

Knowledge Combination to Learn Rotated Detection Without Rotated Annotation

Tianyu Zhu^{1,2,*} Bryce Ferenczi² Pulak Purkait¹ Tom Drummond²
Hamid RezaTofighi² Anton van den Hengel¹
¹Amazon IML. ²Monash University

tianyuzhu52@gmail.com {alanzty, purkaitp, hengelah}@amazon.com
{tianyu.zhu, bryce.ferenczi, tom.drummond, hamid.rezatofighi}@monash.edu

Abstract

Rotated bounding boxes drastically reduce output ambiguity of elongated objects, making it superior to axis-aligned bounding boxes. Despite the effectiveness, rotated detectors are not widely employed. Annotating rotated bounding boxes is such a laborious process that they are not provided in many detection datasets where axis-aligned annotations are used instead. In this paper, we propose a framework that allows the model to predict precise rotated boxes only requiring cheaper axis-aligned annotation of the target dataset.

To achieve this, we leverage the fact that neural networks are capable of learning richer representation of the target domain than what is utilized by the task. The under-utilized representation can be exploited to address a more detailed task. Our framework combines task knowledge of an out-of-domain source dataset with stronger annotation and domain knowledge of the target dataset with weaker annotation. A novel assignment process and projection loss are used to enable the co-training on the source and target datasets. As a result, the model is able to solve the more detailed task in the target domain, without additional computation overhead during inference. We extensively evaluate the method on various target datasets including fresh-produce dataset, HRSC2016 and SSDD. Results show that the proposed method consistently performs on par with the fully supervised approach.

1. Introduction

Rotated detectors introduced in recent works [17, 20, 32] have received attention due to their outstanding performance for top view images [15, 33, 34]. They reduce the output ambiguity of elongated objects for downstream tasks making them superior to axis-aligned detectors in dense scenes with severe occlusions [18]. However, the rotated annotation is more expensive compared to axis-aligned annotation. Furthermore, popular 2D annotation tools such as Sagemaker




Figure 1. KCR combines the task knowledge of a source dataset with stronger rotated annotation and the domain knowledge of the target dataset with weaker axis-aligned annotation, which enables the model to predict rotated detection on the target domain.

Groundtruth¹ and VGG app² do not support rotated bounding box annotations. As a result, many popular detection datasets only have axis-aligned annotations [3, 4, 11]. These problems reduce the potential scope of the implementation of rotated detectors. In this work, we introduce **Knowledge Combination** to learn **Rotated** object detection, a training scheme that only requires cheaper axis-aligned annotation for the target dataset in order to predict rotated boxes.

Neural networks encode data into a latent space, which is then decoded to optimize the given task. The latent embedding is an abstract representation of the data, containing much richer information than the output [29]. Early works in deep learning show that the model implicitly learns to detect image features such as edges and corners [10, 12], which can be used for more detailed tasks if decoded properly. We believe decoding to a more precise task on the target domain can be learnt via co-optimizing with a strongly labelled source dataset. We design a framework that combines task knowledge of rotated detection from a source dataset, and the domain knowledge of a disjoint class of objects in the target dataset with only axis-aligned annotation, as shown in Figure 1. This approach combines the advantage of both weakly-supervised learning and transfer learning.

We follow a design principal that the framework should maximize the target domain knowledge learnt by the model

*Corresponding author.

¹<https://aws.amazon.com/sagemaker/data-labeling/>

²<https://www.robots.ox.ac.uk/vgg/software/via/>

while minimizing the negative impact caused by weaker labels. This is achieved by co-training the source and target dataset with projection losses and a novel assignment process. The design choices are validated through ablation studies. We conduct extensive experiments to demonstrate that our framework is robust to a large domain gap between source and target dataset. Therefore, box orientation can practically be learnt for free with KCR, due to the availability of free public source datasets such as DOTA [32] with rotated annotations. We show the efficacy of this method on a fresh-produce dataset with high density of objects and severe occlusions. The performance (AP50) gap between the proposed method, learning from weak axis-aligned boxes, and the fully-supervised model learning from strong rotated annotation, reduces to only 3.2% for the challenging cucumber dataset. We apply the same framework to HRSC2016 [16] and SSDD [27] datasets to show that our method consistently performs on par with fully supervised models. The performance gap reduces to 1.0% for SSDD. We believe our approach can greatly increase the usage and impact of rotated object detectors. The source code will be publicly available for the community to save future annotation cost. In summary, our main contributions are as follows:

- 1) We introduce a framework that combines task knowledge of a strongly labelled source dataset and domain knowledge of a weakly labelled target dataset.
- 2) We apply this method in 2D rotated detection task, enabling the model to predict rotated bounding box with only axis-aligned annotation and verify the generality of the method with several datasets.
- 3) We demonstrate robustness of the framework to various domain gaps between source and target datasets. Hence, box orientation can be learnt with no additional annotation cost in practical applications.

