

FoundYou: A Unified Model for Personalized Segmentation and Retrieval

Gabriele Trivigno^{*}  Marcos Alfaro^{*1,2}  Claudia Cuttano^{*1} 
Gabriele Berton^{} Luis Payá^{2,3}  Carlo Masone¹ 

¹ Politecnico di Torino ² Miguel Hernández University of Elche

³ Valencian Graduate School and Research Network of Artificial Intelligence

<https://gali13o.github.io/FoundYou/>

Abstract. Personalized segmentation and personalized retrieval both aim to identify the same physical object across different images. While the former localizes the object within a target image, the latter retrieves images where it appears. Despite this shared instance-level objective, the two tasks have largely evolved separately and are addressed with distinct solutions. In this work, we introduce FoundYou, a unified framework built on the observation that Segment Anything 2 (SAM 2), trained to preserve object identity across video frames, inherently captures instance-level cues. We leverage this property to match objects across independent images, enabling segmentation and retrieval to emerge as two outcomes of the same instance alignment process. This unified view unlocks new capabilities beyond traditional benchmarks, including few-shot personalized retrieval and promptable personalized segmentation with flexible prompts. Extensive experiments show consistent gains over unified and task-specific methods, including +18.4 mIoU on PerMIS and +17.8 mAP on ILIAS. Performance scales with additional references and remains robust to weaker prompts. Beyond personalization, FoundYou achieves state-of-the-art results on category-level retrieval benchmarks. Notably, our approach keeps the SAM 2-*small* model entirely frozen and adds only 5.9 M trainable parameters, yielding a 52 M-parameter model that is over $75\times$ faster and $20\times$ smaller than the only prior unified solution.

1 Introduction

Suppose you provide a model with an example of *your handbag* and ask it to segment the same handbag in completely different photos, in unrelated scenes that may contain other similar items (*cf.* Fig. 1-a). Personalized segmentation [74] formalizes this setting: given a visual reference identifying an object, the goal is to segment the same physical object in new images. Unlike few-shot segmentation [12, 33, 75], the objective is not to recognize a category (*e.g.*, *any handbag*), but to preserve identity, *i.e.* distinguish a particular instance from

^{*} Equal contribution.




Fig. 1: FoundYou: a Unified Model for Personalized Segmentation & Retrieval. Given a visual prompt of a user’s handbag, FoundYou segments the *same physical instance* in a new image (a) and retrieves it from a gallery by ranking images via instance-level similarity (b). We compare performance vs. speed against the only prior unified solution, PDM [49], and the strongest task-specific methods: retrieval performance is shown in (c) as mAP vs. FPS, and segmentation in (d) as mIoU vs. FPS. FoundYou improves over PDM by +17.8 mAP and +18.4 mIoU while being over 75× faster (note: x-axis in log scale). It outperforms also task-specific methods, while being the fastest in segmentation and matching the speed of the fastest retrieval competitor.

visually similar ones. Now consider a related objective: instead of segmenting your *handbag* in a given image, you want to retrieve all the photos in a personal collection where it appears (*cf.* Fig. 1-b). Personalized retrieval [49] captures this setting: given the same reference, the task becomes ranking images according to whether they depict that exact instance. Albeit producing different outputs, both tasks fundamentally require instance-level matching, from which a segmentation mask or image-level score is derived. This capability is central to applications, from robotics and embodied perception [5, 28, 48] to cultural heritage digitization [52, 71], personal media organization [1, 27], and creative workflows [11, 37].

Despite this shared objective, personalized segmentation and retrieval have evolved largely independently. Segmentation-focused approaches, such as PerSAM [74], solve the task via correspondence followed by mask prediction with SAM [26], often relying on external models such as DINOv2 [32, 38] or diffusion encoders [45, 57]. Conversely, retrieval-specific methods learn instance-aware representations [9, 53, 56], but do not provide dense localization. PDM [49] represents the closest attempt to bridge the two tasks. However, its reliance on Stable Diffusion leads to prohibitively slow inference (*cf.* Fig. 1-c,d). Additionally, PDM relies on features extracted at different denoising timesteps for segmentation and retrieval, and still relies on prompting SAM for mask generation.

In this work, we take a different perspective by observing that Segment Anything 2 (SAM 2) [43], trained to track objects across video frames, implicitly learns fine grained instance-aware representations. Building on this, we propose

FoundYou, a unified framework for personalized segmentation and retrieval that repurposes SAM 2 memory attention mechanism: instead of propagating identity across temporally adjacent frames, we use it to match a specific instance across *independent* images. This shift, however, entails two fundamental challenges. First, because SAM 2 is trained on natural videos, its representations inherit a temporal continuity bias: inter-frame variations are small and scene context remains stable [12, 55]. In personalized segmentation and retrieval, however, matching occurs across unrelated images without positional or contextual continuity. To address this, we introduce lightweight feature adaptation layers trained for identity discrimination across distinct images, together with a self-distillation objective that preserves spatial structure of SAM 2 representations, ensuring high-quality segmentation masks. Second, SAM 2 lacks a mechanism to assign an image-level similarity score. We therefore introduce a lightweight retrieval head that aggregates dense correspondences from memory attention into a global score, while reusing the original mask decoder for segmentation. Segmentation and retrieval thus emerge from the same instance alignment process.

Since instance matching is realized through memory attention over a set of reference features, FoundYou naturally supports multiple references at inference time. We therefore introduce the task of *few-shot personalized retrieval*, reflecting practical cases where users provide several examples of the same object, for example, taking multiple photos of a *handbag* to search for it more reliably. Additionally, FoundYou supports promptable personalized segmentation with lightweight annotations (*e.g.*, points or boxes), removing the need for full masks. Experiments show that FoundYou consistently outperforms both *i*) the only prior unified solution and *ii*) task-specific methods (*cf.* Fig. 1-c,d). Compared to PDM [49], we improve by +18.4 mIoU on PerMIS (segmentation) and +17.8 mAP on ILIAS (retrieval), while also surpassing *specialized* approaches such as PerSAM-F [74] (+14.8 mIoU) and AMES [53] (+6.1 mAP). Performance scales with additional references and remains robust to weaker annotations such as point prompts. Beyond personalization, FoundYou sets a new state-of-the-art on *category-level* retrieval across 6 domains. Notably, we keep SAM 2 frozen and introduce only **5.9 M trainable parameters**, resulting in a compact 52 M-parameter model 75× faster and 20× smaller than prior unified methods [49].

2 Related Work

Segment Anything 2 for downstream tasks. SAM 2 [43] introduces a unified architecture for image and video segmentation, pretrained to segment objects and propagate masks across frames. Building on this, several works explore SAM 2 for downstream tasks. A first line of research relies on external models such as DINOv2 [38] or BEiT [61] to derive task-specific prompts, enabling few-shot [36, 69], open-vocabulary [68] and referring segmentation [46], also in remote sensing applications [47]. Related to ours, a second line of work operates directly on SAM 2. SAMWISE [13] adapts SAM 2 features through motion- and language-aware conditioning, while GeoSAM2 [14] introduces geometric struc-

ture for 3D segmentation. SANSa [12] and FS-SAM2 [19] are close in spirit to our work. Both leverage SAM 2 to guide segmentation in new images. However, their objective is to generalize matching to broader semantic categories. Conversely, our goal is to *reinforce instance-level discrimination*. Rather than abstracting instance-specific cues, we aim to sharpen them, enabling reliable matching of the same physical object across independent images.

Personalized Segmentation aims to segment user-designated concepts in new images. Recent approaches address this task by first *localizing* the instance in a target image via feature matching with a reference, and then *segmenting* it by prompting SAM [26]. A key challenge lies in the representation used for matching. Early approaches such as PerSAM [74] rely on SAM features, which lack sufficient semantic structure for reliable correspondence. Matcher [32] addresses this by using DINOv2 [38], whose strong semantic representations improve matching but tend to collapse multiple similar instances of the same semantic category. To recover instance-level cues, PDM [49] explores diffusion-based features, while [57] further combines them with DINOv2. Despite improved localization, these diffusion-based pipelines remain computationally expensive, require careful timestep selection, and still rely on SAM for mask generation. In this work, we propose to leverage the instance-aware representations learned by SAM 2 [43] to jointly perform feature matching and segmentation within a unified framework.

Instance Retrieval. Early retrieval works primarily focus on class-level discrimination [18, 35, 50, 63], while instance-level retrieval has been explored in domain-specific settings such as landmarks [9, 66], place recognition [2, 8, 25, 59], and product or fashion collections [3, 20]. ILIAS [27] extends this setting to diverse object types, exposing challenges such as clutter and visually similar distractors, where global embeddings often lack sufficient instance-level discrimination [9, 53]. To address this, prior works refine retrieval candidates using local descriptors and geometric verification [53, 60, 76], including image-matching models such as [17, 21, 54]. While effective for matching covisible scene structures, improving performance in landmark-style retrieval [39, 51], these approaches prioritize geometric consistency and contextual alignment rather than object-centric identity. Specifically designed for re-ranking, AMES [53] learns a transformer on top of frozen DINOv2 features, but remains grounded in representations that emphasize high-level semantics. In contrast, SAM 2 [43], pretrained for object tracking in videos, inherently learns instance-aware representations. We leverage this property with a lightweight retrieval head, achieving strong instance-level discrimination while retaining segmentation capability. Furthermore, we introduce a few-shot retrieval setting, where multiple reference examples are provided at test time, naturally supported by our framework without retraining.

3 Method

Task Definition. A user provides a *visual prompt* $V_r = (\mathcal{I}_r, \mathcal{A}_r)$ where $\mathcal{I}_r \in \mathbb{R}^{H \times W \times 3}$ is a reference image, and $\mathcal{A}_r$ specifies the target instance within the




Fig. 2: Overview of FoundYou. We build on a frozen SAM 2, reinterpreting its memory mechanism to match instances across independent images rather than across temporally adjacent video frames. Given an object specified in the reference image I_r , the Memory Encoder stores its representation S_r in the Memory Bank. Target features F_t are matched with S_r through Memory Attention, producing dense memory features F_t^{mem} , visualized via 3D PCA. From these, the frozen Segmentation Decoder produces a mask, while our lightweight Retrieval Decoder aggregates them into an image-level score, enabling personalized segmentation and retrieval within a single framework.

image. A_r can be a binary mask (e.g. for personalized segmentation [49, 73, 74]), a bounding box (e.g. for personalized retrieval [27, 40, 42]), or a point. When the reference image contains a single object and no box is provided, the box coincides with the image boundaries. Given a target image $I_t \in \mathbb{R}^{H \times W \times 3}$, the goal of **personalized segmentation** is to predict a binary mask $M_t \in \{0, 1\}^{H \times W}$ corresponding to the *same physical instance* indicated by V_r . We define the segmentation function as $M_t = f_{\text{seg}}(V_r, I_t)$. In **personalized retrieval**, given a gallery $\mathcal{G} = \{I_i\}_{i=1}^N$, the objective is to rank database images according to *instance-level relevance*, i.e., whether they depict the same physical object instance as specified by the visual prompt V_r . We define a similarity function $s_i = f_{\text{ret}}(V_r, I_i) \in \mathbb{R}$, which induces a ranking over $\mathcal{G}$ based on decreasing similarity.

Overview. An overview of our method is in Fig. 2. In this work, we observe that SAM 2, unlike other vision foundation models, is *explicitly* optimized to preserve object identity during mask propagation across video frames, implicitly learning instance-aware features. However, such representations are learned under a spatio-temporal continuity prior [12, 55], where identity is maintained within the same scene and under limited positional variation. Personalized segmentation and retrieval, by contrast, require discriminating and matching the same physical instance across *independent* images, where scene context and spatial configuration may vary arbitrarily. Furthermore, SAM 2 lacks a mechanism to assign image-level similarity scores between the reference instance and candidate images, which is necessary to rank results in personalized retrieval. To bridge

this gap, we introduce a lightweight adaptation to remove the implicit tracking prior in SAM 2 features, while preserving and enhancing instance awareness (Sec. 3.1). The adapted representations establish dense correspondences between a visual prompt and a target image without relying on shared scene context. To produce outputs (Sec. 3.2): *i*) for segmentation, we reuse the frozen SAM 2 mask decoder to segment the matched instance; *ii*) for retrieval, we introduce a lightweight decoder that aggregates dense correspondences into an image-level score. We train the adapters and retrieval head on pairs of positive and negative images, and add a self-distillation loss to preserve the spatial structure in the adapted features, required for accurate segmentation (Sec. 3.3). This formulation casts personalized segmentation and retrieval as two outcomes of the same instance-level alignment process, enabling both tasks within a single framework.

3.1 Personalizing SAM 2

SAM 2 [43] is designed to preserve object identity across frames by encoding a reference instance into a **Memory Bank** and matching new frames through attention. We reinterpret this mechanism such that it can work *across independent images*. The reference image $\mathcal{I}_r$ is processed by the **Image Encoder**, producing features $\mathcal{F}_r \in \mathbb{R}^{\frac{H}{16} \times \frac{W}{16} \times D}$. To construct a representation of the reference instance, the **Memory Encoder** fuses visual features with a reference mask $\hat{\mathcal{M}}_r$, which is either: *i*) provided, in personalized segmentation, *i.e.*, $\hat{\mathcal{M}}_r = \mathcal{A}_r$ or *ii*) predicted from a bounding box in retrieval, *i.e.*, $\hat{\mathcal{M}}_r = \text{Mask Decoder}(\mathcal{F}_r, \mathcal{A}_r)$. Formally, the **Memory Encoder** computes the memory representation $\mathcal{S}_r \in \mathbb{R}^{\frac{H}{16} \times \frac{W}{16} \times D}$ as:

$$\mathcal{S}_r = \text{conv}(\mathcal{F}_r) + \text{conv}(\hat{\mathcal{M}}_r) \quad (1)$$

This representation $\mathcal{S}_r$ is stored in the **Memory Bank** (see Fig. 2)¹. Given a target image $\mathcal{I}_t$, the **Image Encoder** extract features $\mathcal{F}_t \in \mathbb{R}^{\frac{H}{16} \times \frac{W}{16} \times D}$. The target features are then matched to the representation $\mathcal{S}_r$ through **Memory Attention**:

$$\mathcal{F}_t^{\text{mem}} = \text{MHCA}(Q(\mathcal{F}_t), K(\mathcal{S}_r), V(\mathcal{S}_r)) \quad (2)$$

where MHCA is a Multi-Head Cross Attention.

This produces memory-conditioned features $\mathcal{F}_t^{\text{mem}} \in \mathbb{R}^{\frac{H}{16} \times \frac{W}{16} \times D}$, which encode dense correspondences between the target image and the reference instance.

However, SAM 2 is trained to preserve instance identity across consecutive video frames, where appearance and position change smoothly over time. In our setting, identity must instead be matched across independent images, where no such continuity exists. To bridge this gap, we introduce a lightweight feature adaptation through AdaptFormer [10] blocks inserted in the last two layers of the **Image Encoder**. Given down- and up- projection matrices $\mathbf{W}_{\text{down}} \in \mathbb{R}^{d \times \tilde{d}}$, $\mathbf{W}_{\text{up}} \in \mathbb{R}^{\tilde{d} \times d}$, an AdaptFormer block operates token-wise:

$$A(x) = \sigma(x \cdot \mathbf{W}_{\text{down}}) \cdot \mathbf{W}_{\text{up}}, \quad (3)$$

¹ If multiple reference images are available, their memory representations are concatenated; with slight abuse of notation, we denote the result as $\mathcal{S}_r$. See Supplementary.

where σ is a ReLU and $\tilde{d} < d$ is the bottleneck dimensionality. The adapted features are summed in a residual fashion in the backbone transformer blocks:

$$\begin{aligned} x_{\text{self}} &= \text{MHSA}(x), \\ x' &= x_{\text{self}} + \text{MLP}(x_{\text{self}}) + A(x_{\text{self}}), \end{aligned} \quad (4)$$

where MHSA is a Multi-Head Self Attention.

The backbone weights are frozen and we only train projections $\mathbf{W}_{\text{down}}$ and $\mathbf{W}_{\text{up}}$.

3.2 Task-Specific Decoders

The resulting memory-conditioned features $\mathcal{F}_t^{\text{mem}}$ are then consumed by task-specific heads: *i*) a Segmentation Decoder, for personalized segmentation or *ii*) a Retrieval Decoder, for personalized retrieval.

Segmentation Decoder. For segmentation, we directly reuse the frozen SAM 2 mask decoder. A learned output token and the memory-conditioned features $\mathcal{F}_t^{\text{mem}}$ iteratively interact through bi-directional cross attention. The features are then progressively upsampled and fused with higher-resolution features from the **Image Encoder** through a top-down pathway. The final mask $\mathcal{M}_t \in \mathbb{R}^{H \times W}$ is obtained by computing a point-wise similarity between the refined features and the output token, followed by thresholding.

Retrieval Decoder. Among the outputs of its decoder, SAM 2 introduces an *occlusion score*, predicted from the memory-conditioned features, to determine whether the tracked object is present in a frame. This mechanism is defined under temporal continuity, and does not transfer directly to independent images (*cf.* ablation in Sec. 4.6). Nevertheless, because this score is computed directly from the memory-conditioned features, it implies that these representations already encode information about *instance presence*. Building on this observation, we propose a Retrieval Decoder that aggregates the $\mathcal{F}_t^{\text{mem}}$ features into an image-level similarity score *across independent images*. Specifically, we derive a simplified variant of the original decoder: *i*) we remove upsampling layers, FPN fusion and tokens required for mask prediction; *ii*) we introduce a learnable retrieval token; *iii*) we train this token to interact with the memory features through cross attention, aggregating spatial correspondences into a compact global representation. Formally, let $\mathcal{F}_t^{\text{mem}}$ be reshaped into a sequence $\mathbf{F}^{(0)} \in \mathbb{R}^{L \times D}$ with $L = \frac{H}{16} \cdot \frac{W}{16}$, and let $\mathbf{z}^{(0)} \in \mathbb{R}^{1 \times D}$ be a learnable retrieval token. We perform two rounds of bidirectional interaction between token and memory features:

$$\begin{aligned} \mathbf{F}^{(k)} &= \text{MHCA}(\mathbf{F}^{(k-1)}, \mathbf{z}^{(k-1)}), \\ \mathbf{z}^{(k)} &= \text{MLP}(\text{MHCA}(\mathbf{z}^{(k-1)}, \mathbf{F}^{(k)})) \quad k \in \{1, 2\}. \end{aligned} \quad (5)$$

The final retrieval score is obtained as

$$s_t = \text{sigmoid}(\text{MLP}_{D \rightarrow 1}(\mathbf{z}^{(2)})), \quad (6)$$

where $s \in [0, 1]$ estimates the probability that the reference instance is present. Qualitative examples shown in Fig. 3 illustrate how scores reflect instance similarity, assigning high confidence to true matches and low confidence to outliers.

3.3 Personalization Loss

Training is formulated as a binary instance-matching problem. For a given reference image $\mathcal{I}_r$, we construct a candidate set $\mathcal{C}_r = \mathcal{P}_r \cup \mathcal{N}_r$, where $\mathcal{P}_r$ (positives) depicts the same physical instance and $\mathcal{N}_r$ (negatives) contains different instances. The reference is encoded into the **Memory Bank**, and each candidate image $I_j \in \mathcal{C}_r$ is processed independently producing a retrieval score $s_j \in [0, 1]$. We assign binary labels and optimize a binary cross-entropy loss:

$$\mathcal{L}_{\text{ret}} = - \sum_{I_j \in \mathcal{C}_r} [y_j \log(s_j) + (1 - y_j) \log(1 - s_j)]. \quad (7)$$

Since SAM 2 already encodes strong instance-aware representations, randomly sampled negatives would result in trivial pairs (*e.g.*, a coffee mug versus a car), which provide limited supervision and slow convergence [24, 67]. The primary challenge instead lies in distinguishing between visually similar objects within the same semantic category, as shown in Fig. 3, *e.g.* two *different coffee mugs*. A common strategy in retrieval is to mine hard negatives during training [2, 7, 34, 67]. However, online mining introduces significant computational overhead, as it must be repeated periodically while the model evolves [15, 44, 65]. Conversely, we leverage the semantic grouping properties of a frozen DINOv2, *i.e.* $\phi(\cdot)$ to obtain semantically related examples (*e.g.*, other visually similar mug for a reference coffee mug). This decouples the mining procedure from the model itself, and can thus be carried out offline only once. We compute global embeddings and select, for each reference $\mathcal{I}_r$, the most similar non-matching images:

$$\mathcal{N}_r = \operatorname{argmax}_{\mathcal{I}_k \notin \mathcal{P}_r} \operatorname{topK} (\phi(\mathcal{I}_r)^\top \cdot \phi(\mathcal{I}_k)). \quad (8)$$

This results in candidate sets emphasizing intra-class instance discrimination.

Mask-based self-distillation. Our adaptation relaxes the spatio-temporal continuity prior implicitly encoded in SAM 2, enabling instance matching across independent images where scene context varies. However, training the adapter with an image-level objective degrades the dense representations required for accurate mask prediction (*cf.* ablations in Sec. 4.6). To preserve this spatial structure, we introduce a self-distillation objective on that regularizes the adapted model to maintain the segmentation capabilities of the frozen backbone. During training, we compute teacher (*i.e.* SAM 2) and student (*i.e.* FoundYou) mask logits on a given image using *i)* the frozen SAM 2 model and *ii)* our adapted model, respectively. We denote with $P^T \in \mathbb{R}^{H \times W}$ and $P^S \in \mathbb{R}^{H \times W}$ the corresponding foreground probability maps obtained by applying a sigmoid to the predicted mask logits. We minimize a pixel-wise Kullback–Leibler divergence:

$$\mathcal{L}_{\text{dist}} = \frac{1}{HW} \sum_{h,w} \text{KL}(P_r^T(h, w) \parallel P_r^S(h, w)). \quad (9)$$

The final training objective combines the personalization and distillation losses:

$$\mathcal{L} = \mathcal{L}_{\text{ret}} + \lambda_{\text{dist}} \mathcal{L}_{\text{dist}}. \quad (10)$$




Fig. 3: Personalized retrieval and segmentation examples. For each reference image (row 1), we show top-2 retrieved candidates (rows 2–3), along with hard negatives examples (rows 4–5) of visually similar but non-matching objects. The retrieval score s_t is reported for each candidate. The last row shows personalized segmentation results on the top-1 retrieved image, using the same reference (row 1) as visual prompt.

4 Experiments

We evaluate FoundYou against unified approaches and task-specific baselines on:

- (1) **Personalized Segmentation.** Given a visual prompt and a reference mask of an instance, the model is tasked with *segmenting* the same object in a target image. We further introduce a *promptable* setting in which we relax the mask assumption, and rely on box or point prompts to refer the object of interest.

- (2) **Personalized Retrieval.** Given a visual prompt depicting an instance or a landmark, and a gallery of candidates, the model is tasked with *retrieving*, *i.e.*, ranking the candidates according to the probability that they contain the object of interest. Furthermore, we propose a *few-shot* retrieval setting, in which multiple examples are provided to the model.

Datasets. For (1) **Personalized Segmentation**, we evaluate on *PerSeg* [74] and *PerMIS* [49]. *PerSeg* contains 40 personal objects captured under substantial variations in pose, scale, and background. *PerMIS* extends this setting with over 432 query–target pairs, captured in more challenging crowded scenes. For (2) **Personalized Retrieval**, we evaluate on *PerMIR* [49] and *ILIAS* [27]. *Per-*

MIR provides a baseline benchmark, but remains limited in scale, with a gallery of 432 images and 216 queries. *ILIAS* introduces a significantly more challenging setting, comprising 1,232 instances with curated positives, alongside 100M distractors from YFCC100M [58], spanning diverse domains (*e.g.*, landmarks, fashion, products, artworks). To assess generalization beyond strict instance matching, we evaluate on class-level retrieval benchmarks, *Cars196*, *iNaturalist*, *RP2K*, *Stanford Online Products*, *Food2K*, *MET*, and the landmark datasets *Google Landmarks v2*, *Revisited Paris*, and *Oxford*. Such datasets provide a complementary setting, where the objective shifts from retrieving the *same physical instance* to retrieving the *same fine-grained category*. More details in the Supp.

Implementation Details. We use the Small version of SAM 2 (46 M parameters), and keep it entirely frozen. We add and train only 5.9 M parameters, *i.e.* the adapter layers and retrieval head. The model is trained with Adam (LR 1×10^{-4} , batch size 16) for 150k iterations, using 4 positive and 8 negative examples per query. λ_{dist} is set to 0.5. We train FoundYou on UnED [70], following [27, 70].

Evaluation protocol. For segmentation, we follow prior work [49, 74] and report mIoU (%) and bIoU (%). On PerMIR, we perform retrieval over the entire gallery of 432 images. For ILIAS, which contains a gallery of 100M images, we follow [27]: the top-1k most similar candidates are first retrieved from the gallery using SigLIP global similarity [72], and then re-ranked by each method. Following [27, 49, 53, 56, 70], we evaluate retrieval performance using mAP (%).

4.1 Unified Personalized Segmentation and Retrieval.

Baselines. In Tab. 1 we compare against state-of-the-art methods for personalized segmentation and retrieval. For segmentation, we compare with generalist models such as SEEM [77], SegGPT [62], GF-SAM [73] and personalized specialist PerSAM [74], of which we also consider fine-tuned version (PerSAM-F). Concerning retrieval, the most relevant comparison are image matching methods [54], which can be used to verify and re-rank retrieved candidates [4, 6, 56, 76]. Among them we consider recent methods such as RoMav2 [16], MatchAnything [21], LightGlue [30] and MAST3R [29]. We use the number of inliers as similarity score, following standard practice [4, 51]. We also evaluate re-ranking specialist methods such as RRT [56] and the recent state-of-the-art AMES [53]. Moreover, we compare with the only prior work that supports both tasks, PDM [49], which relies on StableDiffusion [45] features, and SAM [26] for producing masks.

Discussion. In personalized segmentation, FoundYou achieves the best performance across datasets. On the challenging PerMIS benchmark, which contains crowded scenes with visually similar distractors, FoundYou surpasses PerSAM-F by +14.8 mIoU and PDM by +18.4 mIoU. The improvement in bIoU is even larger (+22.9 over PerSAM-F and +19.4 over PDM), indicating accurate boundary localization. On PerSeg, FoundYou improves over PerSAM-F by +1.1 mIoU and PDM by +6.9 mIoU, showing gains also in simpler scenes. For retrieval, FoundYou achieves 92.1 mAP on PerMIR, outperforming PDM by +26.8. On the

Table 1: Personalized segmentation and retrieval (mIoU, bIoU, and mAP in %, $\uparrow$). For image-matching methods, we use the open-source toolbox from [6]. For PDM, a training-free method, we use the official code. On ILIAS, we use each method to re-rank the same top-1k candidates retrieved by SigLIP; SigLIP alone achieves 19.6 mAP. FoundYou achieves the best performance across benchmarks while using a compact 52 M-parameter model running at 90 FPS on an RTX 4090. It is $75\times$ faster than PDM, faster than segmentation methods, and comparable to the fastest retrieval model.

Model	Segmentation				Retrieval		Efficiency	
	PerSeg		PerMIS		PerMIR	ILIAS	FPS	Params
	mIoU	bIoU	mIoU	bIoU	mAP	mAP	img/s $\uparrow$	M $\downarrow$
<i>Segmentation-only</i>								
SEEM [77] NeurIPS’23	87.1	55.7	34.2	31.8	<i>n.a.</i>	<i>n.a.</i>	8.5	341
SegGPT [62] ICCV’23	94.3	76.5	38.7	35.5	<i>n.a.</i>	<i>n.a.</i>	12.7	354
PerSAM [74] ICLR’24	89.3	71.7	42.1	34.4	<i>n.a.</i>	<i>n.a.</i>	2.5	945
PerSAM-F [74] ICLR’24	<u>95.3</u>	<u>77.9</u>	<u>47.8</u>	34.5	<i>n.a.</i>	<i>n.a.</i>	2.5	945
GF-SAM [73] NeurIPS’24	91.7	74.2	45.6	<u>38.1</u>	<i>n.a.</i>	<i>n.a.</i>	1.5	945
<i>Retrieval-only</i>								
RRT [56] ICCV’21	<i>n.a.</i>	<i>n.a.</i>	<i>n.a.</i>	<i>n.a.</i>	41.7	9.2	31.0	50
MatchAnything [21] ArXiv’25	<i>n.a.</i>	<i>n.a.</i>	<i>n.a.</i>	<i>n.a.</i>	19.7	23.2	3.4	111
MASt3R [29] ECCV’24	<i>n.a.</i>	<i>n.a.</i>	<i>n.a.</i>	<i>n.a.</i>	37.7	25.3	1.7	688
RoMav2 [16] ArXiv’25	<i>n.a.</i>	<i>n.a.</i>	<i>n.a.</i>	<i>n.a.</i>	32.3	<u>26.6</u>	2.7	425
AMES [53] ECCV’24	<i>n.a.</i>	<i>n.a.</i>	<i>n.a.</i>	<i>n.a.</i>	38.2	26.4	90.9	89
<i>Segmentation + Retrieval</i>								
PDM [49] NeurIPS’24	89.5	72.0	44.2	38.0	65.3	14.7	1.2	1200
FoundYou (ours)	96.4	85.6	62.6	57.4	92.1	32.5	<u>90.2</u>	<u>52</u>

ILIAS benchmark, FoundYou reaches 32.5 mAP@1k, exceeding PDM by +17.8 and surpassing strong task-specific methods by +6.1 mAP@1k (AMES). Importantly, FoundYou achieves these results with a single 52 M-parameter model running at 90 FPS, over $75\times$ faster than PDM, faster than specialized segmentation methods, and comparable to the fastest retrieval model (AMES), while keeping SAM 2 frozen and training only 5.9 M additional parameters.

4.2 Few-shot Retrieval

We evaluate FoundYou in the proposed **few-shot retrieval** setting, which reflects realistic scenarios in which users can provide multiple examples of the same instance at inference. FoundYou naturally supports this extension: each reference is stored in the memory bank, allowing candidate images to attend jointly to multiple instance representations via memory attention, enabling aggregation of complementary evidence across views. Competing methods lack a unified multi-view representation; we therefore extend them by computing similarity for each reference and averaging the scores. Results in Table 2 show that increasing the number of references consistently improves performance, progressively reducing the gap to the Oracle (38.9 mAP). Notably, FoundYou benefits substantially more from additional views than prior approaches, improving from 30.1 mAP

Table 2: Proposed few-shot retrieval on ILIAS, mAP@1k. Multiple reference instance examples are provided at inference time. All methods re-rank the same top-1k candidates retrieved with SigLIP [72].

Model	N-shot inference			
	1	2	3	4
Baseline	22.9	22.9	22.9	22.9
XFeat [41]	15.5	16.5	17.4	17.1
RoMa [17]	22.0	21.9	22.1	21.9
SP-LightGlue [30]	22.9	23.0	23.3	22.4
ELoFTR [64]	23.8	24.6	25.7	25.0
MASt3R [29]	24.3	26.0	27.2	27.5
MatchAnything [21]	25.2	26.3	27.2	26.8
RoMa v2 [16]	27.5	<u>28.2</u>	28.4	27.7
AMES [53]	<u>27.6</u>	<u>28.2</u>	<u>29.0</u>	<u>30.2</u>
FoundYou (ours)	30.1	31.8	32.5	34.2
Oracle	38.9	38.9	38.9	38.9




Fig. 4: Re-ranking vs. candidate pool size (ILIAS), mAP (%), ↑. Increasing K enlarges the candidate pool with additional positives and distractors. FoundYou scales more effectively than prior methods, with a widening margin as K grows.

in the 1-shot setting to 34.2 mAP with four references (+4.1). In comparison, AMES improves by +2.6 mAP, whereas RoMa v2 yields marginal gains (+0.2).

4.3 Scaling the Re-Ranking Set

In Fig. 4, we analyze how re-ranking performance evolves as the candidate pool grows on ILIAS. Larger K introduces more positives but also many more distractors. FoundYou improves from 23.4 mAP at $K=10$ to 32.5 at $K=1000$ (+9.1), showing strong scalability to large candidate pools. In contrast, RoMa v2 peaks at 27.3 (at $K=100$) and drops to 26.9 at $K=1000$, a known issue with dense matchers that notoriously suffer from perceptual aliasing [51]. AMES increases more gradually (22.1→26.4), remaining below FoundYou across all K . Crucially, the advantage of FoundYou amplifies with K : the margin over RoMa v2 grows from +0.3 mAP at $K=10$ to +6.5 at $K=1000$. This trend shows that FoundYou effectively scales with the candidate pool, remaining **the most robust to the increasing prevalence of distractors as the candidate set grows**.

4.4 Category-Level Retrieval

Table 3 reports performance across diverse domain-specific retrieval benchmarks. Unlike personalized retrieval, which targets the same physical instance, these benchmarks evaluate retrieval of fine-grained semantic categories, *e.g.*, cars of the same model or birds of the same species. This setting can be viewed as a superset of instance retrieval, as images of the same instance necessarily belong to the same category. Since FoundYou ranks candidates by visual similarity with the query, it naturally prioritizes objects from the same category. Even though

Table 3: Comparison of FoundYou on category-level retrieval and landmark recognition datasets, grouped by domain. Most methods underperform the baseline (*i.e.* no re-ranking), whereas FoundYou, despite being optimized for instance retrieval, always outperforms the baseline and is competitive across domains. For RParis/ROxford, we follow [42], re-ranking the top-100 candidates and compute mAP (% , $\uparrow$). On the remaining datasets we follow [70], re-rank the top-10 candidates, and evaluate with mAP@5 (% , $\uparrow$). The retrieval baseline is SigLIP, following [27].

Method	Products		Cars	Animals	Food	Art	Landmarks		
	RP2K	SOP	Cars196	iNat	Food2K	MET	GLDv2	RPar	ROxf
Baseline	94.8	65.1	<u>90.5</u>	63.9	93.4	77.7	87.5	74.4	39.2
XFeat [41]	96.2	54.9	89.0	58.5	92.6	65.1	84.9	74.7	43.6
SP-LightGlue [30]	95.3	58.6	89.8	60.1	92.9	70.8	88.4	75.5	48.6
ELoFTR [64]	96.2	58.9	89.3	60.0	92.6	74.3	89.0	<u>75.9</u>	<u>49.1</u>
MatchAnything [21]	95.0	56.3	89.2	58.4	92.3	73.1	88.0	75.7	45.6
RoMa v2 [16]	96.1	58.9	90.0	59.2	92.4	78.6	88.0	<u>75.9</u>	46.2
MASt3R [29]	96.1	62.5	89.7	59.6	92.4	80.3	89.3	75.8	45.2
AMES [53]	<u>96.5</u>	<u>66.0</u>	89.8	<u>64.1</u>	<u>93.6</u>	80.8	91.8	75.6	48.7
FoundYou (ours)	97.4	66.2	91.6	65.2	94.4	<u>80.4</u>	<u>91.5</u>	76.1	49.8
Oracle	98.0	76.3	98.3	81.2	98.2	88.4	92.9	76.3	50.3

category-level retrieval is not the primary focus of our model, FoundYou performs on par or better than prior methods. On Cars196 we improve the baseline by +1.1. On iNat, FoundYou reaches 65.2 (+1.1 over AMES). On product benchmarks, we obtain 97.4 on RP2K (+0.9) and 66.2 on SOP (+0.2). For landmarks, FoundYou achieves 76.1 on RParis (+0.2) and 49.8 on ROxford (+0.7). Overall, results highlight the generalization of FoundYou to category-level retrieval.

4.5 Personalized Segmentation with Flexible Prompts

By keeping the SAM 2 decoder frozen, our framework retains the ability to generate masks from flexible prompts, *i.e.* a simple point or bounding box suffices to specify the reference object. Tab. 4-left reports results under different prompts. On PerSeg, FoundYou is nearly unaffected by weaker supervision, achieving 96.4 mIoU with masks, 95.9 with boxes, and 95.2 with points. On the more challenging PerMIS benchmark, the gap between mask and point prompts is 4.0 mIoU for FoundYou (62.6 $\rightarrow$ 58.6), smaller than PerSAM-F’s 6.2 (47.8 $\rightarrow$ 41.6) and PDM’s 9.4 (44.2 $\rightarrow$ 34.8). Qualitative examples in Tab. 4-right illustrate consistent segmentation across prompt types. Our unified formulation enables feature reuse, lowering annotation cost while improving practicality for real-world deployment.

4.6 Ablation Studies

We first ablate the adaptation strategy, adapter design, and adapter placement on personalized segmentation (PerMIS) and retrieval (ILIAS) in Tab. 5. We then analyze how adaptation changes the underlying representation in Fig. 5. Additional ablations on the distillation loss weight λ_{dist} and the number of positive and negative training examples are provided in the Supplementary Material.

Table 4: Promptable personalized segmentation. *Left:* mIoU (% , $\uparrow$) for different prompts types (point, box, mask). *Right:* qualitative results of FoundYou.

Prompt Method		PerSeg PerMIS			
Point	PDM [49]	79.7	34.8	Point	
	PerSAM-F [74]	89.8	41.6		
	FoundYou (ours)	95.2	58.6		
Box	PDM [49]	84.1	39.7	Box	
	PerSAM-F [74]	93.9	46.2		
	FoundYou (ours)	95.9	61.5		
Mask	PDM [49]	89.5	44.2	Mask	
	PerSAM-F [74]	95.3	47.8		
	FoundYou (ours)	96.4	62.6		

Adaptation vs. Fine-tuning. As shown in Tab. 5, applying SAM 2 yields strong segmentation (60.2 mIoU) but poor retrieval performance (3.5 mAP), when using the *occlusion score* as the retrieval score. Full or partial fine-tuning shows limited gains, suggesting disruption of the representations learned during pretraining. Lightweight adaptation achieves the best trade-off, reaching 62.6 mIoU and 32.5 mAP. Notably, self-distillation preserves segmentation quality (+5.9 w.r.t. Adapt w/o Distill) while maintaining strong retrieval performance.

Adapter Design and Placement. Among adapter variants, AdaptFormer performs best, with a $0.5\times$ bottleneck achieving 32.5 mAP. Either smaller or larger bottlenecks reduce performance (-2.0 mAP and -2.7 mAP, respectively). Adapting only the final two stages of the Image Encoder (blocks 2–3) yields the highest accuracy. This suggests that instance discrimination benefits from adapting high-level semantic features while preserving low-level spatial representations.

Mitigating the spatio-temporal continuity bias. SAM 2 is pretrained to preserve object identity across video frames, encouraging its features to maintain spatio-temporal continuity. As a result, these features may be sensitive to changes in an object position and surrounding context. In Fig. 5, we analyze this behavior on the Object-Placement dataset [31] by measuring same-instance fea-

Table 5: Ablations. **Left:** parameter-efficient adaptation outperforms full and partial fine-tuning (FT). **Middle:** AdaptFormer [10] performs best with a $0.5\times$ bottleneck. **Right:** adapting only the last two Image Encoder blocks yields the best performance. We report mIoU (% , $\uparrow$) on PerMIS and mAP (% , $\uparrow$) on ILIAS.

Fine-tuning strategies			Adaptation strategies		Adapter placement		
Method	mIoU	mAP	Method	mAP	Stages		mAP
Frozen SAM 2	60.2	3.5	LoRA [23]	28.9	None	–	20.4
Full FT	27.8	13.8	Adapter [22]	30.2			
Backbone FT	57.2	20.5	AdaptFormer [10]			Early	0–1
QKV FT	59.8	27.2	$0.8\times$ channel dim	30.5		Middle	1–2
Adapt	62.6	32.5	$0.5\times$ channel dim	32.5		Late	2–3
Adapt w/o Distill	56.7	32.7	$0.3\times$ channel dim	29.8	All	0–3	31.8




Fig. 5: Adaptation mitigates SAM 2 spatio-temporal continuity bias. Same-instance feature similarity under (a) object translation, (b) background replacement, and (c) both changes. FoundYou features remain more invariant across these shifts.

ture similarity under object translations, background changes, and their combination. Frozen SAM 2 features become progressively less similar as these changes increase, whereas the features adapted by FoundYou remain substantially more invariant. This shows that our adaptation reduces the continuity bias, yielding features that are more reliable for matching objects across independent images.

5 Conclusion

We presented FoundYou, a unified framework for personalized segmentation and retrieval that extends SAM 2 from video tracking to cross-image instance matching. With lightweight adaptation layers and a compact retrieval head, while keeping SAM 2 entirely frozen, we remove its temporal continuity bias and enable discrimination of a specific physical instance across independent images. The resulting 52 M-parameter model outperforms prior unified and task-specific methods, scales effectively with additional references, and remains robust to weaker prompts. More broadly, our results suggest that large-scale video pre-training in SAM 2 induces identity-aware representations that extend beyond tracking, providing a strong foundation for instance-level reasoning across tasks.

Acknowledgements. The Ministry of Science, Innovation and Universities of Spain has supported this work through FPU23/00587 (Marcos Alfaro). This research work is part of the projects PID2023-149575OB-I00, funded by MICIU/AEI/10.13039/501100011033 and by FEDER, UE, and CIPROM/2024/8, funded by Generalitat Valenciana (program PROMETEO 2025). Claudia Cuttano was supported by the Sustainable Mobility Center (CNMS), which received funding from European Union Next Generation EU (Piano Nazionale di Ripresa e Resilienza (PNRR), Missione 4 Componente 2 Investimento 1.4), grant agreement no. CN_00000023. We acknowledge the CINECA award under the ISCRA initiative, for availability of high performance computing resources.

References

1. Alaluf, Y., Richardson, E., Tulyakov, S., Aberman, K., Cohen-Or, D.: MyVLM: Personalizing VLMs for user-specific queries. In: ECCV. vol. 15071, pp. 73–91 (2024)
2. Arandjelović, R., Gronat, P., Torii, A., Pajdla, T., Sivic, J.: NetVLAD: CNN architecture for weakly supervised place recognition. *IEEE Trans. Pattern Anal. Mach. Intell.* **40**(6), 1437–1451 (2018)
3. Bai, Y., Chen, