

# SAM-6D: Segment Anything Model Meets Zero-Shot 6D Object Pose Estimation

Jiehong Lin<sup>1,2\*</sup>

Lihua Liu<sup>3\*</sup>

Dekun Lu<sup>3</sup>

Kui Jia<sup>2†</sup>

<sup>1</sup>DexForce Co. Ltd., Shenzhen

<sup>2</sup>School of Data Science, The Chinese University of Hong Kong, Shenzhen

<sup>3</sup>South China University of Technology, Guangzhou

<https://github.com/JiehongLin/SAM-6D>

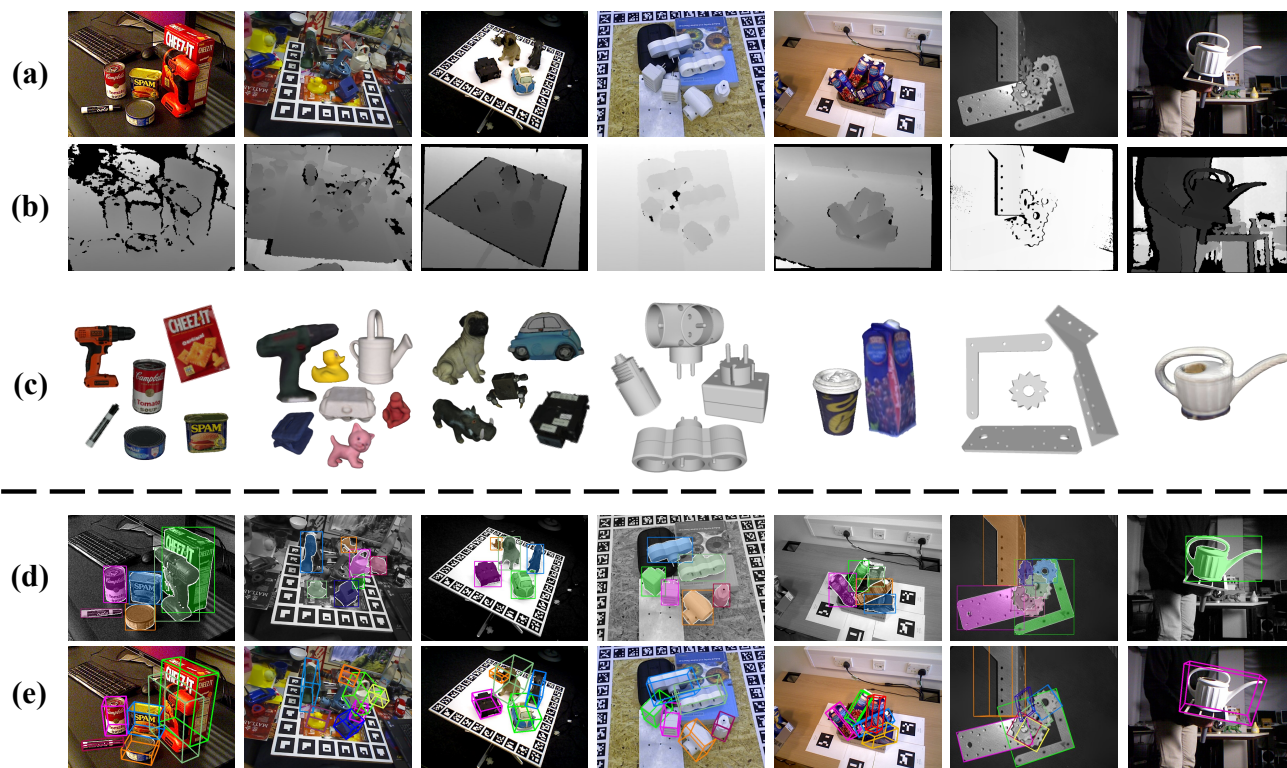


Figure 1. We present SAM-6D for zero-shot 6D object pose estimation. SAM-6D takes an RGB image (a) and a depth map (b) of a cluttered scene as inputs, and performs instance segmentation (d) and pose estimation (e) for novel objects (c). We present the qualitative results of SAM-6D on the seven core datasets of the BOP benchmark [54], including YCB-V, LM-O, HB, T-LESS, IC-BIN, ITODD and TUD-L, arranged from left to right. Best view in the electronic version.

## Abstract

Zero-shot 6D object pose estimation involves the detection of novel objects with their 6D poses in cluttered scenes, presenting significant challenges for model generalizability. Fortunately, the recent Segment Anything Model (SAM) has showcased remarkable zero-shot transfer performance, which provides a promising solution to tackle this task. Motivated by this, we introduce SAM-6D, a novel framework designed to realize the task through two steps, including in-

stance segmentation and pose estimation. Given the target objects, SAM-6D employs two dedicated sub-networks, namely Instance Segmentation Model (ISM) and Pose Estimation Model (PEM), to perform these steps on cluttered RGB-D images. ISM takes SAM as an advanced starting point to generate all possible object proposals and selectively preserves valid ones through meticulously crafted object matching scores in terms of semantics, appearance and geometry. By treating pose estimation as a partial-to-partial point matching problem, PEM performs a two-stage point matching process featuring a novel design of background tokens to construct dense 3D-3D correspondence,

\* Equal contribution. † Corresponding author <kuijia@gmail.com>.

ultimately yielding the pose estimates. Without bells and whistles, SAM-6D outperforms the existing methods on the seven core datasets of the BOP Benchmark for both instance segmentation and pose estimation of novel objects.

## 1. Introduction

Object pose estimation is fundamental in many real-world applications, such as robotic manipulation and augmented reality. Its evolution has been significantly influenced by the emergence of deep learning models. The most studied task in this field is *Instance-level 6D Pose Estimation* [18, 19, 51, 58, 60, 63], which demands annotated training images of the target objects, thereby making the deep models object-specific. Recently, the research emphasis gradually shifts towards the task of *Category-level 6D Pose Estimation* [7, 29–32, 56, 61] for handling unseen objects, yet provided they belong to certain categories of interest. In this paper, we thus delve into a broader task setting of *Zero-shot 6D Object Pose Estimation* [5, 28], which aspires to detect all instances of novel objects, unseen during training, and estimate their 6D poses. Despite its significance, this zero-shot setting presents considerable challenges in both object detection and pose estimation.

Recently, Segment Anything Model (SAM) [26] has garnered attention due to its remarkable zero-shot segmentation performance, which enables prompt segmentation with a variety of prompts, *e.g.*, points, boxes, texts or masks. By prompting SAM with evenly sampled 2D grid points, one can generate potential class-agnostic object proposals, which may be highly beneficial for zero-shot 6D object pose estimation. To this end, we propose a novel framework, named **SAM-6D**, which employs SAM as an advanced starting point for the focused zero-shot task. Fig. 2 gives an overview illustration of SAM-6D. Specifically, SAM-6D employs an **Instance Segmentation Model (ISM)** to realize instance segmentation of novel objects by enhancing SAM with a carefully crafted object matching score, and a **Pose Estimation Model (PEM)** to solve object poses through a two-stage process of partial-to-partial point matching.

The Instance Segmentation Model (ISM) is developed using SAM to take advantage of its zero-shot abilities for generating all possible class-agnostic proposals, and then assigns a meticulously calculated object matching score to each proposal for ascertaining whether it aligns with a given novel object. In contrast to methods that solely focus on object semantics [5, 40], we design the object matching scores considering three terms, including semantics, appearance and geometry. For each proposal, the first term assesses its semantic matching degree to the rendered templates of the object, while the second one further evaluates its appearance similarities to the best-matched template. The final term considers the matching degree based on ge-

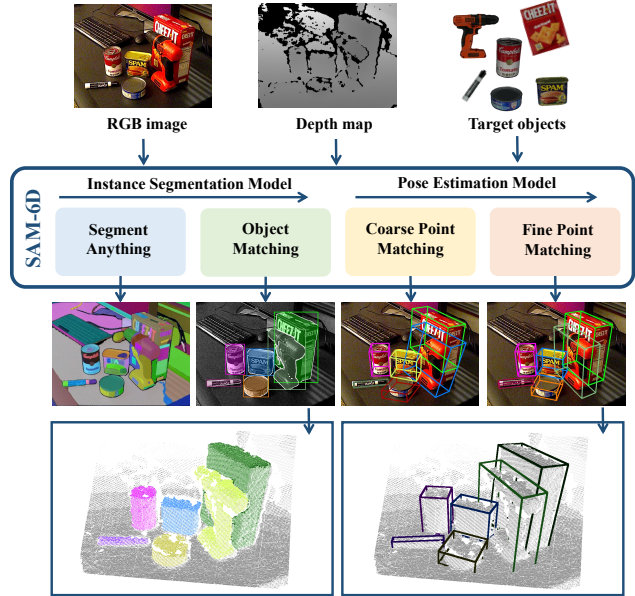


Figure 2. An overview of our proposed SAM-6D, which consists of an Instance Segmentation Model (ISM) and a Pose Estimation Model (PEM) for joint instance segmentation and pose estimation of novel objects in RGB-D images. ISM leverages the Segment Anything Model (SAM) [26] to generate all possible proposals and selectively retains valid ones based on object matching scores. PEM involves two stages of point matching, from coarse to fine, to establish 3D-3D correspondence and calculate object poses for all valid proposals. Best view in the electronic version.

ometry, such as object shape and size, by calculating the Intersection-over-Union (IoU) value between the bounding boxes of the proposal and the 2D projection of the object transformed by a rough pose estimate.