2. Related Work

Rotated Detection Task. Rotated object detection requires the model to predict minimum area rectangles with five degrees of freedom, namely rotated bounding boxes, enclosing objects of interests [32]. In axis-aligned object detection, the output rectangles have four degrees of freedom, which are aligned with the image axes [13]. The rotated boxes occupy a much smaller area when estimating the location status of diagonally positioned elongated objects as shown in Figure 1. Rotated detection is strictly superior to axis-aligned detection as there is less background within the box, and the orientation can potentially convey object pose information. However, there are only a handful of datasets with rotated annotations [16, 27, 32] comparing to large number of readily available large-scale axis-aligned datasets [4, 11, 13, 19]. A potential contributor to this phenomenon is the ease of annotating axis-aligned boxes by simple click-and-drag with current labelling tools. Rotated boxes require much more

effort to tightly enclose the objects with an extra degree of freedom. Popular annotation tools such as AWS Sagemaker and VGG app do not support rotated boxes. In order to acquire tighter rotated annotations, users must pay for instance segmentation, which is significantly more expensive and unnecessary for the final task. Therefore, we propose this work to address the shortcomings of the rotated object detection task pipeline by making orientation free to learn.

Weakly-Supervised Learning. Weakly supervised learning sits between fully supervised learning and unsupervised learning, in the sense that only weak labels are available. The labels are weak either because they are incomplete, inexact or noisy [1, 6, 9, 23, 30, 31]. In computer vision, popular weakly-supervised learning tasks include object detection with only image level annotation [5, 35] and instance segmentation with only box annotation [2, 21]. Due to the difficulty of pixel-wise prediction, weakly-supervised instance segmentation still falls significantly behind a fully-supervised model [28] on novel objects. In this paper, we focus on learning rotated detection requiring five parameters with only axis-aligned annotation during training which provides four parameters. The one parameter difference makes our approach weakly-supervised learning. To the best of our knowledge, such problem has only been attempted on specific category of objects [8] but not approached generally. We believe that for novel elongated objects, solving this problem is more appropriate than solving weakly-supervised instance segmentation directly.

Transfer Learning. Computer vision models are frequently initialised with backbones [7, 14, 22] pretrained on ImageNet or COCO [13, 26], which is a basic form of transfer learning and a common practice. Transfer learning is highly effective when the target dataset has small sample size such as medical imaging [24]. Despite the fact that conventional transfer learning reduces the number of data samples required for a specific task, strong annotation of the target domain is still required for the model to learn the detailed task. In this paper, we utilize co-training strategy to transfer the ability to make more detailed predictions from the source to the target dataset with weaker annotations.

3. KCR

The goal of our work is to learn to predict rotated bounding boxes on a *target* dataset of which we only have axis-aligned bounding boxes in training examples. We develop a co-training scheme that utilizes an out-of-domain but strongly labelled *source* dataset to learn accurate rotations of elongated objects. An example of a potential source and target dataset pair are a satellite imagery dataset, DOTA [32], and a fresh-produce dataset, as shown in Figure 1. In this section, we first briefly describe the workflow of the detector following by our training scheme which enables weakly-supervised learning and knowledge transfer.

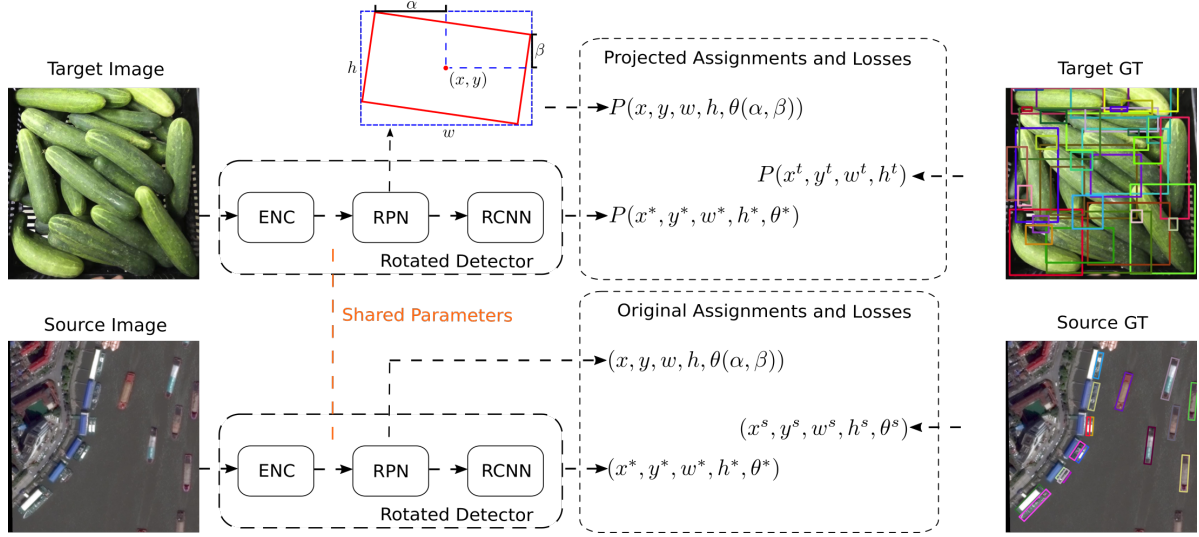


Figure 2. Overall framework of KCR, which learns the knowledge of rotated detection from source dataset combined with domain knowledge of axis-aligned target dataset to infer rotated bounding boxes of target objects. The rotated detector [33] takes in an image encoded by a CNN and generates *first-stage* proposals $(x, y, w, h, \alpha, \beta, p)$ and *second-stage* refinements $(x^*, y^*, w^*, h^*, \theta^*, p^*)$. We use a function P to either project or enlarge the box representation during the assignment processes and loss functions.

