

# Swiss DINO: Efficient and Versatile Vision Framework for On-device Personal Object Search

Kirill Paramonov<sup>1</sup>, Jia-Xing Zhong<sup>1,2,\*</sup>, Umberto Michieli<sup>1</sup>, Jijoong Moon<sup>3</sup>, Mete Ozay<sup>1</sup>

**Abstract**—In this paper, we address a recent trend in robotic home appliances to include vision systems on personal devices, capable of personalizing the appliances on the fly. In particular, we formulate and address an important technical task of personal object search, which involves localization and identification of personal items of interest on images captured by robotic appliances, with each item referenced only by a few annotated images. The task is crucial for robotic home appliances and mobile systems, which need to process personal visual scenes or to operate with particular personal objects (e.g., for grasping or navigation). In practice, personal object search presents two main technical challenges. First, a robot vision system needs to be able to distinguish between many fine-grained classes, in the presence of occlusions and clutter. Second, the strict resource requirements for the on-device system restrict the usage of most state-of-the-art methods for few-shot learning and often prevent on-device adaptation. In this work, we propose Swiss DINO: a simple yet effective framework for one-shot personal object search based on the recent DINOv2 transformer model, which was shown to have strong zero-shot generalization properties. Swiss DINO handles challenging on-device personalized scene understanding requirements and does not require any adaptation training. We show significant improvement (up to 55%) in segmentation and recognition accuracy compared to the common lightweight solutions, and significant footprint reduction of backbone inference time (up to 100×) and GPU consumption (up to 10×) compared to the heavy transformer-based solutions<sup>1</sup>.

## I. INTRODUCTION

Computer vision has a pivotal role in mobile systems and home appliances to understand the surroundings and navigate in complex environments. Scene understanding deep neural networks have obtained outstanding results and have been successfully deployed to mass-accessible personal devices: for example industrial or domestic service robots (e.g., vacuum cleaners), assistive robots, and smartphones.

Recently, increasing attention [1]–[3] has been devoted to the personalization of on-device AI vision models to tackle a variety of practical use cases. In this work, we focus on personal item search, whereby we want robot vision systems to localize and recognize personal user classes (or **fine-grained classes**, e.g. *my dog Archie*, *her dog Bruno*, *my favorite cup*, *your favorite flower*, etc.) on scenes. Specifically, a user provides a small number of reference images with location

annotations (either a segmentation map or a bounding box) for each personal item. Then, given a new scene, a visual system needs to *i*) determine which of the personal objects are present in the scene, and *ii*) provide the location (in the form of a segmentation map or bounding box) for each of the personal objects present in the scene. This task has found significant applications for personal assistants and service robots: for navigation (e.g., *reach my white sofa*), HRI (e.g., *find my dog Archie*), grasping (e.g., *bring me my phone*), etc.

Previous works have focused on different aspects of the task. The closest comparison task to ours is the few-shot semantic segmentation [4], which aims to segment an object on the scene given a provided reference image and mask.

In this paper, we aim to address the following limitations of existing few-shot semantic segmentation methods. First, existing solutions only evaluate the IoU metric for the mask corresponding to the ground truth class on the image, thus not accounting for the multi-class scenario. Second, they require adaptation training on coarse datasets, thus making fine-grained classes indistinguishable in the feature space (which is part of the effect known as the **neural collapse** [5]). Third, current transformer-based solutions rely on large foundation models (e.g., SAM), which may be too costly for on-device implementation.

In this work, we develop a problem statement and metrics for the personal object search task, that is closely related to the practical scenarios. We develop a novel method for the task, which does not rely on coarse dataset training and is very lightweight, allowing seamless implementation on device.

Inspired by works showing the great versatility of the DINOv2 model [6] for downstream tasks [6], [7], we employ DINOv2 as our backbone. Our system is called Swiss DINO, after the Swiss Army Knife, for its incredible versatility and adaptability. Fig. 1 shows our approach and its novelty compared to existing solutions.

Our evaluation focuses on multi-instance personalization (i.e., adaptation to multiple personal objects) via one-shot transfer on multiple tasks (image classification, object detection and semantic segmentation) and datasets (iCubWorld [8] and PerSeg [9]). For one-shot segmentation task, Swiss DINO improves the memory usage by up to 10× and backbone inference time by up to 100× compared to few-shot semantic segmentation competitors based on foundational models, while maintaining similar segmentation accuracy, as well as improving segmentation accuracy of lightweight solutions on cluttered scenes by 46%. To evaluate multi-instance identification accuracy, we adopt metrics from open-

\*Research completed during internship at Samsung R&D Institute UK.

<sup>1</sup> Samsung R&D Institute UK (SRUK), Communications House, South St, Staines, Surrey, United Kingdom {n.surname}@samsung.com

<sup>2</sup> University of Oxford, Wellington Square, Oxford, Oxfordshire, United Kingdom jiaxing.zhong@cs.ox.ac.uk

<sup>3</sup>Samsung Research Korea, Seoul R&D Campus, 56, Seongchon-gil, Seocho-gu, Seoul, Rep. of Korea jijoong.moon@samsung.com

<sup>1</sup>Code is available at:

<https://github.com/SamsungLabs/SwissDINO>.

