

DETR Does Not Need Multi-Scale or Locality Design

Yutong Lin^{1†} Yuhui Yuan^{2†} Zheng Zhang^{2†} Chen Li¹ Nanning Zheng¹ Han Hu^{2†}

¹National Key Laboratory of Human-Machine Hybrid Augmented Intelligence,
National Engineering Research Center for Visual Information and Applications,
and Institute of Artificial Intelligence and Robotics, Xi’an Jiaotong University

²Microsoft Research Asia

Abstract

This paper presents an improved DETR detector that maintains a “plain” nature: using a single-scale feature map and global cross-attention calculations without specific locality constraints, in contrast to previous leading DETR-based detectors that reintroduce architectural inductive biases of multi-scale and locality into the decoder. We show that two simple technologies are surprisingly effective within a plain design to compensate for the lack of multi-scale feature maps and locality constraints. The first is a box-to-pixel relative position bias (BoxRPB) term added to the cross-attention formulation, which well guides each query to attend to the corresponding object region while also providing encoding flexibility. The second is masked image modeling (MIM)-based backbone pre-training which helps learn representation with fine-grained localization ability and proves crucial for remedying dependencies on the multi-scale feature maps. By incorporating these technologies and recent advancements in training and problem formulation, the improved “plain” DETR showed exceptional improvements over the original DETR detector. By leveraging the Object365 dataset for pre-training, it achieved 63.9 mAP accuracy using a Swin-L backbone, which is highly competitive with state-of-the-art detectors which all heavily rely on multi-scale feature maps and region-based feature extraction. Code will be available at <https://github.com/impiga/Plain-DETR>.

1. Introduction

The recent revolutionary advancements in natural language processing highlight the importance of keeping task-specific heads or decoders as general, simple, and lightweight as possible, and shifting main efforts towards building more powerful large-scale foundation models [37, 11, 2]. However, the computer vision community often

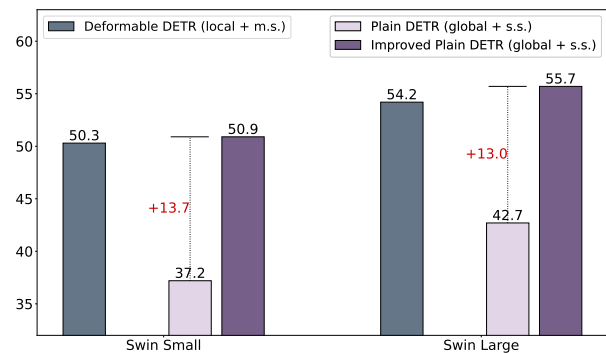


Figure 1: We improve the plain DETR detectors, which rely on global cross-attention calculation and single-scale (s.s.) feature maps, by huge margins, using both Swin-S and Swin-L backbones. It makes plain DETRs as competitive as the present leading DETR detectors based on local cross-attention and multi-scale (m.s.) feature maps.

continues to focus heavily on the tuning and complexity of task-specific heads, resulting in designs that are increasingly heavy and complex.

The development of DETR-based object detection methods follows this trajectory. The original DETR approach [4] is impressive in that it discarded complex and domain-specific designs such as multi-scale feature maps and region-based feature extraction that require a dedicated understanding of the specific object detection problem. Yet, subsequent developments [55, 54] in the field have reintroduced these designs, which do improve training speed and accuracy but also contravene the principle of “fewer inductive biases” [13].

In this work, we aim to improve upon the original DETR detector, while preserving its “plain” nature: *no multi-scale feature maps, no locality design for cross-attention calculation*. This is challenging as object detectors need to handle objects of varying scales and locations. Despite the latest improvements in training and problem formulation, as shown in Table 1, the plain DETR method still lags greatly

[†]Equal contribution. ✉ {yuhui.yuan, hanhu}@microsoft.com

behind state-of-the-art detectors that utilize multi-scale feature maps and regional feature extraction design.

So, how can we compensate for these architectural “inductive biases” in addressing multi-scale and arbitrarily located objects? Our exploration found that two simple technologies, though not entirely new, were surprisingly effective in this context: box-to-pixel relative position bias (BoxRPB) and masked image modeling (MIM) pre-training. BoxRPB is inspired by the relative position bias (RPB) term in vision Transformers [34, 33] which encodes the geometric relationship between pixels and enhances translation invariance. BoxRPB extends RPB to encode the geometric relationship between $4d$ -boxes and $2d$ -pixels. We also present an axial decomposition approach for efficient computation, with no loss of accuracy compared to using the full term. Our experiments show that the BoxRPB term can well guide the cross-attention computation to be well dedicated to individual objects (see Figure 4, and it dramatically improves detection accuracy by +8.9 mAP over a plain DETR baseline of 37.2 mAP on the COCO benchmark (see Table 2).

The utilization of MIM pre-training is another crucial technology in enhancing the performance of plain DETR. Our results demonstrate also a significant improvement of +7.4 mAP over the plain DETR baseline (see Table 2), which may be attributed to its fine-grained localization capability [49]. While MIM pre-training has been shown to moderately improve the performance of other detectors [20, 50], its impact in plain settings is profound. Furthermore, the technology has proven to be a key factor in eliminating the necessity of using multi-scale feature maps from the backbones, thereby expanding the findings in [28, 15] to detectors that utilize hierarchical backbones or single-scale heads.

By incorporating these technologies and the latest improvements in both training and problem formulation, our improved “plain” DETR has demonstrated exceptional improvements over the original DETR detector, as illustrated in Figure 1. Furthermore, our method achieved an accuracy of 63.9 mAP when utilizing the Object365 dataset for pre-training, making it highly competitive with state-of-the-art object detectors that rely on multi-scale feature maps and region-based feature extraction techniques, such as cascade R-CNN [33] and DINO [54], among others.

Beyond these outcomes, our methodology exemplifies how to minimize the architectural “inductive bias” when designing an effective task-specific head or decoder, as opposed to relying on detection-specific multi-scale and localized designs. Our study hopes to inspire future research on using generic plain decoders, such as that of DETR, for a wider range of visual problems with minimal effort, thus allowing the field to shift more energy to developing large foundation visual models, similar to what occurs in the field

of natural language processing.

