 - \mu_m)(\mu_1 - \mu_m)^T) + \pi_2(\Sigma_2 + (\mu_2 - \mu_m)(\mu_2 - \mu_m)^T)$$
$$= \frac{1}{2}\Sigma_1 + \frac{1}{2}\Sigma_2 + \frac{1}{4}(\mu_1 - \mu_2)(\mu_1 - \mu_2)^T$$
$$=: \begin{pmatrix} \alpha & \gamma \\ \gamma & \beta \end{pmatrix}$$

STEP 2

Approximate the Gaussian mixture by the Gaussian distribution $N(\mu_m, \Sigma_m)$

aka ***"Moment-Matched Gaussian Approximation"***

Proposed Method: ProbIoU+ Loss (2/3)

STEP 1

Compute the mean and the covariance matrix of the Gaussian Mixture of GBB1 and GBB2

Mean

$$\mu_m = \pi_1\mu_1 + \pi_2\mu_2$$
$$= \tfrac{1}{2}(\mu_1 + \mu_2)$$

Covariance Matrix

$$\Sigma_m = \pi_1(\Sigma_1 + (\mu_1 - \mu_m)(\mu_1 - \mu_m)^T) + \pi_2(\Sigma_2 + (\mu_2 - \mu_m)(\mu_2 - \mu_m)^T)$$
$$= \frac{1}{2}\Sigma_1 + \frac{1}{2}\Sigma_2 + \frac{1}{4}(\mu_1 - \mu_2)(\mu_1 - \mu_2)^T$$
$$=: \begin{pmatrix} \alpha & \gamma \\ \gamma & \beta \end{pmatrix}$$

STEP 2

Approximate the Gaussian mixture by the Gaussian distribution $N(\mu_m, \Sigma_m)$

*aka **"Moment-Matched Gaussian Approximation"***

STEP 3

Set **s** to the major-axis length of the Gaussian's elliptical contour at confidence level p

$s = 2\sqrt{k\psi_{\max}}$ where $k = \chi^2_{2,p}$ and $\psi_{\max}$ is the **maximum eigenvalue** of Σ_m

Proposed Method: ProbIoU+ Loss (3/3)

RESULTING EXPRESSION

$$\mathcal{L}_{ProbIoU^{+}} = \mathcal{L}_{ProbIoU} + \frac{\|\mu_1 - \mu_2\|_2^2}{s^2}$$

$$\Downarrow$$

$$\mathcal{L}_{ProbIoU^{+}} = \mathcal{L}_{ProbIoU} + \frac{\|\mu_1 - \mu_2\|_2^2}{4k\psi_{\max}} \quad \text{where} \quad \psi_{\max} = \frac{(\alpha+\beta) + \sqrt{(\alpha-\beta)^2 + 4\gamma^2}}{2}$$

$$\Downarrow$$

$$\mathcal{L}_{ProbIoU^{+}} = \mathcal{L}_{ProbIoU} + \frac{(x_1 - x_2)^2 + (y_1 - y_2)^2}{2k\left((\alpha+\beta) + \sqrt{(\alpha-\beta)^2 + 4\gamma^2}\right)}$$

Simulation Study

Table 1: Simulation parameters for target and anchor boxes.

Box type	Position	Area (scale)	Aspect ratio (w/h)	Angle (rad)
Target boxes	1 fixed position at the center of a circular region $\mathcal{D}((15, 15), 5)$	1 (unit area)	1, 2, 3, 4	$-\pi/2, -\pi/3, -\pi/6, 0, \pi/6, \pi/3$
Anchor boxes	5000 positions uniformly sampled within $\mathcal{D}$	0.5, 0.75, 1, 1.5, 2	1, 2, 3, 4	$-\pi/2, -\pi/3, -\pi/6, 0, \pi/6, \pi/3$