3.1. Rotated Detection Overview

In this work, we employ the same architectural choices as [33] and the flow is briefly demonstrated in Figure 2. The forward propagation comes in two stages: the oriented RPN followed by oriented R-CNN, where both methods contain a classification head and a regression head. The RPN takes an image and generates N oriented region proposals $\{R_i\}_{i=1}^N$ which each take the form $R_i = (x_i, y_i, w_i, h_i, \alpha, \beta, p_i)$, where (x_i, y_i) denotes the center, w_i and h_i are the width and height of the tightest axis-aligned external box. α and β are the offsets relative to the midpoints of the top and right sides of the external rectangle. p_i is an object score. The proposed rotated region will then be cropped by rotated roi [33] in feature space. The proposal is then fed to Oriented R-CNN which is a CNN followed by another classification head and a bounding box regression head rectifying the spatial location. We denote the output of second stage is $\{R_i^*\}_{i=1}^N$ and $R_i^* = (x_i^*, y_i^*, w_i^*, h_i^*, \theta_i^*, p_i^*, c_i^*)$ where (x_i^*, y_i^*) denotes the center, w_i^* and h_i^* are the width and height and θ_i^* the rotation angle of the final predicted rotated box. p_i^* is the second stage object score and c_i^* represents the classification score.

3.2. Learning Rotated Region Proposal

We develop a co-training scheme that utilizes both source and target datasets. We denote $\{B_j^s\}_{j=1}^m$ to be the m rotated bounding boxes of a source image where $B_j^s = (x_j^s, y_j^s, w_j^s, h_j^s, \theta_j^s)$ is a rotated box. For the target dataset, we have n axis-aligned boxes $\{B_j^t\}_{j=1}^n$ on a target image

where $B_j^t = (x_j^t, y_j^t, w_j^t, h_j^t)$ is a axis-aligned box.

In the first stage, a set of N oriented regions are proposed where each proposal $R_i = (x_i, y_i, w_i, h_i, \alpha, \beta, p_i)$ is assigned a ground truth label $B_{\sigma(i)}$ based on intersection over union (IoU) matching, where the assignment $\sigma(i)$ and matching score $\tau(i)$ are

$$\sigma(i) = \arg \max_{j \in \{1, \dots, m\}} iou((x_i, y_i, w_i, h_i), P(B_j)), \quad (1)$$

$$\tau(i) = \max_{j \in \{1, \dots, m\}} iou((x_i, y_i, w_i, h_i), P(B_j)). \quad (2)$$

Since in first stage, (x_i, y_i, w_i, h_i) represents the tightest external axis-aligned box instead of the rotated region itself, we need a transformation function P to project the ground truth to axis-aligned box for source rotated dataset. It is important to notice that $P(B_j)$ is strictly equal or larger than canonical axis-aligned box of the object. Therefore, for target axis-aligned dataset, we need to enlarge B_j^t . Here we formally define the transformation function

$$P(B) = \begin{cases} (x_{min}, y_{min}, x_{max}, y_{max}) \\ \quad \text{for } B = (x, y, w, h, \theta), \\ (x_{min} - \gamma w_i, y_{min} - \gamma h_i, x_{max} + \gamma w_i, \\ \quad y_{max} + \gamma h_i) \text{ for } B = (x, y, w, h), \end{cases} \quad (3)$$

where γ is simply an enlargement factor that we can tune.

The loss for this RPN is defined as follows

$$L_S = \frac{1}{N} \sum_{i=1}^N -\mathbf{1}_{\tau(i) \geq 0.5} \log(p_i) + l_1(R_i, B_{\sigma(i)}^s) \quad (4)$$

$$L_{\mathcal{T}} = \frac{1}{N} \sum_{i=1}^N -\mathbf{1}_{\tau(i) \geq 0.5} \log(p_i) + l_1((x_i, y_i, w_i, h_i), P(B_{\sigma(i)}^t)),$$

$$L = L_S + L_{\mathcal{T}}, \quad (5)$$

where the loss for both target and source examples are composed by a binary cross entropy (BCE) loss and an l_1 regression loss on spatial outputs. The label for the BCE is determined by whether the proposal has an overlap larger than 0.5 with any projected ground truth. For the source dataset, we can compute ground truth (α^s, β^s) from θ^s which is the regression for the rotation representation. However, for the target dataset, we only compute regression loss for (x_i, y_i, w_i, h_i) . The model must learn the rotation knowledge from the source dataset. To simplify the mathematical notation, we omit classification loss as it is less relevant to our contribution than object score.

3.3. Learning Rotated R-CNN

In the second stage, the Oriented R-CNN takes in a subset of proposals generated in first stage and make final prediction $R_i^* = (x_i^*, y_i^*, w_i^*, h_i^*, \theta_i^*, p_i^*, c_i^*)$. The bounding box regression is less important because of two reasons. Firstly, the cropped region proposal is already an approximate detection. The goal of regression of this stage is to fine-tune. Secondly, we choose to use class-agnostic rotated bounding box regression here as the model is able to learn general bounding box regression from the source dataset. The classification of the second stage is also straight forward to train because canonical axis-aligned detector fundamentally identical in that respect. The most important and challenging aspect is to produce an accurate object score p_i^* , which is a direct result of ground truth assignment process. We first formulate the assignment process for the source dataset as

$$\sigma_s^*(i) = \arg \max_{j \in \{1, \dots, m\}} iou((x_i, y_i, w_i, h_i, \theta_i), B_j^s), \quad (6)$$

$$\tau_s^*(i) = \max_{j \in \{1, \dots, m\}} iou((x_i, y_i, w_i, h_i, \theta_i), B_j^s). \quad (7)$$

Note that the assignment process is based on the proposal instead of the final refined prediction. The difference from the first stage is that here we use accurate rotated ground truth B_j^s against R_i instead of the external enclosing axis-aligned box. We compute θ of the first stage with (α, β) inline with [33]. This works appropriately with the source

dataset. However, it creates a challenge for the target dataset which lacks rotation information. If we follow the identical assignment process as the source dataset, then a close-to-correct proposal might be classified as negative example if the overlap between the rotated bounding box and axis-aligned bounding box is low. Consequently, false negative rate will be increased, reducing the recall of such training scheme. Therefore, we propose two additional strategies, projection assignment and heuristic selection.

Projection Assignment. To avoid the aforementioned increased false negative problem, we can project the predicted $(x_i, y_i, w_i, h_i, \theta_i)$ as an axis-aligned box and then compute the overlap between projected box and axis-aligned ground truth. Then we have,

$$\sigma_t^*(i) = \arg \max_{j \in \{1, \dots, m\}} iou(P(x_i, y_i, w_i, h_i, \theta_i), P(B_j^t)), \quad (8)$$

$$\tau_t^*(i) = \max_{j \in \{1, \dots, m\}} iou(P(x_i, y_i, w_i, h_i, \theta_i), P(B_j^t)). \quad (9)$$

where P is the transformation defined in previous section. By doing this, we reduce the false negative rate as rotated boxes can be correctly assigned to the axis-aligned counterpart and increase the recall of the model. However, consider a rotated box that has a rotation angle of $\pi - \theta$ or $-\theta$. It will have the same projected axis-aligned box as the one with θ and these two are completely different boxes. Therefore, this process might misclassify negative examples as positive, which results in reduction in precision. The losses are:

$$L_S^* = \frac{1}{N} \sum_{i=1}^N -\mathbf{1}_{\tau_s^*(i) \geq 0.5} \log(p_i^*) + l_1(R_i^*, B_{\sigma(i)}^s), \quad (10)$$

$$L_{\mathcal{T}}^* = \frac{1}{N} \sum_{i=1}^N -\mathbf{1}_{\tau_t^*(i) \geq 0.5} \log(p_i^*) \quad (11)$$

$$L^* = L_S^* + L_{\mathcal{T}}^*. \quad (12)$$

Heuristic Selection. Aside from the projection assignment, we also propose a parallel strategy. We use the same assignment rule as the source dataset 7. However, we evaluate how reliable each assignment is based on the aspect ratio and area size of axis-aligned ground truth and try to learn more from the reliable ground truth. For a particular axis-aligned ground truth box, we can compute its aspect ratio r_j defined by $r_j = \max(w_j, h_j) / \min(w_j, h_j)$. Larger aspect ratio implies the ground truth box is more reliable. In addition, we can find an area threshold $a_{threshold}$ a minimum un-occluded stereotypical object area on a target dataset. For $w_j * h_j < a_{threshold}$, it is likely an occluded box. Following these heuristics, we define a binary reliability switch and

loss for target dataset as

$$g_i = \begin{cases} 1 & \text{if } (r_{\sigma^*(i)} > 3 \text{ or } w_{\sigma^*(i)} * h_{\sigma^*(i)} < a_{threshold}) \\ 0 & \text{otherwise} \end{cases} \quad (13)$$

$$L_{\mathcal{T}}^* = \frac{1}{N} \sum_{i=1}^N -g_i \mathbf{1}_{\tau_i^*(i) \geq 0.5} \log(p_i^*). \quad (14)$$

and we choose to mask unreliable examples. Despite heuristic selection is less general than projection assignment, this can be beneficial for confined industry applications where the aspect ratio of the target object is known such as detecting a bottle on the conveyor belt.

4. Experiments and Discussion

In this section, we first outline our implementation details and experimental setup including the datasets. Then we show the effectiveness of our method through ablation studies followed by the comparison with different target and source datasets. Finally, we include some qualitative results.

4.1. Experimental setup

Source Datasets. Source datasets, in this work, are datasets strongly labelled with box orientations. The model has access to the source datasets to learn the task of rotated box detection. We choose a variety of source datasets including DOTA [32], COCO [13] and Fresh-produce dataset 1 where DOTA [32] is a popular satellite imagery dataset with rotated annotations and COCO [13] is a popular general object detection dataset. In COCO, rotated bounding box ground truth can be generated by finding the minimum enclosing rectangle on each instance segmentation mask ground truth. Fresh-produce dataset 1 is a challenging dataset with high object density and heavy occlusion. The number of images and instances are shown in table 1. There are three long-shape subsets including banana, cucumber and carrot.

Target Datasets. The training subset of the target datasets only contains axis-aligned bounding box ground truth which is used for the model to learn. Alternatively, the validation and test subsets contain rotated ground truth to evaluate the performance of the model. We select HRSC2016 [16], SSDD [27], cucumber and carrot datasets 1 as our target datasets to cover a variety of domain permutations including satellite, single channel, natural and various object density.

