



Article Multi-Level Alignment Network for Cross-Domain Ship Detection

Chujie Xu^{1,2}, Xiangtao Zheng^{1,*} and Xiaoqiang Lu¹

- ¹ Key Laboratory of Spectral Imaging Technology CAS, Xi'an Institute of Optics and Precision Mechanics, Chinese Academy of Sciences, Xi'an 710119, China; xuchujie19@mails.ucas.ac.cn (C.X.); luxiaoqiang@opt.ac.cn (X.L.)
- ² University of Chinese Academy of Sciences, Beijing 100049, China
- * Correspondence: zhengxiangtao@opt.cn

Abstract: Ship detection is an important research topic in the field of remote sensing. Compared with optical detection methods, Synthetic Aperture Radar (SAR) ship detection can penetrate clouds to detect hidden ships in all-day and all-weather. Currently, the state-of-the-art methods exploit convolutional neural networks to train ship detectors, which require a considerable labeled dataset. However, it is difficult to label the SAR images because of expensive labor and well-trained experts. To address the above limitations, this paper explores a cross-domain ship detection task, which adapts the detector from labeled optical images to unlabeled SAR images. There is a significant visual difference between SAR images and optical images. To achieve cross-domain detection, the multi-level alignment network, which includes image-level, convolution-level, and instance-level, is proposed to reduce the large domain shift. First, image-level alignment exploits generative adversarial networks to generate SAR images from the optical images. Then, the generated SAR images and the real SAR images are used to train the detector. To further minimize domain distribution shift, the detector integrates convolution-level alignment and instance-level alignment. Convolutionlevel alignment trains the domain classifier on each activation of the convolutional features, which minimizes the domain distance to learn domain-invariant features. Instance-level alignment reduces domain distribution shift on the features extracted from the region proposals. The entire multi-level alignment network is trained end-to-end and its effectiveness is proved on multiple cross-domain ship detection datasets.

Keywords: ship detection; domain adaptation; convolutional neural network; synthetic aperture radar

1. Introduction

Ships are marine transportation carriers, which are important for both military and civilian fields such as marine surveillance, shipping management, and maritime rescue. With the development of remote sensing technology, ship detection attempts to locate all ship instances from the wide-range sea surface. However, it is difficult to identify ships from the complex background due to orientation, scale variability, motion blurring, speckle noises, and clutter disturbance [1].

To monitor ships, two detection techniques have received extensive attention, including optical ship detection and Synthetic Aperture Radar (SAR) ship detection. Optical ship detection uses optical remote sensing images to distinguish ships, which is conducive to human visual interpretation [2]. Early optical ship detection methods exploit some geometric elements or artificially designed features to locate ship targets [3]. With the development of deep learning, Convolutional Neural Networks (CNNs) become the mainstream method [4–6].

Although optical remote sensing images provide detailed spatial information, the ship targets may be obscured by clouds, mists, or shadows [7]. In contrast, SAR can penetrate clouds and partial occlusions to detect hidden targets. Furthermore, SAR has the ability



Citation: Xu, C.; Zheng, X.; Lu, X. Multi-Level Alignment Network for Cross-Domain Ship Detection. *Remote Sens.* 2022, *14*, 2389. https:// doi.org/10.3390/rs14102389

Academic Editor: Giampaolo Ferraioli

Received: 14 March 2022 Accepted: 12 May 2022 Published: 16 May 2022

Publisher's Note: MDPI stays neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Copyright: © 2022 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (https:// creativecommons.org/licenses/by/ 4.0/). to work in all-day and all-weather. SAR ship detection methods arise active research by exploiting the backscatter characteristics [8], polarization characteristics [9], and geometric characteristics [10]. These methods extract hand-crafted features to describe the ship targets, which lack robustness in the wide-range sea surface such as open sea, offshore area, and harbor. Current state-of-the-art methods detect SAR ships by learning a CNN from a large dataset [11,12]. However, it is difficult to label the SAR images, which requires expensive labor and well-trained experts [13].

To address the above limitations, this paper explores a cross-domain ship detection task, which learns the detector from labeled optical images to unlabeled SAR images. The SAR images have a considerable domain shift compared with the optical images due to speckle noise, geometric distortion, lack of color information, and other factors. However, ship targets share the same shape characteristics in both optical and SAR domains, as shown in Figure 1. This observation provides a solution to transfer the CNN from the easy-to-obtain labeled optical domain to the SAR domain. This paper focuses on cross-domain ship detection: full supervision is given in the optical domain while no supervision is available in the SAR domain, as shown in Figure 2.



(b)

Figure 1. Comparison of optical images and SAR images. There is a large domain shift between the optical images and the SAR images, but ship targets present the same shape characteristics in the optical and SAR images. (**a**) Optical images. (**b**) SAR images.

Existing domain adaptation methods consider feature adaptation. For example, Chen et al. [14] proposed Domain Adaptive Faster R-CNN (DA), which considered the alignment of image features and instance features. Saito et al. [15] proposed Strong-Weak Distribution Alignment (SWDA), which considered the alignment of global features and local features. In this paper, an end-to-end Multi-level Alignment Network (MAN) is proposed. To reduce the large domain shift between the optical domain and the SAR domain, MAN considers three levels of alignment: image-level, convolution-level, and instance-level. First, image-level

alignment exploits generative adversarial networks to generate SAR images from the optical images. Then, the generated SAR images and the real SAR images are used to train the detector. To further narrow the domain distribution shift, the detector embeds convolution-level alignment and instance-level alignment. Convolution-level alignment trains the domain classifier on each activation of the convolutional features, which minimizes the domain distribution shift on the features extracted from the region proposals. The entire multi-level alignment network is trained end-to-end and its effectiveness is proved on multiple cross-domain ship detection datasets.





Figure 2. The cross-domain ship detection task. Full supervision is given in the optical domain while no supervision is available in the SAR domain. The detector is trained in the optical domain but is adapted for the SAR domain.

The contributions can be summarized as follows:

- 1. The cross-domain ship detection task is considered in this paper, which adapts the detector from labeled optical images to unlabeled SAR images. Compared with other cross-domain tasks, the cross-domain ship detection between the optical domain and the SAR domain is more challenging and more realistic;
- 2. The Multi-level Alignment Network (MAN) is proposed to reduce the large domain shift from the optical domain to the SAR domain, which achieves cross-domain alignment at the image-level, convolution-level, and instance-level;
- 3. The multi-level alignment mechanism is embedded into Faster R-CNN, and the entire detector is trained end-to-end without increasing inference time.

