

MBZUAI

Digital.Commons@MBZUAI

Machine Learning Faculty Publications

Scholarly Works

7-28-2022

Semantic-Aligned Matching for Enhanced DETR Convergence and Multi-Scale Feature Fusion

Gongjie Zhang

Zhipeng Luo

Yingchen Yu

Jiaxing Huang

Kaiwen Cui

See next page for additional authors

Follow this and additional works at: <https://dclibrary.mbzuai.ac.ae/mlfp>



Part of the [Artificial Intelligence and Robotics Commons](#)

Preprint: arXiv

Archived with thanks to arXiv

Preprint License: CC by NC ND 4.0

Uploaded 24 August 2022

Authors

Gongjie Zhang, Zhipeng Luo, Yingchen Yu, Jiaying Huang, Kaiwen Cui, Shijian Lu, and Eric Xing

Semantic-Aligned Matching for Enhanced DETR Convergence and Multi-Scale Feature Fusion

Gongjie Zhang[†], Zhipeng Luo[†], Yingchen Yu, Jiaxing Huang, Kaiwen Cui, Shijian Lu[✉], and Eric P. Xing

Abstract—The recently proposed DEtection TRansformer (DETR) has established a fully end-to-end paradigm for object detection. However, DETR suffers from slow training convergence, which hinders its applicability to various detection tasks. We observe that DETR’s slow convergence is largely attributed to the difficulty in matching object queries to relevant regions due to the unaligned semantics between object queries and encoded image features. With this observation, we design *Semantic-Aligned-Matching DETR++* (SAM-DETR++) to accelerate DETR’s convergence and improve detection performance. The core of SAM-DETR++ is a plug-and-play module that projects object queries and encoded image features into the same feature embedding space, where each object query can be easily matched to relevant regions with similar semantics. Besides, SAM-DETR++ searches for multiple representative keypoints and exploits their features for semantic-aligned matching with enhanced representation capacity. Furthermore, SAM-DETR++ can effectively fuse multi-scale features in a coarse-to-fine manner on the basis of the designed semantic-aligned matching. Extensive experiments show that the proposed SAM-DETR++ achieves superior convergence speed and competitive detection accuracy. Additionally, as a plug-and-play method, SAM-DETR++ can complement existing DETR convergence solutions with even better performance, achieving 44.8% AP with merely 12 training epochs and 49.1% AP with 50 training epochs on COCO val2017 with ResNet-50. Codes are available at <https://github.com/ZhangGongjie/SAM-DETR>.

Index Terms—Computer Vision, Object Detection, Vision Transformer, DETR, Model Convergence, Multi-Scale Representation.

1 INTRODUCTION

OBJECT detection [1] is a fundamental computer vision task and has experienced remarkable progress with the recent development of deep learning and convolutional neural networks (ConvNets). However, most modern ConvNet-based object detectors (e.g., Faster R-CNN [2], YOLO [3], FCOS [4]) still heavily rely on a series of hand-crafted components, such as anchors, non-maximum suppression (NMS), rule-based training target assignment, etc., which lead to complex detection pipelines and sub-optimal performance. Recently, the emergence of DEtection TRansformer (DETR) [5] has revolutionized the paradigm for object detection. DETR adopts a simple Transformer encoder-decoder pipeline [6] and removes the need for those hand-crafted components, achieving a fully end-to-end framework for object detection. However, despite its simplicity and promising performance, DETR suffers from severely slow convergence on training, requiring 500 epochs to fully converge on the COCO dataset [7], while most other ConvNet-based object detectors [2, 3, 8, 4] only requires 12~36 training epochs instead. DETR’s slow convergence significantly increases its training cost and thus hinders the wide application of DETR.

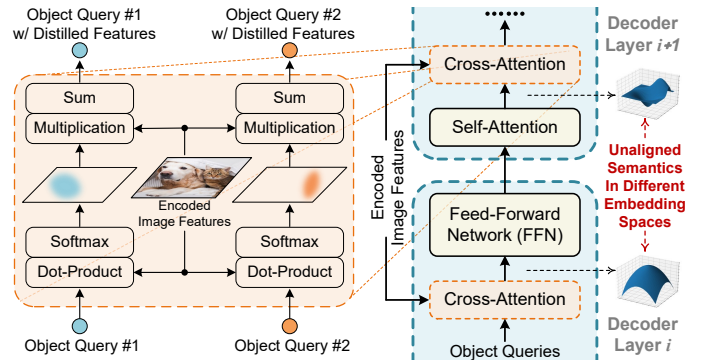


Fig. 1. The analysis for the root of DETR’s slow convergence. **Left:** The cross-attention module in DETR’s decoder layers can be interpreted as a ‘matching and feature distillation’ process. Each object query first matches its particular relevant regions in encoded image features via ‘Dot-Product and Softmax’, and then distills instance-level features from the matched regions for subsequent prediction. **Right:** However, modules between cross-attentions may project object queries and encoded image features into different feature embedding spaces, leading to the unaligned semantics between them. Such unaligned semantics imposes difficulty in cross-attention’s matching process and thus hinders the convergence of Transformer-based object detection frameworks.

- Gongjie Zhang, Zhipeng Luo, Yingchen Yu, Jiaxing Huang, Kaiwen Cui, and Shijian Lu are with the School of Computer Science and Engineering, Nanyang Technological University, Singapore.
- Eric P. Xing is with the School of Computer Science, Carnegie Mellon University. He also serves as the president of Mohamed bin Zayed University of Artificial Intelligence.
- E-mail: GongjieZhang@ntu.edu.sg, Zhipeng001@e.ntu.edu.sg, Yingchen001@e.ntu.edu.sg, Jiaxing.Huang@ntu.edu.sg, Kaiwen001@e.ntu.edu.sg, Shijian.Lu@ntu.edu.sg, Eric.Xing@mbzuai.ac.ae.
- [†] denotes equal contribution; [✉] denotes corresponding author.

Pre-print version. All rights reserved by the authors.

predictions as well as to repeat the subsequent ‘matching and feature distillation’ processes for refined predictions. However, as pointed out in [9, 10, 11, 12, 13], it is difficult for object queries to learn to match appropriate regions. As illustrated in Fig. 1 (right), we observe that the matching difficulty is largely attributed to the unaligned semantics between object queries and encoded image features. Concretely, the modules between cross-attentions project object queries into different feature embedding spaces, in which object queries have different feature semantics from encoded image features. This leads to the complexity of matching object queries with relevant regions and further DETR’s slow training convergence.

An intuitive and promising direction to mitigate the matching difficulty caused by unaligned semantics has been explored in Siamese-based architectures, which adopt identical sub-networks to produce comparable output feature vectors for similarity computation. The effectiveness of Siamese-based architectures has been extensively verified in various matching-involved vision tasks, including object tracking [14, 15, 16, 17, 18, 19], re-identification [20, 21, 22, 23, 24], and few-shot recognition [25, 26, 27, 28, 29]. In light of the success of Siamese-based architectures in matching-involved tasks, we follow the similar philosophy to address the matching difficulty in the cross-attention module of DETR’s decoder.

With these motivations, we propose *Semantic-Aligned-Matching DETR++* (SAM-DETR++) that accelerates the convergence of DETR via a semantic-aligned matching mechanism. Concretely, SAM-DETR++ appends a plug-and-play module ahead of the cross-attention modules in DETR’s decoder layers, with which object queries and encoded image features can be projected into the same semantics-aligned feature embedding spaces and thus be matched efficiently. The aligned semantics imposes a strong prior for each object query to focus on those semantically similar regions in encoded image features. SAM-DETR++ also explicitly identifies multiple representative keypoints for each object query and exploits their features for semantic-aligned matching, which can naturally fit into the original multi-head attention mechanism [6] for enhanced representation capacity. In addition, we extend the semantic-aligned matching mechanism to incorporate multi-scale features that are inherently unaligned in feature semantics. This enables SAM-DETR++ to represent objects at different scales in a ‘divide and conquer’ manner and significantly alleviates the representation complexity, yielding faster convergence and improved accuracy. Further, as SAM-DETR++ works as a plug-in to the original DETR [5] with little modification to the rest operations, it can be easily integrated with existing DETR convergence strategies [11, 30] in a complementary manner, boosting detection performance to a greater extent.

The contributions of this work are summarized below. (i) We propose SAM-DETR++, which accelerates DETR’s convergence with a plug-and-play module that enables semantic-aligned matching between object queries and encoded image features. (ii) We propose to explicitly search for objects’ representative keypoints and leverage their features for semantic-aligned matching, which further strengthens the representation capacity of the introduced semantic-aligned matching mechanism. (iii) We introduce

a multi-scale design into the semantic-aligned matching mechanism to effectively fuse multi-scale features in a coarse-to-fine manner, which enables adaptive representation of objects at different scales and achieves faster convergence as well as superior detection performance. (iv) Our approach offers a unique perspective in mitigating DETR’s slow convergence issue with simply a plug-and-play module, thus can be easily integrated with existing convergence solutions in a complementary manner. Experiments show that with just 12 training epochs, our fully-fledged SAM-DETR++ surpasses the original DETR [5] trained for 500 epochs on the COCO benchmark [7], and achieves the state-of-the-art performance among DETR-based detectors.

This paper is an extension of our previous paper [13] published at the CVPR’2022 conference. Compared with its conference version [13], this paper incorporates the following new contributions. (i) We extend the proposed semantic-aligned matching mechanism to effectively incorporate the multi-scale features that are inherently unaligned in semantics. This enables the adaptive representation of objects of different sizes and further improves the convergence speed and detection accuracy significantly. (ii) We examine the compatibility of SAM-DETR++ with more recently proposed DETR convergence solutions, demonstrating its superior robustness and achieving further improvement in detection accuracy. (iii) A minor tweak to remove the dropout in the Transformer is adopted for superior performance at no computational cost. (iv) We conduct more comprehensive analysis of SAM-DETR++, including visualization, illustration, and experimentation, for a clearer explanation and a better understanding of our method.