2. A Modernized Plain DETR Baseline

2.1. A Review of the Original DETR

The original DETR detector [4] consists of 3 sub-networks:

- *A backbone network \mathcal{F}_b* to extract image features from an image. We denote the input image as $\mathbf{I} \in \mathbb{R}^{H \times W \times 3}$. The backbone network can provide multi-scale feature maps $\mathbf{C}^2, \mathbf{C}^3, \mathbf{C}^4, \mathbf{C}^5$, if a convolutional ConvNet is used, i.e., ResNet [22]. The spatial resolutions are typically $1/4^2, 1/8^2, 1/16^2$, and $1/32^2$ of the input image. The original DETR detector used the mainstream backbone architecture at the time, ResNet, as its backbone network, and either an original ResNet or a variant with a dilated stage 5 network is used. Now the mainstream backbone network has evolved to vision Transformers, which will be used in our experiments, e.g., Swin Transformer [34].
- *A Transformer encoder \mathcal{F}_e* to enhance the image features. It applies on $\mathbf{P}^5 \in \mathbb{R}^{\frac{HW}{32^2} \times C}$ ($C=256$), obtained via a linear projection on \mathbf{C}^5 . The Transformer encoder usually consists of several stacking Transformer blocks, i.e., 6 in the original DETR.
- *A global Transformer decoder \mathcal{F}_d* to decode object bounding boxes from the image feature map using a set of randomly initialized object queries $\mathbf{Q} = \{\mathbf{q}_0, \mathbf{q}_1, \dots, \mathbf{q}_n\}$. The Transformer decoder also usually consists of multiple layers, with each layer including a self-attention block, a cross-attention block, and a feed-forward block. Each of the decoder layers will produce a set of objects with labels and bounding boxes, driven by a set matching loss.

The DETR framework possesses several merits, including: 1) Conceptually straightforward and generic in applicability. It views object detection as a pixel-to-object “translation” task, with a generic notion of decoding image pixels into problem targets. 2) Requiring minimal domain knowledge, such as custom label assignments and hand-designed non-maximum suppression, due to the use of an end-to-end set matching loss. 3) Being plain, avoiding domain-specific multi-scale feature maps and region-based feature extraction.

In the following, we will first build an enhanced DETR-based detector by incorporating recent advancements regarding both training and problem formulation, while maintaining the above nice merits.

2.2. An Enhanced Plain DETR Baseline

Basic setup. Our basic setup mostly follows the original DETR framework, except for the following adaptations: 1) We use a stronger Swin-T backbone, instead of the original ResNet50 backbone; 2) We create a feature map of P_4 from C_5 by deconvolution, instead of adding dilation operations to the last stage of the backbone, for simplicity purpose. 3) We set the number of queries as 300, and the dropout ratio of the Transformer decoder as 0. 4) We use $1\times$ scheduler settings (12 epochs) for efficient ablation study. As shown in Table 1, this basic setup produces a 22.5 mAP on COCO val.

In the following, we incorporate some recent advancements in training and problem formulation into the basic setup, and gradually improve the detection accuracy to 37.2 mAP, as shown in Table 1.

Merging Transformer encoder into the backbone. The backbone network and Transformer encoder serve the same purpose of encoding image features. We discovered that by utilizing a Vision Transformer backbone, we are able to consolidate the computation budget of the Transformer encoder into the backbone, with slight improvement, probably because more parameters are pre-trained. Specifically, we employed a Swin-S backbone and removed the Transformer encoder. This method resulted in similar computation FLOPs compared to the original Swin-T plus 6-layer Transformer encoder. This approach simplifies the overall DETR framework to consist of only a backbone (encoder) and a decoder network.

Focal loss for better classification. We follow [55] to utilize focal loss [30] to replace the default cross-entropy loss, which improves the detection accuracy significantly from 23.1 mAP to 31.6 mAP.

Iterative refinement. We follow the iterative refinement scheme [43, 55, 3] to make each decoder layer predict the box delta over the latest bounding box produced by the previous decoder layer, unlike the original DETR that uses independent predictions within each Transformer decoder layer. This strategy improves the detection accuracy by +1.5 mAP to reach 33.1 mAP.

Content-related query. We follow [55] to generate object queries based on image content. The top 300 most confident predictions are selected as queries for the subsequent decoding process. A set matching loss is used for object query generation, thereby maintaining the merit of no domain-specific label assignment strategy. This modification resulted in a +0.9 mAP improvement in detection accuracy, reaching 34.0 mAP.

Look forward twice. We incorporate the look forward twice scheme [54, 26] to take advantage of the refined box information from previous Transformer decoder layers,

| MTE | FL | IR | TS | LFT | MQS | HM | AP |
|-----|----|----|----|-----|-----|----|------|
| ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | 22.5 |
| ✓ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | 23.1 |
| ✓ | ✓ | ✗ | ✗ | ✗ | ✗ | ✗ | 31.6 |
| ✓ | ✓ | ✓ | ✗ | ✗ | ✗ | ✗ | 33.1 |
| ✓ | ✓ | ✓ | ✓ | ✗ | ✗ | ✗ | 34.0 |
| ✓ | ✓ | ✓ | ✓ | ✓ | ✗ | ✗ | 34.8 |
| ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✗ | 35.2 |
| ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | 37.2 |

Table 1: Preliminary ablation results on the effect of each factor that is used to modernize plain DETR. MTE: merging the Transformer encoder. FL: classification loss as a focal loss. IR: Iterative refinement. TS: two-stage. LFT: look forward twice. MQS: mixed query selection. HM: hybrid matching.

thereby more effectively optimizing the parameters across adjacent Transformer decoder layers. This modification yields +0.8 mAP improvements.

Mixed query selection. This method [54] combines the static content queries with image-adaptive position queries to form better query representations. It yields +0.4 mAP improvements.

Hybrid matching. The original one-to-one set matching is less efficacy in training positive samples. There have been several methods to improve the efficacy through an auxiliary one-to-many set matching loss [26, 6, 27]. We opted for the hybrid matching approach [26], as it preserves the advantage of not requiring additional manual labeling noise or assignment designs. This modification resulted in a +2.0 mAP improvement in detection accuracy, achieving a final 37.2 mAP.

3. Box-to-Pixel Relative Position Bias

In this section, we introduce a simple technology, box-to-pixel relative position bias (BoxRPB), that proves critical to compensate for the lack of multi-scale features and the explicit local cross-attention calculations.

The original DETR decoder adopts a standard cross-attention computation:

$$\mathbf{O} = \text{Softmax}(\mathbf{QK}^T)\mathbf{V} + \mathbf{X}, \quad (1)$$

where X and O are the input and output features of each object query, respectively; Q , K and V are query, key, and value features, respectively.

As will be shown in Figure 4, the original cross-attention formulation often attends to irrelevant image areas within a plain DETR framework. We conjecture that this may be a reason for its much lower accuracy than that with multi-scale and explicit locality designs. Inspired by the success of pixel-to-pixel relative position bias for vision Transformer architectures [34, 33], we explore the use of box-

to-pixel relative position bias (BoxRPB) for cross-attention calculation:

$$\mathbf{O} = \text{Softmax}(\mathbf{Q}\mathbf{K}^T + \mathbf{B})\mathbf{V} + \mathbf{X}, \quad (2)$$

where \mathbf{B} is the relative position bias determined by the geometric relationship between boxes and pixels.

Different from the original relative position bias (RPB) which is defined on $2d$ - relative positions, the BoxRPB needs to handle a larger geometric space of $4d$. In the following, we introduce two implementation variants.

A Naive BoxRPB implementation. We adapt the continuous RPB method [33] to compute the $4d$ - box-to-pixel relative position bias. The original continuous RPB method [33] produces the bias term for each relative position configuration by a meta-network applied on the corresponding $2d$ - relative coordinates. When extending this method for BoxRPB, we use the top-left and bottom-right corners to represent a box and use the relative positions between these corner points and the image pixel point as input to the meta-network. Denote the relative coordinates as $(\Delta\mathbf{x}_1, \Delta\mathbf{y}_1) \in \mathbb{R}^{K \times H \times W \times 2}$ and $(\Delta\mathbf{x}_2, \Delta\mathbf{y}_2) \in \mathbb{R}^{K \times H \times W \times 2}$, the box-to-pixel relative position bias can be defined as:

$$\mathbf{B} = \text{MLP}(\Delta\mathbf{x}_1, \Delta\mathbf{y}_1, \Delta\mathbf{x}_2, \Delta\mathbf{y}_2), \quad (3)$$

where \mathbf{B} is in a shape of $\mathbb{R}^{K \times W \times H \times M}$, with M denoting the number of attention heads, K denoting the number of predicted bounding boxes, W , H denoting the width and height of the output feature maps; the MLP network consists of two linear layers: Linear \rightarrow ReLU \rightarrow Linear. The input/output shapes of these two linear layers are: $K \times H \times W \times 4 \rightarrow K \times H \times W \times 256$ and $K \times H \times W \times 256 \rightarrow K \times H \times W \times M$, respectively.

Our experiments show that this naive implementation already performs very effectively, as shown in Table 3a. However, it will consume a lot of GPU computation and memory budget and thus is not practical.

A decomposed BoxRPB implementation. Now, we present a more efficient implementation of BoxRPB. Instead of directly computing the bias term for a $4d$ - input, we consider decomposing the bias computation into two terms:

$$\mathbf{B} = \text{unsqueeze}(\mathbf{B}_x, 1) + \text{unsqueeze}(\mathbf{B}_y, 2), \quad (4)$$

where $\mathbf{B}_x \in \mathbb{R}^{K \times W \times M}$ and $\mathbf{B}_y \in \mathbb{R}^{K \times H \times M}$ are the biases regarding x - axis and y - axis, respectively. They are computed as:

$$\mathbf{B}_x = \text{MLP}_1(\Delta\mathbf{x}_1, \Delta\mathbf{x}_2), \quad \mathbf{B}_y = \text{MLP}_2(\Delta\mathbf{y}_1, \Delta\mathbf{y}_2), \quad (5)$$

The overall process of the decomposed BoxRPB implementation is also illustrated in Figure 2. The input/output shapes of the two linear layers within MLP_1 are:

$K \times W \times 2 \rightarrow K \times W \times 256$ and $K \times W \times 256 \rightarrow K \times W \times M$, respectively. Similarly, the input/output shapes for the two linear layers within MLP_2 follow the same pattern.

Through decomposition, both the computation FLOPs and memory consumption are significantly reduced, while the accuracy almost keeps, as shown in Table 3a. This decomposition-based implementation is used default in our experiments.

Figure 4 shows the effect of this additional BoxRPB term for cross-attention computation. In general, the BoxRPB term makes the attention focused more on the objects and box boundaries, while the cross-attention without the BoxRPB may attend to many irrelevant areas. This may explain the significantly improved accuracy (+8.9 mAP) by the BoxRPB term, as shown in Table 2.

4. More Improvements

In this section, we introduce two other technologies that can additionally improve the plain DETR framework.

MIM pre-training. We leverage the recent advances of masked image modeling pre-training[1, 20, 51, 28] which have shown better locality[49]. Specifically, we initialize the Swin transformer backbone with SimMIM pre-trained weights that are learned on ImageNet without labels as in[51].

As shown in Table 2, the MIM pre-training brings +7.4 mAP improvements over our plain DETR baseline. The profound gains of MIM pre-training on the plain DETR framework than on other detectors may highlight the importance of the learned localization ability for a plain DETR framework. On a higher baseline where BoxRPB has been involved, the MIM pre-training can still yield +2.6 mAP gains, reaching 48.7 mAP. Moreover, we note that MIM pre-training is also crucial for enabling us abandon the multi-scale backbone features with almost no loss of accuracy, as shown by Table 5b and 5c.

Bounding box regression with re-parameterization. Another improvement we would like to highlight is the bounding box re-parameterization when performing bounding box regression.

The original DETR framework [4] and most of its variants directly scale the box centers and sizes to $[0,1]$. It will face difficulty in detecting small objects due to the large objects dominating the loss computation. Instead, we re-parameterize the box centers and sizes of l -th decoder layer as:

$$\begin{aligned} t_x^l &= (g_x - p_x^{l-1})/p_w^{l-1}, \\ t_y^l &= (g_y - p_y^{l-1})/p_h^{l-1}, \\ t_w^l &= \log(g_w/p_w^{l-1}), \\ t_h^l &= \log(g_h/p_h^{l-1}) \end{aligned} \quad (6)$$

where $p_x^{l-1}/p_y^{l-1}/p_w^{l-1}/p_h^{l-1}$ are the predicted unnormalized




Figure 2: **Illustrating the details of the proposed BoxRPB scheme.** (Left): The black grid represents an input image. The blue sketch region represents a predicted bounding box. We mark the top-left and right-down corners of the box with red stars. (Middle): Our BoxRPB calculates the offsets between all positions and the two corners along both x -axis and y -axis. Then, we concatenate the offset vectors along each axis to form $(\Delta x_1, \Delta x_2)$ and $(\Delta y_1, \Delta y_2)$ and apply an independent MLP to obtain the relative position bias terms B_x and B_y . (Right): We broadcast and add B_x to B_y to get the 2D relative bias term B . We color the positions with higher attention values with red color and blue color otherwise.

| BoxRPB | MIM | reparam. | AP | AP ₅₀ | AP ₇₅ | AP _S | AP _M | AP _L |
|--------|-----|----------|-------------|------------------|------------------|-----------------|-----------------|-----------------|
| ✗ | ✗ | ✗ | 37.2 | 63.7 | 37.6 | 17.8 | 40.5 | 55.6 |
| ✓ | ✗ | ✗ | 46.1 | 67.6 | 49.1 | 27.2 | 50.5 | 64.9 |
| ✗ | ✓ | ✗ | 44.6 | 67.0 | 48.3 | 26.9 | 49.1 | 59.1 |
| ✗ | ✓ | ✓ | 46.3 | 68.2 | 51.1 | 30.7 | 51.0 | 58.4 |
| ✓ | ✓ | ✗ | 48.7 | 67.7 | 53.0 | 31.3 | 53.1 | 63.0 |
| ✓ | ✓ | ✓ | 50.9 | 69.3 | 55.5 | 34.2 | 55.1 | 65.5 |

Table 2: **Core ablation results of the proposed components.** Equipped with these components, a plain DETR could achieve competitive performance.

box positions and sizes of $(l-1)$ -th decoder layer.

Table 2 shows that this modification can enhance the overall detection performance by +2.2 AP. Especially, it achieves a larger +2.9 AP improvements on small objects.

5. Ablation Study and Analysis

5.1. The importance of box relative position bias

In Table 3, we study the effect of each factor within our BoxRPB scheme and report the detailed comparison results in the following discussion.

Effect of axial decomposition. Modeling the 2D relative position without any decomposition is a naive baseline compared with our axial decomposition schema, and it can be parameterized as $(\Delta x_1, \Delta y_1, \Delta x_2, \Delta y_2) \in \mathbb{R}^{k \times H \times W \times 4}$. This baseline requires a quadratic computation overhead and memory consumption while the decomposed one decreases the cost to linear complexity. In Table 3a, we compared the two approaches and find that the axial decomposition scheme achieves comparable performance (50.9 vs. 50.8) while it requires a much lower memory footprint (9.5G vs. 26.8G) and smaller computation overhead (5.8G FLOPs vs. 265.4G FLOPs).

Effect of box points. Table 3b shows the comparison of using only the center points or the two corner points. We

find that applying only the center points improves the baseline (fourth row of Table 2) by +1.7 AP. However, its performance is worse than that of using two corner points. In particular, while the two methods achieve comparable AP₅₀ results, utilizing corner points could boost AP₇₅ by +2.2. This shows that not only the position (center) but also the scale (height and width) of the query box are important to precisely model relative position bias.

Effect of hidden dimension. We study the effect of the hidden dimension in Equation 5. As shown in Table 3c, a smaller hidden dimension of 128 would lead to a performance drop of 0.5, indicating that the position relation is non-trivial and requires a higher dimension space to model.

Comparison with other methods. We study the effect of choosing other schemes to compute the modulation term B in Equation 2. We compared to several representative methods as follows: (i) Conditional cross-attention scheme [35], which computes the modulation term based on the inner product between the conditional spatial (position) query embedding and the spatial key embedding. (ii) DAB cross-attention scheme [31], which builds on conditional cross-attention and further modulates the positional attention map using the box width and height information. (iii) Spatially modulated cross-attention scheme (SMCA) [16], which designs handcrafted query spatial priors, implemented with a 2D Gaussian-like weight map, to constrain the attended features to be around the object queries' initial estimations.

Table 3d reports the detailed comparison results. Our approach achieves the best performance among all the methods. Specifically, the conditional cross-attention module achieves similar performance with our center-only setting (first row of Table 3b). DAB cross-attention and SMCA are slightly better than the conditional cross-attention module, but they still lag behind the BoxRPB by a gap of 2.5 AP and 2.2 AP, respectively.

We also compare BoxRPB with DAB cross-attention based on its official open-source code. Replacing DAB positional module with BoxRPB achieves a +1.8 mAP perfor-

| decomp. | mem. | GFLOPs | AP | | | box points | AP | | | hidden dim. | method | | | | | | |
|--------------|-------|--------|-------------|------------------|------------------|--------------------|-------------|------------------|------------------|-------------|----------------------|-------------|------------------|-------------------------|-------------|-------------|-------------|
| | | | AP | AP ₅₀ | AP ₇₅ | | AP | AP ₅₀ | AP ₇₅ | | standard cross attn. | AP | AP ₅₀ | AP ₇₅ | | | |
| \times | 26.8G | 265.4 | 50.8 | 69.3 | 55.4 | center | 48.0 | 69.0 | 53.3 | 128 | 50.4 | 69.1 | 55.1 | conditional cross attn. | 46.3 | 68.2 | 51.1 |
| \checkmark | 9.5G | 5.8 | 50.9 | 69.3 | 55.5 | 2 \times corners | 50.9 | 69.3 | 55.5 | 256 | 50.9 | 69.4 | 55.4 | DAB cross attn. | 48.4 | 68.9 | 53.4 |
| | | | | | | | | | | 512 | 50.9 | 69.3 | 55.5 | SMCA cross attn. | 48.7 | 69.2 | 53.6 |
| | | | | | | | | | | | | | | ours | 50.9 | 69.3 | 55.5 |

(a) axial decomposition.

(b) box points.

(c) hidden dim.

(d) cross-attention modulation.

Table 3: **Ablation of box relative position bias scheme.** (a) Axial decomposition can significantly decrease the computation overhead and GPU memory footprint. (b) The corner points perform better than the center point. (c) The higher the hidden dimension, the better performance. (d) Our approach performs much better than other related methods designed to modulate the cross-attention maps.

| method | AP | AP ₅₀ | AP ₇₅ | AP _S | AP _M | AP _L |
|------------------------|-------------|------------------|------------------|-----------------|-----------------|-----------------|
| deformable cross attn. | 50.2 | 68.5 | 54.8 | 34.1 | 54.4 | 63.3 |
| RoIAlign | 49.6 | 68.3 | 54.1 | 31.9 | 54.2 | 63.5 |
| RoI Sampling | 49.3 | 68.2 | 53.8 | 33.1 | 53.2 | 63.0 |
| Box Mask | 48.6 | 68.7 | 52.9 | 31.8 | 52.7 | 63.0 |
| Ours | 50.9 | 69.3 | 55.5 | 34.2 | 55.1 | 65.5 |

Table 4: **Comparison with local cross-attention scheme.** Global cross-attention with BoxRPB outperforms all the local cross-attention counterparts and have a significant gain on large objects.

mance gain.

5.2. Comparison with local attention scheme

In this section, we compared our global attention schema with other representative local cross-attention mechanisms, including deformable cross-attention [55], RoIAlign [21], RoI Sampling (sampling fixed points inside the Region of Interest), and box mask inspired by [7]. We illustrate the key differences between those methods in the supplementary material.

As shown in Table 4, our method surpasses all the local cross-attention variants. In addition, we observed that large objects have larger improvements for our method. A similar observation is also reported in DETR [4], it may be due to more effective long-range context modeling based on the global attention scheme.

5.3. On MIM pre-training

We explore different ways of using the backbone and decoder feature maps with or without MIM pre-training. We evaluate the performance of three different architecture configurations, which are illustrated in Figure 3. We discuss and analyze the results as follows.

MIM pre-training brings consistent gains. By comparing the experimental results under the same architecture configuration, we found that using MIM pre-training consistently achieves better performance. For example, as shown in Table 5, using MIM pre-training outperforms using supervised

| backbone \rightarrow decoder | MIM | AP | AP ₅₀ | AP ₇₅ | AP _S | AP _M | AP _L |
|-----------------------------------------------|--------------|-------------|------------------|------------------|-----------------|-----------------|-----------------|
| $(C^3, C^4, C^5) \rightarrow (P^3, P^4, P^5)$ | \times | 49.6 | 69.2 | 53.8 | 31.5 | 53.4 | 65.2 |
| $(C^3, C^4, C^5) \rightarrow (P^3, P^4, P^5)$ | \checkmark | 51.1 | 69.3 | 56.0 | 34.8 | 55.4 | 65.2 |

(a)

| backbone \rightarrow decoder | MIM | AP | AP ₅₀ | AP ₇₅ | AP _S | AP _M | AP _L |
|-----------------------------------|--------------|-------------|------------------|------------------|-----------------|-----------------|-----------------|
| $(C^3, C^4, C^5) \rightarrow P^5$ | \times | 47.0 | 68.2 | 50.4 | 28.0 | 51.5 | 64.2 |
| $(C^3, C^4, C^5) \rightarrow P^4$ | \times | 49.6 | 69.8 | 53.4 | 31.4 | 53.7 | 65.5 |
| $(C^3, C^4, C^5) \rightarrow P^3$ | \times | 49.7 | 69.8 | 53.9 | 32.7 | 53.5 | 65.2 |
| $(C^3, C^4, C^5) \rightarrow P^5$ | \checkmark | 50.3 | 69.3 | 54.9 | 33.4 | 54.7 | 64.9 |
| $(C^3, C^4, C^5) \rightarrow P^4$ | \checkmark | 51.0 | 69.4 | 55.7 | 34.5 | 55.1 | 65.2 |
| $(C^3, C^4, C^5) \rightarrow P^3$ | \checkmark | 50.9 | 69.2 | 55.4 | 34.4 | 55.0 | 64.5 |

(b)

| backbone \rightarrow decoder | MIM | AP | AP ₅₀ | AP ₇₅ | AP _S | AP _M | AP _L |
|--------------------------------|--------------|-------------|------------------|------------------|-----------------|-----------------|-----------------|
| $C^5 \rightarrow P^5$ | \times | 46.4 | 67.7 | 49.7 | 26.9 | 50.5 | 64.4 |
| $C^5 \rightarrow P^4$ | \times | 48.0 | 68.7 | 51.8 | 30.4 | 52.2 | 64.4 |
| $C^5 \rightarrow P^3$ | \times | 48.7 | 69.1 | 52.6 | 30.7 | 52.9 | 64.9 |
| $C^5 \rightarrow P^5$ | \checkmark | 50.2 | 69.1 | 55.0 | 33.5 | 54.5 | 64.6 |
| $C^5 \rightarrow P^4$ | \checkmark | 50.9 | 69.3 | 55.5 | 34.2 | 55.1 | 65.5 |
| $C^5 \rightarrow P^3$ | \checkmark | 50.9 | 69.2 | 55.7 | 34.6 | 54.9 | 65.0 |

(c)

Table 5: **Ablation of MIM pre-training.** (a) multi-scale feature maps output by the backbone + multi-scale feature maps for the Transformer decoder. (b) multi-scale feature maps output by the backbone + single-scale feature map for the Transformer decoder. (c) single-scale feature map output by the backbone + single-scale feature map for the Transformer decoder.

pre-training by 1.5 AP in the $(C^3, C^4, C^5) \rightarrow (P^3, P^4, P^5)$ configuration and 2.9 AP in the $C^5 \rightarrow P^4$ configuration.

Multi-scale feature maps for the decoder can be removed. By comparing the results between Table 5a and Table 5b, we found that using high-resolution feature maps can match or even surpass the performance of using multi-scale feature maps. For example, $(C^3, C^4, C^5) \rightarrow P^3$ achieves comparable performance with $(C^3, C^4, C^5) \rightarrow (P^3, P^4, P^5)$ with or without using MIM pre-training. This observation is not trivial as most existing detection heads still require multi-scale features as input, and it makes building a competitive single-scale plain DETR possible. We hope this finding could ease the design of future de-




Figure 3: We compare the architecture designs when using different feature maps output by the backbone and sent to the Transformer decoder. From (a) to (b), we simplify the dependency on sending multi-scale feature maps to the Transformer decoder. From (b) to (c), we remove the dependency on fusing multi-scale feature output by the backbone. We adopt (c) as our default architecture setting.

| method | AP | AP ₅₀ | AP ₇₅ | AP _S | AP _M | AP _L |
|-----------------------|-------------|------------------|------------------|-----------------|-----------------|-----------------|
| Cascade Mask R-CNN[3] | 53.7 | 71.9 | 58.7 | 36.9 | 57.4 | 69.1 |
| Ours | 53.8 | 73.4 | 58.9 | 35.9 | 57.0 | 68.9 |

Table 6: Comparison of the improved plain DETR and Cascade Mask R-CNN with a MIM pre-trained ViT-Base backbone. Our plain DETR with global cross-attention is slightly better than the region-based, multi-scaled Cascade Mask R-CNN.

tection frameworks.

Multi-scale feature maps from the backbone are non-necessary. We analyze the effect of removing the multi-scale feature maps from the backbone by comparing the results of Table 5b and Table 5c. When using a supervised pre-trained backbone, adopting only the last feature map C^5 from the backbone would hurt the performance. For example, when using the supervised pre-trained backbone, the $C^5 \rightarrow P^5$ reaches 46.4 AP, which is worse than $(C^3, C^4, C^5) \rightarrow P^5$ (47.0 AP) by 0.6 AP. However, when using the MIM pre-trained backbone, $C^5 \rightarrow P^5$ reaches 50.2 mAP, which is comparable with the performance of $(C^3, C^4, C^5) \rightarrow P^5$ (50.3 AP). These results show that MIM pre-training can reduce the reliance on multi-scale feature maps.

Single-scale feature map from the backbone and single-scale feature map for the decoder is enough. Based on the above observations, we can reach a surprisingly simple but important conclusion that we can completely eliminate the need for multi-scale feature maps in both the backbone and Transformer decoder by using our proposed BoxRBP scheme and MIM pre-training.

5.4. Application to a plain ViT

In this section, we build a simple and effective fully plain object detection system by applying our approach to the plain ViT [13]. Our system only uses a single-resolution feature map throughout a plain Transformer encoder-decoder architecture, without any multi-scale design or processing. We compare our approach with the state-of-the-art Cascade Mask R-CNN [3, 28] on the COCO

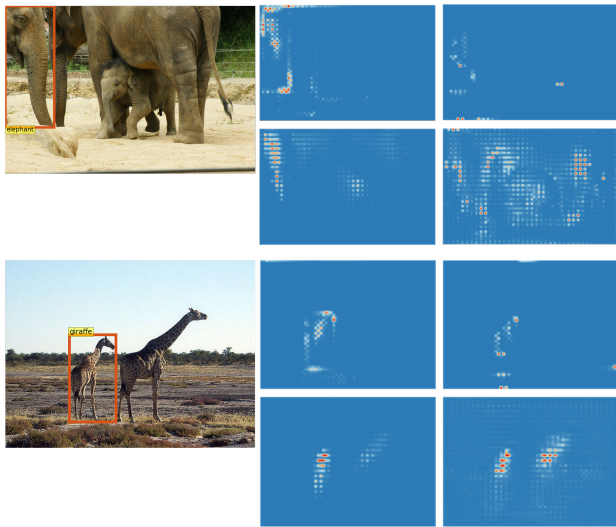


Figure 4: Visualizations of the cross-attention maps of models w. or w/o. BoxRBP. For each group, the first column shows the input image and the object query. The first row presents the attention maps of the model w. BoxRBP, while the second row displays attention maps of the model w/o. BoxRBP. BoxRBP helps to guide the cross-attention to focus on the individual objects.

dataset. For the fair comparison, We use a MAE [20] pre-trained ViT-Base as the backbone and train the object detector for ~ 50 epochs. As shown in Table 6, our method achieves comparable results with Cascade Mask R-CNN which relies on using multi-scale feature maps for better localization across different object scales. Remarkably, our method does not train with instance mask annotations that are usually considered to be beneficial for object detection.

5.5. Visualization of cross-attention maps

Figure 4 shows the cross-attention maps of models with or without BoxRBP. For the model with BoxRBP, the cross-attention concentrate on the individual object. In the contrary, the cross-attention of model without BoxRBP attend to multiple objects that have similar appearance.

| method | framework | extra data | #params | #epoch | AP | AP ₅₀ | AP ₇₅ | AP _S | AP _M | AP _L |
|-----------------|-----------|------------|---------|---------|-------------|------------------|------------------|-----------------|-----------------|-----------------|
| Swin [34] | HTC | | 284M | 72 | 57.7 | 76.2 | 63.1 | 33.4 | 52.9 | 64.0 |
| DETA [36] | DETR | | 218M | 24 | 58.5 | 76.5 | 64.4 | 38.5 | 62.6 | 73.8 |
| DINO-DETR [54] | DETR | | 218M | 36 | 58.6 | 76.9 | 64.1 | 39.4 | 61.6 | 73.2 |
| Ours* | DETR | | 228M | 36 | 60.0 | 78.9 | 66.4 | 42.8 | 62.7 | 73.7 |
| DETA [36] | DETR | O365 | 218M | 24 + 24 | 63.5 | 80.4 | 70.2 | 46.1 | 66.9 | 76.9 |
| DINO-DETR [54]* | DETR | O365 | 218M | 26 + 18 | 63.3 | – | – | – | – | – |
| Ours* | DETR | O365 | 228M | 24 + 24 | 63.9 | 82.1 | 70.7 | 48.2 | 66.8 | 76.7 |

Table 7: System-level comparisons with the state-of-the-art results on COCO test-dev. All methods adopt the Swin-Large backbone. The * marks the results with test time augmentation.

6. System-level Results

We compare our method with other state-of-the-art methods in this section. Table 7 shows results, where all experiments reported in this table utilize a Swin-Large as the backbone. As other works usually apply an encoder to enhance the backbone features, we also stack 12 window-based single-scale transformer layers (with a feature dimension of 256) on top of the backbone for a fair comparison.

With the 36 training epochs, our model achieves 60.0 AP on the COCO test-dev set, which outperforms DINO-DETR by 1.4 AP. Further introducing the Objects365 [40] as the pre-training dataset, our method reaches 63.9 AP on the test-dev set, which is better than DINO-DETR and DETA by a notable margin. These strong results verify that the plain DETR architecture does not have intrinsic drawbacks to prevent it from achieving high performance.

7. Related work

DETR-based object detection. DETR [4] has impressed the field for its several merits, including the conceptually straightforward and generic in applicability, requiring minimal domain knowledge that avoids customized label assignments and non-maximum suppression, and being plain. While the original DETR maintains a plain design, it also suffers from slow convergence rate and lower detection accuracy. There have been many follow-up works including [35, 16, 9, 47, 55, 53, 52, 17, 54], and now many top object detectors have been built upon this line of works, thanks to the reintroduction of multi-scale and locality designs [54, 14, 46]. Unlike these leading works, we aim for an improved DETR framework that maintains a “plain” nature without multi-scale features and local cross-attention computation.

Region-based object detection. Prior to the DETR framework, the object detectors were usually built in a region-based fashion: the algorithms analyze every region of the entire image locally, and the object detections are obtained

by ranking and filtering the results of each region. Due to the locality nature, it’s hard for them to flexibly leverage global information for object detection. Moreover, while some early attempts use single scale feature map on the head [19, 38, 18, 39, 32], later, the leading methods are almost all built by multi-scale features such as FPN [29], BiFPN [42], Cascade R-CNN [3], and HTC [5], etc. We expect our strong plain DETR detector may also inspire research in exploring single-scale feature map for region-based detection.

Position encoding. This paper is also related to position encoding techniques. The original Transformer [45] uses absolute position encoding. Early vision Transformers [4, 12, 44] inherit this absolute position encoding setting. Swin Transformers [34, 33] highlight the importance of relative position bias for Transformer-based visual recognition, where some early variants can be found in both language and vision domains [23, 41, 24, 10, 25, 8, 48]. This paper extends the relative position bias for box-to-pixel pairs, instead of previous pixel-to-pixel pairs. It also reveals that the RPB can effect even more critical in the context of plain DETR detectors.

Pre-training. The pre-training methods [20, 51, 1] that follow the path of masked image modeling have drawn increasing attention due to their strong performance on various core vision tasks such as object detection and semantic segmentation. Although some recent works [28, 49] have revealed some possible reasons why MIM outperforms the conventional supervised pre-training and confirmed that FPN can be simplified, few works attempt to build a fully plain object detection head based on MIM pre-trained backbones. Our experiment results show that MIM pre-training is a key factor in fully plain object detection architecture design.

8. Conclusion

This paper has present an improved plain DETR detector which achieves exceptional improvements over the original plain model, and achieves a 63.9 mAP accuracy using a Swin-L backbone, which is highly competitive with state-of-the-art detectors that have been heavily tuned using multi-scale feature maps and region-based feature extraction. We highlighted the importance of two technologies of BoxRPN and MIM-based pre-training for this improved plain DETR framework. We hope the effective detector empowered by minimal architectural “inductive bias” can encourage future research to explore generic plain decoders in other vision problems.

References

- [1] H. Bao, L. Dong, S. Piao, and F. Wei. Beit: Bert pre-training of image transformers. *arXiv preprint arXiv:2106.08254*, 2021. 4, 8
- [2] T. Brown, B. Mann, N. Ryder, M. Subbiah, J. D. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell, et al. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020. 1
- [3] Z. Cai and N. Vasconcelos. Cascade r-cnn: Delving into high quality object detection. In *CVPR*, pages 6154–6162, 2018. 3, 7, 8
- [4] N. Carion, F. Massa, G. Synnaeve, N. Usunier, A. Kirillov, and S. Zagoruyko. End-to-end object detection with transformers. In *European conference on computer vision*, pages 213–229. Springer, 2020. 1, 2, 4, 6, 8
- [5] K. Chen, J. Pang, J. Wang, Y. Xiong, X. Li, S. Sun, W. Feng, Z. Liu, J. Shi, W. Ouyang, et al. Hybrid task cascade for instance segmentation. In *CVPR*, pages 4974–4983, 2019. 8
- [6] Q. Chen, J. Wang, C. Han, S. Zhang, Z. Li, X. Chen, J. Chen, X. Wang, S. Han, G. Zhang, et al. Group detr v2: Strong object detector with encoder-decoder pretraining. *arXiv preprint arXiv:2211.03594*, 2022. 3
- [7] B. Cheng, I. Misra, A. G. Schwing, A. Kirillov, and R. Girshick. Masked-attention mask transformer for universal image segmentation. *arXiv preprint arXiv:2112.01527*, 2021. 6
- [8] X. Chu, B. Zhang, Z. Tian, X. Wei, and H. Xia. Do we really need explicit position encodings for vision transformers. *arXiv preprint arXiv:2102.10882*, 3(8), 2021. 8
- [9] X. Dai, Y. Chen, J. Yang, P. Zhang, L. Yuan, and L. Zhang. Dynamic detr: End-to-end object detection with dynamic attention. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2988–2997, 2021. 8
- [10] Z. Dai, Z. Yang, Y. Yang, J. Carbonell, Q. V. Le, and R. Salakhutdinov. Transformer-xl: Attentive language models beyond a fixed-length context. *arXiv preprint arXiv:1901.02860*, 2019. 8
- [11] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018. 1
- [12] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020. 8
- [13] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly, J. Uszkoreit, and N. Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. In *ICLR*, 2021. 1, 7
- [14] Y. Fang, W. Wang, B. Xie, Q. Sun, L. Wu, X. Wang, T. Huang, X. Wang, and Y. Cao. Eva: Exploring the limits of masked visual representation learning at scale. *arXiv preprint arXiv:2211.07636*, 2022. 8
- [15] Y. Fang, S. Yang, S. Wang, Y. Ge, Y. Shan, and X. Wang. Unleashing vanilla vision transformer with masked image modeling for object detection. *arXiv preprint arXiv:2204.02964*, 2022. 2
- [16] P. Gao, M. Zheng, X. Wang, J. Dai, and H. Li. Fast convergence of detr with spatially modulated co-attention. In *ICCV*, pages 3621–3630, 2021. 5, 8
- [17] Z. Gao, L. Wang, B. Han, and S. Guo. Adamixer: A fast-converging query-based object detector. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5364–5373, 2022. 8
- [18] R. Girshick. Fast r-cnn. In *ICCV*, pages 1440–1448, 2015. 8
- [19] R. Girshick, J. Donahue, T. Darrell, and J. Malik. Rich feature hierarchies for accurate object detection and semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 580–587, 2014. 8
- [20] K. He, X. Chen, S. Xie, Y. Li, P. Dollár, and R. Girshick. Masked autoencoders are scalable vision learners. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16000–16009, 2022. 2, 4, 7, 8
- [21] K. He, G. Gkioxari, P. Dollár, and R. Girshick. Mask R-CNN. In *ICCV*, 2017. 6
- [22] K. He, X. Zhang, S. Ren, and J. Sun. Deep Residual Learning for Image Recognition. In *CVPR*, 2016. 2
- [23] H. Hu, J. Gu, Z. Zhang, J. Dai, and Y. Wei. Relation networks for object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3588–3597, 2018. 8
- [24] H. Hu, Z. Zhang, Z. Xie, and S. Lin. Local relation networks for image recognition. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3464–3473, 2019. 8
- [25] Z. Huang, D. Liang, P. Xu, and B. Xiang. Improve transformer models with better relative position embeddings. *arXiv preprint arXiv:2009.13658*, 2020. 8
- [26] D. Jia, Y. Yuan, H. He, X. Wu, H. Yu, W. Lin, L. Sun, C. Zhang, and H. Hu. Dets with hybrid matching. *arXiv preprint arXiv:2207.13080*, 2022. 3
- [27] F. Li, H. Zhang, S. Liu, J. Guo, L. M. Ni, and L. Zhang. Dn-detr: Accelerate detr training by introducing query denoising. *arXiv preprint arXiv:2203.01305*, 2022. 3

- [28] Y. Li, H. Mao, R. Girshick, and K. He. Exploring plain vision transformer backbones for object detection. *arXiv preprint arXiv:2203.16527*, 2022. 2, 4, 7, 8
- [29] T.-Y. Lin, P. Dollár, R. Girshick, K. He, B. Hariharan, and S. Belongie. Feature pyramid networks for object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2117–2125, 2017. 8
- [30] T.-Y. Lin, P. Goyal, R. Girshick, K. He, and P. Dollár. Focal loss for dense object detection. In *Proceedings of the IEEE international conference on computer vision*, pages 2980–2988, 2017. 3
- [31] S. Liu, F. Li, H. Zhang, X. Yang, X. Qi, H. Su, J. Zhu, and L. Zhang. Dab-detr: Dynamic anchor boxes are better queries for detr. *arXiv preprint arXiv:2201.12329*, 2022. 5
- [32] W. Liu, D. Anguelov, D. Erhan, C. Szegedy, S. Reed, C.-Y. Fu, and A. C. Berg. Ssd: Single shot multibox detector. In *ECCV*, pages 21–37. Springer, 2016. 8
- [33] Z. Liu, H. Hu, Y. Lin, Z. Yao, Z. Xie, Y. Wei, J. Ning, Y. Cao, Z. Zhang, L. Dong, et al. Swin transformer v2: Scaling up capacity and resolution. In *CVPR*, pages 12009–12019, 2022. 2, 3, 4, 8
- [34] Z. Liu, Y. Lin, Y. Cao, H. Hu, Y. Wei, Z. Zhang, S. Lin, and B. Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *ICCV*, pages 10012–10022, 2021. 2, 3, 8
- [35] D. Meng, X. Chen, Z. Fan, G. Zeng, H. Li, Y. Yuan, L. Sun, and J. Wang. Conditional detr for fast training convergence. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2021. 5, 8
- [36] J. Ouyang-Zhang, J. H. Cho, X. Zhou, and P. Krähenbühl. Nms strikes back. *arXiv preprint arXiv:2212.06137*, 2022. 8
- [37] A. Radford, K. Narasimhan, T. Salimans, I. Sutskever, et al. Improving language understanding by generative pre-training. 2018. 1
- [38] J. Redmon, S. Divvala, R. Girshick, and A. Farhadi. You only look once: Unified, real-time object detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016. 8
- [39] S. Ren, K. He, R. Girshick, and J. Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. *Advances in neural information processing systems*, 28, 2015. 8
- [40] S. Shao, Z. Li, T. Zhang, C. Peng, G. Yu, X. Zhang, J. Li, and J. Sun. Objects365: A large-scale, high-quality dataset for object detection. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 8430–8439, 2019. 8
- [41] P. Shaw, J. Uszkoreit, and A. Vaswani. Self-attention with relative position representations. *arXiv preprint arXiv:1803.02155*, 2018. 8
- [42] M. Tan, R. Pang, and Q. V. Le. Efficientdet: Scalable and efficient object detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10781–10790, 2020. 8
- [43] Z. Teed and J. Deng. Raft: Recurrent all-pairs field transforms for optical flow. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part II 16*, pages 402–419. Springer, 2020. 3
- [44] H. Touvron, M. Cord, M. Douze, F. Massa, A. Sablayrolles, and H. Jégou. Training data-efficient image transformers & distillation through attention. In *International conference on machine learning*, pages 10347–10357. PMLR, 2021. 8
- [45] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017. 8
- [46] W. Wang, J. Dai, Z. Chen, Z. Huang, Z. Li, X. Zhu, X. Hu, T. Lu, L. Lu, H. Li, et al. Internimage: Exploring large-scale vision foundation models with deformable convolutions. *arXiv preprint arXiv:2211.05778*, 2022. 8
- [47] Y. Wang, X. Zhang, T. Yang, and J. Sun. Anchor detr: Query design for transformer-based detector, 2021. 8
- [48] K. Wu, H. Peng, M. Chen, J. Fu, and H. Chao. Rethinking and improving relative position encoding for vision transformer. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10033–10041, 2021. 8
- [49] Z. Xie, Z. Geng, J. Hu, Z. Zhang, H. Hu, and Y. Cao. Revealing the dark secrets of masked image modeling. *arXiv preprint arXiv:2205.13543*, 2022. 2, 4, 8
- [50] Z. Xie, Z. Zhang, Y. Cao, Y. Lin, J. Bao, Z. Yao, Q. Dai, and H. Hu. Simmim: A simple framework for masked image modeling. *arXiv preprint arXiv:2111.09886*, 2021. 2
- [51] Z. Xie, Z. Zhang, Y. Cao, Y. Lin, J. Bao, Z. Yao, Q. Dai, and H. Hu. Simmim: A simple framework for masked image modeling. In *CVPR*, pages 9653–9663, 2022. 4, 8
- [52] G. Zhang, Z. Luo, Y. Yu, J. Huang, K. Cui, S. Lu, and E. P. Xing. Semantic-aligned matching for enhanced detr convergence and multi-scale feature fusion. *arXiv preprint arXiv:2207.14172*, 2022. 8
- [53] G. Zhang, Z. Luo, Y. Yu, Z. Tian, J. Zhang, and S. Lu. Towards efficient use of multi-scale features in transformer-based object detectors. *arXiv preprint arXiv:2208.11356*, 2022. 8
- [54] H. Zhang, F. Li, S. Liu, L. Zhang, H. Su, J. Zhu, L. M. Ni, and H.-Y. Shum. Dino: Detr with improved denoising anchor boxes for end-to-end object detection. *arXiv preprint arXiv:2203.03605*, 2022. 1, 2, 3, 8
- [55] X. Zhu, W. Su, L. Lu, B. Li, X. Wang, and J. Dai. Deformable detr: Deformable transformers for end-to-end object detection. *arXiv preprint arXiv:2010.04159*, 2020. 1, 3, 6, 8