TOTAL

14.4M

Regression cases

Initial Setup

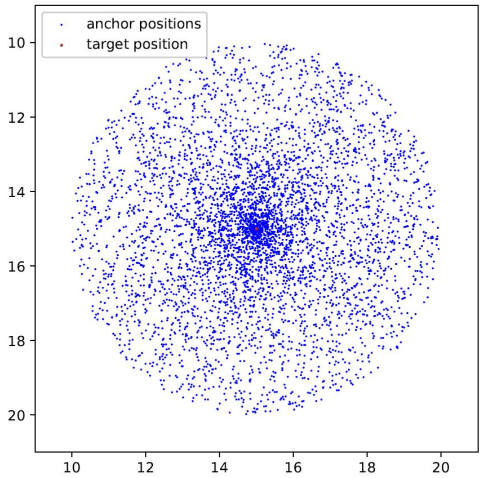

Performance results

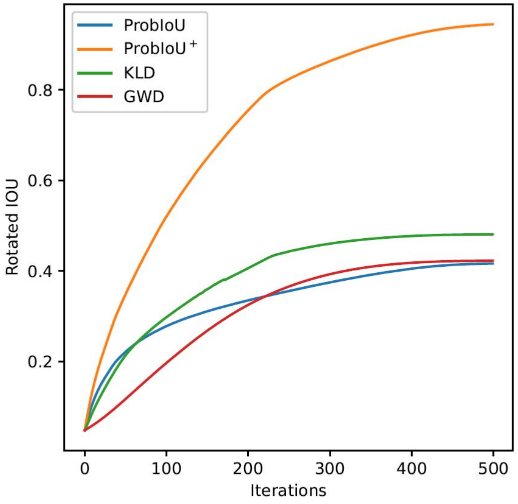

Experiments (1/3)

Dataset	Task	Classes	Images	Train	Test
HRSC2016[ref.11]	Ship Detection	Ships	1070	626	444

Table 2: Evaluation of different localization losses on HRSC2016 dataset

ref.12

Dataset	Detector	Loss	AP_{50}	AP_{60}	AP_{75}	AP_{85}	mAP
HRSC2016	RetinaNet	Smooth l_1	83.09	76.28	50.90	14.28	47.52
		GWD	87.82	82.84	55.56	14.61	51.40
		KLD	89.07	85.55	58.65	12.85	52.29
		ProbIoU	88.76	85.54	60.81	14.33	53.11
		ProbIoU^{+}	**91.11**	**87.70**	**65.28**	**16.53**	**55.90**
	R^3Det	Smooth l_1	90.46	86.65	63.60	19.39	55.22
		GWD	94.06	90.46	62.87	16.12	55.86
		KLD	93.68	90.97	67.92	17.54	57.88
		ProbIoU	**95.24**	92.87	71.81	26.78	61.40
		ProbIoU^{+}	94.99	**92.91**	**73.32**	**29.57**	**62.11**

AP_{75} +4.47%

AP_{75} +1.51%

ref.11: *Liu et al., ICPRAM 2017 (HRSC)* **ref.12:** *Zeng et al., TGRS 2024 (ARS-DETR)*

Experiments (2/3)

Dataset	Task	Classes	Images	Train	Test
ICDAR2015[ref.13]	Scene Text Detection	Text	1500	1000	500

Table 3: Evaluation of different localization losses on ICDAR2015 dataset

Dataset	Detector	Loss	AP_{50}	AP_{60}	AP_{75}	AP_{85}	mAP
ICDAR2015	RetinaNet	Smooth l_1	71.04	62.76	32.52	**5.85**	35.78
		GWD	72.06	61.92	30.97	4.99	35.42
		KLD	71.25	61.34	28.34	3.79	33.98
		ProbIoU	71.73	62.13	31.72	5.23	35.47
		ProbIoU$^+$	**72.75**	**64.61**	**32.91**	5.07	**36.42**
	R^3Det	Smooth l_1	73.26	63.53	28.81	3.28	35.16
		GWD	72.96	62.57	27.81	3.01	34.51
		KLD	73.66	62.05	25.59	2.67	33.94
		ProbIoU	**73.72**	63.59	29.38	3.49	35.53
		ProbIoU$^+$	73.70	**64.74**	**31.68**	**4.00**	**36.05**

AP_{75} +1.19%

AP_{75} +2.30%

ref.13: Karatzas et al., ICDAR 2015

Experiments (3/3)

Dataset	Task	Classes	Images	Train	Test
DOTA-v1.0[ref.1]	Aerial Object Detection	15 Categories	2806	1869	937

