

# Semantic-Aware Transformation-Invariant RoI Align

Guo-Ye Yang<sup>1</sup>, George Kiyohiro Nakayama<sup>2</sup>, Zi-Kai Xiao<sup>1</sup>, Tai-Jiang Mu<sup>1\*</sup>, Xiaolei Huang<sup>3</sup>,  
Shi-Min Hu<sup>1 †</sup>

<sup>1</sup>BNRist Department of Computer Science and Technology, Tsinghua University

<sup>2</sup>Stanford University

<sup>3</sup>College of Information Sciences and Technology, Pennsylvania State University  
yanggy19@mails.tsinghua.edu.cn, w4756677@stanford.edu, xzk23@mails.tsinghua.edu.cn,  
taijiang@tsinghua.edu.cn, suh972@psu.edu, shimin@tsinghua.edu.cn

## Abstract

Great progress has been made in learning-based object detection methods in the last decade. Two-stage detectors often have higher detection accuracy than one-stage detectors, due to the use of region of interest (RoI) feature extractors which extract transformation-invariant RoI features for different RoI proposals, making refinement of bounding boxes and prediction of object categories more robust and accurate. However, previous RoI feature extractors can only extract invariant features under limited transformations. In this paper, we propose a novel RoI feature extractor, termed Semantic RoI Align (SRA), which is capable of extracting invariant RoI features under a variety of transformations for two-stage detectors. Specifically, we propose a semantic attention module to adaptively determine different sampling areas by leveraging the global and local semantic relationship within the RoI. We also propose a Dynamic Feature Sampler which dynamically samples features based on the RoI aspect ratio to enhance the efficiency of SRA, and a new position embedding, *i.e.*, Area Embedding, to provide more accurate position information for SRA through an improved sampling area representation. Experiments show that our model significantly outperforms baseline models with slight computational overhead. In addition, it shows excellent generalization ability and can be used to improve performance with various state-of-the-art backbones and detection methods. The code is available at <https://github.com/cxjyxxme/SemanticRoIAlign>.

## Introduction

As a fundamental computer vision task, object detection aims to locate and recognize objects of interest in input images. In the last decade, great progress has been made in learning-based object detection methods, making them widely useful in our daily uses, such as face recognition, text detection, pedestrian detection, among others.

Most existing detection methods can be grouped into two categories, *i.e.*, one-stage detectors (Liu et al. 2016; Redmon et al. 2016) and two-stage detectors (Ren et al. 2015;

\*Corresponding author is Tai-Jiang Mu.

† Arxiv version with appendices: <https://arxiv.org/abs/2312.09609>

Copyright © 2024, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

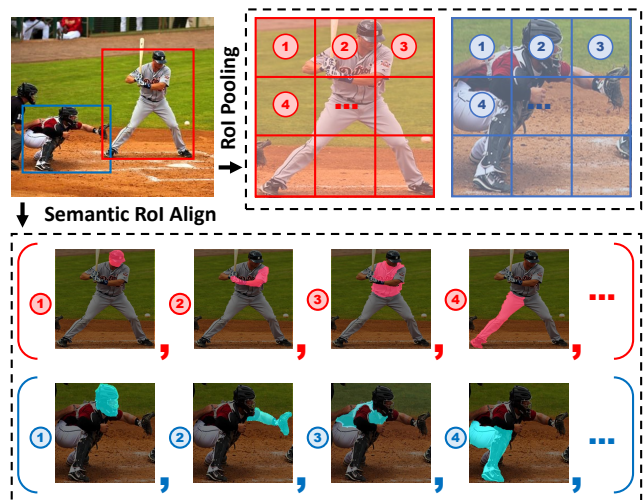


Figure 1: Previous RoI feature extractor versus the proposed Semantic RoI Align (SRA). Top: RoI Pooling samples each feature in some specific positions, making extracted RoI features sensitive to object poses. Bottom: SRA samples each feature from different semantic regions, making it capable of extracting invariant RoI features under various transformations including object pose transformation.

He et al. 2017). One-stage detectors directly predict objects with a single neural network in an end-to-end manner. In contrast, two-stage detectors first propose a list of object proposals and then predict the proposals' labels and refine bounding boxes with extracted RoI features of each proposal. RoI features provide better transformation-invariance for different proposal regions and using them can thus better refine bounding boxes and predict the category of each proposal in the second stage of a two-stage detector. In this paper, we mainly focus on improving the RoI feature extractor for two-stage detectors.

Similar objects may show great appearance differences in images due to different environmental conditions, object poses, *etc.*, making detection difficult to be generalizable under various transformations (Girshick 2015). Therefore,

many RoI feature extractors aim to extract transformation-invariant object features. RoI Pooling (Girshick 2015) is the pioneering work for RoI feature extraction, which pools features in a fixed number of sub-regions of the RoI and obtains scale-invariant features. RoI Align (He et al. 2017) further improves the positional accuracy of RoI Pooling via bilinear interpolation. RoI Transformer (Ding et al. 2019) extracts rotation-invariant features by rotating sampling positions with a regressive rotation angle. However, previous methods cannot extract invariant features for more complex transformations like perspective transformation and object pose transformation. Though there exist works like Deformable RoI Pooling (DRoIPooling) (Dai et al. 2017; Zhu et al. 2019) that can extract invariant features under some complex transformations by adaptively adding a regressive offset to each sampling position, experiments show that it only achieves invariance under scaling transformation and cannot easily be extended to handle others such as rotation. This is because the sampling position offsets are regressed with convolutional networks, which need to be trained with transformed data (Krizhevsky, Sutskever, and Hinton 2012) and different transformations would require different kernels since the same position of a convolutional kernel may correspond to different object regions when the object is transformed under different transformations.

In this paper, we regard different transformations like perspective and pose transformations as being comprised of spatial transformations of different semantic parts, while also considering that high-level features of semantic regions are more stable under varying transformations. From this perspective, we propose a Semantic RoI Align (SRA) to extract transformation-invariant RoI features by sampling features from different semantic regions. RoI Pooling samples features at specific locations in a RoI. A limitation of such sampling can be seen from the example shown in Figure 1 (top row); the 4-th sampling location extracts features of the background in the red RoI, whereas that location extracts features of a player’s leg in the blue RoI. Such sampling loses invariance under pose transformations. In contrast, in our proposed SRA, we design a semantic attention module to obtain different semantic regions by leveraging the global and local semantic relationship within the RoI. Then we sample features from the semantic regions, and concatenate the sampled features as the RoI feature which is semantic-aware and transformation-invariant.

Since the computational efficiency of SRA determined by the feature sampling resolution, we propose a Dynamic Feature Sampler to dynamically sample features according to the aspect ratio of different RoIs, which speeds up SRA while minimizing the impact on accuracy. Furthermore, previous positional embedding methods (Zhao, Jia, and Koltun 2020) only encode information of the sampling center, which cannot accurately represent regional information. We thus propose a new positional embedding, namely Area Embedding, which embeds positions in a sampling area into a fixed-length vector, providing more accurate position information. SRA can replace the RoI feature extractor in most two-stage detectors and brings higher detection accuracy with a slight overhead in the number of network

parameters and computation.

By using SRA as RoI feature extractor for Faster R-CNN (Ren et al. 2015), our method achieves 1.7% higher mAP in the COCO object detection task with only additional 0.2M parameters and 1.1% FLOPs compared to the baseline model. Meanwhile, it also exceeds other RoI feature extractors with less computational overhead. To verify the generalizability of SRA, we equip it to various state-of-the-art backbones and detection methods. Results show that SRA can consistently boost their detection accuracy.

In summary, our contributions are:

- a novel RoI feature extractor, *i.e.*, Semantic RoI Align, which is able to extract transformation-invariant RoI features and can be plugged into most two-stage detectors to improve detection accuracy with little extra cost,
- a Dynamic Feature Sampler which makes SRA implementation efficient, and an Area Embedding which provides more comprehensive and accurate information of sampled positions.
- Extensive experiments that demonstrate the superiority of the SRA and its great generalizability to various state-of-the-art backbones and detection methods.

## Related Work

### Object Detection and RoI Feature Extractors

In recent years, deep learning techniques are dominant in object detection. Most deep object detection methods can be categorized into two types: one-stage detectors (Redmon et al. 2016) and two-stage detectors (Girshick 2015). Faster R-CNN (Ren et al. 2015) is a two-stage network with a Regional Proposal Network (RPN) predicting multiple RoI proposals; RoI features are then extracted by an RoI feature extractor to predict object bounding boxes and categories in the second-stage network. Mask R-CNN (He et al. 2017) proposed a general framework for object instance segmentation tasks. Dynamic head (Dai et al. 2021) proposed to use scale, spatial, and task-aware attention mechanisms to improve detection accuracy.

RoI feature extractors are used to extract transformation-invariant features in two-stage detectors, so that the second-stage network can refine the bounding boxes and predict object categories more accurately. RoI Pooling (Girshick 2015) performs scale-invariant feature extraction by dividing the RoI into a fixed number of bins, pooling the features in each bin, and concatenating them into a vector of fixed size. RoI Align (He et al. 2017) uses bilinear interpolation to more accurately extract features. RoI Transformer (Ding et al. 2019) extracts rotation-invariant features by correcting the extracted features using a learned rigid transformation supervised by ground-truth oriented bounding boxes. However, these methods only model features invariant to rigid transformations, while ignoring non-rigid transformations. Deformable RoI Pooling (Dai et al. 2017; Zhu et al. 2019) extracts features by adding a regressive offset to each sampling position of RoI Pooling. Our experiments show that it can only extract invariant features under scale transformation which could be due to its learning the regressive offset by convolution, making it hard to generalize to other

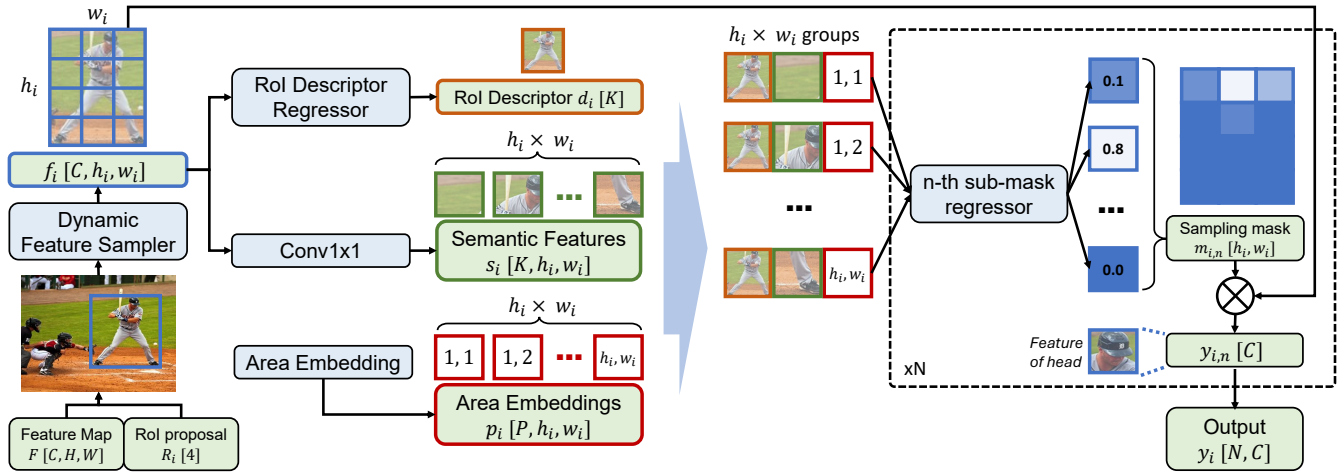


Figure 2: The network architecture of our Semantic RoI Align.  $\otimes$  means matrix multiplication, and  $\oplus$  means concatenation.