The rest of this paper is organized as follows: Section 2 presents related work; Section 3 briefly reviews the architecture of DETR [5]; Section 4 describes our proposed method in detail; Section 5 presents experiment results and our analysis; Section 6 draws the concluding remarks.

2 RELATED WORK

2.1 ConvNet-Based Object Detection

Most modern object detectors are based on ConvNets and have experienced remarkable progress with the development of deep learning [1]. Most of these ConvNet-based detectors can be divided into two categories: two-stage and one-stage detectors. Two-stage detectors mainly involve Faster R-CNN [2] and its extensions [31, 32, 33, 34, 35, 36, 37, 38, 39, 40], which use a region proposal network (RPN) to first produce region proposals and then perform region-wise predictions over the proposals. These two-stage detectors are generally more accurate and have achieved state-of-the-art performance in various object detection challenges. Differently, one-stage detectors [41, 3, 8, 4, 42, 43, 44, 45, 46, 47, 48, 49, 50] remove the proposal generation stage and directly predict over the densely placed shape priors, which achieve better inference speed.

Even with promising results, these ConvNet-based object detectors are still not optimal. They perform object detection by defining and solving surrogate regression and classification tasks, and rely on many hand-crafted components, such as non-maximum suppression (NMS), anchors, and rule-based training target assignment. Therefore, these

ConvNet-based detectors' architectures are rather complicated, hyper-parameter-intensive, and not fully end-to-end, which leads to sub-optimal detection performance.

2.2 Transformer-Based Object Detection

Distinct from those ConvNet-based detectors, the recently proposed DETR [5] has revolutionized the object detection paradigm using a Transformer architecture [6] optimized by a set-based global loss. DETR eliminates the need for many hand-designed components and achieves the first fully end-to-end object detector with competitive performance. However, DETR suffers from extremely slow convergence and needs extra-long training to achieve good performance compared with those ConvNet-based object detectors. A few recent works have been proposed to mitigate this issue. Deformable DETR [9], Efficient DETR [51], Sparse DETR [52], and ViDT [53] replace the original dense attention with sparse attention mechanisms. PnP-DETR [54] proposes a poll-and-pool sampling strategy in its attention mechanism. Besides, Conditional DETR [10], SMCA-DETR [11], Anchor DETR [12], and DAB-DETR [55] make substantial modifications to the attention mechanism, aiming to add spatial constraints to the original cross-attention to better focus on prominent regions. Furthermore, the recently proposed DN-DETR [30] designs a novel de-noising training strategy to speed up DETR's training procedure, which also achieves very promising results.

In this work, we also aim to improve DETR's convergence and performance, but from a distinctive perspective. Our method does not modify the original attention mechanism in DETR, nor does it change the training strategy for DETR. Our method notices the unaligned semantics between object queries and encoded image features, which causes the difficulty of matching object queries to relevant regions and further DETR's slow convergence. To address this issue, our method only appends a plug-and-play module into the DETR architecture and thus can naturally work with existing convergence solutions for DETR [11, 30] in complementary. Besides, our method can be easily extended to fuse multi-scale features, which further reduces the complexity in representing objects of different sizes, thus accelerating the convergence to a greater extent.

2.3 Siamese Architectures for Matching-Involved Tasks

Matching is a common task in computer vision, of which the core idea is to predict the similarity between a pair of inputs. The concept of matching is commonly referred to in contrastive tasks, like face recognition [56, 57], object tracking [14, 58, 15, 16, 17, 18, 19, 59, 60, 61, 62], re-identification [20, 21, 22, 23, 24, 63], and few-shot recognition [25, 26, 27, 28, 29]. Empirical results have shown that Siamese-based architectures, which project both inputs into the same embedding space using identical sub-networks, perform extraordinarily well on these matching tasks. This is because Siamese-based architectures can produce comparable feature vectors with aligned semantics for the inputs, thus facilitating the computation of similarities among them.

Our work is motivated by the success of Siamese-based architectures in various matching-involved tasks. We interpret DETR's cross-attention as a 'matching and feature

distillation' process and leverage the philosophy of Siamese networks to facilitate the matching procedure. We believe that it is essential to impose aligned semantics between object queries and encoded image features so that the similarity between them can be efficiently and accurately computed for the accelerated convergence of DETR.

3 A BRIEF REVIEW OF DETR

Since our proposed SAM-DETR++ is developed on top of DETR [5] for its accelerated convergence and superior detection performance, we first briefly review the basic architecture of DETR [5] before introducing our method.

Unlike ConvNet-based object detectors [2, 41, 3, 4, 47] that address object detection by solving surrogate classification and regression tasks, DETR [5] directly formulates object detection as a set prediction problem. The pipeline of DETR is simple: a backbone network, a Transformer encoder, and a Transformer decoder. Given an input image $\mathbf{I} \in \mathbb{R}^{H_0 \times W_0 \times 3}$, the backbone network and the Transformer encoder produce the encoded features for the input image $\mathbf{F} \in \mathbb{R}^{HW \times d}$, in which d denotes the number of feature channels, and H_0 , W_0 and H , W denote the spatial sizes of the image and the encoded features, respectively. After that, the Transformer decoder takes the encoded image features \mathbf{F} and a small set of object queries $\mathbf{Q} \in \mathbb{R}^{N \times d}$ as input, and then produces the detection results. Here, N denotes the number of object queries, which is typically set to 100~300 [5, 9, 10, 11, 64, 12, 55, 13, 65, 66, 67, 68].

The Transformer decoder consists of multiple decoder layers, in which object queries are sequentially processed by a self-attention module, a cross-attention module, and a feed-forward network (FFN) to produce the outputs. The object queries output by each decoder layer are further fed into the subsequent layers and go through a Multi-Layer Perceptron (MLP) to produce detection predictions. The cross-attention module is the key element in the Transformer decoder, in which object queries interact with the encoded image features. As discussed in Section 1 and illustrated in Fig. 1, the cross-attention module can be interpreted as a 'matching and feature distillation' process: object queries first search for the relevant regions to match, then distill instance-level features from the matched regions to generate detection predictions. We formulate and interpret cross-attention as:

$$\mathbf{Q}' = \text{Softmax}\left(\underbrace{\frac{(\mathbf{Q}\mathbf{W}_q)(\mathbf{F}\mathbf{W}_k)^T}{\sqrt{d}}}_{\text{to match relevant regions}}\right)(\mathbf{F}\mathbf{W}_v), \quad (1)$$

to distill features from relevant regions

where \mathbf{W}_q , \mathbf{W}_k , and \mathbf{W}_v denote the linear projections for query, key, and value, respectively, in the Transformer attention mechanism, and $\mathbf{Q}' \in \mathbb{R}^{N \times d}$ denotes the cross-attention's generated object queries.

Preferably, the cross-attention's output object queries \mathbf{Q}' should contain instance-level features distilled from the relevant regions, which are used to produce detection predictions. However, as discussed before and also verified in [9, 10, 11, 12, 55, 13], the object queries are initially equally matched to all spatial locations in the encoded image features and are very challenging to learn to focus on specific