Implementation Details. We pretrain our detector, Oriented-RCNN [33] using DOTA [32]. The ship class and the images with the ships are removed from DOTA for this paper for pretraining and its role as a source dataset, due to class overlap with some target datasets. The main statistic we use in the paper to evaluate our method is AP50. It calculates the average precision with IOU threshold of 0.5. Average precision (AP) is the area under precision and recall curve.

Table 1. Fresh-produce datasets.

Class	Number of images	Number of instances
banana	158	7391
cucumber	48	2036
carrot	47	4647

We utilise the mmrotate framework [36] and [33] for training. We use a batch of 2 images from target dataset and a batch of 2 images from the source dataset for a combined mini-batch of 4. The losses and forward propagation for two batches are computed independently, then the losses are added and backpropagated through the network. We train the models up to 50 epochs of target dataset. During inference, we set the non-maximum-suppression threshold to 0.5 instead of 0.1 used generally for aerial datasets, as the higher object density fresh-produce dataset contains far more object overlapping cases. We conduct the experiments using one 2080ti GPU with 11GB of memory. The training time depends on the dataset size, ranging from 30 minutes to 2 hours to train. The test can be done within 5 minutes for each dataset with speed of 15 FPS, which is the same as the original detector.

4.2. Ablation Studies

To tackle the weakly-supervised learning of rotated bounding box given only axis-aligned ground truth, we build our approach progressively. In this section, we show how the approach evolves using cucumber as our axis-aligned training target dataset and DOTA [32] as our source dataset. We choose this particular pair because their domain gap is large in terms object appearance, density and occlusion severity. It is more convincing if the model is able to learn rotation with such a large domain gap between source and target.

We first establish a baseline by training a rotated detector using our axis-aligned cucumber training dataset and test it on rotated cucumber test set. This baseline is the first row of table 2 with AP50 of 0.491. Then we introduce the cotraining strategy with the source dataset without any modification of the training scheme. That means we simply treat both source and target datasets equally and completely follow the training scheme of [33]. This gives us AP50 of 0.542 which is an incremental improvement over the baseline.

RPN projection 5 improves the performance to 0.581 (Table 2). Instead of learning the wrong rotation supervision from the axis-aligned dataset, it chooses not to learn. After that, we try the single class strategy: only use one class label for all classes in source dataset. This effectively improves the performance to 0.633 because the model focuses on learning the object score and rotation angle instead of the classification problem of the source dataset. If the target dataset is multi-class, we can apply the single class strategy to the source dataset because the classification is not important.

Table 2. Ablation Studies of the proposed methods. We use DOTA as source and fresh cucumber as target dataset. This table decoupled the effectiveness of rotated cotraining, rpn projection loss, single class, heuristic selection, projection assignment and axis-aligned pretraining using other objects.

Rotated Cotraining	RPN Proj.	Single Class	H. Select.	P. Assign.	Pretraining	AP50↑
-	-	-	-	-	-	0.491
✓	-	-	-	-	-	0.542
✓	✓	-	-	-	-	0.581
✓	✓	✓	-	-	-	0.633
✓	✓	✓	✓	-	-	0.666
✓	✓	✓	-	✓	-	0.664
✓	✓	✓	-	✓	✓	0.683

Table 3. Analysis for enlargement factor γ . We use DOTA as source and HRSC as target dataset for this analysis.

γ	1.00	1.05	1.10	1.15	1.20	1.25
AP50↑	0.791	0.763	0.756	0.733	0.718	0.650

Table 4. We evaluate KCR with dense cucumber and carrot datasets with severe occlusion as target datasets. Tight rotated ground truth is used to evaluate AP50 performance of the model in the test time. We compare KCR method with original Oriented-RCNN [33] with weak axis-aligned annotation for training and fully-supervised models trained with strong rotated ground truth. Results show that KCR enables the model to perform on par with fully-supervised model and works reasonably well under large domain gap.

	Method	Source	Target Train	AP50↑
Cucumber	original [33]	-	axis-aligned	0.491
	KCR(ours)	DOTA [32]	axis-aligned	0.683
	KCR(ours)	banana	axis-aligned	0.788
	fully-supervised	-	rotated	0.820
Carrot	original [33]	-	axis-aligned	0.560
	KCR(ours)	DOTA [32]	axis-aligned	0.614
	KCR(ours)	banana	axis-aligned	0.723
	fully-supervised	-	rotated	0.765

The projection assignment and heuristic selection methods are implemented separately. The projection assignment strategy as described in section 3.3 improves the performance to 0.664 since it reduces the false negative rate. It cannot further improve performance due to the false positive assignment problem. Heuristic selection strategy improves the model performance to 0.666. Finally, we pretrain this model on other axis-aligned objects and improve this performance to 0.683 which is 0.192 higher than the baseline. We take the last row of table 2 as our final approach and apply it to permutations of target and source datasets in section 4.2.

We also provide an ablation of the enlargement factor $\gamma \geq 1$. As we can see from Table 3, $\gamma = 1.00$ is the best, which is what we have used throughout all the experiments.