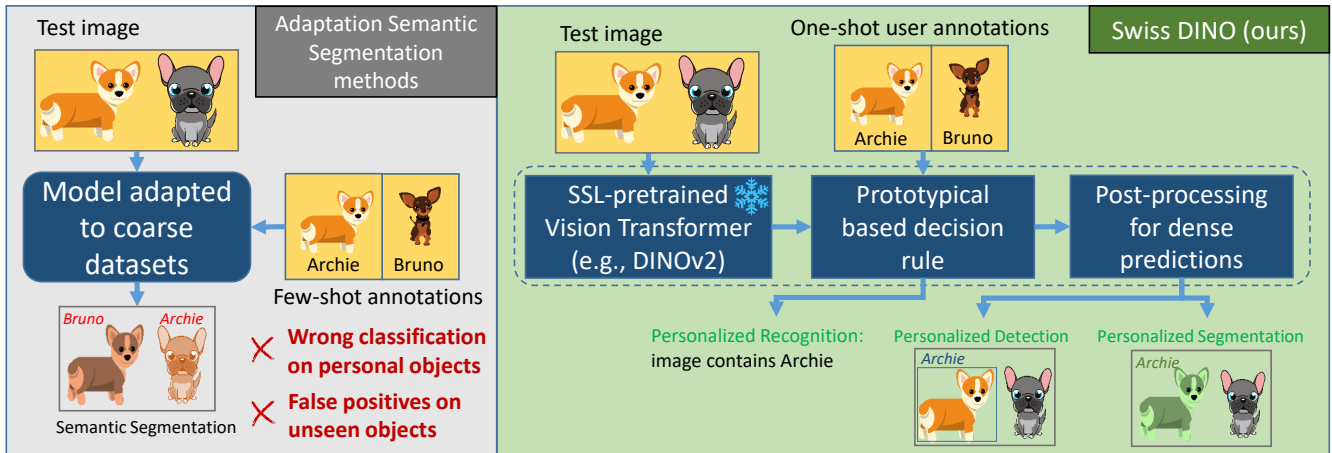


Fig. 1: Comparison with semantic segmentation methods. Left: common adaptive semantic segmentation methods are adapting models to coarse datasets and do not account for multiple personal objects or unseen personal objects on a scene, thus generating false positive errors. Right: our Swiss DINO avoids false positive errors by performing open-set classification on parts of the image prior to generating segmentation masks.

set recognition task [10]. We adapt existing segmentation methods for the multi-instance setup and show that compared to lightweight competitors pre-trained on coarse classes, Swiss DINO achieves 55% identification improvement on simple scenes, and 42% improvement on cluttered scenes.

The remainder of the paper is organized as follows: Sec. II positions our paper in the current landscape of personalized scene understanding, Sec. III formalizes our problem setup, Sec. IV presents the details of our method, Sec. V shows the results on several benchmark, finally Sec. VI draws the conclusion of our work.

## II. RELATED WORKS

*a) Few-Shot Semantic Segmentation:* While early works on few-shot semantic segmentation resorted to fine-tuning large parts of models [11]–[14], recent approaches are based on sparse feature matching [15] or training adaptation layers with the prototypical loss [16]–[20]. These latter approaches compute class prototypes as the average embedding of all images of a class. The label of a new (query) image is predicted by identifying the nearest prototype vector computed from the training (support) set. The training and evaluation are usually performed on popular segmentation datasets with **coarse-level classes** (e.g. *person, cat, car, chair*), namely PASCAL – 5<sup>i</sup> [21] and COCO – 20<sup>i</sup> [22].

Recent advancements in large vision models have led to novel few-shot scene understanding works, especially applied to semantic segmentation, such as PerSAM [9], and Matcher [7]. PerSAM, a training-free approach, uses a single image with a reference mask to localize and segment target concepts. Matcher, utilizing off-the-shelf vision foundation models, can showcase impressive generalization across tasks. On the other hand, both approaches are computationally expensive and not applicable on low-resource devices.

*b) Object Detection Datasets for Robotic Applications:* Object detection and fine-grained identification are crucial

tasks for robotic manipulators [23]. To boost the development of object detection methods, several datasets have been introduced. In particular, iCubWorld [8] is a collection of images recording the visual experience of the iCub humanoid robot observing personal user objects in its typical environment, such as a laboratory or an office. CORE50 [24] further enriches the field, offering a new benchmark for continuous object recognition, designed specifically for real-world applications such as fine-grained object detection in robot vision systems. These datasets align with our task as they represent scenarios where few-shot personalization can be used to enhance the robot’s ability to recognize new or fine-grained objects, serving as practical representations of the use cases where our method can be applied. It was shown that the common classification architectures trained on coarse-level datasets have low accuracy on aforementioned datasets with fine-grained classes [25] when applied out-of-the-box. This is due to the fact that the fine-grained classes become indistinguishable in the feature space after long training on coarse class classification, part of the effect known as **neural collapse** [5]. Therefore, fine-tuning [25] or adaptation [16], [19] methods are often employed to separate the feature vectors for fine-grained datasets.

*c) Pre-trained DINOv2 as An All-purpose Backbone:* In self-supervised learning (SSL), significant contributions have been made to the development of pre-trained models, such as DINO [26] and DINOv2 [6]. These models have demonstrated remarkable capabilities in feature extraction and object localization, making them highly transferable to our task of few-shot personalization. Siméoni et al. presents a method named LOST [27] to leverage pre-trained vision transformer features for unsupervised object localization. Melas-Kyriazi et al. [28] reframed image decomposition as a graph partitioning problem, using eigenvectors from self-supervised networks to segment images and localize objects. These methods not only provide a strong foundation for

our few-shot personalization method but also highlight the potential of SSL transformer backbones in overcoming the challenge of neural collapse.

### III. PROBLEM STATEMENT

In this section, we present the problem formulation of personal object search and notation for each of the three stages involved.

1) *Pre-training Stage*: The first stage is to pre-train a backbone model on a large dataset. The backbone should provide localization information of objects on an image, and have a strong ability to transfer to new personal classes, in particular avoiding neural collapse of generated features.

2) *On-device Personalization Stage*: After the system is implemented on a mobile or robotic device (e.g., a robot vacuum cleaner or service robot), it is shown a few images of personal objects, together with their label (e.g., *dog Archie*, *dog Bruno*, *my mug*, etc.), and a prompt indicating the location of the object on the image, in the form of a bounding box or a segmentation map. Those images are also known in the few-shot literature as *support* images.

Although our setup can be applied to any number of support images per personal object, to simplify evaluation and notations, for the rest of the paper, we consider the most challenging one-shot setup, i.e., we get a single support image  $S_c$  for each personal object index  $c = 1, 2, \dots, C$ , where  $C$  is the number of personal objects.

3) *On-device Open-set Personal Object Segmentation, Detection, and Recognition*: During the on-device inference stage, we are given a new test image  $Q$  (also known in the literature as the *query* image). For this image, we need to determine: *i*) which personal objects, if any, are present on the image; *ii*) for each of the personal objects present on the image, find its location in the form of segmentation map or a bounding box.

More formally, we define a personal object search method POS by

$$\text{POS}(Q) := (\text{oloc}_1(Q), \dots, \text{oloc}_C(Q)), \quad (1)$$

$$\text{oloc}_c := \begin{cases} \text{loc}_c(Q), & \text{if the object } c \text{ is present} \\ \text{None}, & \text{otherwise} \end{cases} \quad (2)$$

where  $\text{loc}_c(Q)$  can take the form of a bounding box or a segmentation map for the object  $c$  on the image  $Q$ .

### IV. METHODOLOGY

Our Swiss DINO system consists of three main components: *i*) patch-level feature map extraction; *ii*) support feature map processing; *iii*) query feature map processing. An overview of our Swiss DINO is shown in Fig. 2.

#### A. Patch-level feature map extractor

We utilize a pre-trained transformer-based patch-level feature extractor. Inspired by the previous work [6] on the DINOv2 model, making use of its localization and fine-grained separation capabilities of its feature map, we choose DINOv2 as our transformer backbone (for a comparison between different backbone models, see Section V-F.1).

The backbone  $\mathcal{B}$  takes an image  $X$  as an input and produces *i*) patch-wise feature map  $X^F = (X_{1,1}^F, \dots, X_{N_P, N_P}^F)$ , where  $N_P$  is the number of patches along each side of the image, and  $X_{i,j}^F$  is the  $D$ -dimensional vector corresponding to the  $(i, j)$ -th spatial patch on the image; and *ii*) a  $D$ -dimensional class token  $X^C$ , such that  $\mathcal{B}(X) = (X^F, X^C)$ .

Given support images  $\{S_c\}_{c=1}^C$  for each personal-level class and a query image  $Q$ , we compute the corresponding feature maps  $\{S_c^F\}_{c=1}^C$  and  $Q^F$ .

#### B. Support feature map processing

For each personal class  $c = 1, 2, \dots, C$ , we apply the same processing steps to the feature map  $S_c^F$ . In the following, we drop the index  $c$  to make the notation less cluttered.

1) *(optional) Bounding box into segmentation map*: If we are given the ground truth bounding box  $b$  for the support image  $S$ , we consider the union of all patches  $P_{i,j}$  that have non-empty intersection with  $b$ , denoted by  $b^P$ , as well as patches bordering  $b^P$ , denoted by  $\partial b^P$ . We then partition a set of corresponding feature vectors  $\{S_{i,j}^F | P_{i,j} \in b^P \cup \partial b^P\}$  into  $k_S$  clusters using the  $k$ -means method (in our implementation,  $k_S$  was empirically chosen to be  $k_S = 5$ ), denoting the set of patches in each cluster as  $\mathcal{K}_r^P$ , with  $r = 1, 2, \dots, k_S$ .

Given that the patches from  $\partial b^P$  are outside of the bounding box, and thus do not belong to the object of interest, we filter out the patch clusters which contain those ‘negative’ patches, thus resulting in an (approximate) segmentation map:

$$\text{seg} = \bigcup_r \{\mathcal{K}_r^P \mid \partial b^P \cap \mathcal{K}_r^P = \emptyset\}. \quad (3)$$

This process allows us to separate the object of interest within the bounding box from the background.

2) *Patch pooling from segmentation map*: Given a (ground truth or approximate) segmentation map  $\text{seg}$  of the support image, we pick the patches  $P_{i,j}$  which partially intersect with  $\text{seg}$ , denoting the set of those patches as  $\text{seg}^P$ .

We compute the patch prototype with simple average over patches in  $\text{seg}^P$  by

$$\text{proto} := \text{avg}(S_{i,j}^F | P_{i,j} \in \text{seg}^P). \quad (4)$$

3) *Adaptive threshold for class prototype*: To pick the patches of interest from the query object, we choose a feature distance metric and a threshold to determine patches of interest on the query image. As a distance metric between feature vectors, we pick the widely used cosine similarity metric. To determine the distance threshold, we use the information about positive and negative patches on the support image.

More concretely, we denote  $\text{seg}^P$  to be a set of patches that have non-empty intersection with the segmentation map  $\text{seg}$ , and  $n\text{seg}^P$  to be a set of patches that have empty intersection with  $\text{seg}$ . We compute the set of positive patch distances and the set of negative patch distances

$$\text{pd} = \{\text{dist}(S_{i,j}^F, \text{proto}) \mid P_{i,j} \in \text{seg}^P\}, \quad (5)$$

$$\text{nd} = \{\text{dist}(S_{i,j}^F, \text{proto}) \mid P_{i,j} \in n\text{seg}^P\}. \quad (6)$$

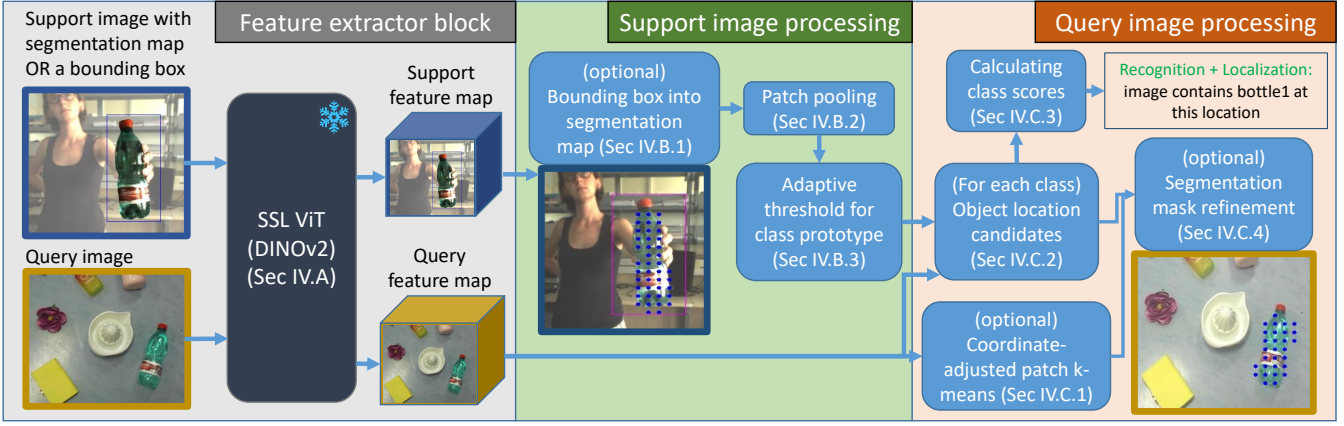


Fig. 2: High-level overview of our Swiss DINO system.

We also remove some possible patch outliers (by removing the highest 5 percent from  $pd$  and the lowest 5 percent from  $nd$ ) to find positive and negative thresholds by  $ptr = \text{percentile}(pd, 95)$  and  $ntr = \text{percentile}(nd, 5)$ . The final adaptive threshold is taken as the minimum of positive and negative thresholds by  $tr := \min(ptr, ntr)$ .

### C. Query feature map processing

Given a tuple  $(proto_c, tr_c)$  for each personal class  $c = 1, \dots, C$  and query feature map  $Q^F$ , we use the following steps to find the patches belonging to the objects of interest.

1) (optional for refined segmentation map) *Coordinate-adjusted patch k-means*: First, agnostic to the set of support classes, we perform a pre-processing step on the query feature map  $Q^F$ . To group together the patches corresponding to the same object, we apply  $k$ -means to the patch feature vectors. In addition, we augment feature vectors with spatial information to reinforce the connectivity of patch clusters:

$$Q_{i,j}^{F,aug} := \text{concat}(Q_{i,j}^F, \alpha_{co}i/N_P, \alpha_{co}j/N_P), \quad (7)$$

where  $\alpha_{co}$  is a coordinate scaling factor aiming to control the effect of spatial information on the resulting clusters.

We cluster the augmented patch features  $Q_{i,j}^{F,aug}$  into  $k_Q$  clusters  $\mathcal{K}_r^Q$ ,  $r = 1, \dots, k_Q$  and save those clusters for the segmentation map refinement step later.

2) *Object location candidates*: For each personal class  $c = 1, \dots, C$ , we find patches in  $Q^F$  which are close enough to the class prototype, resulting in a set of patches we denote as  $pseg_c^{raw}$ :

$$pseg_c^{raw} := \{P_{i,j} \mid \text{dist}(Q_{i,j}^F, proto_c) < tr_c\}, \quad (8)$$

where  $\text{dist}$  is the cosine similarity between feature vectors.

If  $pseg_c^{raw}$  is empty, we choose the patch which is closest to the prototype:  $pseg_c^{raw} = \arg \min_{P_{i,j}} \text{dist}(Q_{i,j}^F, proto_c)$ .

To account for cluttered scenes with similar objects, we split  $pseg_c^{raw}$  into  $L$  connected subsets  $(pseg_c^1, \dots, pseg_c^L)$ , thus generating  $L$  candidates for the location of the object  $c$  on the image.

3) *Calculating class scores*: For each candidate set of patches  $pseg_c^l$ ,  $l = 1, \dots, L$ , we find the class score via patch prototype distance with support image:

$$score_c^l := \text{dist}(\text{avg}(Q_{i,j}^F \mid P_{i,j} \in pseg_c^l), proto_c). \quad (9)$$

We then choose the candidate  $l_{max}$  with the maximum class score as the predicted segmentation map

$$pseg_c := \bigcup_{i,j} \{P_{i,j} \mid P_{i,j} \in pseg_c^{l_{max}}\}, \quad (10)$$

and classification score  $score_c := \max_l score_c^l$ .

From the class score  $score_c$ , we determine whether object  $c$  is on the image. Similar to other score-based approaches in open-set classification [10], a classification threshold needs to be selected for a given dataset to control which predicted masks  $pseg_c$  we would accept, and which ones we would reject. In the actual implementation, the classification threshold needs to be selected for each scenario empirically, while in this work we measure the capability of the method to separate positive examples from negative examples via score precision metric (see Section V-B).

4) (optional) *Segmentation map refinement*: While we can use the  $pseg_c$  as a segmentation map for the object of interest, the map usually covers only part of the object or contains holes. To capture the whole object, we refine the patches from  $pseg_c$  with clusters  $\mathcal{K}_r^Q$  obtained from the previous step using  $k$ -means by

$$pseg_c^{ref} := \bigcup_r \{\mathcal{K}_r^Q \mid \mathcal{K}_r^Q \cap pseg_c \neq \emptyset\}. \quad (11)$$

5) (optional) *Bounding box from segmentation map*: We can also generate a detection bounding box using the refined segmentation map  $pseg_c^{ref}$  by taking the extreme coordinates of the segmentation map.

## V. EXPERIMENTS

### A. Datasets

In this section, we describe the datasets used for the evaluation of our framework. We specifically choose datasets which *i*) contain images of personal objects with different

position/scale/background/lighting variations with fine-level class annotations, and *ii*) include either a segmentation map or bounding box annotations.

1) *PerSEG*: The PerSEG dataset [9] is a convenient choice for one-shot segmentation tasks due to a collection of 40 personalized classes and high-quality segmentation maps. The images contain salient objects that take a large part of the image and with simple, non-cluttered backgrounds, making the segmentation and classification task easier compared to other, noisier, datasets. For few-shot evaluation, we take the first image for each class as a reference image, and test one-shot open-set classification and segmentation on the rest of the images in the class, following [9].

2) *iCubWorld*: The iCubWorld dataset [8] is aimed specifically for robotics application of fine-grained object identification. The dataset contains images from several sessions where a single object is being moved in hand across the scene, and several additional sessions where various objects are filmed in cluttered environments.

We take the subset of the sessions within dataset that contain bounding box annotations, namely *i*) MIX sessions where 50 personal objects are captured in various poses, scales, and lighting conditions, one session per object, and *ii*) TABLE, FLOOR1, FLOOR2, SHELF sessions where a subset of personal objects are scattered on the same scene (altogether 19 personal objects are included in those sessions). For the evaluation of our framework, we take the first image of the MIX session as the support image. We evaluate detection and open-set classification accuracy separately on the collection of MIX sessions (called iCW-single here) and the collection of cluttered sessions TABLE, FLOOR1, FLOOR2, SHELF (called iCW-cluttered here).

## B. Metrics

In this section, we present metrics for the personal object search task, which include a localization metric and two open-set identification metrics.

The precise definitions of the metrics are to follow.

*i*) To measure localization, we employ the common mIoU metric [29] between ground truth localization and predicted localization for a given personal item on the image:

$$\text{mIoU} := \text{avg}_i \left( \text{IoU}(\text{pseg}_{gtc_i}^{ref}(Q_i), \text{gtseg}_i) \right), \quad (12)$$

where  $(Q_i, gtc_i, \text{gtseg}_i)$  are triplets of reference image, index of a personal object on the image, and ground truth localization (segmentation map for PerSEG and bounding box for iCubWorld) of the object on the image, respectively (for cluttered scenes, different personal objects on the same image would correspond to different tuples with same  $Q_i$ ).

*iii*) To measure identification accuracy (denoted by ACC), we check that the predicted score for the ground truth class is the highest among candidate locations of comparison classes near the ground truth location:

$$\text{ACC} := \text{avg}_i \left( \text{acc}(\arg \max_c (\text{score}_c^{loc}(Q_i))) \right), \quad (13)$$

where  $\text{score}_c^{loc}(Q_i)$  is the score of the location candidate of personal class  $c$  with the highest intersection with ground

truth map  $\text{gtseg}_i$  (the score is 0 if there is no intersecting candidate).

*iii*) To measure open set identification accuracy, we employ the Average Precision metric across class scores (denoted by cPREC), which measures how well the class scores for positive examples are separated from the scores from negative examples:

$$\text{cPREC} := \text{avg}_c \left( \text{AP}_i(\text{score}_c^{loc}(Q_i)) \right). \quad (14)$$

We also compare the footprints of the methods in the form of *i*) inference time: how much time does it take for a backbone to process a single image; and *ii*) GPU memory consumption (vRAM): how much GPU memory is required to pass a single image through the backbone, without gradients.

Since the pre- and post-processing steps are done on the CPU, the timings for those steps depend on I/O throughput and specific implementation of those steps. However, since the k-means pre-processing step takes a considerable amount of time in Swiss DINO, we discuss the impact of k-means on the time footprint in Section V-E.3.

## C. Experimental setup

For our experiments, we use DINOv2 backbone (version without registers), with input resized to 448x448 resolution, patch size 14, and patch number  $N_P = 32$ .

To measure the footprints, we use NVIDIA A40 single GPU, with batch size 1 during inference.

For segmentation refinement hyperparameters from Section IV-C.1, we empirically chose  $k_Q = 30, \alpha_{co} = 200$  for the iCubWorld dataset, and  $k_Q = 150, \alpha_{co} = 200$  for the PerSEG dataset. We employ efficient k-means++ method [30] to speed up the clustering step.

## D. Comparison methods

To compare our method against the existing solutions, we focus primarily on the training-free methods for semantic segmentation or detection. To be able to adapt semantic segmentation methods for few-shot prototype-based identification task, we choose methods that provide a feature map or a prototype vector for each predicted segmentation mask or bounding box.

1) *Matcher / PerSAM / DINOv2+M2F*: Matcher [7] and PerSAM [9] are state-of-the-art training-free methods for one-shot semantic segmentation. The methods are based on prompt engineering for a large SAM [33] segmentation model either based on positive-negative pairs [9] or on DINOv2’s features [7]. We include Matcher and PerSAM with default SAM ViT-h backbones, as well as PerSAM with smaller SAM ViT-b backbone to compare the footprint efficiency of those methods. We also consider DINOv2+Mask2Former, using an M2F [32] segmentation head on top of DINOv2 [6] specifically trained for segmentation on the ADE20k dataset [34] with coarse-level classes.

To adapt and evaluate PerSAM and Matcher for multi-class identification, we *i*) extract the feature map from the respective backbones (DINOv2 for Matcher, SAM for

TABLE I: Results on the iCubWorld dataset for the object detection task.

Method	Backbone	mIoU $\uparrow$		cPREC $\uparrow$		ACC $\uparrow$		Time (ms) $\downarrow$	vRAM (MB) $\downarrow$
		single	cluttered	single	cluttered	single	cluttered		
YOLOv8-seg [31]	YOLOv8-s	54.2	6.5	8.1	10.6	9.5	11.6	7.8	390
	YOLOv8-m	56.0	8.2	10.8	13.4	11.2	17.0	12.6	520
	YOLOv8-l	53.1	7.6	10.8	10.4	9.6	6.0	12.2	676
Swiss DINO (ours)	DINOv2 (ViT-s)	65.7	49.8	61.1	54.8	46.8	67.3	<b>7.3</b>	<b>152</b>
	DINOv2 (ViT-b)	68.7	50.3	62.5	<b>55.3</b>	65.1	<b>69.1</b>	<b>7.3</b>	444
	DINOv2 (ViT-l)	<b>69.9</b>	<b>53.4</b>	<b>65.7</b>	52.3	<b>68.2</b>	68.7	14.6	1250
DINOv2 bbox oracle (upper bound)	DINOv2 (ViT-s)	-	-	68.9	96.0	70.8	93.0	<b>7.3</b>	152
	DINOv2 (ViT-b)	-	-	68.5	96.7	72.4	92.9	<b>7.3</b>	444
	DINOv2 (ViT-l)	-	-	70.6	94.4	74.6	94.2	14.6	1250

TABLE II: Results on the PerSEG dataset for the semantic segmentation task.

Method	Backbone	mIoU $\uparrow$	cPREC $\uparrow$	ACC $\uparrow$	Time (ms) $\downarrow$	vRAM (MB) $\downarrow$
YOLOv8-seg	YOLOv8-s	85.6	29.9	29.0	7.8	390
	YOLOv8-b	88.3	40.8	34.5	12.6	520
	YOLOv8-l	87.4	33.6	32.3	12.2	676
DINOv2+M2F [6], [32]	DINOv2 (ViT-g)+M2F	68.5	46.3	37.7	1415	17980
PerSAM [9]	SAM (ViT-b)	86.1	89.5	84.3	758	1674
PerSAM [9]	SAM (ViT-h)	<b>89.3</b>	91.8	85.6	1001	6874
Matcher [7]	DINOv2+SAM (ViT-h)	76.6	<b>91.9</b>	<b>86.7</b>	3787	8670
Swiss DINO (ours)	DINOv2 (ViT-s)	83.5	91.4	82.0	<b>7.3</b>	<b>152</b>
	DINOv2 (ViT-b)	83.5	90.4	81.5	<b>7.3</b>	444
	DINOv2 (ViT-l)	82.4	89.5	81.3	14.6	1250
DINOv2 bbox oracle (upper bound)	DINOv2 (ViT-s)	-	99.9	97.6	7.3	152
	DINOv2 (ViT-b)	-	98.8	96.0	7.3	444
	DINOv2 (ViT-l)	-	98.8	98.0	14.6	1250

PerSAM); *ii*) average the features over the support and predicted query masks to get the prototypes; and *iii*) apply cosine distance to calculate  $score_c(Q)$  for each query image.

To adapt DINOv2+M2F for personal object search task, we use pre-softmax feature vectors of M2F head. We then perform the same steps as our method, but using DINOv2+M2F as an alternative backbone.

The methods above require a precise segmentation map extracted from the reference image, without good extension to bounding box annotations. Therefore, we don’t consider those methods for the iCubWorld dataset, which only provides bounding box annotation for the reference images of personal objects.

2) *YOLOv8-seg*: YOLOv8-seg is an instance segmentation model based on the state-of-the-art YOLOv8 [31] lightweight detection method, and pre-trained on COCO [35] dataset. It is particularly convenient for us, since the model outputs feature vectors for each of the candidate bounding boxes and masks.

To adapt the model for the personal object search task, we *i*) extract the support prototype vector for the bounding box with the highest IoU and the ground truth bounding box; *ii*) find the query bounding box with the closest prototype on the query image to get IoU score; and *iii*) extract the query prototype from the ground-truth query bounding box to calculate cPREC and ACC. We apply YOLOv8-seg on the iCubWorld dataset for detection and on the PerSEG dataset for segmentation.

3) *DINOv2 bounding box oracle*: To have an upper bound reference for the identification metrics cPREC and ACC, we assume that the ground truth location of the personal object on the test image is known. We call this method “DINOv2 bbox oracle” which has similar computational resources.

Knowing ground truth bounding boxes, we crop support images  $S_c$  and a query image  $Q$  into  $S_c^{bb}$  and  $Q^{bb}$  respectively. We then use DINOv2 class tokens  $(S_c^{bb})^C$  and  $(Q^{bb})^C$  as prototypes and compute the class score  $score_c(Q)$  as the cosine distance between corresponding prototypes. Then, we can calculate cPREC and ACC metrics using  $score_c(Q)$  as utilized before.

### E. Main results and discussion

1) *Results on iCubWorld*: As we see from Table I, Swiss DINO significantly outperforms the lightweight comparison method YOLOv8-seg on personal object detection on the iCubWorld dataset. Swiss DINO achieves 16%/46% IOU improvement on single-object and cluttered scenes, respectively. Swiss DINO also shows significant improvement in personal object identification, showing 55%/40% cPREC open-set score improvement, and 57%/52% classification accuracy improvement on single and cluttered scenes respectively.

Significant mIOU gap in cluttered scenes compared to YOLOv8-seg method is caused by the wrong bounding box picked as the prediction, and the large gap in cPREC and ACC metrics is caused by a large number of false positive

predictions near the ground truth location of the object. Given that our adaptation of YOLOv8-seg chooses the bounding box with the feature vector closest to the ground truth class prototype, this shows the poor separation of the feature vectors for fine-grained classes.

Compared to the bounding box-oracle method, on the iCubWorld-single dataset the cPREC and ACC scores of our approach are only about 5% below the upper bound, while on the iCubWorld-cluttered dataset, we observe a more significant 25% accuracy gap, likely due to smaller object scale and presence of similar objects on the cluttered images.

2) *Results on PerSEG*: From Table II we see that Swiss DINO also outperforms YOLOv8-seg on semantic segmentation on the PerSEG dataset in terms of personal object identification metrics (50% cPREC improvement) and classification accuracy (48% improvement), while maintaining similar computational footprint and slightly smaller IoU on segmentation maps.

Compared to DINOv2+M2F, Swiss DINO shows 25% IoU, 45% cPREC, and 49% ACC improvement, while using a much smaller backbone. This again shows how fine-tuning of segmentation head on a coarse-level dataset harms discriminative properties of the features in personalized scenarios. Compared to heavy SAM-based Matcher and PerSAM-b/h, Swiss DINO achieves 100× backbone inference time speedup and 10× improved GPU memory usage while maintaining competitive segmentation and identification scores.

Overall, the results show outstanding capabilities for zero-shot transfer of DINOv2 feature maps onto new tasks (i.e., segmentation and detection) and personalized classes compared to other backbones trained on large datasets, namely: CNN-based YOLOv8 architecture, specialized Mask2Former segmentation head and SAM foundation model.

3) *Impact of k-means*: In the on-server implementation of our method, most of the inference time is spent on the k-means pre-processing step for the query images (Step IV-C.1), which is done on the CPU. This is because k-means is performed on  $N_P^2 = 1024$  feature vectors of  $D = 384/768/1024$  dimensions from vit-s/b/l backbones respectively.

The k-means method takes 0.19/0.23/0.25 seconds per query image for vit-s/b/l backbones respectively, which is about 95% of the overall inference time, compared to about 10 milliseconds spent on query feature map extraction (Step IV-A), and 1 millisecond spent on the rest of the query post-processing (Steps IV-C.2-IV-C.5). Our experiments were done on 64 Intel(R) Xeon(R) Gold 5218 CPU @ 2.30GHz cores.

Note that the k-means refinement step for query images is optional and does not affect identification metrics cPREC and ACC, since the class score  $score_c(Q)$  is computed using the non-refined segmentation map. Therefore, the non-refined version is perfect for applications that only need partial localization information (e.g., a single point on the object of interest). For mIOU scores for segmentation maps without refinement, see Table IV.

Also note that even with costly refinement step, the overall

TABLE III: Ablation on the choice of the vision transformer backbone on the iCW-single dataset. The number in parenthesis is the number of patches along the image side.

Backbone	mIoU ↑	cPREC ↑	Time (ms) ↓	vRAM (MB) ↓
OpenCLIP-ViT-b (14)	29.7	52.3	8.1	374
OpenCLIP-ViT-l (16)	35.4	61.9	16.6	934
DeiT-s (14)	39.6	50.3	6.4	112
DeiT-b (14)	37.5	60.3	6.7	388
DINO-ViT-s (32)	58.6	54.7	8.5	182
DINO-ViT-b (32)	56.3	54.0	8.6	554
DINOv2-ViT-s (32)	65.7	61.1	7.3	152
DINOv2-ViT-b (32)	68.7	62.5	7.3	444
DINOv2-ViT-l (32)	69.9	65.7	14.6	1250

TABLE IV: Ablation on the number of clusters  $k_Q$  and coordinate scaling  $\alpha_{co}$ . All experiments are done with the DINOv2 (ViT-s) backbone.

$k_Q$	$\alpha_{co}$	iCW-single mIoU	iCW-clut mIoU	PerSEG mIoU
$\infty$	0	37.2	7.3	81.1
150	200	51.2	32.0	<b>83.5</b>
60	200	60.4	47.4	80.3
30	200	<b>65.7</b>	<b>49.8</b>	75.7
30	0	68.7	43.2	<b>79.5</b>
30	50	<b>69.2</b>	45.3	78.5
30	200	65.7	<b>49.8</b>	75.7

inference time is still 2× lower compared to heavy-weight methods like PerSAM and Matcher, while still maintaining significant gains by 10× on vRAM footprint.

### E. Ablations

1) *ViT backbone ablations*: In this section, we motivate the choice of DINOv2 backbone by comparing the accuracy of the method using popular self-supervised vision transformer backbones: DINO [26], DeiT [36] and OpenCLIP [37].

As we see from Table III, DINOv2 outperforms other backbones by 7-11% in IoU and 1-5% in cPREC metrics, while maintaining similar or better computational footprints.

2) *Hyperparameter ablations*: In this section, we analyze the effect of hyperparameters. In particular, we see how  $k_Q$  and  $\alpha_{co}$  used in segmentation mask refinement affect IoU score for the iCW-single, iCW-cluttered, and PerSEG datasets. From Table IV, we see that reducing the number of clusters  $K_Q$  improves the IoU score by 32%/42% on the iCW single and cluttered respectively, while the number of clusters needs to be high on the cleaner PerSEG dataset to exclude the false positive patches. Coordinate scaling  $\alpha_{co}$  does not affect the accuracy metrics much (yielding 5% improvement on cluttered scenes), however, we see from qualitative observations that the inclusion of coordinate scaling makes the results more robust to scene variation. Following this result, we choose  $k_Q = 30$ ,  $\alpha_{co} = 200$  for the iCubWorld dataset, and  $k_Q = 150$ ,  $\alpha_{co} = 200$  for the PerSEG dataset.

## VI. CONCLUSION

In this work, we have introduced a novel problem and metrics for the personal object search task, which is directly related to practical robot vision tasks performed by mobile and robotic systems such as home appliances and robotic manipulators, in which the system needs to localize all present objects of interest in a cluttered scene, where each object is only referenced by a few images.

To address this task, we have introduced Swiss DINO that leverages SSL-pretrained DINOv2's feature maps having strong discriminative and localization properties. Swiss DINO presents novel clustering-based segmentation/detection mechanisms to alleviate the need for additional specialized modules for such dense prediction tasks.

We compare our framework to common lightweight solutions, as well as heavy transformer-based solutions. We show significant improvement (up to 55%) of segmentation and recognition accuracy compared to the former methods, and significant footprint reduction of backbone inference time (up to 100×) and GPU consumption (up to 10×) compared to the latter methods, allowing seamless implementation on robotic devices.

Altogether, this work shows the power and versatility of self-supervised transformer models on personal object search and various downstream tasks. In future work, we plan to extend Swiss DINO for continually learning new generic as well as new personal objects.

## REFERENCES

- [1] Francesco Barbato, Umberto Michieli, Jijoong Moon, Pietro Zanuttigh, and Mete Ozay, "Cross-architecture auxiliary feature space translation for efficient few-shot personalized object detection," in *IROS*, 2024.
- [2] Tyler L Hayes and Christopher Kanan, "Online Continual Learning for Embedded Devices," *CoLLAs*, 2022.
- [3] Umberto Michieli and Mete Ozay, "Online continual learning for robust indoor object recognition," in *IROS*. IEEE, 2023.
- [4] N. Catalano and M. Matteucci, "Few shot semantic segmentation: a review of methodologies and open challenges," *arXiv:2304.05832*, 2023.
- [5] Vardan Papyan, X. Y. Han, and David L. Donoho, "Prevalence of neural collapse during the terminal phase of deep learning training," *PNAS*, vol. 117, no. 40, pp. 24652–24663, 2020.
- [6] 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:2304.07193*, 2023.
- [7] Yang Liu, Muzhi Zhu, Hengtao Li, Hao Chen, Xinlong Wang, and Chunhua Shen, "Matcher: Segment anything with one shot using all-purpose feature matching," in *ICLR*, 2024.
- [8] Sean Ryan Fanello, Carlo Ciliberto, Matteo Santoro, Lorenzo Natale, Giorgio Metta, Lorenzo Rosasco, and Francesca Odone, "iCub World: Friendly Robots Help Building Good Vision Data-Sets," in *CVPRW*, 2013.
- [9] Renrui Zhang, Zhengkai Jiang, Ziyu Guo, Shilin Yan, Junting Pan, Hao Dong, Yu Qiao, Peng Gao, and Hongsheng Li, "Personalize segment anything model with one shot," in *ICLR*, 2024.
- [10] Lawrence Neal, Matthew Olson, Xiaoli Fern, Weng-Keen Wong, and Fuxin Li, "Open set learning with counterfactual images," in *ECCV*, 2018, pp. 613–628.
- [11] Umberto Michieli and Pietro Zanuttigh, "Incremental learning techniques for semantic segmentation," in *CVPRW*, 2019.
- [12] Qi She, Fan Feng, Xinyue Hao, Qihan Yang, Chuanlin Lan, Vincenzo Lomonaco, Xuesong Shi, Zhengwei Wang, Yao Guo, Yimin Zhang, et al., "Openloris-object: A robotic vision dataset and benchmark for lifelong deep learning," in *ICRA*. IEEE, 2020, pp. 4767–4773.
- [13] Jonas Frey, Hermann Blum, Francesco Milano, Roland Siegwart, and Cesar Cadena, "Continual adaptation of semantic segmentation using complementary 2d-3d data representations," *IEEE RA-L*, vol. 7, no. 4, pp. 11665–11672, 2022.
- [14] Umberto Michieli and Pietro Zanuttigh, "Continual semantic segmentation via repulsion-attraction of sparse and disentangled latent representations," in *CVPR*, 2021, pp. 1114–1124.
- [15] Kuan Xu, Chen Wang, Chao Chen, Wei Wu, and Sebastian Scherer, "Aircode: A robust object encoding method," *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 1816–1823, 2022.
- [16] Fabio Cermelli, Massimiliano Mancini, Yongqin Xian, Zeynep Akata, and Barbara Caputo, "Prototype-based incremental few-shot segmentation," in *BMVC*, 2021.
- [17] Jiacheng Chen, Bin-Bin Gao, Zongqing Lu, Jing-Hao Xue, Chengjie Wang, and Qingmin Liao, "Apanet: Adaptive prototypes alignment network for few-shot semantic segmentation," *IEEE T-MM*, 2022.
- [18] Nanqing Dong and Eric P Xing, "Few-shot semantic segmentation with prototype learning," in *BMVC*, 2018, vol. 3.
- [19] Gen Li, Varun Jampani, Laura Sevilla-Lara, Deqing Sun, Jonghyun Kim, and Joongkyu Kim, "Adaptive prototype learning and allocation for few-shot segmentation," in *CVPR*, 2021, pp. 8334–8343.
- [20] Kaixin Wang, Jun Hao Liew, Yingtian Zou, Daquan Zhou, and Jiashi Feng, "Panet: Few-shot image semantic segmentation with prototype alignment," in *ICCV*, 2019, pp. 9197–9206.
- [21] Amirreza Shaban, Shray Bansal, Zhen Liu, Irfan Essa, and Byron Boots, "One-shot learning for semantic segmentation," 2017.
- [22] Khoi Nguyen and Sinisa Todorovic, "Feature weighting and boosting for few-shot segmentation," in *ICCV*, 2019.
- [23] Chaitanya Mitash, Fan Wang, Shiyang Lu, Vikedo Terhija, Tyler Garaas, Felipe Polido, and Manikantan Nambi, "Armbench: An object-centric benchmark dataset for robotic manipulation," in *ICRA*. IEEE, 2023, pp. 9132–9139.
- [24] Vincenzo Lomonaco and Davide Maltoni, "Core50: a new dataset and benchmark for continuous object recognition," in *CoRL*, 2017.
- [25] Giulia Pasquale, Carlo Ciliberto, Francesca Odone, Lorenzo Rosasco, and Lorenzo Natale, "Are we done with object recognition? the icub robot's perspective," *RAS*, vol. 112, pp. 260–281, 2019.
- [26] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin, "Emerging properties in self-supervised vision transformers," in *ICCV*, 2021, pp. 9650–9660.
- [27] Oriane Siméoni, Gilles Puy, Huy V Vo, Simon Roburin, Spyros Gidaris, Andrei Bursuc, Patrick Pérez, Renaud Marlet, and Jean Ponce, "Localizing objects with self-supervised transformers and no labels," *arXiv:2109.14279*, 2021.
- [28] Luke Melas-Kyriazi, Christian Rupprecht, Iro Laina, and Andrea Vedaldi, "Deep spectral methods: A surprisingly strong baseline for unsupervised semantic segmentation and localization," in *CVPR*, 2022.
- [29] Gabriela Csurka, Diane Larlus, Florent Perronnin, and France Meylan, "What is a good evaluation measure for semantic segmentation?," in *BMVC*, 2013, vol. 27, pp. 10–5244.
- [30] David Arthur, Sergei Vassilvitskii, et al., "k-means++: The advantages of careful seeding," in *Soda*, 2007, vol. 7, pp. 1027–1035.
- [31] Glenn Jocher, Ayush Chaurasia, and Jing Qiu, "Ultralytics YOLOv8," 2023.
- [32] Bowen Cheng, Ishan Misra, Alexander G Schwing, Alexander Kirillov, and Rohit Girdhar, "Masked-attention mask transformer for universal image segmentation," in *CVPR*, 2022, pp. 1290–1299.
- [33] 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 anything," *ICCV*, 2023.
- [34] Bolei Zhou, Hang Zhao, Xavier Puig, Sanja Fidler, Adela Barriuso, and Antonio Torralba, "Scene parsing through ade20k dataset," in *CVPR*, 2017, pp. 633–641.
- [35] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick, "Microsoft coco: Common objects in context," in *ECCV*, 2014, pp. 740–755.
- [36] Hugo Touvron, Matthieu Cord, Matthijs Douze, Francisco Massa, Alexandre Sablayrolles, and Hervé Jégou, "Training data-efficient image transformers & distillation through attention," in *International conference on machine learning*. PMLR, 2021, pp. 10347–10357.
- [37] Mehdi Cherti, Romain Beaumont, Ross Wightman, Mitchell Wortsman, Gabriel Ilharco, Cade Gordon, Christoph Schuhmann, Ludwig Schmidt, and Jenia Jitsev, "Reproducible scaling laws for contrastive language-image learning," in *CVPR*, 2023, pp. 2818–2829.