Table 4: Evaluation of different localization losses on DOTA-v1.0 dataset

Detector	Loss	PL	BD	BR	GTF	SV	LV	SH	TC	BC	ST	SBF	RA	HA	SP	HC	mAP
RetinaNet	Smooth l_1	88.74	73.68	42.49	67.80	77.26	58.93	76.40	90.88	76.95	82.60	56.78	63.01	53.73	63.53	44.01	67.79
	GWD	88.31	74.09	38.77	69.60	74.94	57.27	76.94	90.85	79.52	82.57	55.50	64.81	51.96	66.70	30.58	66.83
	KLD	88.95	75.97	39.75	70.17	74.99	59.91	77.23	90.85	76.40	82.66	52.26	61.69	52.06	65.64	38.56	67.14
	ProbIoU	88.64	75.86	39.28	69.73	78.01	62.92	77.65	90.86	80.77	82.56	53.83	60.49	56.98	67.29	35.48	68.02
	ProbIoU$^+$	88.84	77.64	42.90	66.72	78.25	63.32	77.83	90.86	79.83	83.41	55.65	65.90	59.60	68.01	47.75	69.77
R^{3}Det	Smooth l_1	89.32	72.73	47.32	70.02	77.86	71.88	79.10	90.82	79.55	84.67	55.04	62.72	62.76	63.29	39.59	69.78
	GWD	89.46	73.69	42.40	70.52	75.46	70.59	79.15	90.86	79.80	83.92	54.97	62.03	54.97	67.05	48.28	69.54
	KLD	89.10	74.06	43.38	69.63	75.30	71.58	79.24	90.86	76.50	83.17	54.34	63.28	55.27	57.48	44.92	68.54
	ProbIoU	89.12	74.07	48.51	71.85	77.63	74.15	84.95	90.87	78.65	83.70	59.82	63.74	65.08	67.57	45.39	71.67
	ProbIoU$^+$	89.39	77.28	48.16	72.31	78.29	73.98	84.97	90.86	79.03	83.74	55.52	64.44	65.18	65.51	51.65	72.02

mAP +1.75%

mAP +0.35%

Top 2 performances are highlighted in red and blue, respectively

ref.1: *Xia et al., CVPR 2018 (DOTA)*

Conclusion

Main Contributions

01 Generalized OBB→GBB conversion for oriented object detection

02 **ProbIoU+** loss: a novel localization loss that enhances ProbIoU using Gaussian distribution properties

03 Comprehensive evaluation: Data- and model-agnostic simulation & Experiments across public datasets

Thank You

ICPR 2026
International Conference on Pattern Recognition LYON

Detailed bibliography

[1] Xia et al., “DOTA: A Large-Scale Dataset for Object Detection in Aerial Images,” CVPR 2018.

[2] Girshick, R., “Fast R-CNN,” ICCV 2015.

[3] Yu et al., “UnitBox: An Advanced Object Detection Network,” ACM MM 2016.

[4] Rezatofighi et al., “Generalized Intersection over Union: A Metric and a Loss for Bounding Box Regression,” CVPR 2019.

[5] Zheng et al., “Distance-IoU Loss: Faster and Better Learning for Bounding Box Regression,” AAAI 2020.

[6] Zhou et al., “IoU Loss for 2D/3D Object Detection,” 3DV 2019.

[7] Yang et al., “Learning High-Precision Bounding Box for Rotated Object Detection via Kullback-Leibler Divergence,” NeurIPS 2021.

[8] Dai et al., “AO2-DETR: Arbitrary-Oriented Object Detection Transformer,” IEEE TCSVT 2023.

[9] Yang et al., “Rethinking Rotated Object Detection with Gaussian Wasserstein Distance Loss,” ICML 2021.

[10] Murrugarra-Llerena et al., “Probabilistic Intersection-Over-Union for Training and Evaluation of Oriented Object Detectors,” IEEE TIP 2024.

[11] Liu et al., “A High Resolution Optical Satellite Image Dataset for Ship Recognition,” ICPRAM, 2017.

[12] Zeng et al., “ARS-DETR: Aspect Ratio-Sensitive Detection Transformer for Aerial Oriented Object Detection,” TGRS, 2024.

[13] Karatzas et al., “ICDAR 2015 Competition on Robust Reading,” ICDAR 2015.