Table 5. Evaluation of KCR method with HRSC2016 [16] and SSDD [27] as target datasets against tight rotated ground truth in the test. Influence of source dataset is investigated using COCO [13], COCO with rotation augmentation and DOTA [32]. We also compare KCR against original Oriented-RCNN [33], Grabcut [25] with weak axis-aligned annotation for training and fully-supervised model trained with strong rotated ground truth. We use projection assignment methods for KCR except for the ones with HS, which stands for heuristic selection.

	Method	Source	Target Train	AP50↑
HRSC2016 [16]	original [33]	-	axis-aligned	0.175
	Grabcut [25]	-	axis-aligned	0.240
	Grabcut Train	-	axis-aligned	0.629
	KCR(ours)	COCO [13]	axis-aligned	0.579
	KCR(ours)	COCO Aug	axis-aligned	0.783
	KCR(ours)	DOTA [32]	axis-aligned	0.791
	KCR-HS(ours)	DOTA [32]	axis-aligned	0.778
	fully-supervised	-	rotated	0.903
SSDD [27]	original [33]	-	axis-aligned	0.432
	Grabcut Train	-	axis-aligned	0.585
	KCR(Ours)	COCO [13]	axis-aligned	0.888
	KCR(Ours)	COCO Aug	axis-aligned	0.874
	KCR(Ours)	DOTA [32]	axis-aligned	0.881
	KCR-HS(Ours)	DOTA [32]	axis-aligned	0.874
	fully-supervised	-	rotated	0.898

4.3. Main results

In this section, we apply KCR framework to various combinations of source and target datasets to investigate the generality of the method under different domain gaps. We also compare KCR with a popular foreground segmentor [25].

Fresh-produce dataset. In Table 4 we can see results progressively improve when information that is closer to desired result is introduced. We initially evaluate our method where no knowledge combination is used, that is, for the cucumber and carrot, we use the axis-aligned bounding boxes and allow the network to learn incorrect α, β and θ values. In the next iteration, we use DOTA [32] aerial satellite images for knowledge combination. The domain gap be-




Figure 3. Visualization of KCR performance against original Oriented-Rcnn trained with weak axis-aligned annotation. The images are from test sets of HRSC2016 [16] and SSDD [27]. We use the COCO as our source dataset which has large domain gap from the target. The model trained with KCR methods learnt to predict accurate rotated bounding box, which is much more precise than the original model.

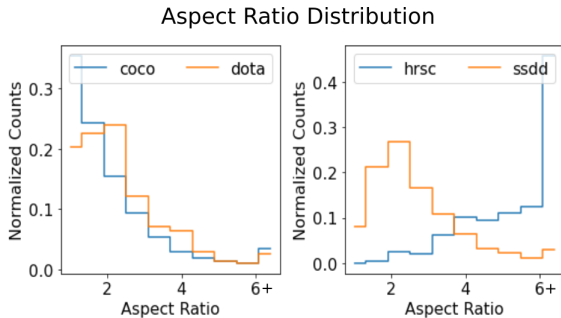


Figure 4. Distribution of Aspect ratios on 4 datasets.

tween these images and the cucumber and carrots is large because of the sparsity of objects in DOTA. Nonetheless, introducing rotated-bounding box information, even with dissimilar dataset, results in a large performance improvement for both cucumber and carrot as seen in the second row of Table 4. The next progression combines knowledge with bananas annotated with rotated bounding-box. This further improves the results and provides an insight into task similarity strengthening knowledge combination. The banana dataset offers the opportunity for the model to learn rotated detection with severe occlusion. For reference, the last row in Table 4 is the result when training the model directly on ground truth rotated bounding box, i.e. fully supervised training. We note that the performance gap between our approach and fully supervised training is small.

HRSC2016. HRSC2016 [16] is a rotated satellite imagery dataset that focuses on ships. The objects in HRSC2016 typically have a larger aspect ratio and occupies larger area of the image in comparison DOTA [32] targets. To prevent class overlap, images with ships have been removed from DOTA when using as a source dataset. We establish our baseline by using original training regime and only axis-aligned ground truth of HRSC2016 yielding AP50 of 0.175, which is 72.8% lower than a fully supervised result as shown in Table 5.

We investigate the performance of KCR with COCO [13], containing only natural images, as source dataset because the domain gap between COCO and HRSC2016 is visually large. We generate rotated bounding box ground truth from instance segmentation masks. The result significantly improves to 0.579 when we use KCR framework. Although the rotated bounding box is used in source dataset, objects in COCO are typically axis-aligned, hence produce a weak rotation training signal. We therefore rotate images by 0 – 180 degrees followed by horizontal and diagonal flip augmentation to increase the number of rotated examples in the source dataset. As a result, AP50 rises to 0.783 with strongly augmented COCO. Thus we show that KCR enables the transfer learning of extra parameter under large domain gap. At last, we use DOTA dataset as source and achieves AP50 of 0.791, which is only 11.2% lower than the fully supervised model. The performance gap is potentially due to lack of equally long objects with large aspect ratio in the source dataset.

SSDD. SSDD [27] is SAR dataset which also focuses on detection of ships. Images in SSDD are single channel and vary in frequency depending on the sensor used for acquisition. These images are commonly low resolution and contain high frequency noise, which results in a visually large domain gap from any source datasets. We follow the same knowledge combination strategy as before. Using KCR, a trained rotated box model performs almost equally well from three different source datasets (DOTA, coco and coco augmented). The final AP50 of 0.888 is 45.6% better than the baseline and only 1% lower than a fully supervised model.