transformations. RoIAttn (Liang and Song 2022) proposes to enhance the RoI features by passing them through multiple self-attention layers. However, simply doing so is limited w.r.t. the ability to obtain invariant RoI features for two reasons: 1) performing self-attention on RoI Align extracted features has more limited flexibility than sampling on the original feature map, 2) the regression ability of typical self-attention is insufficient for identifying specific semantic regions under different transformations. Our SRA obtains different semantic regions with a novel semantic attention structure by leveraging the global and local semantic relationship within the RoI. We then sample the RoI features from the semantic regions, which makes it easy to achieve invariance under more diverse transformations, and thus obtain higher detection accuracy than existing methods.

### Attention Mechanism

In computer vision, attention can be regarded as an adaptive process, which mimics the human visual system’s ability to focus on important regions. RAM (Mnih et al. 2014) is the pioneering work to introduce the attention concept in computer vision. After that, there have been some works (Hu, Shen, and Sun 2018; Zhao, Jia, and Koltun 2020) exploring the use of attention mechanisms for different computer vision tasks. Recently, transformer networks, which have achieved great success in natural language processing (Vaswani et al. 2017), are explored in computer vision and have shown great potential. ViT (Dosovitskiy et al. 2020) is the first work to bring transformer into computer vision by regarding a  $16 \times 16$  pixel region as a word and an image as a sentence. Due to the strong modeling capability of visual transformer networks, they have been applied to various vision tasks such as image recognition (Liu et al. 2021; Guo et al. 2022), object detection (Carion et al. 2020; Yang et al. 2021), etc. In this paper, we introduce a novel semantic attention mechanism to capture invariant RoI features under more variety of transformations.

### Methodology

In this section, we will first introduce the general architecture of the proposed Semantic RoI Align (SRA). Next, we will detail how the semantic masks of SRA are obtained. We will then present a dynamic feature sampling method to dynamically sample features for SRA according to different RoI aspect ratios, which improves model accuracy and efficiency. Finally, the proposed Area Embedding is introduced to replace the previous position embedding, so as to provide more accurate position information for the model.

### Semantic RoI Align

The pipeline of the proposed Semantic RoI Align (SRA) is shown in Figure 2. The SRA extracted RoI feature of an object consists of  $N$  sub-features, each of which is sampled in a specific semantic region, making the sampling position adaptive to image transformations, and thereby improving the transformation-invariance. In Figure 3, we visualize partial semantic masks (left 5 columns) produced by our SRA for 3 RoI proposals (one for each row). The semantic samplings of SRA can sample on the same semantic parts for the object under different perspective transformations such as rotation (top row and middle row) and object pose transformations (top row and bottom row), giving the extracted RoI feature better transformation-invariance, and thus being beneficial to bounding boxes regression and semantic labels prediction in the second-stage network.

The inputs of SRA are a feature map  $F$  with shape  $(C, H, W)$  where  $C$ ,  $H$  and  $W$  represent the number of feature channels, height, and width of the feature map, respectively, and a list of RoI proposals  $R = \{R_i\}$  where  $R_i = \{x_{i,0}, y_{i,0}, x_{i,1}, y_{i,1}\}$  indicates a bounding box in the feature map with  $(x_{i,0}, y_{i,0})$  and  $(x_{i,1}, y_{i,1})$  being the coordinates of the top left corner and the bottom right corner, respectively. For each RoI proposal  $R_i$ , SRA first exploits the Dynamic Feature Sampler to sample a feature map  $f_i$  from the input feature map  $F$  with the bounding box  $R_i$ . We then obtain  $N$  semantic masks  $m_i = \{m_{i,n}\}, 1 \leq n \leq N$ , which

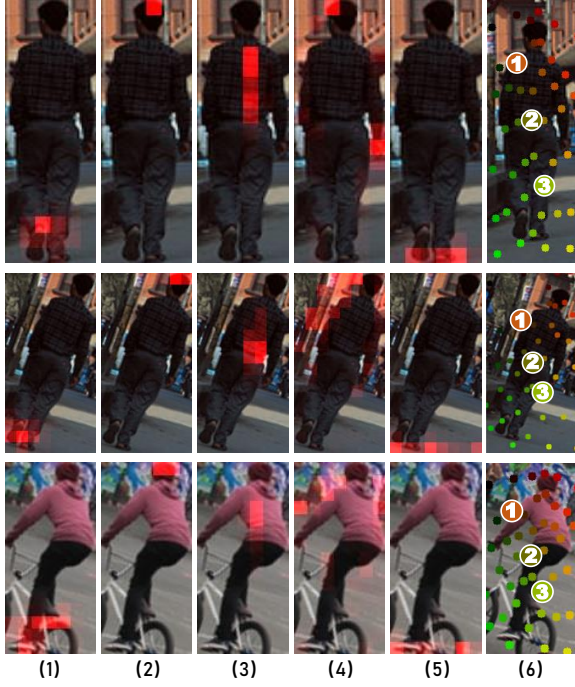


Figure 3: Semantic sampling masks of SRA (columns 1 to 5), sampling locations of DRoIPooling (Zhu et al. 2019) (column 6), and sampling mask of directly passing features extracted by ROI Align through a standard self-attention layer (column 7). Each row represents an ROI in an image. The object in the middle row is rotated by 30 degrees from the top row, and the bottom row shows another object of the same class as the top row. Red and yellow sampling masks are overlaid on images and the sampling locations of DRoIPooling are indicated with different colors.

have the same size as  $f_i$ . The output transformation-invariant ROI feature  $y_i$  of our SRA is finally obtained by sampling  $f_i$  using the semantic masks  $m_i$ . More specifically,  $y_i$  is calculated as the weighted sum of  $f_i$  elements using  $m_i$  elements as weights:

$$y_i(n, c) = \sum_{j=1}^{h_i} \sum_{k=1}^{w_i} f_i(c, j, k) \cdot m_{i,n}(j, k), \quad (1)$$

$$\text{for all } n \in \{1, \dots, N\}, c \in \{1, \dots, C\},$$

where  $h_i, w_i$  are the height and width of feature map  $f_i$ . Next, we will introduce how the semantic masks  $m_i$  are estimated in the SRA.

### Obtaining Semantic Masks

The pipeline of obtaining semantic masks of SRA is also shown in Figure 2. The goal of SRA is to generate  $N$  separable semantic part masks  $m_i$  for the input ROI proposal  $R_i$  and the sampled feature map  $f_i$ . To achieve this, we want the value of  $m_{i,n}(j, k)$  to be positively correlated with the likelihood that position  $(j, k)$  in  $R_i$  belongs to the  $n$ -th semantic part of the object in  $R_i$ . Let us denote that likelihood as  $m'_i(n, j, k)$ . The likelihood is related to two factors,

namely, 1) what it is in  $R_i$ , and 2) what it is at the position  $(j, k)$  of  $R_i$ . The former is expressed by a  $K$ -dimensional ROI descriptor  $d_i$  representing the overall features of  $R_i$ , and the latter is characterized by a semantic feature map  $s_i$  with shape  $(K, h_i, w_i)$  meaning the semantic feature at different positions in  $R_i$ . To make the final ROI feature computed based on Eq. 1 transformation-invariant, the  $m_i$  should transform accordingly when the object transforms. Therefore, we obtain the likelihood  $m'_i$  by using a same regressor at different positions  $(j, k)$ :

$$m'_{i,n}(j, k) = \xi_n([d_i, s_i(:, j, k)]), \quad (2)$$

$$\text{for all } n \in \{1, \dots, N\}, j \in \{1, \dots, h_i\}, k \in \{1, \dots, w_i\},$$

where the  $\xi$  are  $N$  learnable sub-mask regressors, each composed of two lots of Norm-ReLU-Linear,  $[\cdot, \cdot]$  means the concatenate operation, and  $s_i(:, j, k)$  represents the semantic feature in the position  $(j, k)$ . By doing so, if some transformation of the object causes the feature at position  $(j, k)$  to move to position  $(j', k')$ , the transformed masks  $\hat{m}_{i,n}(j', k') = \xi_n([\hat{d}_i, \hat{s}_i(:, j', k')])$  will have similar value with  $m_{i,n}(j, k)$ , since the transformed  $\hat{d}_i$  and  $\hat{s}_i(:, j', k')$  are similar to  $d_i$  and  $s_i(:, j, k)$ , respectively. This means the semantic masks transform accordingly with the transformation. We obtained  $s_i$  by performing a  $1 \times 1$  convolution on  $f_i$ , and we explored various forms of the ROI Descriptor Regressor to obtain  $d_i$ :

$$\text{- Concatenation: } d_i = \psi(\text{Flatten}(f_i))$$

$$\text{- Maximum: } d_i = \psi\left(\max_{j=1}^{h_i} \max_{k=1}^{w_i} f_{i,(:,j,k)}\right) \quad (3)$$

$$\text{- Average: } d_i = \psi\left(\frac{1}{h_i \times w_i} \sum_{j=1}^{h_i} \sum_{k=1}^{w_i} f_{i,(:,j,k)}\right)$$

Here  $\psi$  is a linear layer with  $K$  output channels.

Sampling features based on semantic masks may cause the model to lose position information, which is important for the object detection task. We thus use a position embedding  $p_i$  (see (Zhao, Jia, and Koltun 2020)) to provide position information for the model, and use positional embedded  $m'_{i,n}(j, k) = \xi_n([d_i, s_i(:, j, k), p_i(:, j, k)])$  instead of Eq. 2. The  $p_i$  is obtained by performing a  $1 \times 1$  convolution with output channels of  $P$  on  $p'_i$ , where  $p'_i$  with shape  $(2, h_i, w_i)$  is the relative position of each position in the ROI, and is normalized to  $[-1, 1]$ :

$$p'_i(1, j, k) = j/h_i \times 2 - 1, \quad (4)$$

$$p'_i(2, j, k) = k/w_i \times 2 - 1,$$

The semantic masks  $m_i$  is then obtained by  $m_i = \text{softmax}(m'_i \cdot \gamma)$ , where  $\gamma$  is an amplification factor that amplifies the backpropagation response of masks, and the softmax acts on the last two dimensions to ensure the sum of each mask to 1. Finally, the output ROI feature  $y_i$  is obtained by summing up elements of  $f_i$  weighted by the  $N$  semantic masks  $m_i$  as shown in Eq. 1.

### Dynamic Feature Sampler

In SRA, the semantic masks  $m_i$  are estimated via the sub-sampled feature map  $f_i$ , and thus the computational overhead of SRA is proportional to the input size of  $f_i$ , i.e.



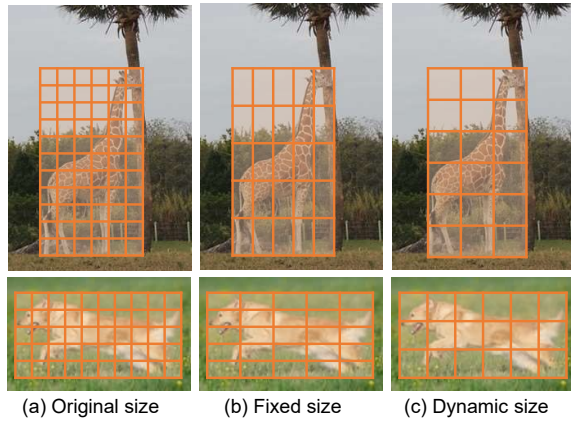


Figure 4: Different methods to determine the size of the sub-sampled feature map. (a) Using a size that is the same resolution as the original feature map is costly as it results in too many samples. (b) Using a fixed size may cause the aspect ratio of the region represented by each sub-sampled feature to be inconsistent for different RoIs, which we believe is harmful to the model, e.g. a ratio of approximately 2 for the upper RoI and 0.5 for the lower RoI. (c) Our Dynamic Feature Sampler overcomes the limitations of the above two methods, yielding both consistent and limited samples.

$h_i \times w_i$ . The size of feature map  $f_i$  can be set to different values for different RoIs. A straightforward solution is to set the size to the original resolution of  $R_i$ . However, this would lead to a large computational cost for some large RoIs. Another way is to set the size to a fixed value. However, as shown in Figure 4(b), this may cause the aspect ratio of the region represented by each feature of  $f_i$  to be inconsistent for different RoIs, and experiments show that this will lead to loss of accuracy. To balance sampling quality and computational efficiency, we propose a Dynamic Feature Sampler to select the size of  $f_i$  for each proposal, which keeps the aspect ratio of each sub-sampled region close to 1, and has a limited size. Specifically, for each RoI  $R_i$ , we pick the size that has the closest aspect ratio to  $R_i$  while not exceeding a maximal area  $M$ . Mathematically, this can be formulated as:

$$h_i, w_i = \arg \min_{\substack{(h', w') \in \mathbb{Z}^2 \\ \text{with } h' \cdot w' \leq M}} \left| \frac{h'}{w'} - \frac{x_{i,1} - x_{i,0}}{y_{i,1} - y_{i,0}} \right|. \quad (5)$$

The sub-sampled feature map  $f_i$  is then obtained by dividing  $R_i$  region of  $F$  into  $h_i \times w_i$  blocks and averaging the feature values in each block. With the Dynamic Feature Sampler, our SRA yields a good performance with a small computational overhead.

### Area Embedding

In Eq. 4, we use position embedding of each grid center in the mask to provide sampling position information to the model. However, as shown in Figure 5(a), since we use a dynamic way to determine the size of the sub-sampled feature map, the same center position may represent different sampling areas. We thus propose Area Embedding to encode the




Figure 5: Schematic diagram comparison of traditional Position Embedding and proposed Area Embedding. (a) Position embedding only embeds the coordinates of the sampling center, while the same center position may represent different sampling areas. (b) Our Area Embedding embeds the entire sampling area.

sampled area of each point in the output feature with two fixed-length vectors, each representing both the position and the coverage on the horizontal and vertical axes. We set the length of this vector to  $M$ , which is the maximal number of samples per axis. For each point  $(j, k) \in \mathbb{Z}^2$  sampled by SRA, we calculate  $p'_i$  by:

$$\begin{aligned} p'_i(1 \cdots M, j, k) &= \text{Upsample}(\text{OneHot}(j; h_i); M), \\ p'_i((M + 1) \cdots 2M, j, k) &= \text{Upsample}(\text{OneHot}(k; w_i); M). \end{aligned} \quad (6)$$

where the  $\text{OneHot}(b; a)$  operator takes an integer  $b$  less or equal to  $a$  as input and produces the one hot embedding of  $b$  within a vector of length  $a$ , and  $\text{Upsample}(\mathbf{v}; M)$  upsamples vector  $\mathbf{v}$  to a  $M$ -sized vector. The upsampling method can vary; in Figure 5, we use nearest sampling for convenience of illustration, while in our experiments we use linear sampling for higher accuracy. The Area Embedding provides the model more accurate sampling position information and experiments show that it improves the accuracy of the model.

## Experiments

We conduct our experiments on the MS COCO dataset (Lin et al. 2014), and use the train2017 for training and use the val2017 and the test2017 for testing. We report the standard COCO evaluation metrics including mean Average Precision (AP) under different Intersection over Union (IoU) thresholds and at different object scales, denoted as  $AP$  for the object detection task and  $AP^m$  for the instance segmentation task. Our model is implemented based on Jitter (Hu et al. 2020) and JDet<sup>1</sup> library. The implementation details of our model are given in the supplementary material.

### Ablation Studies

We first conduct a series of ablation experiments to verify the effectiveness of each part of the proposed model. The ablation experiments are conducted on the object detection task using the MS COCO validation set. We couple our proposed feature extractor with Faster-RCNN using ResNet-50 as the backbone. RoI Align (He et al. 2017) is used as our baseline model, if not specifically mentioned.

<sup>1</sup><https://github.com/Jitter/JDet>

Method	AP	AP <sub>50</sub>	AP <sub>75</sub>	AP <sub>S</sub>	AP <sub>M</sub>	AP <sub>L</sub>	Params	FLOPs
RoI Align	37.5	58.2	40.8	21.8	41.1	48.1	41.8M	340.9G
DRoIPooling	37.9	59.4	41.8	22.4	41.4	49.3	149.4M	349.0G
w/ Conv.	36.4	57.4	39.6	21.5	40.0	46.9	71.9M	350.1G
w/ SA	37.6	58.0	40.8	20.9	41.1	48.7	42.7M	348.1G
SRA (Ours)	<b>39.2</b>	<b>59.6</b>	<b>42.6</b>	<b>22.5</b>	<b>42.6</b>	<b>51.9</b>	42.0M	344.2G

Table 1: Results of comparative and ablation experiments between our SRA and baseline models.

**The Effectiveness of SRA.** To verify the effectiveness of the proposed SRA, we replaced it with RoI Align and Modulated Deformable RoI Pooling (DRoIPooling) (Zhu et al. 2019). The results in Table 1 show that, our model outperforms the baseline model by 1.7% AP with a minor computational and parameters cost, and also outperforms the DRoIPooling by 1.3% AP with a much smaller model.

We also compared SRA with two other baselines, by performing some simple operations on the features extracted by RoI Align: applying a convolutional layer on the features to obtain the sampling masks and re-sampling the features with these masks (denoted as “w/ Conv.”), as well as directly passing the features through a standard multi-head self-attention layer (Vaswani et al. 2017) (denoted as “w/ SA”). The results in Table 1 show that our model achieves a gain of 2.8% and 1.6% in AP, respectively, with a smaller number of parameters and FLOPs. This improvement can be attributed to its enhanced capability in identifying consistent semantic parts across diverse transformations, leading to better transformation-invariance.

We also visualize some samplings of SRA, DRoIPooling, and “w/ SA” in Figure 3. Rows 1 & 2 together show how samplings respond to different transformations of the same object, while rows 1 & 3 together indicate how samplings respond to different objects of the same class. We found the sampling masks of SRA (columns 1 to 5) can be divided into two classes. The first class samples on different semantic parts. For example, columns 1-4 show the samples on the human’s feet, head, and body, and around the human, respectively. The second class of sampling is for positioning, which is only activated in certain positions. For example, the 5th column is only activated on the bottom position of the RoI. Our semantic samplings can sample on the same semantic parts for the object under different transformations, which gives the extracted RoI feature better transformation invariance. In comparison, the sampling locations of DRoIPooling (column 6 of Figure 3) are distributed mostly inside the object as the object transforms, however, they will not vary accordingly with object pose transformations. Taking 3 samplings (in column 6) as an example, in the top row and middle row, the circles numbered from 1 to 3 do not rotate with the object, which means DRoIPooling can not always achieve transformation-invariance under some complex transformations like rotation. Also, simply passing features through a self-attention layer (w/ SA, column 7 of Figure 3) cannot ensure sampling on the same semantic parts,

D	S	A	AP	AP <sub>50</sub>	AP <sub>75</sub>	AP <sub>S</sub>	AP <sub>M</sub>	AP <sub>L</sub>
✓	✗	✗	31.4	52.2	32.7	17.8	35.1	40.0
✗	✓	✗	36.2	57.5	38.6	21.0	39.8	46.6
✗	✗	✓	37.3	58.3	40.7	21.5	41.0	48.1
✗	✓	✓	38.9	59.3	42.3	22.5	42.5	51.3
✓	✗	✓	37.4	58.4	40.4	22.1	41.0	48.4
✓	✓	✗	36.4	57.3	39.0	20.9	39.9	46.7
✓	✓	PE	38.8	59.2	42.5	<b>22.6</b>	42.3	50.7
✓	✓	✓	<b>39.2</b>	<b>59.6</b>	<b>42.6</b>	22.5	<b>42.6</b>	<b>51.9</b>

Table 2: Ablation study on the effectiveness of each module in our SRA. D, S, and A denote RoI descriptor, semantic feature map, and Area Embedding respectively.

Setting	AP	AP <sub>50</sub>	AP <sub>75</sub>	AP <sub>S</sub>	AP <sub>M</sub>	AP <sub>L</sub>
DR=Con.	38.9	59.3	41.9	<b>22.5</b>	42.3	51.1
DR=Max.	39.0	59.3	<b>42.9</b>	<b>22.5</b>	<b>42.6</b>	51.1
DR=Avg.	<b>39.2</b>	<b>59.6</b>	42.6	<b>22.5</b>	<b>42.6</b>	<b>51.9</b>
$\gamma = 1$	38.3	58.7	41.8	22.2	41.8	50.2
$\gamma = 5$	38.7	59.2	42.0	22.1	42.5	51.0
$\gamma = 50$	<b>39.2</b>	<b>59.6</b>	<b>42.6</b>	<b>22.5</b>	<b>42.6</b>	<b>51.9</b>
$\gamma = 500$	36.8	57.6	39.9	21.2	40.2	48.1

Table 3: Experiments on different module settings. DR denotes RoI Descriptor Regressor, and  $\gamma$  denotes the amplification factor.

Setting	AP	AP <sub>50</sub>	AP <sub>75</sub>	AP <sub>S</sub>	AP <sub>M</sub>	AP <sub>L</sub>	Params	FLOPs
$N = 9$	38.0	58.2	41.3	21.8	41.3	49.9	31.5M	340.7G
$N = 25$	38.6	59.0	41.9	22.8	42.3	50.5	35.7M	342.1G
$N = 49$	39.2	59.6	42.6	22.5	<b>42.6</b>	<b>51.9</b>	42.0M	344.2G
$N = 100$	<b>39.4</b>	<b>59.9</b>	<b>42.9</b>	<b>22.9</b>	<b>42.6</b>	51.7	55.4M	348.7G
$size = 3^2$	36.3	57.3	39.0	20.6	39.9	46.4	31.3M	337.8G
$size = 5^2$	37.2	58.0	40.7	21.4	40.8	47.8	35.5M	339.0G
$size = 7^2$	37.3	58.3	40.7	21.5	41.0	48.1	41.8M	340.9G
$size = 10^2$	37.6	58.4	40.8	22.0	41.3	48.8	55.1M	344.9G

Table 4: Comparison between SRA with a different number of masks (N) and the baseline model with comparable RoI sizes.

thus failing to obtain transformation-invariant RoI features.

**Structure of SRA and Area Embedding.** We also conduct experiments to verify the effectiveness of different components in the SRA by controlling whether to concatenate the RoI descriptor (D), semantic feature map (S), and Area Embedding (A) when regressing the masks. The results are listed in Table 2. Comparing the last row with the 5th row in the table, our model with semantic feature map obtains a gain of 1.8% in AP, as our model determines the masks for sampling based on semantic features,

Setting	AP	AP <sub>50</sub>	AP <sub>75</sub>	Avg. S	Params	FLOPs
fixed	38.8	59.2	42.4	64	42.0M	344.1G
$M = 32$	38.3	59.1	42.0	16	42.0M	341.7G
$M = 64$	38.8	59.5	42.5	32	42.0M	342.5G
$M = 128$	<b>39.2</b>	<b>59.6</b>	<b>42.6</b>	64	42.0M	344.2G
$M = 256$	39.1	59.5	42.5	128	42.0M	347.9G

Table 5: Experiments on the Dynamic Feature Sampler.  $M$  denotes the dynamic feature map size limit.

Method	AP	AP <sub>50</sub>	AP <sub>75</sub>	AP <sub>S</sub>	AP <sub>M</sub>	AP <sub>L</sub>
RoI Pooling	37.2	58.9	40.3	21.5	40.2	46.0
RoI Align	37.7	58.9	40.6	21.9	40.7	46.4
DRoIPooling	38.1	60.0	42.0	22.0	41.2	47.2
Ada. RoI Align	37.7	58.8	40.7	21.8	40.7	46.5
Pr. RoI Align	37.8	58.9	40.9	22.1	40.8	46.7
RoIAttn	38.0	59.3	40.9	22.4	41.1	46.9
SRA	<b>39.2</b>	<b>59.8</b>	<b>42.6</b>	<b>22.6</b>	<b>42.1</b>	<b>49.0</b>

Table 6: Comparison with different RoI extractors on the MS COCO detection test-dev set.

which makes the sampled features invariant under a variety of transformations, thus achieving better performance. We also tested our model with or without Area Embedding (8th row and 6th row respectively) and replaced the Area Embedding with position embedding (denoted as PE, 7th row). The results show that the model with AE obtains 2.8% higher AP than without AE, and 0.4% higher AP than with PE, which demonstrates AE can describe more accurately the sampling information of Dynamic Feature Sampler and provide better position information for the model.

**Choices of the RoI Descriptor Regressor.** We tested various choices of the RoI Descriptor Regressor in Eq. 3, denoted as DR=Con., Max., and Avg., respectively, where the choice of *concatenation* is tested on the  $8 \times 8$  fixed size feature sampler as it cannot be adapted to the Dynamic Feature Sampler. The results are shown in Table 3. Though the choice of DR=Con. shows slightly better performance than  $8 \times 8$  fixed size feature sampler with the choice of *average* (38.8% AP in Table 5), considering that it cannot be adapted to the Dynamic Feature Sampler and will lead to a larger amount of calculation and more parameters, we finally use the *average* RoI Descriptor Regressor in our model.

**Parameters Setting.** The number of masks  $N$  determines the size of the RoI features extracted by our model. We tested different settings of  $N$  (denoted as “ $N = x$ ”) and compared them with the baseline model with different settings of RoI Align output size (denoted as “*size* =  $x$ ”). The results in Table 4 show that, with the same RoI feature size (comparing the 1st row with the 5th row, the 2nd row with the 6th row, etc.), our model has a 1.4%-1.9% higher AP, which proves that the transformation-invariant features extracted by our model contain richer information under the

same feature length and are more conducive to object detection. Considering the balance between model parameters and accuracy, we finally set  $N = 49$ . We also tested different settings of the amplification factor  $\gamma$ , denoted as “ $\gamma = x$ ” in Table 3. Results show that setting it to an appropriate value is beneficial to the regression of the semantic mask, so we set  $\gamma = 50$  according to the experiment.

**Dynamic Feature Sampler.** To evaluate the effectiveness of the Dynamic Feature Sampler, we compared  $8 \times 8$  fixed size feature sampler (denoted as fixed) with Dynamic Feature Sampler ( $M = 128$ ), which has the same number of average samplings and similar FLOPs. As shown in the 1st and 4th row of Table 5, the Dynamic Feature Sampler obtained better results as its sub-sampled feature represented region has a more consistent aspect ratio. We also tested different dynamic feature map size limit  $M$ . A larger  $M$  brings a higher feature sampling resolution. In general, the accuracy increases with the increment in the resolution; however, the accuracy improvement brought by the increase in sampling resolution is limited by the resolution of the original feature map, and will gradually tend to zero. So, we choose  $M = 128$  based on the experiment.

## Comparison with Other Methods

We also compared our model with other methods on the COCO test set. We first compared ours with different RoI feature extractors: RoI Pooling (Girshick 2015), RoI Align (He et al. 2017), Adaptive RoI Align (Jung et al. 2018), Precise RoI Align (Jiang et al. 2018), DRoIPooling (Zhu et al. 2019), and RoIAttn (Liang and Song 2022) on Faster R-CNN with ResNet50 as backbone, trained for 12 epochs. The results are shown in Table 6. One can see that our model achieves the best performance. In particular, compared to the RoIAttn which incorporates several self-attention layers on the RoI Align extracted feature, our model obtains 1.2% higher AP. The results demonstrate the advantage of our SRA method for extracting RoI features that are invariant to more types of transformations.

To verify the generalizability of our method, we also examined the performance gain by using SRA across different detection methods and backbones. Please refer to the supplementary material for more details. Our model improves the accuracy of various detection methods and backbone networks, demonstrating the generalizability of our model.

## Conclusions

In this paper, we proposed SRA, a transformation-invariant RoI feature extractor. It regresses semantic masks based on a novel semantic attention structure, and obtains RoI features by sampling the feature map with these semantic masks, making it invariant under more diverse transformations. We further proposed the Dynamic Feature Sampler to speed up the process while minimizing the impact on accuracy, and proposed Area Embedding to provide more accurate sampling area information. Benefiting from the capability and generalizability of SRA, experiments show that its utilization can bring significant performance improvement to various baselines and state-of-the-art models with a small computational and parameter overhead.

## Acknowledgments

This work was supported by the National Key Research and Development Program of China under Grant 2021ZD0112902, the National Natural Science Foundation of China under Grant 62220106003, and the Research Grant of Beijing Higher Institution Engineering Research Center and Tsinghua-Tencent Joint Laboratory for Internet Innovation Technology.

## References

- Carion, N.; Massa, F.; Synnaeve, G.; Usunier, N.; Kirillov, A.; and Zagoruyko, S. 2020. End-to-End Object Detection with Transformers.
- Dai, J.; Qi, H.; Xiong, Y.; Li, Y.; Zhang, G.; Hu, H.; and Wei, Y. 2017. Deformable convolutional networks. In *Proceedings of the IEEE international conference on computer vision*, 764–773.
- Dai, X.; Chen, Y.; Xiao, B.; Chen, D.; Liu, M.; Yuan, L.; and Zhang, L. 2021. Dynamic head: Unifying object detection heads with attentions. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 7373–7382.
- Ding, J.; Xue, N.; Long, Y.; Xia, G.-S.; and Lu, Q. 2019. Learning RoI transformer for oriented object detection in aerial images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2849–2858.
- Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissenborn, D.; Zhai, X.; Unterthiner, T.; Dehghani, M.; Minderer, M.; Heigold, G.; Gelly, S.; Uszkoreit, J.; and Housley, N. 2020. An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale.
- Girshick, R. 2015. Fast r-cnn. In *Proceedings of the IEEE international conference on computer vision*, 1440–1448.
- Guo, M.-H.; Lu, C.-Z.; Liu, Z.-N.; Cheng, M.-M.; and Hu, S.-M. 2022. Visual attention network. *arXiv preprint arXiv:2202.09741*.
- He, K.; Gkioxari, G.; Dollár, P.; and Girshick, R. 2017. Mask r-cnn. In *Proceedings of the IEEE international conference on computer vision*, 2961–2969.
- Hu, J.; Shen, L.; and Sun, G. 2018. Squeeze-and-excitation networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 7132–7141.
- Hu, S.-M.; Liang, D.; Yang, G.-Y.; Yang, G.-W.; and Zhou, W.-Y. 2020. Jittor: a novel deep learning framework with meta-operators and unified graph execution. *Science China Information Sciences*, 63: 1–21.
- Jiang, B.; Luo, R.; Mao, J.; Xiao, T.; and Jiang, Y. 2018. Acquisition of localization confidence for accurate object detection. In *Proceedings of the European conference on computer vision (ECCV)*, 784–799.
- Jung, I.; Son, J.; Baek, M.; and Han, B. 2018. Real-time md-net. In *Proceedings of the European conference on computer vision (ECCV)*, 83–98.
- Krizhevsky, A.; Sutskever, I.; and Hinton, G. E. 2012. Imagenet classification with deep convolutional neural networks. *Advances in neural information processing systems*, 25: 1097–1105.
- Liang, X.; and Song, P. 2022. Excavating roi attention for underwater object detection. In *2022 IEEE International Conference on Image Processing (ICIP)*, 2651–2655. IEEE.
- Lin, T.-Y.; Maire, M.; Belongie, S.; Hays, J.; Perona, P.; Ramanan, D.; Dollár, P.; and Zitnick, C. L. 2014. Microsoft coco: Common objects in context. In *European conference on computer vision*, 740–755. Springer.
- Liu, W.; Anguelov, D.; Erhan, D.; Szegedy, C.; Reed, S.; Fu, C.-Y.; and Berg, A. C. 2016. Ssd: Single shot multibox detector. In *European conference on computer vision*, 21–37. Springer.
- Liu, Z.; Lin, Y.; Cao, Y.; Hu, H.; Wei, Y.; Zhang, Z.; Lin, S.; and Guo, B. 2021. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 10012–10022.
- Mnih, V.; Heess, N.; Graves, A.; et al. 2014. Recurrent models of visual attention. *Advances in neural information processing systems*, 27.
- Redmon, J.; Divvala, S.; Girshick, R.; and Farhadi, A. 2016. You only look once: Unified, real-time object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 779–788.
- Ren, S.; He, K.; Girshick, R.; and Sun, J. 2015. Faster r-cnn: Towards real-time object detection with region proposal networks. *Advances in neural information processing systems*, 28.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.
- Yang, G.-Y.; Li, X.-L.; Martin, R. R.; and Hu, S.-M. 2021. Sampling equivariant self-attention networks for object detection in aerial images. *arXiv preprint arXiv:2111.03420*.
- Zhao, H.; Jia, J.; and Koltun, V. 2020. Exploring self-attention for image recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 10076–10085.
- Zhu, X.; Hu, H.; Lin, S.; and Dai, J. 2019. Deformable convnets v2: More deformable, better results. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 9308–9316.