The remainder of this paper is organized as follows. Section 2 reviews ship detection methods. Section 3 describes the proposed MAN in detail. Section 4 provides full experimental results. Finally, Section 6 concludes this paper.

2. Related Works

2.1. Optical Ship Detection

For optical ship detection, early methods exploit some geometric elements or artificially designed features to locate ship targets [2,3,16]. Zhu et al. [17] proposed a hierarchical

complete and operational optical ship detection approach based on shape and texture features, which is considered a sequential coarse-to-fine elimination process of false alarms. Li et al. [16] proposed ship head classification and body boundary determination for inshore ship detection. Yang et al. [3] proposed the saliency segmentation and the local binary pattern descriptor combined with ship structure for ship detection. Since artificially designed features can only utilize low-level information with poor generalization ability, these methods are often affected by complex backgrounds, resulting in false or missed detections.

In recent years, with the development of deep learning technology, CNN has demonstrated its powerful feature representation ability. The optical ship detection methods based on CNN have been extensively studied and have achieved great success. At present, CNN-based ship detection methods are mainly divided into two categories. One is the two-stage detectors represented by Faster R-CNN [18] and Mask R-CNN [19]. The other is single-stage detectors represented by You Only Look Once (YOLO) [20], CenterNet [21], etc. For example, Nie et al. [22] proposed to achieve inshore ship detection based on Mask R-CNN. Soft-non-maximum suppression is introduced to improve the robustness of nearby inshore ships. Shamsolmoali et al. [23] proposed a multipatch feature pyramid network, which integrated automatic patch selection, feature aggregation, and semantic domain projection. Shamsolmoali et al. [24] proposed a rotation equivariant feature image pyramid network to deal with the complicated object deformation. More detectors directly use rotated bounding boxes to locate objects [5,6,25]. Due to the limitation of the imaging mechanism, the performance of optical ship detection methods will drop sharply in the case of severe weather such as clouds and fog.

2.2. SAR Ship Detection

Compared with optical ship detection, SAR ship detection has attracted more attention because of its penetrating ability and all-day, all-weather working ability. The widely used traditional SAR ship detection methods include the Constant False Alarm Rate (CFAR) algorithms [26–28]. They adjust the detection threshold adaptively according to the established statistical distribution model of background clutter. Wang et al. [29] improved the CFAR algorithm by fusing intensity and spatial information. However, due to changes in the environment, it is difficult to establish an accurate statistical model.

In recent years, CNN-based SAR ship detection methods have also been extensively studied. Lin et al. [30] proposed Squeeze and Excitation Rank Faster R-CNN (SER Faster R-CNN) to further improve the detection performance by using the squeeze and excitation mechanism. Cui et al. [31] proposed a dense attention pyramid network, which extracts abundant features containing resolution and semantic information. Li et al. [12] developed a multidimensional domain deep learning network to exploit the spatial and frequency-domain complementary features. Cui et al. [32] proposed the spatial shuffle-group enhance attention module in CenterNet to extract stronger semantic features and reduce missed detections. Yang et al. [1] proposed a coordinate attention module to obtain stronger semantic features and a receptive field increased module to capture multi-scale contextual information. Ma et al. [33] proposed an anchor-free framework with skip connections and aggregation nodes, which is designed to fuse multi-resolution features and detect multi-scale ship targets. And some works [34,35] considered the orientation of the ships. However, these methods still require full supervision in the SAR domain and cannot avoid the high cost of manual annotation of SAR images.

2.3. Cross-Domain Object Detection

Domain adaptation, which addresses the problem of label scarcity in new domains, has been actively studied in classification tasks [36]. For general object detection, most methods [14,15,37] learn domain-invariant features by adversarial training with the help of domain discriminators. Chen et al. [14] improved the cross-domain robustness of object detection through the image features adaptation and the instance features adaptation. Saito et al. [15] proposed the weak global feature alignment and strong local feature align-

ment. Zhu et al. [38] proposed to selectively align features of source and target domain data by mining the discriminative regions. Hsu et al. [39] proposed to achieve center-aware alignment by paying more attention to foreground pixel features. Xu et al. [40] proposed image-level classification regularization and categorical consistency regularization to match crucial image regions and important instances across domains. VS et al. [41] proposed memory-guided attention for category-aware domain adaptation. Wang et al. [42] proposed to align the sequence feature distributions extracted by Transformer detectors. These methods often only align on certain levels of features, and it is difficult to cope with the huge domain shift between the optical domain and the SAR domain. There are also works [43–45] that utilize the detector trained in the source domain to generate pseudo-labels in the target domain for further training. However, the inaccuracy of pseudo-labels limits the adaptability of the detector in the target domain.

For ship targets, Li et al. [46] considered cross-domain ship detection between optical images in different weather conditions. Zhao et al. [47] considered cross-domain detection between SAR images captured by different satellites. However, no related work considered cross-domain detection from the optical domain to the SAR domain.

3. Method

The labeled optical domain is defined as (X_o, Y_o) , where X_o denotes the optical images and Y_o denotes their corresponding bounding box coordinates. Cross-domain ship detection adapts the learned ship detector to an unlabeled target SAR domain (X_s, Y_s) , where X_s denotes SAR images, but the corresponding labels Y_s are not available. The overall framework of the proposed MAN is shown in Figure 3. Faster R-CNN is used as the base detector and a multi-level alignment mechanism is used to transfer the knowledge learned from the optical domain to the SAR domain.



Figure 3. The overall framework of the proposed MAN. MAN reduces the domain shift on different levels: image-level, convolution-level, and instance-level. First, image-level alignment exploits generative adversarial networks to generate SAR images from the optical images. Then, the generated SAR images and the real SAR images are used to train the detector. To further minimize domain distribution shift, the detector integrates convolution-level alignment and instance-level alignment. Convolution-level alignment trains the domain classifier on each activation of the convolutional feature, which minimizes the domain distance to learn domain-invariant features. Instance-level alignment reduces domain distribution shift on the features extracted from the region proposals.