Rotated Boxes with Grabcut. GrabCut [25] is a computer vision algorithm that predicts a foreground initialised with a region-of-interest. This foreground segmentation can then be used for predicting a rotated box with a variety of heuristic algorithms. This predicted rotation can be used at inference time as a post-processing step, transforming the axis aligned boundary box to a rotated boundary box. At test time, this




Figure 5. Visualization of KCR performance trained with axis-aligned only dataset. The images are chosen from an unlabeled set to ensure generality. The model is clearly capable of predicting tight rotated bounding box even with high object density and occlusion.

improves the baseline by a minimal 6.5%, however reduces inference throughput to 0.3 FPS. Alternatively, this algorithm can be performed offline on the axis-aligned ground truth to produce a noisy rotated boundary box ground truth for training. This method is significantly more effective, improving test time performance on rotated ground truth to 0.629 AP50 for HRSC and 0.585 for SSDD. However, KCR outperforms training on noisy ground truth by 16.2% on HRSC and 30.3% on SSDD.

Analysis on aspect ratios. We show histograms of aspect ratio in Fig 4. The biggest difference between an axis-aligned and rotated box happens when the instance has high aspect ratio **and** rotated. Most COCO objects are neither long **nor** rotated while DOTA has a fair distribution of long and rotated objects. Performance gap between original COCO and DOTA as source dataset is bigger for HRSC than SSDD because HRSC has a more tail heavy aspect ratio distribution.

4.4. Qualitative Results

We visualize performance of KCR on HRSC2016 [16] and SSDD [27] in Figure 3. The model has successfully learnt to predict accurate rotated bounding box with weak axis-aligned annotation. The source dataset we use to gain rotation knowledge is COCO [13], which has a large domain gap distant from the target dataset. The rotated prediction from our framework produces boxes much tighter than a model which was trained on and predicts axis-aligned boxes.

Fresh-produce dataset is more challenging due to higher density of objects with severe occlusion. We visualize the

performance of KCR with cucumber and carrot as target datasets and a banana dataset as the source dataset supplying knowledge of rotation. The depicted images are from a separate unlabeled set to ensure generality. As shown in Figure 5, the model is clearly capable of predicting tight rotated bounding boxes in a challenging scenario which is core contribution of this paper. The model is able to complete the task with high precision and recall in a scene with frequent and extremely occlusions, various lighting conditions and different object sizes. For the two images in the left column, the model detects almost every object.

5. Conclusion

Rotated detection improves the performance of downstream tasks by reducing the overall area of the enclosing box, improving the foreground to background ratio. This is particularly important for scenes with high density of targets and complex occlusions. However, most existing datasets only provide axis-aligned annotation, with the lack of the capability to annotate rotated boxes. In this paper, we address this problem by proposing KCR, which is a novel knowledge combination training scheme that only requires axis-aligned annotation for the target object class to train the model. At inference time, the model predicts accurate rotated bounding boxes on par with fully-supervised approach. This approach will enable the detector to predict an extra but crucial parameter. We believe this work can greatly extend the use case of rotated object detection by reducing annotation costs.

References

- [1] Jiwoon Ahn and Suha Kwak. Learning pixel-level semantic affinity with image-level supervision for weakly supervised semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4981–4990, 2018. 2
- [2] Jifeng Dai, Kaiming He, and Jian Sun. Boxesup: Exploiting bounding boxes to supervise convolutional networks for semantic segmentation. In *Proceedings of the IEEE international conference on computer vision*, pages 1635–1643, 2015. 2
- [3] Patrick Dendorfer, Hamid Rezaatofghi, Anton Milan, Javen Shi, Daniel Cremers, Ian Reid, Stefan Roth, Konrad Schindler, and Laura Leal-Taixé. Mot20: A benchmark for multi object tracking in crowded scenes. *arXiv preprint arXiv:2003.09003*, 2020. 1
- [4] Mark Everingham, Luc Van Gool, Christopher KI Williams, John Winn, and Andrew Zisserman. The pascal visual object classes (voc) challenge. *International journal of computer vision*, 88(2):303–338, 2010. 1, 2
- [5] Xiaoxu Feng, Xiwen Yao, Gong Cheng, and Junwei Han. Weakly supervised rotation-invariant aerial object detection network. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14146–14155, 2022. 2
- [6] Ramazan Gokberk Cinbis, Jakob Verbeek, and Cordelia Schmid. Multi-fold mil training for weakly supervised object localization. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2409–2416, 2014. 2
- [7] 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. 2
- [8] Javed Iqbal, Muhammad Akhtar Munir, Arif Mahmood, Afsheen Rafaqat Ali, and Mohsen Ali. Leveraging orientation for weakly supervised object detection with application to firearm localization. *Neurocomputing*, 440:310–320, 2021. 2
- [9] Alexander Kolesnikov and Christoph H Lampert. Seed, expand and constrain: Three principles for weakly-supervised image segmentation. In *European conference on computer vision*, pages 695–711. Springer, 2016. 2
- [10] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. *Communications of the ACM*, 60(6):84–90, 2017. 1
- [11] Alina Kuznetsova, Hassan Rom, Neil Alldrin, Jasper Uijlings, Ivan Krasin, Jordi Pont-Tuset, Shahab Kamali, Stefan Popov, Matteo Mallocci, Alexander Kolesnikov, et al. The open images dataset v4. *International Journal of Computer Vision*, 128(7):1956–1981, 2020. 1, 2
- [12] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998. 1
- [13] 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*, 2014. 2, 5, 6, 7, 8
- [14] 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 *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10012–10022, 2021. 2
- [15] Zikun Liu, Jingao Hu, Lubin Weng, and Yiping Yang. Rotated region based cnn for ship detection. In *2017 IEEE International Conference on Image Processing (ICIP)*, pages 900–904. IEEE, 2017. 1
- [16] Zikun Liu, Hongzhen Wang, Lubin Weng, and Yiping Yang. Ship rotated bounding box space for ship extraction from high-resolution optical satellite images with complex backgrounds. *IEEE Geoscience and Remote Sensing Letters*, 13(8):1074–1078, 2016. 2, 5, 6, 7, 8
- [17] Zikun Liu, Liu Yuan, Lubin Weng, and Yiping Yang. A high resolution optical satellite image dataset for ship recognition and some new baselines. In *International conference on pattern recognition applications and methods*, volume 2, pages 324–331. SciTePress, 2017. 1
- [18] Shijie Looi. Rotated mask r-cnn: From bounding boxes to rotated bounding boxes. https://github.com/mrlooi/rotated_maskrcnn, 2019. 1
- [19] Anton Milan, Laura Leal-Taixé, Ian Reid, Stefan Roth, and Konrad Schindler. Mot16: A benchmark for multi-object tracking. *arXiv preprint arXiv:1603.00831*, 2016. 2