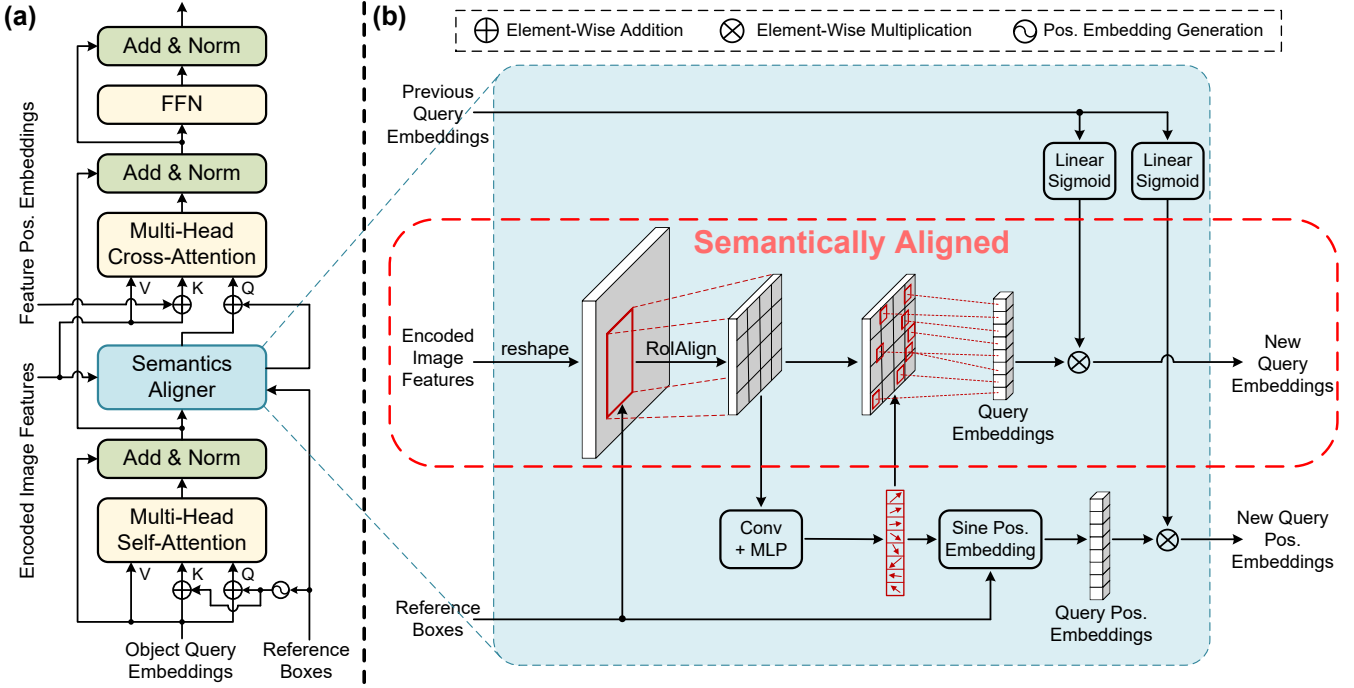


Fig. 2. (a) The overview of one Transformer decoder layer of the proposed SAM-DETR++. It models a learnable *reference box* for each object query, whose center location is used to generate the corresponding positional embeddings. It also appends a *Semantics Aligner* ahead of the cross-attention module, which generates new object queries that are semantically aligned with the encoded image features, thus facilitating their subsequent matching process within cross-attention. (b) The architecture of the proposed *Semantics Aligner*. Only one object query is presented for concise illustration. Semantics Aligner first extracts region features from the reference boxes' corresponding regions via RoIAlign, which are used to predict the coordinates of several representative keypoints with the most distinctive features. The features from these representative keypoints are then sampled as the new query embeddings, which are semantically aligned with the encoded image features. Finally, the new query embeddings are further reweighted by the previous query embeddings to incorporate useful information from them.

regions properly. The matching difficulty with unaligned semantics is the key root for DETR's slow convergence.

4 SAM-DETR++

This section presents the proposed SAM-DETR++ in detail, which greatly accelerates the convergence speed of the original DETR [5] with a plug-and-play module to achieve semantic-aligned matching. The proposed SAM-DETR++ can also effectively fuse multi-scale features in a coarse-to-fine manner on the basis of the introduced semantic-aligned matching mechanism. In addition, being like a plug and play, SAM-DETR++ can be integrated with multiple DETR's convergence solutions to achieve even faster convergence and superior detection accuracy.

4.1 Overview

Motivated by the observations discussed before, our proposed SAM-DETR++ aims to address DETR's slow convergence issue by relieving the complexity of the matching process as illustrated in Fig. 1. The core idea is to project both object queries and encoded image features into the same feature embedding space, thus imposing a strong prior for each object query to focus on its relevant region with similar semantics within the cross-attention module. To achieve this, SAM-DETR++ only makes some minor modifications to the decoder of the original DETR [5].

Fig. 2(a) illustrates the overall architecture of the Transformer decoder of SAM-DETR++. Same as the original DETR [5], each decoder layer is repeated six times, with

zeros as input for the first layer and previous layer's outputs as input for the subsequent layers. As shown in Fig. 2(a), in each decoder layer, a plug-and-play module, named *Semantics Aligner*, is appended ahead of the cross-attention module for imposing aligned semantics between object queries and encoded image features. Besides, SAM-DETR++ also models a learnable reference box for each object query instead of directly modeling its query positional embeddings. The learnable reference boxes $\mathbf{R}_{\text{box}} \in \mathbb{R}^{N \times 4}$ are modeled at the first decoder layer, representing object queries' initial locations. With the spatial guidance of these reference boxes, the proposed Semantics Aligner takes the previous object query embeddings \mathbf{Q} and the encoded image features \mathbf{F} as inputs to obtain new object query embeddings \mathbf{Q}^{new} and their corresponding positional embeddings $\mathbf{Q}_{\text{pos}}^{\text{new}}$, and then feed them into the subsequent cross-attention module. In this way, the generated embeddings \mathbf{Q}^{new} lie within the same embedding space as the encoded image features \mathbf{F} , which facilitates the following matching process between them, allowing object queries to quickly and properly attend to relevant regions with similar semantics in the encoded image features. It is worth noting that no modification is made to the other components in DETR's decoder layers, including multi-head self-attention, multi-head cross-attention, and FFN.

4.2 Semantics Aligner

The detailed architecture of the appended *Semantics Aligner* is illustrated in Fig. 2(b).

Semantic-Aligned Matching. As formulated in Eq. 1 and illustrated in Fig. 1 (left), the cross-attention module uses dot-product to produce the attention heatmaps that represent the matching between object queries and encoded image features. It is natural and intuitive to adopt dot-product for generating the attention heatmaps, as dot-product is a good metric for the similarity between two feature vectors, which encourages object queries to have higher attention weights for regions with higher similarities. However, as illustrated in Fig. 1 (right), the modules between cross-attentions project object queries into a different feature embedding space from that of the encoded image features, leading to unaligned semantics between them. The unaligned semantics causes each object query to almost equally match all spatial locations within the encoded image features at initialization, adding substantial complexity for learning a meaningful matching between them.

Motivated by these observations, we design Semantics Aligner to ensure that object query embeddings are within the same feature embedding space as encoded image features before being processed by cross-attention. This guarantees that the dot-product between query embeddings and encoded image features is always a meaningful measurement of similarity without the need to be explicitly learned, which imposes a prior for object queries to match relevant regions with similar semantics and reduces the matching difficulty.

The alignment of semantics is achieved by re-sampling new object query embeddings from the encoded image features. Concretely, as shown in Fig. 2 (b), Semantics Aligner first restores the encoded image features' spatial dimensions from 1D sequences $HW \times d$ to 2D feature maps $H \times W \times d$. Then, Semantics Aligner extracts region features $\mathbf{F}_R \in \mathbb{R}^{N \times 7 \times 7 \times d}$ from the encoded image features via RoIAlign [69] from the corresponding reference boxes. Finally, the new object query embeddings \mathbf{Q}^{new} and new query positional embeddings $\mathbf{Q}_{\text{pos}}^{\text{new}}$ are obtained by re-sampling features from \mathbf{F}_R . Mathematically, Semantics Aligner can be formulated at a high level as:

$$\mathbf{F}_R = \text{RoIAlign}(\mathbf{F}, \mathbf{R}_{\text{box}}), \quad (2)$$

$$\mathbf{Q}^{\text{new}}, \mathbf{Q}_{\text{pos}}^{\text{new}} = \text{Re-Sample}(\mathbf{F}_R, \mathbf{R}_{\text{box}}, \mathbf{Q}). \quad (3)$$

As the re-sampling procedure does not involve any projection (e.g., ConvNet or MLP), the new object query embeddings \mathbf{Q}^{new} always lie within the same feature embedding space as the encoded image features \mathbf{F} , which encourages object queries to focus on semantically similar regions in the following cross-attention module. The design choice for the re-sampling procedure is to be detailed later.

Semantic-Aligned Matching with Multiple Representative Keypoints. The re-sampling process can be easily accomplished by simple operations like applying global average-pooling or global max-pooling on region features \mathbf{F}_R . But instead, we propose a more sophisticated approach to re-sample new object query embeddings, inspired by the prior works [50, 70, 71, 72, 5, 73, 10] that identify the importance of objects' representative keypoints in object detection. Specifically, Semantics Aligner explicitly searches for multiple representative keypoints for each object query




Fig. 3. Each decoder layer in SAM-DETR++ searches multiple representative keypoints (cyan dots) within each reference box (red box), and uses their features for semantic-aligned matching. As detection proceeds, the keypoints gradually fall on salient and semantically meaningful locations, and the attention heatmaps gradually become more precise.

and extracts their features for the aforementioned semantic-aligned matching. Such design can naturally fit in the multi-head attention mechanism [6] without any modification, enabling every attention head to produce different weights to focus on different parts.

Here, we denote the number of attention heads as M , which is set to 8 in most DETR-based object detectors [5, 9, 10, 11, 74, 12, 55, 13]. M is also the number of representative keypoints to search for each object query. As shown in Fig. 2 (b), after retrieving region features \mathbf{F}_R , Semantics Aligner adopts a ConvNet followed by an MLP to predict the spatial locations of the M keypoints for each object query, representing the locations that are crucial for recognizing and localizing the potential objects, which can be formulated as:

$$\mathbf{R}_{\text{SP}} = \text{MLP}(\text{ConvNet}(\mathbf{F}_R)), \quad (4)$$

where $\mathbf{R}_{\text{SP}} \in \mathbb{R}^{N \times M \times 2}$ denotes the coordinates of M keypoints for each of the N object queries. Note that the predicted coordinates are constrained to be inside their corresponding reference boxes. With the predicted \mathbf{R}_{SP} , the features of these representative keypoints can be then sampled from \mathbf{F}_R via bi-linear interpolation. Semantics Aligner finally concatenates the M sampled features vectors corresponding to the M representative keypoints as the new query embeddings, which are fed into the subsequent multi-head cross-attention so that each attention head can focus on features of one representative keypoint. Similarly, the object queries' corresponding positional embeddings can be computed using sinusoidal functions based on the keypoints' image-scale coordinates, and are then also concatenated to feed into the multi-head cross-attention module.

$$\mathbf{Q}^{\text{new}'} = \text{Concat}(\{\mathbf{F}_R[\dots, x, y, \dots] \text{ for } x, y \in \mathbf{R}_{\text{SP}}\}) \quad (5)$$

$$\mathbf{Q}_{\text{pos}}^{\text{new}'} = \text{Concat}(\text{Sinusoidal}(\mathbf{R}_{\text{box}}, \mathbf{R}_{\text{SP}})) \quad (6)$$

Fig. 3 visualizes the searched representative keypoints and the attention heatmaps produced by cross-attention. It can be observed that the search keypoints gradually fall on the salient positions with rich semantics (e.g., head and extremities) as detection proceeds. In addition, the attention heatmaps also gradually become more precise and focus

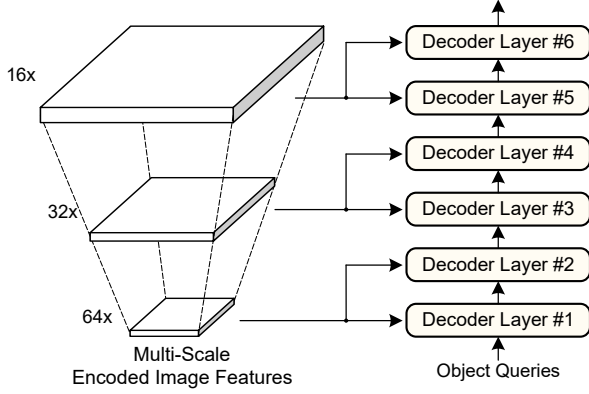


Fig. 4. The proposed SAM-DETR++ can be extended to fuse multi-scale features by simply feeding different feature scales into different decoder layers in a coarse-to-fine manner. Thanks to the introduced semantic-aligned matching mechanism, SAM-DETR++ can effectively fuse multi-scale features that are inherently unaligned in feature semantics.

on those semantically meaningful regions. These results validate the effectiveness of searching and exploiting key-point features for semantic-aligned matching in enhancing its representation capacity.