3.1. Base Detector

Faster R-CNN [18] is the base detector in this paper. It is a two-stage detector mainly composed of three major components: backbone network, Region Proposal Network (RPN), and Detection Heads (DHs). The backbone network first extracts global convolutional

features for each input image. Then, the RPN predicts candidate region proposals based on the global convolutional features. Finally, the DHs consisting of a classifier and a regressor predict the detection results. The classifier filters out the real ships from the candidate regions, while the regressor predicts the bounding boxes of the ships. The training loss of the Faster R-CNN detector L_{det} can be expressed as the sum of RPN loss L_{RPN} and DHs loss L_{DHs} :

$$L_{\rm det} = L_{\rm RPN} + L_{\rm DHs}.$$
 (1)

Both RPN loss L_{RPN} and DHs loss L_{DHs} include classification loss and regression loss. For more details on the architecture and the training process, please refer to [18].

3.2. Image-Level Alignment

Due to the large domain shift between the optical domain and the SAR domain, it is difficult for a detector trained directly on the optical images to have a good application on the SAR images. Inspired by CycleGAN [48], image-level alignment aims to transfer the optical images X_o into the SAR images X_s .

The network structure of image-level alignment is shown in Figure 4. Image-level alignment consists of two generators and two discriminators. Specifically, a generator G_s and a discriminator D_s are added in front of the detector. G_s learns the mapping function $G_s : X_o \rightarrow X_s$ to generate images similar to SAR images, while D_s aims to distinguish between the generated SAR images and the real SAR images. The objective can be expressed as:

$$L_{\text{GAN}}(G_s, D_s) = \mathbb{E}_{x \sim X_s}[\log D_s(x)] + \mathbb{E}_{x \sim X_o}[1 - \log D_s(G_s(x))].$$
(2)

For adversarial training, G_s aims to minimize this objective while D_s tries to maximize it. Their data transfer flows are shown by the solid yellow arrows in Figure 4.

To avoid the model collapse problem caused by optimizing the adversarial objective in isolation, the generator $G_o : X_s \to X_o$ and the discriminator D_o are introduced. Similarly, the generator G_o and the discriminator D_o can be optimized by the objective $L_{\text{GAN}}(G_o, D_o)$, and the solid green arrows in Figure 4 mark their data transfer flows. It is worth noting that, to better show the data flows of MAN, the generator G_o and the discriminator D_o are not drawn in Figure 3.

During training, each batch consists of two images, including one optical image and one SAR image. For an optical image $x \in X_o$, the image $G_o(G_s(x))$ reconstructed by the two generators G_s and G_o should be similar to the original optical image x, i.e., $x \to G_s(x) \to G_o(G_s(x)) \approx x$. Such cycle consistency should also be present for SAR images. Cycle consistency loss is used to motivate this behavior:

$$L_{\rm cyc}^{o} = \mathbb{E}_{x \sim X_{o}}[\|G_{o}(G_{s}(x)) - x\|_{1}], \tag{3}$$

$$L_{\text{cyc}}^{s} = \mathbb{E}_{x \sim X_{s}}[\|G_{s}(G_{o}(x)) - x\|_{1}].$$
(4)

The constraint process of cycle consistency is shown by the dashed arrows in Figure 4. The complete image-level alignment loss can be expressed as:

$$\min_{G_s,G_o} \max_{D_s,D_o} L_{\text{img}} = L_{\text{GAN}}(G_s, D_s) + L_{\text{GAN}}(G_o, D_o) + \lambda \left(L_{\text{cyc}}^o + L_{\text{cyc}}^s \right),$$
(5)

where λ controls the relative importance of the two objectives. When optimizing G_s and G_o , L_{img} needs to be minimized, and when optimizing D_s and D_o , L_{img} needs to be maximized.



Figure 4. Image-level alignment structure. Image-level alignment aims to transfer the optical image into the SAR image.

3.3. Convolution-Level Alignment

Image-level alignment exploits generative adversarial networks to generate SAR images from the optical images. The generated SAR images $X_m = \{G_s(x) | x \in X_o\}$ and the original optical image labels Y_o form the intermediate domain (X_m, Y_o) . The intermediate domain and the SAR domain are used to train the Faster R-CNN detector.

Although the generated images X_m are similar to real SAR images X_s , there are still differences in distribution. To eliminate the domain distribution shift between X_m and X_s , convolution-level alignment trains the domain classifier D_{conv} on each activation of global convolutional features, which minimizes the domain distance to learn domain-invariant feature. Specifically, the backbone network and the domain classifier form a set of adversarial networks. The domain classifier is to distinguish whether the features come from X_m or X_s , but the backbone network tries to extract domain-invariant features so that the domain classifier cannot distinguish them.

For each image, *Z* denotes its domain label, where Z = 0 indicates the intermediate domain (X_m, Y_o) , and Z = 1 indicates the SAR domain (X_s, Y_s) . Let $\hat{Z}(i, j)$ denote the domain label prediction of the convolution-level domain classifier D_{conv} for the activation at the location (i, j) of the convolution features. The convolution-level alignment loss can be written as:

$$L_{\rm conv} = -\sum_{i,j} [Z \log \hat{Z}(i,j) + (1-Z) \log (1 - \hat{Z}(i,j))].$$
(6)

To align the domain distributions, the parameters of the domain classifier are optimized by minimizing the above alignment loss L_{conv} , while the weights of the backbone network are optimized by maximizing L_{conv} . Gradient Reverse Layer (GRL) [49] is added between the backbone network and the domain classifier to achieve adversarial training.

3.4. Instance-Level Alignment

The instance-level feature refers to the Region-of-Interest (RoI) based feature before feeding into the final classifier and regressor, which determine the final prediction of the

detector. Instance-level alignment also trains the domain classifier D_{ins} on local instance features to reduce domain shift.

Let \hat{Z}_k denote the predicted probability of the domain classifier D_{ins} for the *k*-th region proposal domain label, and the instance-level alignment loss is as follows:

$$L_{\rm ins} = -\sum_{k} [Z \log \hat{Z}_k + (1 - Z) \log (1 - \hat{Z}_k)].$$
(7)

GRL is also used to implement adversarial training.