- [20] Nibal Nayef, Fei Yin, Imen Bizid, Hyunsoo Choi, Yuan Feng, Dimosthenis Karatzas, Zhenbo Luo, Uma-pada Pal, Christophe Rigaud, Joseph Chazalon, et al. Icdar2017 robust reading challenge on multi-lingual scene text detection and script identification-rrc-mlt. In *2017 14th IAPR International Conference on Document Analysis and Recognition (ICDAR)*, volume 1, pages 1454–1459. IEEE, 2017. [1](#)
- [21] George Papandreou, Liang-Chieh Chen, Kevin P Murphy, and Alan L Yuille. Weakly-and semi-supervised learning of a deep convolutional network for semantic image segmentation. In *Proceedings of the IEEE international conference on computer vision*, pages 1742–1750, 2015. [2](#)
- [22] Joseph Redmon. Darknet: Open source neural networks in c, 2013. [2](#)
- [23] Zhongzheng Ren, Zhiding Yu, Xiaodong Yang, Ming-Yu Liu, Yong Jae Lee, Alexander G Schwing, and Jan Kautz. Instance-aware, context-focused, and memory-efficient weakly supervised object detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10598–10607, 2020. [2](#)
- [24] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer, 2015. [2](#)
- [25] Carsten Rother, Vladimir Kolmogorov, and Andrew Blake. " grabcut" interactive foreground extraction using iterated graph cuts. *ACM transactions on graphics (TOG)*, 23(3):309–314, 2004. [6](#), [7](#)
- [26] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. Imagenet large scale visual recognition challenge. *International journal of computer vision*, 115(3):211–252, 2015. [2](#)
- [27] J. Li T. Zhang, X. Zhang and X. Xu. Sar ship detection dataset (ssdd): Official release and comprehensive data analysis. *Remote Sensing*, 13(18):3690, Sep. 2021. [2](#), [5](#), [6](#), [7](#), [8](#)
- [28] Zhi Tian, Chunhua Shen, Xinlong Wang, and Hao Chen. Boxinst: High-performance instance segmentation with box annotations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5443–5452, 2021. [2](#)
- [29] Naftali Tishby and Noga Zaslavsky. Deep learning and the information bottleneck principle. In *2015 IEEE information theory workshop (itw)*, pages 1–5. IEEE, 2015. [1](#)
- [30] Yude Wang, Jie Zhang, Meina Kan, Shiguang Shan, and Xilin Chen. Self-supervised equivariant attention mechanism for weakly supervised semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12275–12284, 2020. [2](#)
- [31] Yunchao Wei, Huaxin Xiao, Honghui Shi, Zequn Jie, Jiashi Feng, and Thomas S Huang. Revisiting dilated convolution: A simple approach for weakly-and semi-supervised semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7268–7277, 2018. [2](#)
- [32] 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. [1](#), [2](#), [5](#), [6](#), [7](#)
- [33] Xingxing Xie, Gong Cheng, Jiabao Wang, Xiwen Yao, and Junwei Han. Oriented r-cnn for object detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3520–3529, 2021. [1](#), [3](#), [4](#), [5](#), [6](#)
- [34] 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](#)
- [35] Dingwen Zhang, Junwei Han, Gong Cheng, and Ming-Hsuan Yang. Weakly supervised object localization and detection: A survey. *IEEE transactions on pattern analysis and machine intelligence*, 44(9):5866–5885, 2021. [2](#)
- [36] Yue Zhou, Xue Yang, Gefan Zhang, Jiabao Wang, Yanyi Liu, Liping Hou, Xue Jiang, Xingzhao Liu, Junchi Yan, Chengqi Lyu, et al. Mmrotate: A rotated object detection benchmark using pytorch. *arXiv preprint arXiv:2204.13317*, 2022. [5](#)