Feature Reweighting with Previous Query Embeddings.

So far, Semantics Aligner can generate new object query embeddings $\mathbf{Q}^{\text{new}/}$ with aligned semantics with the encoded image features. However, it also brings one issue: the cross-attention cannot leverage the previous query embeddings \mathbf{Q} that contain valuable information for detection. To mitigate this issue, Semantics Aligner further receives the previous query embeddings \mathbf{Q} as inputs to produce a set of reweighting coefficients via ‘Linear Projection + Sigmoid’. The reweighting coefficients are applied to new query embeddings $\mathbf{Q}^{\text{new}/}$ and their positional embeddings $\mathbf{Q}_{\text{pos}}^{\text{new}/}$ through element-wise multiplication, highlighting those important features. As a result, the useful information from previous query embeddings can be effectively leveraged in cross-attention with the introduced semantic-aligned matching mechanism. Note that feature reweighting does not affect the aligned semantics of query embeddings, as it does not perform any projection on the query embeddings. The described feature reweighting process can be formulated as:

$$\mathbf{Q}^{\text{new}} = \mathbf{Q}^{\text{new}/} \otimes \sigma(\mathbf{Q}\mathbf{W}_{\text{RW1}}), \quad (7)$$

$$\mathbf{Q}_{\text{pos}}^{\text{new}} = \mathbf{Q}_{\text{pos}}^{\text{new}/} \otimes \sigma(\mathbf{Q}\mathbf{W}_{\text{RW2}}), \quad (8)$$

where \mathbf{W}_{RW1} and \mathbf{W}_{RW2} are the learnable parameters for linear projections, $\sigma(\cdot)$ denotes sigmoid function, and \otimes denotes element-wise multiplication.

4.3 Multi-Scale Feature Fusion with Aligned Semantics

Detecting objects of vastly different scales has always been one major challenge in object detection. Modern ConvNet-based object detectors (e.g., Faster R-CNN [2] w/ FPN [75], M2Det [45], EfficientDet [46]) usually incorporate multi-scale features to accommodate this issue by representing objects at different scales in a ‘divide and conquer’ manner, which reduces the representation complexity and achieves superior detection accuracy and faster convergence. With

this motivation, we further extend the proposed semantic-aligned matching strategy to fuse multi-scale features in a coarse-to-fine manner.

As shown in Fig. 4, the proposed method to fuse multi-scale features is simple and concise. Considering the cascade nature of DETR’s decoder, we feed the features of different scales into different stages of the decoder, making the detection pipeline a coarse-to-fine refinement process. Concretely, the first two decoder layers receive the coarsest feature maps to reduce search space for initial localization; the subsequent two layers receive finer feature maps for more precise localization; the last two layers receive high-resolution feature maps for detecting tiny objects. This simple and effective design does not introduce extra parameters but allows SAM-DETR++ to adaptively fuse multi-scale features to represent objects at different scales, thus greatly lowering the learning complexity.

It is worth noting that it is our proposed semantic-aligned matching mechanism that enables this simple approach to fuse multi-scale features effectively. Experiments in Section 5.4 show that without the proposed semantic-aligned matching, directly fusing multi-scale features does not bring a clear performance gain. This is because of the inevitable unaligned semantics across different feature scales, which causes extra complexity in the matching processes between object queries and encoded image features, as discussed before. Our introduced Semantics Aligner alleviates this issue by explicitly aligning the semantics between object query embeddings and encoded image features at all decoder layers, thus enabling the effective fuse of features from different scales.

4.4 Removing Dropout in Transformer

Most existing DETR-like detectors [5, 9, 74, 11, 10, 52], including the CVPR’2022 conference version of SAM-DETR [13], contain dropout [76] in the Transformer encoder-decoder architecture [6]. Dropout [76] is included to mitigate the overfitting issue in natural language processing (NLP) tasks. However, in the task of object detection within images, we empirically find that dropout does not mitigate overfitting but harms the performance of object detection. We conjecture that this is largely attributed to the unique property of object detection within 2D images, where adjacent pixels within feature maps are strongly correlated. Therefore, we incorporate a minor tweak to remove the dropout in the Transformer, which yields better performance at no computational cost. The effectiveness of this modification is verified in Section 5.4.

4.5 Compatibility with Existing Convergence Solutions

As illustrated in Fig. 2(a), SAM-DETR++ only appends a plug-and-play module into the Transformer decoder layer, leaving most other operations unchanged. Besides, SAM-DETR++ speeds up DETR’s training convergence from a distinct perspective from existing convergence solutions. These properties make it easy and effective to integrate SAM-DETR++ with other approaches to achieve even faster convergence and superior detection accuracy. Here, we integrate our method with two recent works to validate the strong compatibility of SAM-DETR++.




Fig. 5. Visualization of the searched representative keypoints and the attention heatmaps of different attention heads in cross-attention from our proposed SAM-DETR++. The searched representative keypoints mostly fall around objects of interest, and typically fall on the positions with the most distinctive features for recognition or localization, like object extremities or central points. Our method’s attention heatmaps are much more focused compared with the original DETR without semantic-aligned matching, which proves the effectiveness of our approach in relieving the complication in the matching processes between object queries and encoded image features, which accelerates DETR’s convergence. Red-colored arrows highlight those fine details in attention heatmaps. Zoom-in may be required to view details.

4.5.1 Compatibility with SMCA-DETR

SMCA-DETR [11] replaces DETR’s original cross-attention module with Spatially Modulated Co-Attention (SMCA), which estimates the position of each object query, and then applies a series of 2D-Gaussian weight maps to constrain the attention responses in different attention heads. Both the center locations and the scales for SMCA’s 2D-Gaussian weight maps are predicted from the corresponding object query embeddings. SMCA [11] effectively accelerates DETR’s convergence by imposing spatial constraints to the SMCA module.

To integrate SMCA [11] into our proposed SAM-DETR++, we make one minor modification to the SMCA mechanism: we adopt the coordinates of the M representative keypoints as the central locations for the 2D Gaussian weight maps. The scales of the weight maps are also predicted from the region features in parallel to the central locations. Experiment results in Section 5.3 validate the complementary effect between our proposed SAM-DETR++ and SMCA [11].

4.5.2 Compatibility with DN-DETR

The recently proposed DN-DETR [30] introduces a novel de-noising training strategy to speed up DETR’s training

procedure, which is also complementary to our approach without any adaptation. Experiment results in Section 5.3 also validate the complementary effect between our proposed SAM-DETR++ and DN-DETR [30].

5 EXPERIMENTS

5.1 Experiment Setup

Dataset and Evaluation Metrics. Following prior works [5, 9, 10, 11, 12, 13, 55], we mainly perform the experiments on the COCO 2017 dataset [7], using the $\sim 117k$ images in the train2017 set for training and the 5k images in the val2017 set for evaluation. We adopt the standard metrics defined by COCO to evaluate the performance of object detection.

Implementation Details. SAM-DETR++’s implementation details mostly align with the original DETR [5] and other prior works [9, 10, 11, 12, 55]. We use ImageNet-pretrained [78] ResNet-50 [79] as the backbone network. All experiments are performed on servers with $8 \times$ Nvidia V100 GPUs. We train our models with AdamW optimizer [80, 81]. The batch size is set to 16 for training, except when ResNet-50-DC5 is used as the backbone, the batch size is set to 8. The initial learning rate is 1×10^{-5} for the backbone parameters

TABLE 1
Object detection performance under the 12-epoch (1x) training schedule on COCO val2017.