The Pose Estimation Model (PEM) is designed to calculate a 6D object pose for each identified proposal that matches the novel object. Initially, we formulate this pose estimation challenge as a partial-to-partial point matching problem between the sampled point sets of the proposal and the target object, considering the factors such as occlusions, segmentation inaccuracies, and sensor noises. To solve this problem, we propose a simple yet effective solution that involves the use of background tokens; specifically, for the two point sets, we learn to align their non-overlapped points with the background tokens in the feature space, and thus effectively establish an assignment matrix to build the necessary correspondence for predicting the object pose. Based on the design of background tokens, we further develop PEM with two point matching stages, *i.e.*, **Coarse Point Matching** and **Fine Point Matching**. The first stage realizes sparse correspondence to derive an initial object pose, which is subsequently used to transform the point set of the proposal, enabling the learning of positional encodings. The second stage incorporates the positional encodings of the two point sets to inject the initial correspondence, and

builds dense correspondence for estimating a more precise object pose. To effectively model dense interactions in the second stage, we propose an innovative design of Sparse-to-Dense Point Transformers, which realize interactions on the sparse versions of the dense features, and subsequently, distribute the enhanced sparse features back to the dense ones using Linear Transformers [12, 24].

For the two models of SAM-6D, ISM, built on SAM, does not require any network re-training or fine-tuning, while PEM is trained on the large-scale synthetic images of ShapeNet-Objects [4] and Google-Scanned-Objects [9] datasets provided by [28]. We evaluate SAM-6D on the seven core datasets of the BOP benchmark [54], including LM-O, T-LESS, TUD-L, IC-BIN, ITODD, HB, and YCB-V. The qualitative results are visualized in Fig. 1. SAM-6D outperforms the existing methods on both tasks of instance segmentation and pose estimation of novel objects, thereby showcasing its robust generalization capabilities.

Our main contributions could be summarized as follows:

- We propose a novel framework of SAM-6D, which realizes joint instance segmentation and pose estimation of novel objects from RGB-D images, and outperforms the existing methods on seven datasets of BOP benchmark.
- We leverage the zero-shot capacities of Segmentation Anything Model (SAM) to generate all possible proposals, and devise a novel object matching score to identify the proposals corresponding to novel objects.
- We approach pose estimation as a partial-to-partial point matching problem with a simple yet effective design of background tokens, and propose a two-stage point matching model for novel objects. The first stage realizes coarse point matching to derive initial object poses, which are then refined in the second stage of fine point matching using newly proposed Sparse-to-Dense Point Transformers.

## 2. Related Work

### 2.1. Segment Anything

Segment Anything (SA) [26], is a promptable segmentation task that focuses on predicting valid masks for various types of prompts, *e.g.*, points, boxes, text, and masks. To tackle this task, the authors propose a powerful segmentation model called Segment Anything Model (SAM), which comprises three components, including an image encoder, a prompt encoder and a mask decoder. SAM has demonstrated remarkable zero-shot transfer segmentation performance in real-world scenarios, including challenging situations such as medical images [36, 37, 71], camouflaged objects [22, 55], and transparent objects [13, 23]. Moreover, SAM has exhibited high versatility across numerous vision applications [69], such as image inpainting [35, 62, 64, 67], object tracking [15, 65, 73], 3D detection and segmentation [2, 66, 70], and 3D reconstruction [3, 49, 59].

Recent studies have also investigated semantically segmenting anything due to the critical role of semantics in vision tasks. Semantic Segment Anything (SSA) [6] is proposed on top of SAM, aiming to assign semantic categories to the masks generated by SAM. Both PerSAM [72] and Matcher [34] employ SAM to segment the object belonging to a specific category in a query image by searching for point prompts with the aid of a reference image containing an object of the same category. CNOS [40] is proposed to segment all instances of a given object model, which firstly generates mask proposals via SAM and subsequently filters out proposals with low feature similarities against object templates rendered from the object model.

For efficiency, FastSAM [74] is proposed by utilizing instance segmentation networks with regular convolutional networks instead of visual transformers used in SAM. Additionally, MobileSAM [68] replaces the heavy encoder of SAM with a lightweight one through decoupled distillation.

### 2.2. Pose Estimation of Novel Objects

**Methods Based on Image Matching** Methods within this group [1, 28, 33, 38, 39, 41, 42, 46, 50] often involve comparing object proposals to templates of the given novel objects, which are rendered with a series of object poses, to retrieve the best-matched object poses. For example, Gen6D [33], OVE6D [1], and GigaPose [41] are designed to select the viewpoint rotations via image matching and then estimate the in-plane rotations to obtain the final estimates. MegaPose [28] employs a coarse estimator to treat image matching as a classification problem, of which the recognized object poses are further updated by a refiner.

**Methods Based on Feature Matching** Methods within this group [5, 10, 11, 17, 20, 53] align the 2D pixels or 3D points of the proposals with the object surface in the feature space [21, 52], thereby building correspondence to compute object poses. OnePose [53] matches the pixel descriptors of the proposals with the aggregated point descriptors of the point sets constructed by Structure from Motion (SfM) for 2D-3D correspondence, while OnePose++ [17] further improves it with a keypoint-free SfM and a sparse-to-dense 2D-3D matching model. ZeroPose [5] realizes 3D-3D matching via geometric structures, and GigaPose [41] establishes 2D-2D correspondence to regress in-plane rotation and 2D scale. Moreover, [11] introduces a zero-shot category-level 6D pose estimation task, along with a self-supervised semantic correspondence learning method. Unlike the above one-stage point matching work, the unique contributions in our Pose Estimation Model are: (a) a two-stage pipeline that boosts performance by incorporating coarse correspondence for finer matching, (b) an efficient design of background tokens to eliminate the need of optimal transport with iterative optimization [48], and (c) a Sparse-to-Dense Point Transformer to effectively model dense relationship.

### 3. Methodology of SAM-6D

We present **SAM-6D** for zero-shot 6D object pose estimation, which aims to detect all instances of a specific novel object, unseen during training, along with their 6D object poses in the RGB-D images. To realize the challenging task, SAM-6D breaks it down into two steps via two dedicated sub-networks, *i.e.*, an **Instance Segmentation Model (ISM)** and a **Pose Estimation Model (PEM)**, to first segment all instances and then individually predict their 6D poses, as shown in Fig. 2. We detail the architectures of ISM and PEM in Sec. 3.1 and Sec. 3.2, respectively.

#### 3.1. Instance Segmentation Model

SAM-6D uses an Instance Segmentation Model (ISM) to segment the instances of a novel object  $\mathcal{O}$ . Given a cluttered scene, represented by an RGB image  $\mathcal{I}$ , ISM leverages the zero-shot transfer capabilities of Segment Anything Model (SAM) [26] to generate all possible proposals  $\mathcal{M}$ . For each proposal  $m \in \mathcal{M}$ , ISM calculates an object matching score  $s_m$  to assess the matching degree between  $m$  and  $\mathcal{O}$  in terms of semantics, appearance, and geometry. The matched instances with  $\mathcal{O}$  can then be identified by simply setting a matching threshold  $\delta_m$ .

In this subsection, we initially provide a brief review of SAM in Sec. 3.1.1 and then explain the computation of the object matching score  $s_m$  in Sec. 3.1.2.

##### 3.1.1 Preliminaries of Segment Anything Model

Given an RGB image  $\mathcal{I}$ , Segment Anything Model (SAM) [26] realizes promptable segmentation with various types of prompts  $\mathcal{P}_r$ , *e.g.*, points, boxes, texts, or masks. Specifically, SAM consists of three modules, including an image encoder  $\Phi_{\text{Image}}$ , a prompt encoder  $\Phi_{\text{Prompt}}$ , and a mask decoder  $\Psi_{\text{Mask}}$ , which could be formulated as follows:

$$\mathcal{M}, \mathcal{C} = \Psi_{\text{Mask}}(\Phi_{\text{Image}}(\mathcal{I}), \Phi_{\text{Prompt}}(\mathcal{P}_r)), \quad (1)$$

where  $\mathcal{M}$  and  $\mathcal{C}$  denote the predicted proposals and the corresponding confidence scores, respectively.

To realize zero-shot transfer, one can prompt SAM with evenly sampled 2D grids to yield all possible proposals, which can then be filtered based on confidence scores, retaining only those with higher scores, and applied to Non-Maximum Suppression to eliminate redundant detections.

##### 3.1.2 Object Matching Score

Given the proposals  $\mathcal{M}$ , the next step is to identify the ones that are matched with a specified object  $\mathcal{O}$  by assigning each proposal  $m \in \mathcal{M}$  with an object matching score  $s_m$ , which comprises three terms, each evaluating the matches in terms of semantics, appearance, and geometry, respectively.

Following [40], we sample  $N_{\mathcal{T}}$  object poses in SE(3) space to render the templates  $\{\mathcal{T}_k\}_{k=1}^{N_{\mathcal{T}}}$  of  $\mathcal{O}$ , which are fed

into a pre-trained visual transformer (ViT) backbone [8] of DINOv2 [45], resulting in the class embedding  $\mathbf{f}_{\mathcal{T}_k}^{cls}$  and

$N_{\mathcal{T}_k}^{patch}$  patch embeddings  $\{\mathbf{f}_{\mathcal{T}_k,i}^{patch}\}_{i=1}^{N_{\mathcal{T}_k}^{patch}}$  of each template  $\mathcal{T}_k$ . For each proposal  $m$ , we crop the detected region out from  $\mathcal{I}$ , and resize it to a fixed resolution. The image crop is denoted as  $\mathcal{I}_m$  and also processed through the same ViT to obtain the class embedding  $\mathbf{f}_{\mathcal{I}_m}^{cls}$  and the patch embeddings  $\{\mathbf{f}_{\mathcal{I}_m,j}^{patch}\}_{j=1}^{N_{\mathcal{I}_m}^{patch}}$ , with  $N_{\mathcal{I}_m}^{patch}$  denoting the number of patches within the object mask. Subsequently, we calculate the values of the individual score terms.

**Semantic Matching Score** We compute a semantic score  $s_{sem}$  through the class embeddings by averaging the top  $K$  values from  $\{\frac{\langle \mathbf{f}_{\mathcal{I}_m}^{cls}, \mathbf{f}_{\mathcal{T}_k}^{cls} \rangle}{|\mathbf{f}_{\mathcal{I}_m}^{cls}| \cdot |\mathbf{f}_{\mathcal{T}_k}^{cls}|}\}_{k=1}^{N_{\mathcal{T}}}$  to establish a robust measure of semantic matching, with  $\langle, \rangle$  denoting an inner product. The template that yields the highest semantic value can be seen as the best-matched template, denoted as  $\mathcal{T}_{best}$ , and is used in the computation of the subsequent two scores.

**Appearance Matching Score** Given  $\mathcal{T}_{best}$ , we compare  $\mathcal{I}_m$  and  $\mathcal{T}_{best}$  in terms of appearance using an appearance score  $s_{appe}$ , based on the patch embeddings, as follows:

$$s_{appe} = \frac{1}{N_{\mathcal{I}_m}^{patch}} \sum_{j=1}^{N_{\mathcal{I}_m}^{patch}} \max_{i=1, \dots, N_{\mathcal{T}_{best}}^{patch}} \frac{\langle \mathbf{f}_{\mathcal{I}_m,j}^{patch}, \mathbf{f}_{\mathcal{T}_{best},i}^{patch} \rangle}{|\mathbf{f}_{\mathcal{I}_m,j}^{patch}| \cdot |\mathbf{f}_{\mathcal{T}_{best},i}^{patch}|}. \quad (2)$$

$s_{appe}$  is utilized to distinguish objects that are semantically similar but differ in appearance.

**Geometric Matching Score** In terms of geometry, we score the proposal  $m$  by considering factors like object shapes and sizes. Utilizing the object rotation from  $\mathcal{T}_{best}$  and the mean location of the cropped points of  $m$ , we have a coarse pose to transform the object  $\mathcal{O}$ , which is then projected onto the image to obtain a compact bounding box  $\mathcal{B}_o$ . Afterwards, the Intersection-over-Union (IoU) value between  $\mathcal{B}_o$  and the bounding box  $\mathcal{B}_m$  of  $m$  is used as the geometric score  $s_{geo}$ :

$$s_{geo} = \frac{\mathcal{B}_m \cap \mathcal{B}_o}{\mathcal{B}_m \cup \mathcal{B}_o}. \quad (3)$$

The reliability of  $s_{geo}$  is easily impacted by occlusions. We thus compute a visible ratio  $r_{vis}$  to evaluate the confidence of  $s_{geo}$ , which is detailed in the supplementary materials.

By combining the above three score terms, the object matching score  $s_m$  could be formulated as follows:

$$s_m = \frac{s_{sem} + s_{appe} + r_{vis} \cdot s_{geo}}{1 + 1 + r_{vis}}. \quad (4)$$

#### 3.2. Pose Estimation Model

SAM-6D uses a Pose Estimation Model (PEM) to predict the 6D poses of the proposals matched with the object  $\mathcal{O}$ .

For each object proposal  $m$ , PEM uses a strategy of point registration to predict the 6D pose w.r.t.  $\mathcal{O}$ . Denoting the sampled point set of  $m$  as  $\mathcal{P}_m \in \mathbb{R}^{N_m \times 3}$  with  $N_m$  points

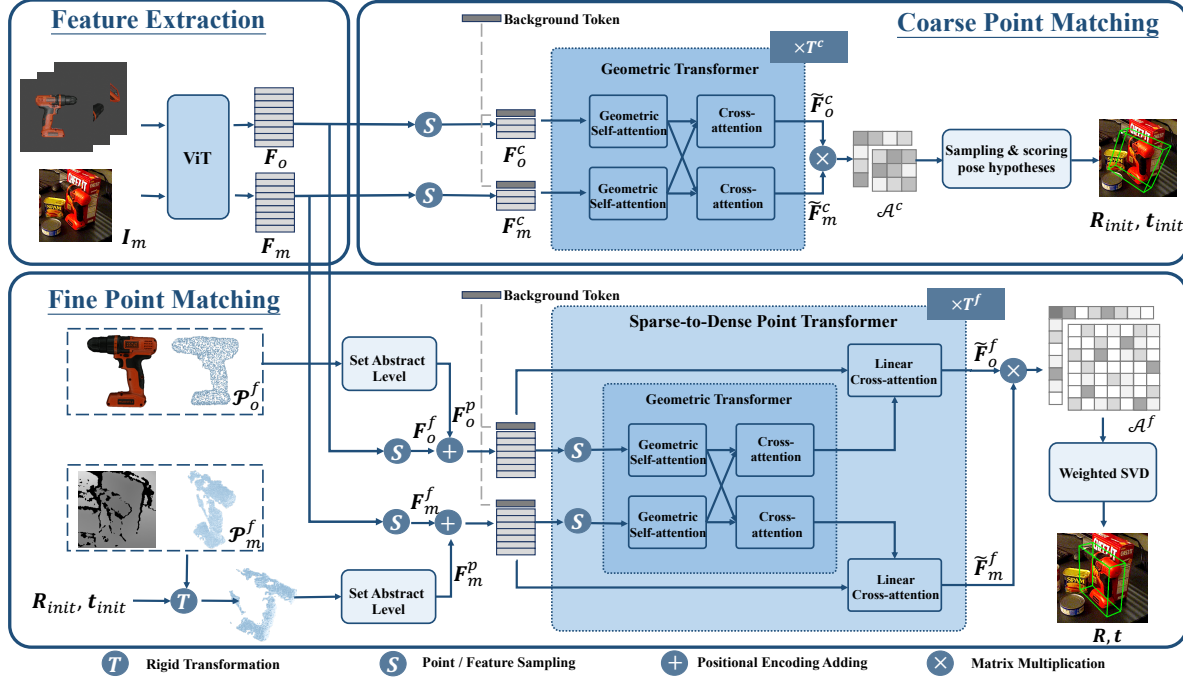


Figure 3. An illustration of Pose Estimation Model (PEM) of SAM-6D.

and that of  $\mathcal{O}$  as  $\mathcal{P}_o \in \mathbb{R}^{N_o \times 3}$  with  $N_o$  points, the goal is to solve an assignment matrix to present the partial-to-partial correspondence between  $\mathcal{P}_m$  and  $\mathcal{P}_o$ . Partial-to-partial correspondence arises as  $\mathcal{P}_o$  only partially matches  $\mathcal{P}_m$  due to occlusions, and  $\mathcal{P}_m$  may partially align with  $\mathcal{P}_o$  due to segmentation inaccuracies and sensor noises. We propose to equip their respective point features  $\mathbf{F}_m \in \mathbb{R}^{N_m \times C}$  and  $\mathbf{F}_o \in \mathbb{R}^{N_o \times C}$  with learnable **Background Tokens**, denoted as  $\mathbf{f}_m^{bg} \in \mathbb{R}^C$  and  $\mathbf{f}_o^{bg} \in \mathbb{R}^C$ , where  $C$  is the number of feature channels. This simple design resolves the assignment problem of non-overlapped points in two point sets, and the partial-to-partial correspondence thus could be effectively built based on feature similarities. Specifically, we can first compute the attention matrix  $\mathcal{A}$  as follows:

$$\mathcal{A} = [\mathbf{f}_m^{bg}, \mathbf{F}_m] \times [\mathbf{f}_o^{bg}, \mathbf{F}_o]^T \in \mathbb{R}^{(N_m+1) \times (N_o+1)}, \quad (5)$$

and then obtain the soft assignment matrix  $\tilde{\mathcal{A}}$

$$\tilde{\mathcal{A}} = \text{Softmax}_{\text{row}}(\mathcal{A}/\tau) \cdot \text{Softmax}_{\text{col}}(\mathcal{A}/\tau), \quad (6)$$

where  $\text{Softmax}_{\text{row}}()$  and  $\text{Softmax}_{\text{col}}()$  denote Softmax operations executed along the row and column of the matrix, respectively.  $\tau$  is a constant temperature. The values in each row of  $\tilde{\mathcal{A}}$ , excluding the first row associated with the background, indicate the matching probabilities of the point  $\mathbf{p}_m \in \mathcal{P}_m$  aligning with background and the points in  $\mathcal{P}_o$ . Specifically, for  $\mathbf{p}_m$ , its corresponding point  $\mathbf{p}_o \in \mathcal{P}_o$  can be identified by locating the index of the maximum score  $\tilde{a} \in \tilde{\mathcal{A}}$  along the row; if this index equals zero, the embedding of  $\mathbf{p}_m$  aligns with the background token, indicating it

has no valid correspondence in  $\mathcal{P}_o$ . Once  $\tilde{\mathcal{A}}$  is obtained, we can gather all the matched pairs  $\{(\mathbf{p}_m, \mathbf{p}_o)\}$ , along with their scores  $\{\tilde{a}\}$ , to compute the pose using weighted SVD.

Building on the above strategy with background tokens, PEM is designed in two point matching stages. For the proposal  $m$  and the target object  $\mathcal{O}$ , the first stage involves **Coarse Point Matching** between their sparse point sets  $\mathcal{P}_m^c$  and  $\mathcal{P}_o^c$ , while the second stage involves **Fine Point Matching** between their dense sets  $\mathcal{P}_m^f$  and  $\mathcal{P}_o^f$ ; we use the upper scripts ‘c’ and ‘f’ to indicate respective variables of these two stages. The aim of the first stage is to derive a coarse pose  $\mathbf{R}_{init}$  and  $\mathbf{t}_{init}$  from sparse correspondence. Then in the second stage, we use the initial pose to transform  $\mathcal{P}_m^f$  for learning the positional encodings, and employ stacked **Sparse-to-Dense Point Transformers** to learn dense correspondence for a final pose  $\mathbf{R}$  and  $\mathbf{t}$ . Prior to two point matching modules, we incorporate a **Feature Extraction** module to learn individual point features of  $m$  and  $\mathcal{O}$ . Fig. 3 gives a detailed illustration of PEM.

### 3.2.1 Feature Extraction

The Feature Extraction module is employed to extract point-wise features  $\mathbf{F}_m$  and  $\mathbf{F}_o$  for the point sets  $\mathcal{P}_m$  and  $\mathcal{P}_o$  of the proposal  $m$  and the given object  $\mathcal{O}$ , respectively.

Rather than directly extracting features from the discretized points  $\mathcal{P}_m$ , we utilize the visual transformer (ViT) backbone [8] on its masked image crop  $\mathcal{I}_m$  to capture patch-wise embeddings, which are then reshaped and interpolated to match the size of  $\mathcal{I}_m$ . Each point in  $\mathcal{P}_m$  is assigned with the corresponding pixel embedding, yielding  $\mathbf{F}_m$ .

We represent the object  $\mathcal{O}$  with templates rendered from different camera views. All visible object pixels (points) are aggregated across views and sampled to create the point set  $\mathcal{P}_o$ . Corresponding pixel embeddings, extracted using the ViT backbone, are then used to form  $F_o$ .

### 3.2.2 Coarse Point Matching

The Coarse Point Matching module is used to initialize a coarse object pose  $R_{init}$  and  $t_{init}$  by estimating a soft assignment matrix  $\tilde{S}^c$  between sparse versions of  $\mathcal{P}_m$  and  $\mathcal{P}_o$ .

As shown in Fig. 3, we first sample a sparse point set  $\mathcal{P}_m^c \in \mathbb{R}^{N_m^c \times 3}$  with  $N_m^c$  points from  $\mathcal{P}_m$ , and  $\mathcal{P}_o^c \in \mathbb{R}^{N_o^c \times 3}$  with  $N_o^c$  points from  $\mathcal{P}_o$ , along with their respective sampled features  $F_m^c$  and  $F_o^c$ . Then we concatenate  $F_m^c$  and  $F_o^c$  with learnable background tokens, and process them through  $T^c$  stacked Geometric Transformers [48], each of which consists of a geometric self-attention for intra-point-set feature learning and a cross-attention for inter-point-set correspondence modeling. The processed features, denoted as  $\tilde{F}_m^c$  and  $\tilde{F}_o^c$ , are subsequently used to compute the soft assignment matrix  $\tilde{A}^c$  based on (5) and (6).

With  $\tilde{A}^c$ , we obtain the matching probabilities between the overlapped points of  $\mathcal{P}_m^c$  and  $\mathcal{P}_o^c$ , which can serve as the distribution to sample multiple triplets of point pairs and compute pose hypotheses [14, 25]. We assign each pose hypothesis  $R_{hyp}$  and  $t_{hyp}$  a pose matching score  $s_{hyp}$  as:

$$s_{hyp} = N_m^c / \sum_{\mathbf{p}_m^c \in \mathcal{P}_m^c} \min_{\mathbf{p}_o^c \in \mathcal{P}_o^c} \|\mathbf{R}_{hyp}^T(\mathbf{p}_o^c - \mathbf{t}_{hyp}) - \mathbf{p}_m^c\|_2. \quad (7)$$

Among the pose hypotheses, the one with the highest pose matching score is chosen as the initial pose  $R_{init}$  and  $t_{init}$  inputted into the next Fine Point Matching module.

### 3.2.3 Fine Point Matching

The Fine Point Matching module is utilized to build dense correspondence and estimate a more precise pose  $R$  and  $t$ .

To build finer correspondence, we sample a dense point set  $\mathcal{P}_m^f \in \mathbb{R}^{N_m^f \times 3}$  with  $N_m^f$  points from  $\mathcal{P}_m$ , and  $\mathcal{P}_o^f \in \mathbb{R}^{N_o^f \times 3}$  with  $N_o^f$  points from  $\mathcal{P}_o$ , along with their respective sampled features  $F_m^f$  and  $F_o^f$ . We then inject the initial correspondence, learned by the coarse point matching, through the inclusion of positional encodings. Specifically, we transform  $\mathcal{P}_m^f$  with the coarse pose  $R_{init}$  and  $t_{init}$  and apply it to a multi-scale Set Abstract Level [47] to learn the positional encodings  $F_m^p$ ; similarly, positional encodings  $F_o^p$  are also learned for  $\mathcal{P}_o^f$ . We then add  $F_m^p$  and  $F_o^p$  to  $F_m^f$  and  $F_o^f$ , concatenate each with a background token, and process them to yield  $\tilde{F}_m^f$  and  $\tilde{F}_o^f$ , resulting in the soft assignment matrix  $\tilde{A}^f$  based on (5) and (6).

However, the commonly used transformers [48, 57] incur a significant computational cost when learning dense point features. The recent Linear Transformers [12, 24], while being more efficient, exhibit less effective modeling of point

interactions, since they implement attentions along the feature dimension. To address this, we propose a novel design of **Sparse-to-Dense Point Transformer** (SDPT), as shown in Fig. 3. Specifically, given two dense point features  $F_m^f$  and  $F_o^f$ , SDPT first samples two sparse features from them and applies a Geometric Transformer [48] to enhance their interactions, resulting in two improved sparse features, denoted as  $F_m^{f'}$  and  $F_o^{f'}$ . SDPT then employs a Linear Cross-attention [12] to spread the information from  $F_m^{f'}$  to  $F_o^{f'}$ , treating the former as the key and value of the transformer, and the latter as the query. The same operations are applied to  $F_o^{f'}$  and  $F_m^{f'}$  to update  $F_o^f$ .

In Fine Point Matching, we stack  $T^f$  SDPTs to model the dense correspondence and learn the soft assignment matrix  $\tilde{A}^f$ . We note that, in each SDPT, the background tokens are consistently maintained in both sparse and dense point features. After obtaining  $\tilde{A}^f$ , we search within  $\mathcal{P}_o^f$  for the corresponding points to all foreground points in  $\mathcal{P}_m^f$  along with the probabilities, building dense correspondence, and compute the final object pose  $R$  and  $t$  via weighted SVD.

## 4. Experiments

In this section, we conduct experiments to evaluate our proposed SAM-6D, which consists of an Instance Segmentation Model (ISM) and a Pose Estimation Model (PEM).

**Datasets** We evaluate our proposed SAM-6D on the seven core datasets of the BOP benchmark [54], including LM-O, T-LESS, TUD-L, IC-BIN, ITODD, HB, and YCB-V. PEM is trained on the large-scale synthetic ShapeNet-Objects [4] and Google-Scanned-Objects [9] datasets provided by [28], with a total of 2,000,000 images across  $\sim 50,000$  objects.

**Implementation Details** For ISM, we follow [40] to utilize the default ViT-H SAM [26] or FastSAM [74] for proposal generation, and the default ViT-L model of DINOv2 [44] to extract class and patch embeddings. For PEM, we set  $N_m^c = N_o^c = 196$  and  $N_m^f = N_o^f = 2048$ , and use InfoNCE loss [43] to supervise the learning of attention matrices (5) for both matching stages. We use ADAM to train PEM with a total of 600,000 iterations; the learning rate is initialized as 0.0001, with a cosine annealing schedule used, and the batch size is set as 28. For each object, we use two rendered templates for training PEM. During evaluation, we follow [40] and use 42 templates for both ISM and PEM.

**Evaluation Metrics** For instance segmentation, we report the mean Average Precision (mAP) scores at different Intersection-over-Union (IoU) thresholds ranging from 0.50 to 0.95 with a step size of 0.05. For pose estimation, we report the mean Average Recall (AR) w.r.t three error functions, *i.e.*, Visible Surface Discrepancy (VSD), Maximum Symmetry-Aware Surface Distance (MSSD) and Maximum Symmetry-Aware Projection Distance (MSPD). For further details about these evaluation metrics, please refer to [54].

Method	Segmentation Model	Object Matching Score			BOP Dataset							Mean
		$s_{sem}$	$s_{appe}$	$s_{geo}$	LM-O	T-LESS	TUD-L	IC-BIN	ITODD	HB	YCB-V	
ZeroPose [5]	SAM [26]	-	-	-	34.4	32.7	41.4	25.1	22.4	47.8	51.9	36.5
CNOS [40]	FastSAM [74]	-	-	-	39.7	37.4	48.0	27.0	25.4	51.1	59.9	41.2
CNOS [40]	SAM [26]	-	-	-	39.6	39.7	39.1	28.4	28.2	48.0	59.5	40.4
SAM-6D (Ours)	FastSAM [74]	✓	×	×	39.5	37.6	48.7	25.7	25.3	51.2	60.2	41.2
		✓	✓	×	40.6	39.3	50.1	27.7	29.0	52.2	60.6	42.8
		✓	×	✓	40.4	41.4	49.7	28.2	30.1	54.0	61.1	43.6
		✓	✓	✓	42.2	42.0	51.7	29.3	31.9	54.8	<b>62.1</b>	44.9
	SAM [26]	✓	×	×	43.4	39.1	48.2	33.3	28.8	55.1	60.3	44.0
		✓	✓	×	44.4	40.8	49.8	34.5	30.0	55.7	59.5	45.0
		✓	×	✓	44.0	44.7	54.8	33.8	31.5	58.3	59.9	46.7
		✓	✓	✓	<b>46.0</b>	<b>45.1</b>	<b>56.9</b>	<b>35.7</b>	<b>33.2</b>	<b>59.3</b>	60.5	<b>48.1</b>

Table 1. Instance segmentation results of different methods on the seven core datasets of the BOP benchmark [54]. We report the mean Average Precision (mAP) scores at different Intersection-over-Union (IoU) values ranging from 0.50 to 0.95 with a step size of 0.05.

Method	Input Type	Detection / Segmentation	BOP Dataset							Mean
			LM-O	T-LESS	TUD-L	IC-BIN	ITODD	HB	YCB-V	
With Supervised Detection / Segmentation										
MegaPose [28]	RGB	MaskRCNN [16]	18.7	19.7	20.5	15.3	8.00	18.6	13.9	16.2
MegaPose <sup>†</sup> [28]	RGB		53.7	62.2	58.4	43.6	30.1	72.9	60.4	54.5
MegaPose <sup>†</sup> * [28]	RGB-D		58.3	54.3	71.2	37.1	40.4	75.7	63.3	57.2
ZeroPose [5]	RGB-D		26.1	24.3	61.1	24.7	26.4	38.2	29.5	32.6
ZeroPose <sup>†</sup> [5]	RGB-D		56.2	53.3	<b>87.2</b>	41.8	<b>43.6</b>	68.2	58.4	58.4
SAM-6D (Ours)	RGB-D		<b>66.5</b>	<b>66.0</b>	80.9	<b>61.9</b>	31.9	<b>81.8</b>	<b>79.6</b>	<b>66.9</b>
With Zero-Shot Detection / Segmentation										
ZeroPose [5]	RGB-D	ZeroPose [5]	26.0	17.8	41.2	17.7	38.0	43.9	25.7	25.7
ZeroPose <sup>†</sup> [5]	RGB-D		49.1	34.0	74.5	39.0	42.9	61.0	57.7	51.2
SAM-6D (Ours)	RGB-D		<b>63.5</b>	<b>43.0</b>	<b>80.2</b>	<b>51.8</b>	<b>48.4</b>	<b>69.1</b>	<b>79.2</b>	<b>62.2</b>
MegaPose* [28]	RGB	CNOS (FastSAM) [40]	22.9	17.7	25.8	15.2	10.8	25.1	28.1	20.8
MegaPose <sup>†</sup> * [28]	RGB		49.9	47.7	65.3	36.7	31.5	65.4	60.1	50.9
MegaPose <sup>†</sup> * [28]	RGB-D		62.6	48.7	<b>85.1</b>	46.7	46.8	73.0	76.4	62.8
ZeroPose <sup>†</sup> * [5]	RGB-D		53.8	40.0	83.5	39.2	52.1	65.3	65.3	57.0
GigaPose [41]	RGB		29.9	27.3	30.2	23.1	18.8	34.8	29.0	27.6
GigaPose <sup>†</sup> [41]	RGB		59.9	<b>57.0</b>	63.5	46.7	39.7	72.2	66.3	57.9
SAM-6D (Ours)	RGB-D		<b>65.1</b>	47.9	82.5	<b>49.7</b>	<b>56.2</b>	<b>73.8</b>	<b>81.5</b>	<b>65.3</b>
SAM-6D (Ours)	RGB-D		SAM-6D (FastSAM)	<b>66.7</b>	<b>48.5</b>	<b>82.9</b>	<b>51.0</b>	<b>57.2</b>	<b>73.6</b>	<b>83.4</b>
SAM-6D (Ours)	RGB-D	SAM-6D (SAM)	<b>69.9</b>	<b>51.5</b>	<b>90.4</b>	<b>58.8</b>	<b>60.2</b>	<b>77.6</b>	<b>84.5</b>	<b>70.4</b>

Table 2. Pose estimation results of different methods on the seven core datasets of the BOP benchmark [54]. We report the mean Average Recall (AR) among VSD, MSSD and MSPD, as introduced in Sec. 4. The symbol ‘†’ denotes the use of pose refinement proposed in [28]. The symbol ‘\*’ denotes the results published on BOP leaderboard. Our used masks of MaskRCNN [16] are provided by CosyPose [27].

#### 4.1. Instance Segmentation of Novel Objects

We compare our ISM of SAM-6D with ZeroPose [5] and CNOS [40], both of which score the object proposals in terms of semantics solely, for instance segmentation of novel objects. The quantitative results are presented in Table 1, demonstrating that our ISM, built on the publicly available foundation models of SAM [26] / FastSAM [74] and ViT (pre-trained by DINOv2 [44]), delivers superior results without the need for network re-training or finetuning. Note that our baseline with only semantic matching score  $s_{sem}$ , whether based on SAM or FastSAM [74], aligns precisely with the method of CNOS; the only difference is that we adjust the hyperparameters of SAM to generate more

proposals for scoring. Further enhancements to our baselines are achieved via the inclusion of appearance and geometry matching scores, *i.e.*,  $s_{appe}$  and  $s_{geo}$ , as verified in Table 1. Qualitative results of ISM are visualized in Fig. 1.

#### 4.2. Pose Estimation of Novel Objects

##### 4.2.1 Comparisons with Existing Methods

We compare our PEM of SAM-6D with the representative methods, including MegaPose [28], ZeroPose [5], and GigaPose [41], for pose estimation of novel objects. Quantitative comparisons, as presented in Table 5, show that our PEM, without the time-intensive render-based refiner [28], outperforms the existing methods under various mask pre-

dictions. Importantly, the mask predictions from our ISM significantly enhance the performance of PEM, compared to other mask predictions, further validating the advantages of ISM. Qualitative results of PEM are visualized in Fig. 1.

### 4.2.2 Ablation Studies and Analyses

We conduct ablation studies on the YCB-V dataset to evaluate the efficacy of individual designs in PEM, with the mask predictions generated by ISM based on SAM.

**Efficacy of Background Tokens** We address the partial-to-partial point matching issue through a simple yet effective design of background tokens. Another existing solution is the use of optimal transport [48] with iterative optimization, which, however, is time-consuming. The two solutions are compared in Table 3, which shows that our PEM with background tokens achieves results comparable to optimal transport, but with a faster inference speed. As the density of points for matching increases, optimal transport requires more time to derive the assignment matrices.

**Efficacy of Two Point Matching Stages** With the background tokens, we design PEM with two stages of point matching via a Coarse Point Matching module and a Fine Point Matching module. Firstly, we validate the effectiveness of the Fine Point Matching module, which effectively improves the results of the coarse module, as verified in Table 4. Further, we evaluate the effectiveness of the Coarse Point Matching module by removing it from PEM. In this case, the point sets of object proposals are not transformed and are directly used to learn the positional encodings in the fine module. The results, presented in Table 4, indicate that the removal of Coarse Point Matching significantly degrades the performance, which may be attributed to the large distance between the sampled point sets of the proposals and target objects, as no initial poses are provided.

**Efficacy of Sparse-to-Dense Point Transformers** We design Sparse-to-Dense Point Transformers (SDPT) in the Fine Point Matching module to manage dense point interactions. Within each SDPT, Geometric Transformers [48] is employed to learn the relationships between sparse point sets, which are then spread to the dense ones via Linear Transformers [24]. We conduct experiments on either Geometric Transformers using sparse point sets with 196 points or Linear Transformers using dense point sets with 2048 points. The results, presented in Table 5, indicate inferior performance compared to using our SDPTs. This is because Geometric Transformers struggle to handle dense point sets due to high computational costs, whereas Linear Transformers prove to be ineffective in modeling dense correspondence with attention along the feature dimension.

### 4.3. Runtime Analysis

We conduct evaluation on a server with a GeForce RTX 3090 GPU, and report in Table 6 the runtime averaged

	AR	Time (s)
PEM with Optimal Transport	81.4	4.31
PEM with Background Tokens	<b>84.5</b>	<b>1.36</b>

Table 3. Quantitative results of Optimal Transport [48] and our design of Background Tokens in the Pose Estimation Model on YCB-V. The reported time is the average per-image processing time of pose estimation across the entire dataset on a server with a GeForce RTX 3090 GPU.

Coarse Point Matching	Fine Point Matching	AR
✓	×	77.6
×	✓	40.2
✓	✓	<b>84.5</b>

Table 4. Ablation studies on the the strategy of two point matching stages in the Pose Estimation Model on YCB-V.

Transformer	#Point	AR
Geometric Transformer [48]	196	81.7
Linear Transformer [24]	2048	78.4
Sparse-to-Dense Point Transformer	196 → 2048	<b>84.5</b>

Table 5. Quantitative comparisons among various types of transformers employed in the Fine Point Matching module of the Pose Estimation Model on YCB-V.

Segmentation Model	Time (s)		
	Instance Segmentation	Pose Estimation	All
FastSAM [74]	0.45	0.98	1.43
SAM [26]	2.80	1.57	4.37

Table 6. Runtime of SAM-6D with different segmentation models. The reported time is the average per-image processing time across the seven core datasets of BOP benchmark on a server with a GeForce RTX 3090 GPU.

on the seven core datasets of BOP benchmark, indicating the efficiency of SAM-6D which avoids the use of time-intensive render-based refiners. We note that SAM-based method takes more time on pose estimation than FastSAM-based one, due to more object proposals generated by SAM.

## 5. Conclusion

In this paper, we take Segment Anything Model (SAM) as an advanced starting point for zero-shot 6D object pose estimation, and present a novel framework, named SAM-6D, which comprises an Instance Segmentation Model (ISM) and a Pose Estimation Model (PEM) to accomplish the task in two steps. ISM utilizes SAM to segment all potential object proposals and assigns each of them an object matching score in terms of semantics, appearance, and geometry. PEM then predicts the object pose for each proposal by solving a partial-to-partial point matching problem through two stages of Coarse Point Matching and Fine Point Matching. The effectiveness of SAM-6D is validated on the seven core datasets of BOP benchmark, where SAM-6D significantly outperforms existing methods.



## References

- [1] Dingding Cai, Janne Heikkilä, and Esa Rahtu. Ove6d: Object viewpoint encoding for depth-based 6d object pose estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6803–6813, 2022. 3
- [2] Jun Cen, Yizheng Wu, Kewei Wang, Xingyi Li, Jingkang Yang, Yixuan Pei, Lingdong Kong, Ziwei Liu, and Qifeng Chen. Sad: Segment any rgb. *arXiv preprint arXiv:2305.14207*, 2023. 3
- [3] Jiazhong Cen, Zanwei Zhou, Jiemin Fang, Wei Shen, Lingxi Xie, Xiaopeng Zhang, and Qi Tian. Segment anything in 3d with nerfs. *arXiv preprint arXiv:2304.12308*, 2023. 3
- [4] Angel X Chang, Thomas Funkhouser, Leonidas Guibas, Pat Hanrahan, Qixing Huang, Zimo Li, Silvio Savarese, Manolis Savva, Shuran Song, Hao Su, et al. Shapenet: An information-rich 3d model repository. *arXiv preprint arXiv:1512.03012*, 2015. 3, 6
- [5] Jianqiu Chen, Mingshan Sun, Tianpeng Bao, Rui Zhao, Liwei Wu, and Zhenyu He. 3d model-based zero-shot pose estimation pipeline. *arXiv preprint arXiv:2305.17934*, 2023. 2, 3, 7
- [6] Jiaqi Chen, Zeyu Yang, and Li Zhang. Semantic segment anything. <https://github.com/fudan-zvg/Semantic-Segment-Anything>, 2023. 3
- [7] Kai Chen and Qi Dou. Sgpa: Structure-guided prior adaptation for category-level 6d object pose estimation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2773–2782, 2021. 2
- [8] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020. 4, 5
- [9] Laura Downs, Anthony Francis, Nate Koenig, Brandon Kinman, Ryan Hickman, Krista Reymann, Thomas B McHugh, and Vincent Vanhoucke. Google scanned objects: A high-quality dataset of 3d scanned household items. In *2022 International Conference on Robotics and Automation (ICRA)*, pages 2553–2560. IEEE, 2022. 3, 6
- [10] Zhiwen Fan, Panwang Pan, Peihao Wang, Yifan Jiang, De-jia Xu, Hanwen Jiang, and Zhangyang Wang. Pope: 6-dof promptable pose estimation of any object, in any scene, with one reference. *arXiv preprint arXiv:2305.15727*, 2023. 3
- [11] Walter Goodwin, Sagar Vaze, Ioannis Havoutis, and Ingmar Posner. Zero-shot category-level object pose estimation. In *European Conference on Computer Vision*, pages 516–532. Springer, 2022. 3
- [12] Dongchen Han, Xuran Pan, Yizeng Han, Shiji Song, and Gao Huang. Flatten transformer: Vision transformer using focused linear attention. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5961–5971, 2023. 3, 6
- [13] Dongsheng Han, Chaoning Zhang, Yu Qiao, Maryam Qamar, Yuna Jung, SeungKyu Lee, Sung-Ho Bae, and Choong Seon Hong. Segment anything model (sam) meets glass: Mirror and transparent objects cannot be easily detected. *arXiv preprint arXiv:2305.00278*, 2023. 3
- [14] Rasmus Laurvig Haugaard and Anders Glent Buch. Surfemb: Dense and continuous correspondence distributions for object pose estimation with learnt surface embeddings. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6749–6758, 2022. 6
- [15] Haibin He, Jing Zhang, Mengyang Xu, Juhua Liu, Bo Du, and Dacheng Tao. Scalable mask annotation for video text spotting. *arXiv preprint arXiv:2305.01443*, 2023. 3
- [16] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *Proceedings of the IEEE international conference on computer vision*, pages 2961–2969, 2017. 7
- [17] Xingyi He, Jiaming Sun, Yuang Wang, Di Huang, Hujun Bao, and XiaoWei Zhou. Onepose++: Keypoint-free one-shot object pose estimation without cad models. *Advances in Neural Information Processing Systems*, 35:35103–35115, 2022. 3
- [18] Yisheng He, Wei Sun, Haibin Huang, Jianran Liu, Haoqiang Fan, and Jian Sun. Pvn3d: A deep point-wise 3d keypoints voting network for 6dof pose estimation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 11632–11641, 2020. 2
- [19] Yisheng He, Haibin Huang, Haoqiang Fan, Qifeng Chen, and Jian Sun. Ffb6d: A full flow bidirectional fusion network for 6d pose estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3003–3013, 2021. 2
- [20] Yisheng He, Yao Wang, Haoqiang Fan, Jian Sun, and Qifeng Chen. Fs6d: Few-shot 6d pose estimation of novel objects. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6814–6824, 2022. 3
- [21] Shengyu Huang, Zan Gojcic, Mikhail Usvyatsov, Andreas Wieser, and Konrad Schindler. Predator: Registration of 3d point clouds with low overlap. In *Proceedings of the IEEE/CVF Conference on computer vision and pattern recognition*, pages 4267–4276, 2021. 3
- [22] Ge-Peng Ji, Deng-Ping Fan, Peng Xu, Ming-Ming Cheng, Bowen Zhou, and Luc Van Gool. Sam struggles in concealed scenes—empirical study on “segment anything”. *arXiv preprint arXiv:2304.06022*, 2023. 3
- [23] Wei Ji, Jingjing Li, Qi Bi, Wenbo Li, and Li Cheng. Segment anything is not always perfect: An investigation of sam on different real-world applications. *arXiv preprint arXiv:2304.05750*, 2023. 3
- [24] Angelos Katharopoulos, Apoorv Vyas, Nikolaos Pappas, and François Fleuret. Transformers are rnns: Fast autoregressive transformers with linear attention. In *International conference on machine learning*, pages 5156–5165. PMLR, 2020. 3, 6, 8
- [25] Tong Ke and Stergios I Roumeliotis. An efficient algebraic solution to the perspective-three-point problem. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7225–7233, 2017. 6
- [26] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C Berg, Wan-Yen Lo, et al. Segment any-

- thing. *arXiv preprint arXiv:2304.02643*, 2023. [2](#), [3](#), [4](#), [6](#), [7](#), [8](#)
- [27] Yann Labbé, Justin Carpentier, Mathieu Aubry, and Josef Sivic. Cosypose: Consistent multi-view multi-object 6d pose estimation. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XVII 16*, pages 574–591. Springer, 2020. [7](#)
- [28] Yann Labbé, Lucas Manuelli, Arsalan Mousavian, Stephen Tyree, Stan Birchfield, Jonathan Tremblay, Justin Carpentier, Mathieu Aubry, Dieter Fox, and Josef Sivic. Megapose: 6d pose estimation of novel objects via render & compare. In *Proceedings of the 6th Conference on Robot Learning (CoRL)*, 2022. [2](#), [3](#), [6](#), [7](#)
- [29] Jiehong Lin, Hongyang Li, Ke Chen, Jiangbo Lu, and Kui Jia. Sparse steerable convolutions: An efficient learning of se (3)-equivariant features for estimation and tracking of object poses in 3d space. *Advances in Neural Information Processing Systems*, 34:16779–16790, 2021. [2](#)
- [30] Jiehong Lin, Zewei Wei, Zhihao Li, Songcen Xu, Kui Jia, and Yuanqing Li. Dualposenet: Category-level 6d object pose and size estimation using dual pose network with refined learning of pose consistency. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3560–3569, 2021.
- [31] Jiehong Lin, Zewei Wei, Changxing Ding, and Kui Jia. Category-level 6d object pose and size estimation using self-supervised deep prior deformation networks. In *European Conference on Computer Vision*, pages 19–34. Springer, 2022.
- [32] Jiehong Lin, Zewei Wei, Yabin Zhang, and Kui Jia. Vi-net: Boosting category-level 6d object pose estimation via learning decoupled rotations on the spherical representations. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 14001–14011, 2023. [2](#)
- [33] Yuan Liu, Yilin Wen, Sida Peng, Cheng Lin, Xiaoxiao Long, Taku Komura, and Wenping Wang. Gen6d: Generalizable model-free 6-dof object pose estimation from rgb images. In *European Conference on Computer Vision*, pages 298–315. Springer, 2022. [3](#)
- [34] Yang Liu, Muzhi Zhu, Hengtao Li, Hao Chen, Xinlong Wang, and Chunhua Shen. Matcher: Segment anything with one shot using all-purpose feature matching. *arXiv preprint arXiv:2305.13310*, 2023. [3](#)
- [35] Zhaoyang Liu, Yanan He, Wenhai Wang, Weiyun Wang, Yi Wang, Shoufa Chen, Qinglong Zhang, Yang Yang, Qingyun Li, Jiashuo Yu, et al. Internchat: Solving vision-centric tasks by interacting with chatbots beyond language. *arXiv preprint arXiv:2305.05662*, 2023. [3](#)
- [36] Jun Ma and Bo Wang. Segment anything in medical images. *arXiv preprint arXiv:2304.12306*, 2023. [3](#)
- [37] Maciej A Mazurowski, Haoyu Dong, Hanxue Gu, Jichen Yang, Nicholas Konz, and Yixin Zhang. Segment anything model for medical image analysis: an experimental study. *Medical Image Analysis*, 89:102918, 2023. [3](#)
- [38] Van Nguyen Nguyen, Yinlin Hu, Yang Xiao, Mathieu Salzmann, and Vincent Lepetit. Templates for 3d object pose estimation revisited: Generalization to new objects and robustness to occlusions. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 6771–6780, 2022. [3](#)
- [39] Van Nguyen Nguyen, Thibault Groueix, Yinlin Hu, Mathieu Salzmann, and Vincent Lepetit. Nope: Novel object pose estimation from a single image. *arXiv preprint arXiv:2303.13612*, 2023. [3](#)
- [40] Van Nguyen Nguyen, Thibault Groueix, Georgy Ponimatkin, Vincent Lepetit, and Tomas Hodan. Cnos: A strong baseline for cad-based novel object segmentation. *arXiv preprint arXiv:2307.11067*, 2023. [2](#), [3](#), [4](#), [6](#), [7](#)
- [41] Van Nguyen Nguyen, Thibault Groueix, Mathieu Salzmann, and Vincent Lepetit. Gigapose: Fast and robust novel object pose estimation via one correspondence. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024. [3](#), [7](#)
- [42] Brian Okorn, Qiao Gu, Martial Hebert, and David Held. Zephyr: Zero-shot pose hypothesis rating. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pages 14141–14148. IEEE, 2021. [3](#)
- [43] Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*, 2018. [6](#)
- [44] Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, et al. Dinov2: Learning robust visual features without supervision. *arXiv preprint arXiv:2304.07193*, 2023. [6](#), [7](#)
- [45] Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, et al. Dinov2: Learning robust visual features without supervision. *arXiv preprint arXiv:2304.07193*, 2023. [4](#)
- [46] Panwang Pan, Zhiwen Fan, Brandon Y Feng, Peihao Wang, Chenxin Li, and Zhangyang Wang. Learning to estimate 6dof pose from limited data: A few-shot, generalizable approach using rgb images. *arXiv preprint arXiv:2306.07598*, 2023. [3](#)
- [47] Charles Ruizhongtai Qi, Li Yi, Hao Su, and Leonidas J Guibas. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. *Advances in neural information processing systems*, 30, 2017. [6](#)
- [48] Zheng Qin, Hao Yu, Changjian Wang, Yulan Guo, Yuxing Peng, and Kai Xu. Geometric transformer for fast and robust point cloud registration. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 11143–11152, 2022. [3](#), [6](#), [8](#)
- [49] Qihong Shen, Xingyi Yang, and Xinchao Wang. Anything-3d: Towards single-view anything reconstruction in the wild. *arXiv preprint arXiv:2304.10261*, 2023. [3](#)
- [50] Ivan Shugurov, Fu Li, Benjamin Busam, and Slobodan Ilic. Osop: A multi-stage one shot object pose estimation framework. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6835–6844, 2022. [3](#)
- [51] Yongzhi Su, Mahdi Saleh, Torben Fetzner, Jason Rambach, Nassir Navab, Benjamin Busam, Didier Stricker, and Fed-

- erico Tombari. ZebraPose: Coarse to fine surface encoding for 6dof object pose estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6738–6748, 2022. 2
- [52] Jiaming Sun, Zehong Shen, Yuang Wang, Hujun Bao, and Xiaowei Zhou. Loftr: Detector-free local feature matching with transformers. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8922–8931, 2021. 3
- [53] Jiaming Sun, Zihao Wang, Siyu Zhang, Xingyi He, Hongcheng Zhao, Guofeng Zhang, and Xiaowei Zhou. OnePose: One-shot object pose estimation without cad models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6825–6834, 2022. 3
- [54] Martin Sundermeyer, Tomáš Hodaň, Yann Labbe, Gu Wang, Eric Brachmann, Bertram Drost, Carsten Rother, and Jiří Matas. Bop challenge 2022 on detection, segmentation and pose estimation of specific rigid objects. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2784–2793, 2023. 1, 3, 6, 7
- [55] Lv Tang, Haoke Xiao, and Bo Li. Can sam segment anything? when sam meets camouflaged object detection. *arXiv preprint arXiv:2304.04709*, 2023. 3
- [56] Meng Tian, Marcelo H Ang, and Gim Hee Lee. Shape prior deformation for categorical 6d object pose and size estimation. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXI 16*, pages 530–546. Springer, 2020. 2
- [57] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017. 6
- [58] Chen Wang, Danfei Xu, Yuke Zhu, Roberto Martín-Martín, Cewu Lu, Li Fei-Fei, and Silvio Savarese. Densefusion: 6d object pose estimation by iterative dense fusion. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 3343–3352, 2019. 2
- [59] Dongqing Wang, Tong Zhang, Alaa Abboud, and Sabine Süsstrunk. Inpaintnerf360: Text-guided 3d inpainting on unbounded neural radiance fields. *arXiv preprint arXiv:2305.15094*, 2023. 3
- [60] Gu Wang, Fabian Manhardt, Federico Tombari, and Xiangyang Ji. Gdr-net: Geometry-guided direct regression network for monocular 6d object pose estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16611–16621, 2021. 2
- [61] He Wang, Srinath Sridhar, Jingwei Huang, Julien Valentin, Shuran Song, and Leonidas J Guibas. Normalized object coordinate space for category-level 6d object pose and size estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2642–2651, 2019. 2
- [62] Qian Wang, Biao Zhang, Michael Birsak, and Peter Wonka. Instructedit: Improving automatic masks for diffusion-based image editing with user instructions. *arXiv preprint arXiv:2305.18047*, 2023. 3
- [63] Yu Xiang, Tanner Schmidt, Venkatraman Narayanan, and Dieter Fox. Posecnn: A convolutional neural network for 6d object pose estimation in cluttered scenes. *arXiv preprint arXiv:1711.00199*, 2017. 2
- [64] Defeng Xie, Ruichen Wang, Jian Ma, Chen Chen, Haonan Lu, Dong Yang, Fobo Shi, and Xiaodong Lin. Edit everything: A text-guided generative system for images editing. *arXiv preprint arXiv:2304.14006*, 2023. 3
- [65] Jinyu Yang, Mingqi Gao, Zhe Li, Shang Gao, Fangjing Wang, and Feng Zheng. Track anything: Segment anything meets videos. *arXiv preprint arXiv:2304.11968*, 2023. 3
- [66] Yunhan Yang, Xiaoyang Wu, Tong He, Hengshuang Zhao, and Xihui Liu. Sam3d: Segment anything in 3d scenes. *arXiv preprint arXiv:2306.03908*, 2023. 3
- [67] Tao Yu, Runseng Feng, Ruoyu Feng, Jinming Liu, Xin Jin, Wenjun Zeng, and Zhibo Chen. Inpaint anything: Segment anything meets image inpainting. *arXiv preprint arXiv:2304.06790*, 2023. 3
- [68] Chaoning Zhang, Dongshen Han, Yu Qiao, Jung Uk Kim, Sung-Ho Bae, Seungkyu Lee, and Choong Seon Hong. Faster segment anything: Towards lightweight sam for mobile applications. *arXiv preprint arXiv:2306.14289*, 2023. 3
- [69] Chaoning Zhang, Sheng Zheng, Chenghao Li, Yu Qiao, Taegoo Kang, Xinru Shan, Chenshuang Zhang, Caiyan Qin, Francois Rameau, Sung-Ho Bae, et al. A survey on segment anything model (sam): Vision foundation model meets prompt engineering. *arXiv preprint arXiv:2306.06211*, 2023. 3
- [70] Dingyuan Zhang, Dingkan Liang, Hongcheng Yang, Zhikang Zou, Xiaoqing Ye, Zhe Liu, and Xiang Bai. Sam3d: Zero-shot 3d object detection via segment anything model. *arXiv preprint arXiv:2306.02245*, 2023. 3
- [71] Haojie Zhang, Yongyi Su, Xun Xu, and Kui Jia. Improving the generalization of segmentation foundation model under distribution shift via weakly supervised adaptation. *arXiv preprint arXiv:2312.03502*, 2023. 3
- [72] Renrui Zhang, Zhengkai Jiang, Ziyu Guo, Shilin Yan, Junting Pan, Hao Dong, Peng Gao, and Hongsheng Li. Personalize segment anything model with one shot. *arXiv preprint arXiv:2305.03048*, 2023. 3
- [73] Zhenghao Zhang, Zhichao Wei, Shengfan Zhang, Zuozhuo Dai, and Siyu Zhu. Uvosam: A mask-free paradigm for unsupervised video object segmentation via segment anything model. *arXiv preprint arXiv:2305.12659*, 2023. 3
- [74] Xu Zhao, Wenchao Ding, Yongqi An, Yinglong Du, Tao Yu, Min Li, Ming Tang, and Jinqiao Wang. Fast segment anything. *arXiv preprint arXiv:2306.12156*, 2023. 3, 6, 7, 8