Both convolution-level and instance-level alignment train domain classifiers to alleviate domain shift. The prediction results of domain classifiers D_{conv} and D_{ins} should be consistent. Enforcing consistency between convolution-level and instance-level alignment can help to learn the cross-domain robustness of target localization and classification. The consistency regularizer can be written as:

$$L_{\rm cst} = \sum_{k} \|\frac{1}{I} \sum_{i,j} \hat{Z}(i,j) - \hat{Z}_{k}\|_{2},$$
(8)

where *I* denotes the total number of activations in a feature map, and $\|\cdot\|_2$ is the L_2 distance. The convolution-level domain classifier produces a domain probability for each activation of the global feature. Hence, the average of all activations in the feature is taken as its global domain probability.

The complete MAN is trained using an end-to-end joint optimization method. For optical images, the overall training loss L_o can be expressed as:

$$L_o = L_{det} + \alpha L_{img} + \beta (L_{conv} + L_{ins} + L_{cst}),$$
(9)

where α and β are trade-off parameters to balance the detection loss with each alignment loss. For SAR images, the overall training loss L_s can be expressed as:

$$L_s = \alpha L_{\rm img} + \beta (L_{\rm conv} + L_{\rm ins} + L_{\rm cst}).$$
(10)

It is worth noting that the alignment modules are only used during training. In the training phase, MAN transfers the learned knowledge from the optical domain to the SAR domain through the multi-level alignment mechanism. Therefore, in the testing phase, all alignment modules can be simply offloaded, and the Faster R-CNN detector is used to predict ships in SAR images.

4. Results

To evaluate the effectiveness of MAN, comprehensive experiments are conducted in this section. Firstly, the cross-domain ship detection datasets and evaluation metric used in this paper are introduced. Next, the implementation details are described in detail. Following that, the performance of MAN is analyzed by ablation experiments. Finally, the effectiveness of MAN is confirmed by comparing it with some state-of-the-art methods.

4.1. Datasets

This is the first work to achieve cross-domain ship detection from the optical domain to the SAR domain. The existing datasets are not suitable for this cross-domain task. Therefore, to verify the effectiveness of MAN, this work constructs two cross-domain datasets based on some commonly used optical object detection datasets and SAR ship datasets. In each cross-domain dataset, the number of optical ships is more than that of SAR ships.

4.1.1. HRRSD \rightarrow SSDD

Optical images are from the HRRSD dataset [50]. It is a high-resolution remote sensing detection dataset, which contains 21,761 images and 55,740 object instances. The HRRSD

dataset contains 13 categories of objects, and only 2165 ship images are selected as the optical domain. The spatial resolution is from 0.15 m to 1.2 m.

SAR images are from the SSDD dataset [51]. It is a SAR ship dataset, which is mainly derived from RadarSat-2, TerraSAR-X, and Sentinel-1 sensors with four polarization modes including HH, VV, VH, and HV. The resolution of SAR images ranges from 1 m to 15 m. It consists of 1160 images and 2456 ships. The average number of ships per image is 2.12.

4.1.2. DIOR \rightarrow HRSID

Optical images are from the DIOR dataset [52], which is an optical remote sensing image dataset for object detection. It consists of 23,463 images and 192,472 object instances. DIOR dataset contains 20 categories of objects, and only 2702 ship images are selected as the optical domain.

SAR images are from the HRSID dataset [53]. It is a high-resolution SAR images dataset, which is constructed from Sentinel-1B, TerraSAR-X, and TanDEMX imageries. The resolution of SAR images is as follows: 0.5 m, 1 m, and 3 m. It contains a total of 5604 SAR images and 16,951 ship instances.

Table 1 lists the relevant statistics of the two cross-domain datasets.

Table 1. Statistics of two cross-domain datasets.

	Optical Domain	SAR Domain
$\mathrm{HRRSD}\to\mathrm{SSDD}$	2165 ship images and 3886 ship instances	1160 ship images and 2459 ship instances
$\text{DIOR} \rightarrow \text{HRSID}$	2702 ship images and nearly 64,000 ship instances	5604 ship images and 16,951 ship instances

4.2. Evaluation Metric

The mean Average Precision (mAP) [54] is used to evaluate the ship detection performance, which is the area under the Precision-Recall (P-R) curve. It is defined as

$$mAP = \int_0^1 P(R) \cdot R dR, \qquad (11)$$

where *P* is Precision, *R* is Recall,

$$P = \frac{\mathrm{TP}}{\mathrm{TP} + \mathrm{FP}},\tag{12}$$

$$R = \frac{\mathrm{TP}}{\mathrm{TP} + \mathrm{FN}}.$$
 (13)

TP denotes the number of correctly detected ship targets, FP denotes the number of false alarms, and FN denotes the number of missed detections.

4.3. Implementation Details

Faster R-CNN [18] with RoIAlign [19] is adopted in the experiments. ResNet-101 [55] is pre-trained on ImageNet and used to initialize the backbone network. The loss balance parameter α is set to 2, β is set to 0.01, and λ is set to 10. MAN is trained with a learning rate of 0.0001 for 100k iterations and then reduces the learning rate to 0.00001 for another 20k iterations. Each batch is composed of two images, one from the optical domain and one from the SAR domain, which are randomly sampled. The Stochastic Gradient Descent (SGD) optimizer with momentum of 0.9 and weight decay of 0.0005 is adopted in the experiments. MAN is implemented in PyTorch, and all experiments are performed on Ubuntu 20.04 system, Intel(R) Xeon(R) Silver 4210 CPU, and a Quadro RTX 6000 GPU.

4.4. Ablation Studies

To verify the effect of each alignment, some ablation studies have been completed on $DIOR \rightarrow HRSID$ dataset. To ensure fairness, other settings are consistent with implementation details when conducting experiments with different strategies.

4.4.1. The Impact of Different Level Alignments

To evaluate the significance of different level alignments in MAN, comprehensive ablation experiments are performed. The experimental results are shown in Table 2.

Table 2. The impact of different level alignments on detection performance. The best result is indicated by bold.