Method	multi-scale	#Epochs	#Params (M)	GFLOPs	AP	AP ₅₀	AP ₇₅	AP _S	AP _M	AP _L
• Backbone: ResNet-50 (Single-Scale Features)										
Faster-R-CNN-R50 [2]		12	34	547	35.7	56.1	38.0	19.2	40.9	48.7
DETR-R50 [5] ‡		12	41	86	22.3	39.5	22.2	6.6	22.8	36.6
Deformable-DETR-R50 [9]		12	34	78	31.8	51.4	33.5	15.0	35.7	44.7
Conditional-DETR-R50 [10]		12	44	90	32.2	52.1	33.4	13.9	34.5	48.7
SMCA-DETR-R50 [11]		12	42	86	31.6	51.7	33.1	14.1	34.4	46.5
SAM-DETR-R50 (Ours)		12	57	107	34.2	55.8	35.3	15.0	37.7	52.5
SAM-DETR-R50 w/ SMCA (Ours)		12	57	107	37.0	58.0	38.5	17.8	40.3	56.1
• Backbone: ResNet-50-DC5 (High-Resolution Features)										
Faster-R-CNN-R50-DC5 [2]		12	166	320	37.3	58.8	39.7	20.1	41.7	50.0
DETR-R50-DC5 [5] ‡		12	41	187	25.9	44.4	26.0	7.9	27.1	41.4
Deformable-DETR-R50-DC5 [9]		12	34	128	34.9	54.3	37.6	19.0	38.9	47.5
Conditional-DETR-R50-DC5 [10]		12	44	195	35.9	55.8	38.2	17.8	38.8	52.0
SMCA-DETR-R50-DC5 [11]		12	42	187	32.5	52.8	33.9	14.2	35.4	48.1
Anchor-DETR-R50-DC5 [12]		12	39	151	37.1	57.8	39.1	19.0	40.8	51.4
DAB-DETR-R50-DC5 [55]		12	44	216	38.0	60.3	39.8	19.2	40.9	55.4
DN-DETR-R50-DC5 [30]		12	44	216	41.7	61.4	44.1	21.2	45.0	60.2
SAM-DETR-R50-DC5 (Ours)		12	57	229	39.1	59.9	41.2	20.9	42.8	55.5
SAM-DETR-R50-DC5 w/ SMCA (Ours)		12	57	229	41.3	61.6	43.6	22.1	44.9	59.2
SAM-DETR-R50-DC5 w/ DN (Ours)		12	57	229	42.3	61.7	45.2	22.8	45.7	60.0
SAM-DETR-R50-DC5 w/ SMCA + DN (Ours)		12	57	229	43.7	63.0	46.8	24.3	47.4	61.4
• Backbone: ResNet-50 (Multi-Scale Features)										
Faster-R-CNN-R50-FPN [2, 75]	✓	12	42	180	37.9	58.8	41.1	22.4	41.1	49.1
Cascade-R-CNN-R50-FPN [31, 75]	✓	12	69	230	40.4	58.9	44.1	22.8	43.7	54.0
FCOS-R50 [4]	✓	12	32	201	38.6	57.2	41.7	23.5	42.8	48.9
Sparse-R-CNN-R50-FPN [77]	✓	12	106	166	40.1	59.4	43.5	22.9	43.6	52.9
Deformable-DETR-R50 [9]	✓	12	40	173	37.2	55.5	40.5	21.1	40.7	50.5
SMCA-DETR-R50 [11]	✓	12	40	152	35.0	54.1	37.8	18.7	37.7	48.1
SAM-DETR++-R50 (Ours)	✓	12	55	203	41.9	60.5	45.3	24.6	45.5	57.4
SAM-DETR++-R50 w/ SMCA (Ours)	✓	12	55	203	43.2	61.5	46.5	25.5	46.5	58.6
SAM-DETR++-R50 w/ SMCA + DN (Ours)	✓	12	55	203	44.8	62.6	47.9	26.7	48.2	60.9

‘‡’ denotes the original DETR baseline [5] with increased number of object query (100→300) and focal loss as the classification loss function.

and 1×10^{-4} for the other parameters. The weight decay is set to 1×10^{-4} . Two training schedules are experimented: (i) the 12-epoch (1x) schedule that is widely adopted in ConvNet-based detectors [2, 75, 8, 4], where the learning rate decays at the 10th epoch; (ii) the 50-epoch schedule that is often used in Transformer-based detectors [9, 10, 11, 12, 55], where the learning rate decays at the 40th epoch. Model-related hyper-parameters (*e.g.*, feature channel dimension, number of encoder and decoder layers) remain the same with DETR [5], except we make two minor modifications following some recent works [9, 74, 10, 11, 12] to improve DETR’s convergence speed: the number of object queries N is increased from 100 to 300; the sigmoid focal loss [8] is adopted as the classification loss instead of the cross-entropy loss. These two modifications are also applied to the original DETR [5] for a fair comparison with the baseline.

The same data augmentation used in prior works [5, 9, 10, 12, 55, 13] is adopted, which includes random resize, horizontal flip, and random crop. We constrain the training images’ longest sides to be less or equal than 1333 pixels and the shortest sides to be larger or equal than 480 pixels.

5.2 Visualization and Analysis

Fig. 5 visualizes the representative keypoints searched by the proposed Semantics Aligner as well as their corresponding attention heatmaps generated from the subsequent multi-head cross-attention module. We also compare the attention heatmaps with the ones generated from the

original DETR [5]. Results are obtained under the 12-epoch (1x) training schedule using ResNet-50 [79].

The visualization shows that the searched representative keypoints mostly fall around the target objects, and typically at those representative positions with the most distinctive features, such as object extremities or central points. The attention response heatmaps generated by the subsequent cross-attention modules also show high responses on those searched representative keypoints accordingly. In addition, compared with the original DETR [5], our method shows clearly more precise and focused responses, which validates that the proposed semantic-aligned matching mechanism successfully facilitates the matching of object queries with appropriate regions for distillation of relevant instance-level features, thus accelerating DETR’s convergence.

5.3 Experiment Results

This subsection presents experiment results under two different training schedules. Here, we denote our proposed method with multi-scale feature fusion included as SAM-DETR++, and denote our method without multi-scale feature fusion as SAM-DETR. It is noteworthy that SAM-DETR is identical to its CVPR’2022 conference version [13], except that we remove dropout [76] as discussed in Section 4.4.

Results under the 12-epoch (1x) Schedule. We first present object detection results under the short 12-epoch

TABLE 2
Comparison with state-of-the-art object detectors on COCO val 2017 under longer training schedules.

Method	#Epochs	#Params (M)	GFLOPs	AP	AP ₅₀	AP ₇₅	AP _S	AP _M	AP _L
• Baseline methods trained for extra-long epochs:									
Faster-R-CNN-R50-FPN [2, 75]	108	42	180	42.0	62.1	45.5	26.6	45.4	53.4
DETR-R50 [5]	500	41	86	42.0	62.4	44.2	20.5	45.8	61.1
DETR-R50-DC5 [5]	500	41	187	43.3	63.1	45.9	22.5	47.3	61.1
• ConvNet-based object detectors:									
Cascade-Mask-RCNN-R50-FPN [31]	36	77	394	44.3	62.2	48.0	26.6	47.7	57.7
TSP-FCOS-R50-FPN [82]	36	52	189	43.1	62.3	47.0	26.6	46.8	55.9
TSP-R-CNN-R50-FPN [82]	36	64	188	43.8	63.3	48.3	28.6	46.9	55.7
TSP-R-CNN-R50-FPN [82]	96	64	188	45.0	64.5	49.6	29.7	47.7	58.0
Sparse-R-CNN-R50-FPN [77]	36	106	166	45.0	63.4	48.2	26.9	47.2	59.5
• Transformer-based object detectors:									
DETR-R50 [5] ‡	50	41	86	34.9	55.5	36.0	14.4	37.2	54.5
DETR-R50-DC5 [5] ‡	50	41	187	36.7	57.6	38.2	15.4	39.8	56.3
UP-DETR-R50 [83]	150	41	86	40.5	60.8	42.6	19.0	44.4	60.0
UP-DETR-R50 [83]	300	41	86	42.8	63.0	45.3	20.8	47.1	61.7
Deformable-DETR-R50 [9]	50	40	173	43.8	62.6	47.7	26.4	47.1	58.0
Deformable-DETR-R50 (two-stage) [9]	50	40	173	46.2	65.2	50.0	28.8	49.2	61.7
SMCA-DETR-R50 [11]	50	40	152	43.7	63.6	47.2	24.2	47.0	60.4
SMCA-DETR-R50 [11]	108	40	152	45.6	65.5	49.1	25.9	49.3	62.6
Conditional-DETR-R50 [10]	50	44	90	40.9	61.8	43.3	20.8	44.6	59.2
Conditional-DETR-R50 [10]	108	44	90	43.0	64.0	45.7	22.7	46.7	61.5
Conditional-DETR-R50-DC5 [10]	50	44	195	43.8	64.4	46.7	24.0	47.6	60.7
Conditional-DETR-R50-DC5 [10]	108	44	195	45.1	65.4	48.5	25.3	49.0	62.2
Anchor-DETR-R50 [12]	50	37	93	42.1	63.1	44.9	22.3	46.2	60.0
Anchor-DETR-R50-DC5 [12]	50	37	172	44.2	64.7	47.5	24.7	48.2	60.6
DAB-DETR-R50 [55]	50	44	94	42.2	63.1	44.7	21.5	45.7	60.3
DAB-DETR-R50-DC5 [55]	50	44	202	44.5	65.1	47.7	25.3	48.2	62.3
DAB-DETR-R50-DC5 (3 Patterns) [55]	50	44	216	45.7	66.2	49.0	26.1	49.4	63.1
Sparse-DETR-R50 [52]	50	41	136	46.3	66.0	50.1	29.0	49.5	60.8
DN-DETR-R50 [30]	50	44	94	44.1	64.4	46.7	22.9	48.0	63.4
DN-DETR-R50-DC5 [30]	50	44	202	46.3	66.4	49.7	26.7	50.0	64.3
SAM-DETR++-R50 (Ours)	50	55	203	47.5	66.5	51.3	29.3	50.8	62.7
SAM-DETR++-R50 w/ SMCA (Ours)	50	55	203	48.0	66.6	52.2	29.9	51.5	64.6
SAM-DETR++-R50 w/ SMCA + DN (Ours)	50	55	203	49.1	67.2	53.2	30.5	52.6	64.7

‡ denotes the original DETR baseline [5] with increased number of object query (100→300) and focal loss as the classification loss function.

(1x) training schedule, which is widely used in conventional ConvNet-based object detectors [2, 3, 8, 43, 4]. As shown in Table 1, when using ResNet-50 or ResNet-50-DC5 as backbones, Faster R-CNN [2] is able to achieve relatively satisfactory detection accuracy, while the original DETR [5] is still heavily under-trained. Other recently proposed DETR-based object detectors [9, 10, 11, 12, 55] obtain clearly better results over the original DETR [5], but still have large gaps with Faster R-CNN [2]. With the proposed semantic-aligned matching mechanism incorporated into DETR [5], the standalone SAM-DETR without multi-scale feature fusion obtains significant performance gain over the original DETR baseline [5] (+11.9% AP w/ R50 and +13.2% AP w/ R50-DC5), achieves comparable detection accuracy as Faster R-CNN [2], and also significantly outperforms other Transformer-based object detectors [9, 10, 11, 12, 55]. Furthermore, since SAM-DETR does not modify the attention mechanism nor does it modify the training strategy, one unique advantage of SAM-DETR is that it can be easily integrated with other DETR’s convergence solutions. As shown in Table 1, our approach can be integrated with SMCA-DETR [11], DN-DETR [30], or both of them. The integrated methods achieve consistent and significant performance gain over the standalone SAM-DETR as well as their respective baselines [11, 30]. Our methods even significantly outperform the known-to-be-fast-converging Faster

R-CNN [2] by very large margins. These results verify the effectiveness of our proposed semantic-aligned matching mechanism and its strong compatibility.

In Table 1, we also present the object detection performance of our proposed SAM-DETR++ that is extended to fuse multi-scale features, as well as the performance of other detectors exploiting multi-scale features. As shown in Table 1, fusing multi-scale features via the introduced semantic-aligned matching mechanism significantly improves the detection accuracy over SAM-DETR-R50-DC5 with even reduced computational cost. In addition, it is noteworthy that SAM-DETR++-R50 w/ SMCA+DN achieves a state-of-the-art detection performance of 44.8% AP with only 12 training epochs, which outperforms the original DETR-R50-DC5 [5] trained for 500 epochs (43.3% AP), reducing the required number of training epochs by more than 97.6%. These results show that fusing multi-scale features via our proposed semantic-aligned matching mechanism further improves the detection accuracy and accelerates the convergence to a greater extent.

Results under Longer Training Schedules. Table 2 further compares SAM-DETR++ with other state-of-the-art object detectors under the longer training schedules. When trained for 50 epochs, our proposed SAM-DETR++ already outperforms the original DETR [5] trained for 500 epochs by large




Fig. 6. The convergence curves of DETR, SAM-DETR (w/o multi-scale feature fusion), and SAM-DETR++ (w/ multi-scale feature fusion). DETR is trained with 500 epochs, with the learning rate dropped at the 400th epoch. SAM-DETR and SAM-DETR++ are trained under the 12-epoch (1x) and 50-epoch learning schedules. Our methods converge much faster and achieve clearly better detection performance over the original DETR.

TABLE 3

Ablation study on the design choices of SAM-DETR++. Results are obtained on COCO val 2017 under the 12-epoch (1x) learning schedule.

Semantics Aligner	Query Re-Sampling Strategy				Feature Reweighting	Remove Dropout	Multi-Scale Feature Fusion	AP	AP ₅₀	AP ₇₅	AP _S	AP _M	AP _L
	AvgPool	MaxPool	Keypoint x1	Keypoints x8									
✓	✓							22.3	39.5	22.2	6.6	22.8	36.6
✓		✓						25.2	48.9	23.3	8.9	26.4	41.3
✓			✓					27.0	50.2	25.8	10.3	28.0	43.9
✓				✓				28.6	50.3	28.1	12.4	31.2	44.4
✓					✓			30.3	52.0	29.8	12.4	32.8	47.3
✓				✓				32.0	53.4	32.8	13.5	35.3	49.2
✓				✓	✓			33.1	54.2	33.7	13.9	36.5	51.7
✓				✓	✓	✓		34.2	55.8	35.3	15.0	37.7	52.5
✓				✓	✓	✓	✓	29.1	53.2	28.6	11.5	31.8	46.0
✓				✓	✓	✓	✓	41.9	60.5	45.3	24.6	45.5	57.4

margins, and also achieves state-of-the-art performance among all Transformer-based object detectors. In addition, as SAM-DETR++ works from a distinct perspective from existing solutions, combining our proposed SAM-DETR++ with SMCA [11] and DN [30] (SAM-DETR++ w/ SMCA and SAM-DETR++ w/ SMCA+DN) brings further performance gains, achieving the state-of-the-art accuracy of 49.1% AP on COCO val 2017 with ResNet-50, without bells and whistles.

Convergence Curves. We also present the convergence curves of the proposed SAM-DETR and SAM-DETR++ in Fig. 6, which show significantly accelerated convergence speed of our methods over the baselines. These experiment results well validate our method’s superior learning efficiency, good detection accuracy, and strong compatibility.

5.4 Ablation Study

We conduct comprehensive ablation experiments to validate the effectiveness of our designs in SAM-DETR++. The ablation experiments are performed with ResNet-50 [79] as the backbone network and under the short 12-epoch (1x) learning schedule.

Effect of Semantic-Aligned Matching. The first row in Table 3 shows the detection result of the original DETR baseline [5]. As shown in Table 3, the proposed Semantic Aligner, together with any query re-sampling strategy, consistently improves the performance over the baseline. We highlight that even with the naive max-pooling re-sampling, AP and AP₅₀ significantly improves by 4.7% and

10.7%, respectively. The results validate our claim that the proposed semantic-aligned matching mechanism effectively eases the matching difficulty between object queries and their corresponding target features, thus accelerating the training convergence of DETR.

Effect of Semantic-Aligned Matching with Representative Keypoints. As shown in Table 3, different object query re-sampling strategies lead to large variance in detection performance. Max-pooling performs clearly better than average-pooling, which suggests that object detection relies more on salient features rather than treating all features within reference boxes equally. This motivates us to explicitly search representative keypoints and employ their features for the introduced semantic-aligned matching mechanism. Results show that searching just one keypoint and re-sampling its features as new object queries outperforms the naive re-sampling strategies (AvgPool and MaxPool). Furthermore, searching multiple representative keypoints can naturally work with the multi-head attention mechanism [6] to further strengthen the representation capability of the re-sampled new object queries produced by Semantic Aligner and achieve superior performance.

Fig. 7 also studies the effect of the number of representative keypoints for each object query. As the figure shows, the performance increases as the number of keypoints increases and saturates at 8. Therefore, we set the number of representative keypoints at 8 by default in our approach.

Searching within Reference Boxes vs. Searching within Images. As introduced in Section 4.2, the representative




Fig. 7. Ablation study on the number of searched representative keypoint(s) for each object query. Results are obtained on COCO val2017 under the 12-epoch (1x) learning schedule without multi-scale feature fusion and removing dropout.

TABLE 4

Ablation study on the search range of representative keypoints. Results are obtained on COCO val2017 under the 12-epoch (1x) learning schedule without multi-scale feature fusion and removing dropout.

Keypoint Search Range		AP	AP _{0.5}	AP _{0.75}
within ref box	within image			
✓		33.1	54.2	33.7
	✓	30.0	52.3	29.2

keypoints are searched within their corresponding reference boxes. We also evaluate the performance when allowing representative keypoints to be outside their corresponding reference boxes by relaxing the search range constraint. As shown in Table 4, searching representative keypoints at the image scale impedes the performance. We suspect the performance drop is due to the increased difficulty of matching given a larger search range. Note that, in the original DETR [5], object queries do not have explicit search ranges, which corresponds to the setup of image-scale searching. While our proposed SAM-DETR++ models learnable reference boxes with interpretable meanings, which can effectively narrow down the search range, leading to faster training convergence.

Effect of Feature Reweighting with Previous Query Embeddings. As discussed in Section 4.2, previous object queries’ embeddings contain helpful information for object detection, which cannot be directly leveraged due to the introduced re-sampling process. As a workaround, we propose to perform feature reweighting on re-sampled query embeddings based on previous query embeddings. This incorporates information from previous query embeddings while does not impede the aligned semantics. As shown in Table 3, the proposed feature reweighting mechanism consistently boosts performance, indicating its effectiveness.

Effect of Removing Dropout in Transformer. As shown in Table 3, the simple tweak of removing dropout [76] in Transformer [6] for object detection increases the performance of SAM-DETR without involving multi-scale feature fusion by 1.1% AP, at no extra computational cost.

Effect of Multi-Scale Feature Fusion with Aligned Semantics. As shown in Table 3, on top of SAM-DETR, incorporating multi-scale feature fusion improves the detection performance by a considerable margin of 7.7% AP. This

verifies that multi-scale feature fusion effectively reduces the complexity of representing objects of different sizes and can adaptively choose appropriate feature scales for object representation, leading to further performance gain.

It is noteworthy that multi-scale feature fusion highly depends on our proposed semantic-aligned matching mechanism. As shown in Table 3, performing multi-scale feature fusion without our proposed semantic-aligned matching only yields a poor performance of 29.1% AP (-12.8% AP compared with SAM-DETR++). This is because there exists inevitable unaligned semantics across different feature scales. Without imposing aligned semantics, directly fusing multi-scale features that are projected into different feature embedding spaces (*i.e.*, with unaligned semantics) causes extra matching difficulty in cross-attention, as explained in Section 1 and Fig. 1.