Image-Level Alignment	Convolution-Level Alignment	Instance-Level Alignment	mAP (%)
×	×	×	31.38
\checkmark	×	×	36.86
\checkmark	\checkmark	×	54.74
✓	\checkmark	\checkmark	57.37

Faster R-CNN is the baseline and does not consider cross-domain alignment. It is trained on optical images and tested directly on SAR images, which achieves 31.38% mAP. Based on Faster R-CNN, each different level alignment is increased.

Image-level alignment attempts to translate an optical image into an image close to the SAR domain. After adding image-level alignment on Faster R-CNN, the detector achieves 36.86% mAP. It facilitates the detector to be adapted to the SAR domain. Figure 5 shows the result of image-level alignment. It can be seen that the generated images are close to SAR images. However, image-level alignment is difficult to fit the speckle noise and geometric distortion of SAR images, and there is still a certain distribution difference between the generated SAR images and the real SAR images.



Figure 5. Image-level alignment examples. (a) The optical images. (b) The generated SAR images.

Convolution-level alignment further greatly improves the detection performance. Convolution-level alignment is beneficial for aligning the feature distributions of optical and SAR images. After adding convolution-level alignment, the detector achieves 54.74% mAP. Figure 6 shows the comparison of features with or without convolution-level alignment. The features are from the last convolutional layer of the backbone network. It can be seen that with the addition of convolution-level alignment, the shores that are more likely to have ships get more activations. This helps the detector improve recall.



Figure 6. Comparison of features with and without convolution-level alignment. (**a**) SAR images with ground-truths. (**b**) Feature maps w/o convolution-level alignment. (**c**) Feature maps w/ convolution-level alignment.

Instance-level alignment is added last to help get better classification and regression features, resulting in a detection performance of 57.37% mAP. In addition, some visual detection examples are compared, and the results are shown in Figure 7. Each alignment is incrementally increased and its detection results are shown separately. Not considering domain alignment can lead to a large number of missed detections. Gradually increasing the alignments at each level improves the recall rate and obtains better detection performance. This also demonstrates the effectiveness of the multi-level alignment mechanism.



(a)

Figure 7. Comparison of detection effects. The red bounding box denotes ground-truth and the green bounding box denotes predicted results. (a) Faster R-CNN. (b) Faster R-CNN w/ Image-level alignment. (c) Faster R-CNN w/ image-level alignment and convolution-level alignment. (d) MAN (ours).

4.4.2. The Impact of Different Training Strategies

In this paper, two training strategies are explored to handle the image-level alignment module and the domain adaptive detector incorporating convolution-level and instancelevel alignments. The step-by-step strategy first trains the image-level alignment module and then trains the domain adaptive detector. This means that the optical images are first translated into SAR images. Then, the generated SAR images input to the domain adaptive detector are invariant. The end-to-end strategy jointly optimizes the image-level alignment module and the domain adaptive detector in end-to-end training. Table 3 lists the results for the two training strategies, where the step-by-step strategy only obtains 39.13% mAP, while the end-to-end strategy obtains 57.37% mAP. The end-to-end training strategy achieves huge gains for two main reasons. First, when training the detector, image-level alignment is being trained at the same time. Therefore, the generated SAR images in different epochs are different. This difference encourages the detector to be more robust. Second, when jointly optimizing, the detection loss can guide not only detector training but also imagelevel alignment training. This means that image-level alignment can generate images in directions more suitable for ship detection.

Training Strategies	mAP (%)
Step-by-step	39.13
End-to-end	57.37

Table 3. The impact of different training strategies on detection performance. The best result is indicated by bold.

4.4.3. The Impact of Hyperparameters Changes

As shown in Equations (9) and (10), hyperparameters α and β control the balance of alignment losses at all levels. The choice of their value is worth exploring. α select 1, 2 and 3 for the experiment; while β select 1.00, 0.100, 0.010, and 0.001 for the experiment. The experimental results are shown in Figure 8. It can be seen that the adjustment of α has a large impact on the detection performance. In contrast, β is not sensitive. When $\alpha = 2$ and $\beta = 0.01$, MAN can be trained to obtain the best performance 57.37% mAP.



Figure 8. Analysis of trade-off parameters α and β . (a) The effect of parameter α on detection performance when β is fixed at 0.01. (b) The effect of parameter β on detection performance when α is fixed at 2.

4.4.4. The Impact of Different Backbone Networks

Different backbone networks are used to test the performance of MAN. ResNet-50 and ResNet-101 are selected for the experiments. The experimental results are shown in Table 4. MAN with ResNet-50 achieves 55.65% mAP, while MAN with ResNet-101 achieves 57.37% mAP.

Table 4. The impact of different backbone networks on detection performance. The best result is indicated by bold.

Backbone Networks	mAP (%)
ResNet-50	55.65
ResNet-101	57.37

4.5. Comparisons with State-of-the-Art Methods

The proposed MAN is compared with Faster R-CNN [18], Domain Adaptive Faster R-CNN (DA) [14], and Strong-Weak Distribution Alignment (SWDA) [15]. Faster R-CNN does not consider domain alignment. It is only trained in the optical domain and tested directly in the SAR domain. DA considers the alignment of convolution-level and instance-level, while SWDA considers the alignment of global and local features. These domain adaptation methods were originally applied to natural scenes, such as normal-to-foggy adaptation, synthetic-to-real adaptation, and cross-camera adaptation. This paper tests their performance on the cross-domain ship detection task. Comparative experiments are

performed on two datasets: HRRSD \rightarrow SSDD and DIOR \rightarrow HRSID. The experimental results are shown in Table 5.

Table 5. Comparison of SAR ship detection performance (mAP) on different cross-domain ship detection datasets. The best results are indicated by bold.

Methods	mAP (%)	
	$HRRSD \rightarrow SSDD$	$\mathbf{DIOR} \to \mathbf{HRSID}$
Faster R-CNN [18]	43.55	31.38
DA [14]	44.24	50.48
SWDA [15]	48.29	42.90
MAN (Ours)	61.92	57.37

Experimental results show that MAN achieves the best results on both cross-domain datasets. It achieves 57.37% mAP on DIOR \rightarrow HRSID and 61.92% mAP on HRRSD \rightarrow SSDD. Such results confirm the effectiveness of the proposed MAN. Meanwhile, Figure 9 shows the P-R curve results of the two datasets. It can also be seen that the area under the P-R curve of MAN is the largest. MAN can achieve higher precision and recall. It is worth noting that MAN does not have any efficiency drop compared to other methods. This is because the multi-level alignment mechanism only works during training and does not require any computation during testing.