5.5 Further Discussions

On the Compatibility among SAM-DETR++, SMCA-DETR [11], and DN-DETR [30]. One of the key advantages of the proposed SAM-DETR++ is its excellent compatibility, which we demonstrate by integrating it with SMCA-DETR [11] and DN-DETR [30] and achieving superior detection performance. The reason behind their excellent compatibility is that each of them effectively accelerates the convergence of DETR from distinct perspectives, thus complementing each other. Concretely, SMCA-DETR [11] accelerates DETR’s convergence by imposing strong spatial constraints for the cross-attention module, in which each object query is limited to attend to a specific region adaptively. SMCA-DETR [11] effectively reduces the search space for each object query in cross-attention, thus improving the training convergence. DN-DETR [30] proposes a de-noising training strategy to mitigate the instability of bipartite graph matching that causes inconsistent optimization goals in DETR’s early training stages. With its proposed de-noising training strategy, the optimization objectives of DETR become consistent even in early training stages, which accelerates DETR’s training convergence. Unlike the above two methods, our proposed SAM-DETR++ aims to reduce the matching difficulty between object query and encoded image features by enforcing aligned semantics, which encourages each object query to attend to those features with similar semantics. We demonstrate that with adequately addressed factors that obstruct convergence, Transformer-based detectors do not fall behind conventional ConvNet-based detectors [2, 4, 46] in terms of convergence speed, with even superior performance and simpler pipelines.

Relevance and Difference with Sparse R-CNN [77]. We encode instance-level information with reference boxes and object queries, which have certain similarities to the proposal boxes and proposal features in Sparse R-CNN [77]. Besides, both methods leverage RoIAlign [69] to pool region features. However, the two methods are fundamentally distinct. As a member of the R-CNN family, Sparse R-CNN directly feeds the pooled features to a heavy R-CNN head to produce region-wise detection results. In contrast, our proposed SAM-DETR++ searches and extracts objects’ salient features from the pooled region features using a lightweight network. The extracted features are fed to the

TABLE 5

Learnable reference boxes in SAM-DETR++ are not sensitive to gaps across datasets.

Ref. Box	initialization trainable?	scratch	pretrained on COCO	random
		✓	×	×
mAP@0.5 (%)		79.6	79.5	79.3

Transformer modules for global predictions with accelerated convergence.

Are Learned Reference Boxes Sensitive to Gaps across Datasets? The learned reference boxes encode statistical information for object distribution of specific datasets. To study whether these reference boxes affect generalization across datasets, we train and evaluate SAM-DETR++ w/ SMCA on Pascal VOC [84] (with notably different statistics from COCO [7]) over three setups for reference boxes: (i) learning from scratch on Pascal VOC, (ii) inheriting from COCO and remaining fixed, and (iii) remaining fixed from random initialization. Except for the reference boxes, all other parameters are trained on Pascal VOC. We use Pascal VOC trainval07+12 for training and use test07 for evaluation. Results in Table 5 show that the learned reference boxes generalize well across datasets. Even with totally random reference boxes, SAM-DETR++ can still deliver satisfactory detection accuracy. This is because (i) the reference boxes are dense enough to cover most image regions, and (ii) SAM-DETR++ involves multiple stages of box adjustment, thus initial reference boxes do not have a clear impact on final predictions.

6 CONCLUSION

This paper presents SAM-DETR++ to accelerate the convergence of DETR. The core of SAM-DETR++ is a plug-and-play module that semantically aligns object queries and encoded image features to facilitate the matching procedure between them. It also explicitly searches multiple representative keypoints with the most discriminative features for semantic-aligned matching. Besides, on the basis of semantic-aligned matching, SAM-DETR++ can further benefit from multi-scale feature fusion in a coarse-to-fine manner. By simply introducing a plug-and-play module, our proposed SAM-DETR++ accelerates DETR’s convergence from a unique perspective, and thus can be easily integrated with existing convergence solutions to boost performance to a greater extent. On the COCO benchmark, the fully-fledged SAM-DETR++ achieves 44.8% AP with only 12 training epochs, outperforming Faster R-CNN by a large margin. It also achieves state-of-the-art detection accuracy among Transformer-based detectors. We hope our work paves the way for more comprehensive research and applications of Transformer-based object detectors.

ACKNOWLEDGMENTS

This study is supported under the RIE 2020 Industry Alignment Fund – Industry Collaboration Projects (IAF-ICP) Funding Initiative, as well as cash and in-kind contribution from the industry partner(s).

REFERENCES

- [1] L. Liu, W. Ouyang, X. Wang, P. Fieguth, J. Chen, X. Liu, and M. Pietikäinen, “Deep learning for generic object detection: A survey,” *International Journal of Computer Vision*, vol. 128, pp. 261–318, 2020.
- [2] S. Ren, K. He, R. B. Girshick, and J. Sun, “Faster R-CNN: Towards real-time object detection with region proposal networks,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 39, pp. 1137–1149, 2015.
- [3] J. Redmon and A. Farhadi, “YOLO 9000: Better, faster, stronger,” in *CVPR*, 2017.
- [4] Z. Tian, C. Shen, H. Chen, and T. He, “FCOS: Fully convolutional one-stage object detection,” in *ICCV*, 2019.
- [5] N. Carion, F. Massa, G. Synnaeve, N. Usunier, A. Kirillov, and S. Zagoruyko, “End-to-end object detection with transformers,” in *ECCV*, 2020.
- [6] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, “Attention is all you need,” in *NeurIPS*, 2017.
- [7] T.-Y. Lin, M. Maire, S. J. Belongie, L. D. Bourdev, R. B. Girshick, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick, “Microsoft COCO: Common objects in context,” in *ECCV*, 2014.
- [8] T.-Y. Lin, P. Goyal, R. Girshick, K. He, and P. Dollár, “Focal loss for dense object detection,” in *ICCV*, 2017.
- [9] X. Zhu, W. Su, L. Lu, B. Li, X. Wang, and J. Dai, “Deformable DETR: Deformable transformers for end-to-end object detection,” in *ICLR*, 2021.
- [10] D. Meng, X. Chen, Z. Fan, G. Zeng, H. Li, Y. Yuan, L. Sun, and J. Wang, “Conditional DETR for fast training convergence,” in *ICCV*, 2021.
- [11] P. Gao, M. Zheng, X. Wang, J. Dai, and H. Li, “Fast convergence of DETR with spatially modulated co-attention,” in *ICCV*, 2021.
- [12] Y. Wang, X. Zhang, T. Yang, and J. Sun, “Anchor DETR: Query design for Transformer-based detector,” in *AAAI*, 2022.
- [13] G. Zhang, Z. Luo, Y. Yu, K. Cui, and S. Lu, “Accelerating DETR convergence via semantic-aligned matching,” in *CVPR*, 2022.
- [14] L. Bertinetto, J. Valmadre, J. F. Henriques, A. Vedaldi, and P. H. Torr, “Fully-convolutional siamese networks for object tracking,” in *ECCV*, 2016.
- [15] B. Li, J. Yan, W. Wu, Z. Zhu, and X. Hu, “High performance visual tracking with siamese region proposal network,” in *CVPR*, 2018.
- [16] B. Li, W. Wu, Q. Wang, F. Zhang, J. Xing, and J. Yan, “SiamRPN++: Evolution of siamese visual tracking with very deep networks,” in *CVPR*, 2019.
- [17] P. Voigtlaender, J. Luiten, P. H. Torr, and B. Leibe, “Siam R-CNN: Visual tracking by re-detection,” in *CVPR*, 2020.
- [18] N. Wang, W. Zhou, J. Wang, and H. Li, “Transformer meets tracker: Exploiting temporal context for robust visual tracking,” in *CVPR*, 2021.
- [19] X. Chen, B. Yan, J. Zhu, D. Wang, X. Yang, and H. Lu, “Transformer tracking,” in *CVPR*, 2021.
- [20] D. Chung, K. Tahboub, and E. J. Delp, “A two stream siamese convolutional neural network for person re-identification,” in *ICCV*, 2017.
- [21] M. Zheng, S. Karanam, Z. Wu, and R. J. Radke, “Re-identification with consistent attentive siamese networks,” in *CVPR*, 2019.
- [22] L. Wu, Y. Wang, J. Gao, and X. Li, “Where-and-when to look: Deep siamese attention networks for video-based person re-identification,” *IEEE Transactions on Multimedia*, vol. 21, no. 6, pp. 1412–1424, 2018.
- [23] C. Shen, Z. Jin, Y. Zhao, Z. Fu, R. Jiang, Y. Chen, and X.-S. Hua, “Deep siamese network with multi-level similarity perception for person re-identification,” in *ACM MM*, 2017.
- [24] Y. Shen, T. Xiao, H. Li, S. Yi, and X. Wang, “Learning deep neural networks for vehicle Re-ID with visual-spatio-temporal path proposals,” in *ICCV*, 2017.
- [25] G. Koch, R. Zemel, and R. Salakhutdinov, “Siamese neural networks for one-shot image recognition,” in *ICML Deep Learning Workshop*, 2015.
- [26] J. Snell, K. Swersky, and R. Zemel, “Prototypical networks for few-shot learning,” in *NeurIPS*, 2017.
- [27] F. Sung, Y. Yang, L. Zhang, T. Xiang, P. H. Torr, and T. M. Hospedales, “Learning to compare: Relation network for few-shot learning,” in *CVPR*, 2018.
- [28] T.-I. Hsieh, Y.-C. Lo, H.-T. Chen, and T.-L. Liu, “One-shot object detection with co-attention and co-excitation,” in *NeurIPS*, 2019.

- [29] G. Zhang, Z. Luo, K. Cui, and S. Lu, "Meta-DETR: Image-level few-shot object detection with inter-class correlation exploitation," *arXiv preprint arXiv:2103.11731*, 2021.
- [30] F. Li, H. Zhang, S. Liu, J. Guo, L. M. Ni, and L. Zhang, "DN-DETR: Accelerate DETR training by introducing query denoising," in *CVPR*, 2022.
- [31] Z. Cai and N. Vasconcelos, "Cascade R-CNN: High quality object detection and instance segmentation," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 43, pp. 1483–1498, 2021.
- [32] J. Pang, K. Chen, J. Shi, H. Feng, W. Ouyang, and D. Lin, "Libra R-CNN: Towards balanced learning for object detection," in *CVPR*, 2019.
- [33] L. Tychsen-Smith and L. Petersson, "Improving object localization with fitness NMS and bounded IoU loss," in *CVPR*, 2018.
- [34] H. Hu, J. Gu, Z. Zhang, J. Dai, and Y. Wei, "Relation networks for object detection," in *CVPR*, 2018.
- [35] G. Zhang, S. Lu, and W. Zhang, "CAD-Net: A context-aware detection network for objects in remote sensing imagery," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 57, no. 12, pp. 10 015–10 024, 2019.
- [36] X. Yan, Z. Chen, A. Xu, X. Wang, X. Liang, and L. Lin, "Meta R-CNN: Towards general solver for instance-level low-shot learning," in *ICCV*, 2019.
- [37] M. Liao, P. Lyu, M. He, C. Yao, W. Wu, and X. Bai, "Mask TextSpotter: An end-to-end trainable neural network for spotting text with arbitrary shapes," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 43, no. 2, pp. 532–548, 2021.
- [38] Y. Xiao and R. Marlet, "Few-shot object detection and viewpoint estimation for objects in the wild," in *ECCV*, 2020.
- [39] Q. Fan, W. Zhuo, C.-K. Tang, and Y.-W. Tai, "Few-shot object detection with attention-RPN and multi-relation detector," in *CVPR*, 2020.
- [40] X. Wang, T. E. Huang, T. Darrell, J. E. Gonzalez, and F. Yu, "Frustratingly simple few-shot object detection," in *ICML*, 2020.
- [41] W. Liu, D. Anguelov, D. Erhan, C. Szegedy, S. Reed, C.-Y. Fu, and A. C. Berg, "SSD: Single shot multibox detector," in *ECCV*, 2016.
- [42] B. Kang, Z. Liu, X. Wang, F. Yu, J. Feng, and T. Darrell, "Few-shot object detection via feature reweighting," in *ICCV*, 2019.
- [43] S. Zhang, L. Wen, X. Bian, Z. Lei, and S. Z. Li, "Single-shot refinement neural network for object detection," in *CVPR*, 2018.
- [44] S. Liu, D. Huang, and Y. Wang, "Receptive field block net for accurate and fast object detection," in *ECCV*, 2018.
- [45] Q. Zhao, T. Sheng, Y. Wang, Z. Tang, Y. Chen, L. Cai, and H. Ling, "M2Det: A single-shot object detector based on multi-level feature pyramid network," in *AAAI*, 2019.
- [46] M. Tan, R. Pang, and Q. V. Le, "EfficientDet: Scalable and efficient object detection," in *CVPR*, 2020.
- [47] X. Zhou, D. Wang, and P. Krähénbühl, "Objects as points," in *arXiv preprint arXiv:1904.07850*, 2019.
- [48] J.-M. Perez-Rua, X. Zhu, T. M. Hospedales, and T. Xiang, "Incremental few-shot object detection," in *CVPR*, 2020.
- [49] G. Zhang, K. Cui, R. Wu, S. Lu, and Y. Tian, "PNPDet: Efficient few-shot detection without forgetting via plug-and-play sub-networks," in *WACV*, 2021.
- [50] X. Zhou, J. Zhuo, and P. Krahenbuhl, "Bottom-up object detection by grouping extreme and center points," in *CVPR*, 2019.
- [51] Z. Yao, J. Ai, B. Li, and C. Zhang, "Efficient DETR: improving end-to-end object detector with dense prior," *arXiv preprint arXiv:2104.01318*, 2021.
- [52] B. Roh, J. Shin, W. Shin, and S. Kim, "Sparse DETR: Efficient end-to-end object detection with learnable sparsity," in *ICLR*, 2022.
- [53] H. Song, D. Sun, S. Chun, V. Jampani, D. Han, B. Heo, W. Kim, and M.-H. Yang, "ViDT: An efficient and effective fully transformer-based object detector," in *ICLR*, 2022.
- [54] T. Wang, L. Yuan, Y. Chen, J. Feng, and S. Yan, "PnP-DETR: Towards efficient visual analysis with Transformers," in *ICCV*, 2021.
- [55] 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," in *ICLR*, 2022.
- [56] F. Schroff, D. Kalenichenko, and J. Philbin, "FaceNet: A unified embedding for face recognition and clustering," in *CVPR*, 2015.
- [57] L. Song, D. Gong, Z. Li, C. Liu, and W. Liu, "Occlusion robust face recognition based on mask learning with pairwise differential siamese network," in *ICCV*, 2019.
- [58] R. Tao, E. Gavves, and A. W. Smeulders, "Siamese instance search for tracking," in *CVPR*, 2016.
- [59] X. Dong and J. Shen, "Triplet loss in siamese network for object tracking," in *ECCV*, 2018.
- [60] A. He, C. Luo, X. Tian, and W. Zeng, "A twofold siamese network for real-time object tracking," in *CVPR*, 2018.
- [61] Z. Zhu, Q. Wang, B. Li, W. Wu, J. Yan, and W. Hu, "Distractor-aware siamese networks for visual object tracking," in *ECCV*, 2018.
- [62] Z. Zhang and H. Peng, "Deeper and wider siamese networks for real-time visual tracking," in *CVPR*, 2019.
- [63] S. He, H. Luo, P. Wang, F. Wang, H. Li, and W. Jiang, "TransReID: Transformer-based object re-identification," in *ICCV*, 2021.
- [64] G. Zhang, Z. Luo, K. Cui, and S. Lu, "Meta-DETR: Few-shot object detection via unified image-level meta-learning," *arXiv preprint arXiv:2103.11731v2*, 2021.
- [65] J. Zhang, J. Huang, Z. Luo, G. Zhang, and S. Lu, "DA-DETR: Domain adaptive detection transformer by hybrid attention," *arXiv preprint arXiv:2103.17084*, 2021.
- [66] X. Cao, P. Yuan, B. Feng, and K. Niu, "CF-DETR: Coarse-to-fine transformers for end-to-end object detection," in *AAAI*, 2022.
- [67] X. Dai, Y. Chen, J. Yang, P. Zhang, L. Yuan, and L. Zhang, "Dynamic DETR: End-to-end object detection with dynamic attention," in *ICCV*, 2021.
- [68] 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.
- [69] K. He, G. Gkioxari, P. Dollár, and R. Girshick, "Mask R-CNN," in *ICCV*, 2017.
- [70] Z. Yang, S. Liu, H. Hu, L. Wang, and S. Lin, "RepPoints: Point set representation for object detection," in *ICCV*, 2019.
- [71] R. Wu, G. Zhang, S. Lu, and T. Chen, "Cascade EF-GAN: Progressive facial expression editing with local focuses," in *CVPR*, 2020.
- [72] Y. Chen, Z. Zhang, Y. Cao, L. Wang, S. Lin, and H. Hu, "RepPoints v2: Verification meets regression for object detection," in *NeurIPS*, 2020.
- [73] G. Zhang, K. Cui, T.-Y. Hung, and S. Lu, "Defect-GAN: High-fidelity defect synthesis for automated defect inspection," in *WACV*, 2021.
- [74] G. Zhang, Z. Luo, K. Cui, and S. Lu, "Meta-DETR: Image-level few-shot object detection with inter-class correlation exploitation," *arXiv preprint arXiv:2103.11731v3*, 2021.
- [75] T.-Y. Lin, P. Dollár, R. Girshick, K. He, B. Hariharan, and S. Belongie, "Feature pyramid networks for object detection," in *CVPR*, 2017.
- [76] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov, "Dropout: a simple way to prevent neural networks from overfitting," *The Journal of Machine Learning Research (JMLR)*, vol. 15, no. 1, pp. 1929–1958, 2014.
- [77] P. Sun, R. Zhang, Y. Jiang, T. Kong, C. Xu, W. Zhan, M. Tomizuka, L. Li, Z. Yuan, C. Wang, and P. Luo, "Sparse R-CNN: End-to-end object detection with learnable proposals," in *CVPR*, 2021.
- [78] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, "ImageNet: A large-scale hierarchical image database," in *CVPR*, 2009.
- [79] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *CVPR*, 2016.
- [80] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," in *ICLR*, 2015.
- [81] I. Loshchilov and F. Hutter, "Decoupled weight decay regularization," in *ICLR*, 2019.
- [82] Z. Sun, S. Cao, Y. Yang, and K. M. Kitani, "Rethinking Transformer-based set prediction for object detection," in *ICCV*, 2021.
- [83] Z. Dai, B. Cai, Y. Lin, and J. Chen, "UP-DETR: Unsupervised pre-training for object detection with transformers," in *CVPR*, 2021.
- [84] M. Everingham, L. Van Gool, C. K. I. Williams, J. Winn, and A. Zisserman, "The Pascal Visual Object Classes (VOC) Challenge," *International Journal of Computer Vision*, vol. 88, no. 2, pp. 303–338, 2010.