Figure 9. P-R curves of different detection methods. (a) HRRSD \rightarrow SSDD dataset. (b) DIOR \rightarrow HRSID dataset.

5. Discussion

As shown in Figure 7, the detection results of our method are mostly correct, but some unexpected results still exist. Figure 10 shows some examples of failed detections. It can be found that in some complex backgrounds, the buildings and ships on shore cannot be effectively distinguished. As shown in Figure 10a, some supervisors with similar reflection characteristics to ships are falsely detected as ships. As shown in Figure 10b, multiple closely-arranged ships are easily missed. In response to such a situation, some more effective cross-domain alignment methods need to be explored. Meanwhile, some prior knowledge of SAR ships should also be considered to help us distinguish ships from backgrounds.



Figure 10. (**a**,**b**) Some cases of unsuccessful results. The red bounding box denotes ground-truth and the green bounding box denotes the predicted results.

6. Conclusions

This paper considers a cross-domain ship detection task: adapting the ship detector from labeled optical images to unlabeled SAR images. In view of the large domain shift between the optical domain and SAR domain, the multi-level alignment network is proposed, which achieves cross-domain alignment at the image-level, convolution-level, and instance-level. Image-level alignment attempts to translate optical images into SAR images, although it is not expected to be the same as SAR images. Convolution-level alignment trains the domain classifier on each activation of the convolutional feature, which minimizes the domain distance to learn domain-invariant features. Instance-level alignment reduces domain distribution shift on the features extracted from the region proposals. The multi-level alignment mechanism is embedded into Faster R-CNN, and the entire detector is trained end-to-end without increasing inference time. The effectiveness of the multi-level alignment mechanism is proved on multiple cross-domain ship detection datasets. MAN can detect most ships, but there are still some unexpected results. The main reason is that some complex backgrounds interfere with cross-domain alignment.

In future work, we will further explore some more effective cross-domain alignment methods. Meanwhile, some prior knowledge of SAR ships should also be considered to help us distinguish ships from backgrounds. In addition, in some practical application scenarios, it is not entirely impossible to obtain a small number of labeled SAR images. Therefore, the semi-supervised cross-domain ship detection is also a meaningful research direction.

Author Contributions: Conceptualization, C.X. and X.Z.; methodology, C.X.; software, C.X.; validation, C.X., X.Z. and X.L.; formal analysis, X.Z.; investigation, C.X.; resources, X.Z.; data curation, C.X.; writing—original draft preparation, C.X.; writing—review and editing, X.Z.; visualization, C.X.; supervision, X.L.; project administration, X.L.; funding acquisition, X.L. All authors have read and agreed to the published version of the manuscript.

Funding: This work was supported in part by the National Science Fund for Distinguished Young Scholars under Grant 61925112, in part by the Innovation Capability Support Program of Shaanxi

under Grant 2020KJXX-091 and Grant 2020TD-015, and in part by the Shaanxi Natural Science Basic Research Program 2022JQ-693.

Data Availability Statement: Not applicable.

Conflicts of Interest: The authors declare no conflict of interest.

References

- Yang, X.; Zhang, X.; Wang, N.; Gao, X. A Robust One-Stage Detector for Multiscale Ship Detection With Complex Background in Massive SAR Images. *IEEE Trans. Geosci. Remote Sens.* 2022, 60, 5217712. [CrossRef]
- Liu, G.; Zhang, Y.; Zheng, X.; Sun, X.; Fu, K.; Wang, H. A New Method on Inshore Ship Detection in High-Resolution Satellite Images Using Shape and Context Information. *IEEE Geosci. Remote Sens. Lett.* 2014, 11, 617–621. [CrossRef]
- 3. Yang, F.; Xu, Q.; Li, B. Ship Detection From Optical Satellite Images Based on Saliency Segmentation and Structure-LBP Feature. *IEEE Geosci. Remote Sens. Lett.* 2017, 14, 602–606. [CrossRef]
- Li, Q.; Mou, L.; Liu, Q.; Wang, Y.; Zhu, X.X. HSF-Net: Multiscale Deep Feature Embedding for Ship Detection in Optical Remote Sensing Imagery. *IEEE Trans. Geosci. Remote Sens.* 2018, 56, 7147–7161. [CrossRef]
- Cui, Z.; Leng, J.; Liu, Y.; Zhang, T.; Quan, P.; Zhao, W. SKNet: Detecting Rotated Ships as Keypoints in Optical Remote Sensing Images. *IEEE Trans. Geosci. Remote Sens.* 2021, 59, 8826–8840. [CrossRef]
- 6. Liu, L.; Bai, Y.; Li, Y. Locality-Aware Rotated Ship Detection in High-Resolution Remote Sensing Imagery Based on Multiscale Convolutional Network. *IEEE Geosci. Remote Sens. Lett.* **2022**, *19*, 3502805. [CrossRef]
- Wang, R.; You, Y.; Zhang, Y.; Zhou, W.; Liu, J. Ship detection in foggy remote sensing image via scene classification R-CNN. In Proceedings of the 2018 International Conference on Network Infrastructure and Digital Content (IC-NIDC), Guiyang, China, 22–24 August 2018; pp. 81–85. [CrossRef]
- 8. Qin, X.; Zhou, S.; Zou, H.; Gao, G. A CFAR Detection Algorithm for Generalized Gamma Distributed Background in High-Resolution SAR Images. *IEEE Geosci. Remote Sens. Lett.* **2013**, *10*, 806–810. [CrossRef]
- 9. Wang, C.; Wang, Z.; Zhang, H.; Zhang, B.; Wu, F. A PolSAR ship detector based on a multi-polarimetric-feature combination using visual attention. *Int. J. Remote Sens.* 2014, *35*, 7763–7774. [CrossRef]
- Zhu, J.; Qiu, X.; Pan, Z.; Zhang, Y.; Lei, B. Projection Shape Template-Based Ship Target Recognition in TerraSAR-X Images. *IEEE Geosci. Remote Sens. Lett.* 2017, 14, 222–226. [CrossRef]
- Li, Y.; Zhang, S.; Wang, W.Q. A Lightweight Faster R-CNN for Ship Detection in SAR Images. *IEEE Geosci. Remote Sens. Lett.* 2022, 19, 4006105. [CrossRef]
- 12. Li, D.; Liang, Q.; Liu, H.; Liu, Q.; Liu, H.; Liao, G. A Novel Multidimensional Domain Deep Learning Network for SAR Ship Detection. *IEEE Trans. Geosci. Remote Sens.* **2022**, *60*, 5203213. [CrossRef]
- Schwegmann, C.; Kleynhans, W.; Salmon, B.; Mdakane, L.; Meyer, R. Very deep learning for ship discrimination in Synthetic Aperture Radar imagery. In Proceedings of the 2016 IEEE International Geoscience and Remote Sensing Symposium (IGARSS), Beijing, China, 10–15 July 2016; pp. 104–107. [CrossRef]
- Chen, Y.; Li, W.; Sakaridis, C.; Dai, D.; Van Gool, L. Domain Adaptive Faster R-CNN for Object Detection in the Wild. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Salt Lake City, UT, USA, 18–23 June 2018; pp. 3339–3348.
- Saito, K.; Ushiku, Y.; Harada, T.; Saenko, K. Strong-Weak Distribution Alignment for Adaptive Object Detection. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), Long Beach, CA, USA, 15–20 June 2019; pp. 6956–6965.
- 16. Li, S.; Zhou, Z.; Wang, B.; Wu, F. A Novel Inshore Ship Detection via Ship Head Classification and Body Boundary Determination. *IEEE Geosci. Remote Sens. Lett.* **2016**, *13*, 1920–1924. [CrossRef]
- 17. Zhu, C.; Zhou, H.; Wang, R.; Guo, J. A Novel Hierarchical Method of Ship Detection from Spaceborne Optical Image Based on Shape and Texture Features. *IEEE Trans. Geosci. Remote Sens.* **2010**, *48*, 3446–3456. [CrossRef]
- Ren, S.; He, K.; Girshick, R.; Sun, J. Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks. In Advances in Neural Information Processing Systems; Cortes, C., Lawrence, N., Lee, D., Sugiyama, M., Garnett, R., Eds.; Curran Associates, Inc.: Red Hook, NY, USA, 2015; Volume 28.
- He, K.; Gkioxari, G.; Dollár, P.; Girshick, R. Mask R-CNN. In Proceedings of the IEEE International Conference on Computer Vision (ICCV), Venice, Italy, 22–29 October 2017; pp. 2961–2969.
- Redmon, J.; Divvala, S.; Girshick, R.; Farhadi, A. You Only Look Once: Unified, Real-Time Object Detection. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Las Vegas, NV, USA, 27–30 June 2016; pp. 779–788.
 Zhou, X.; Wang, D.; Krähenbühl, P. Objects as points. *arXiv* 2019, arXiv:1904.07850.
- 22. Nie, S.; Jiang, Z.; Zhang, H.; Cai, B.; Yao, Y. Inshore Ship Detection Based on Mask R-CNN. In Proceedings of the 2018 IEEE International Geoscience and Remote Sensing Symposium, Valencia, Spain, 22–27 July 2018; pp. 693–696. [CrossRef]
- Shamsolmoali, P.; Chanussot, J.; Zareapoor, M.; Zhou, H.; Yang, J. Multipatch Feature Pyramid Network for Weakly Supervised Object Detection in Optical Remote Sensing Images. *IEEE Trans. Geosci. Remote Sens.* 2022, 60, 5610113. [CrossRef]
- 24. Shamsolmoali, P.; Zareapoor, M.; Chanussot, J.; Zhou, H.; Yang, J. Rotation Equivariant Feature Image Pyramid Network for Object Detection in Optical Remote Sensing Imagery. *IEEE Trans. Geosci. Remote Sens.* **2022**, *60*, 5608614. [CrossRef]

- Wang, T.; Li, Y. Rotation-Invariant Task-Aware Spatial Disentanglement in Rotated Ship Detection Based on the Three-Stage Method. *IEEE Trans. Geosci. Remote Sens.* 2022, 60, 5609112. [CrossRef]
- 26. Robey, F.; Fuhrmann, D.; Kelly, E.; Nitzberg, R. A CFAR adaptive matched filter detector. *IEEE Trans. Aerosp. Electron. Syst.* **1992**, 28, 208–216. [CrossRef]
- 27. Ai, J.; Qi, X.; Yu, W.; Deng, Y.; Liu, F.; Shi, L. A New CFAR Ship Detection Algorithm Based on 2-D Joint Log-Normal Distribution in SAR Images. *IEEE Geosci. Remote Sens. Lett.* 2010, 7, 806–810. [CrossRef]
- Leng, X.; Ji, K.; Yang, K.; Zou, H. A Bilateral CFAR Algorithm for Ship Detection in SAR Images. *IEEE Geosci. Remote Sens. Lett.* 2015, 12, 1536–1540. [CrossRef]
- Wang, C.; Bi, F.; Zhang, W.; Chen, L. An Intensity-Space Domain CFAR Method for Ship Detection in HR SAR Images. *IEEE Geosci. Remote Sens. Lett.* 2017, 14, 529–533. [CrossRef]
- 30. Lin, Z.; Ji, K.; Leng, X.; Kuang, G. Squeeze and Excitation Rank Faster R-CNN for Ship Detection in SAR Images. *IEEE Geosci. Remote Sens. Lett.* **2019**, *16*, 751–755. [CrossRef]
- Cui, Z.; Li, Q.; Cao, Z.; Liu, N. Dense Attention Pyramid Networks for Multi-Scale Ship Detection in SAR Images. *IEEE Trans. Geosci. Remote Sens.* 2019, 57, 8983–8997. [CrossRef]
- 32. Cui, Z.; Wang, X.; Liu, N.; Cao, Z.; Yang, J. Ship Detection in Large-Scale SAR Images Via Spatial Shuffle-Group Enhance Attention. *IEEE Trans. Geosci. Remote Sens.* 2021, 59, 379–391. [CrossRef]
- Ma, X.; Hou, S.; Wang, Y.; Wang, J.; Wang, H. Multi-Scale and Dense Ship Detection in SAR Images Based on Key-Point Estimation and Attention Mechanism. *IEEE Trans. Geosci. Remote Sens.* 2022, 60, 5221111. [CrossRef]
- Zhang, J.; Xing, M.; Sun, G.C.; Li, N. Oriented Gaussian Function-Based Box Boundary-Aware Vectors for Oriented Ship Detection in Multiresolution SAR Imagery. *IEEE Trans. Geosci. Remote Sens.* 2022, 60, 5211015. [CrossRef]
- Sun, Y.; Sun, X.; Wang, Z.; Fu, K. Oriented Ship Detection Based on Strong Scattering Points Network in Large-Scale SAR Images. IEEE Trans. Geosci. Remote Sens. 2022, 60, 5218018. [CrossRef]
- Gong, T.; Zheng, X.; Lu, X. Cross-Domain Scene Classification by Integrating Multiple Incomplete Sources. *IEEE Trans. Geosci. Remote Sens.* 2021, 59, 10035–10046. [CrossRef]
- Chen, C.; Zheng, Z.; Huang, Y.; Ding, X.; Yu, Y. I3Net: Implicit Instance-Invariant Network for Adapting One-Stage Object Detectors. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), Nashville, TN, USA, 20–25 June 2021; pp. 12576–12585.
- Zhu, X.; Pang, J.; Yang, C.; Shi, J.; Lin, D. Adapting Object Detectors via Selective Cross-Domain Alignment. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), Long Beach, CA, USA, 15–20 June 2019; pp. 687–696.
- Hsu, C.C.; Tsai, Y.H.; Lin, Y.Y.; Yang, M.H. Every Pixel Matters: Center-Aware Feature Alignment for Domain Adaptive Object Detector. In *Proceedings of the European Conference on Computer Vision (ECCV)*; Lecture Notes in Computer Science; Springer: Cham, Switzerland, 2020; pp. 733–748.
- Xu, C.D.; Zhao, X.R.; Jin, X.; Wei, X.S. Exploring Categorical Regularization for Domain Adaptive Object Detection. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), Seattle, WA, USA, 13–19 June 2020; pp. 11724–11733.
- VS, V.; Gupta, V.; Oza, P.; Sindagi, V.A.; Patel, V.M. MeGA-CDA: Memory Guided Attention for Category-Aware Unsupervised Domain Adaptive Object Detection. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), Nashville, TN, USA, 20–25 June 2021; pp. 4516–4526.
- Wang, W.; Cao, Y.; Zhang, J.; He, F.; Zha, Z.J.; Wen, Y.; Tao, D. Exploring Sequence Feature Alignment for Domain Adaptive Detection Transformers. In Proceedings of the 29th ACM International Conference on Multimedia, Chengdu, China, 20–24 October 2021; pp. 1730–1738.
- Khodabandeh, M.; Vahdat, A.; Ranjbar, M.; Macready, W.G. A Robust Learning Approach to Domain Adaptive Object Detection. In Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV), Seoul, Korea, 27 October–2 November 2019; pp. 480–490.
- Kim, S.; Choi, J.; Kim, T.; Kim, C. Self-Training and Adversarial Background Regularization for Unsupervised Domain Adaptive One-Stage Object Detection. In Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV), Seoul, Korea, 27 October–2 November 2019.
- RoyChowdhury, A.; Chakrabarty, P.; Singh, A.; Jin, S.; Jiang, H.; Cao, L.; Learned-Miller, E. Automatic Adaptation of Object Detectors to New Domains Using Self-Training. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), Long Beach, CA, USA, 15–20 June 2019; pp. 780–790.
- 46. Li, L.; Zhou, Z.; Wang, B.; Miao, L.; An, Z.; Xiao, X. Domain Adaptive Ship Detection in Optical Remote Sensing Images. *Remote Sensing* **2021**, *13*, 3168. [CrossRef]
- 47. Zhao, S.; Zhang, Z.; Guo, W.; Luo, Y. An Automatic Ship Detection Method Adapting to Different Satellites SAR Images with Feature Alignment and Compensation Loss. *IEEE Trans. Geosci. Remote Sens.* **2022**, *60*, 5225217. [CrossRef]
- Zhu, J.Y.; Park, T.; Isola, P.; Efros, A.A. Unpaired Image-To-Image Translation Using Cycle-Consistent Adversarial Networks. In Proceedings of the IEEE International Conference on Computer Vision (ICCV), Venice, Italy, 22–29 October 2017; pp. 2223–2232.
- 49. Ganin, Y.; Lempitsky, V. Unsupervised Domain Adaptation by Backpropagation. In Proceedings of the 32nd International Conference on Machine Learning, Lille, France, 6–11 June 2015; Volume 37, pp. 1180–1189.

- 50. Zhang, Y.; Yuan, Y.; Feng, Y.; Lu, X. Hierarchical and Robust Convolutional Neural Network for Very High-Resolution Remote Sensing Object Detection. *IEEE Trans. Geosci. Remote Sens.* **2019**, *57*, 5535–5548. [CrossRef]
- 51. Li, J.; Qu, C.; Shao, J. Ship detection in SAR images based on an improved faster R-CNN. In Proceedings of the 2017 SAR in Big Data Era: Models, Methods and Applications (BIGSARDATA), Beijing, China, 13–14 November 2017; pp. 1–6. [CrossRef]
- Li, K.; Wan, G.; Cheng, G.; Meng, L.; Han, J. Object detection in optical remote sensing images: A survey and a new benchmark. ISPRS J. Photogramm. Remote Sens. 2020, 159, 296–307. [CrossRef]
- 53. Wei, S.; Zeng, X.; Qu, Q.; Wang, M.; Su, H.; Shi, J. HRSID: A High-Resolution SAR Images Dataset for Ship Detection and Instance Segmentation. *IEEE Access* 2020, *8*, 120234–120254. [CrossRef]
- 54. Everingham, M.; Van Gool, L.; Williams, C.K.; Winn, J.; Zisserman, A. The pascal visual object classes (voc) challenge. *Int. J. Comput. Vis.* **2010**, *88*, 303–338. [CrossRef]
- He, K.; Zhang, X.; Ren, S.; Sun, J. Deep Residual Learning for Image Recognition. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Las Vegas, NV, USA, 27–30 June 2016; pp. 770–778.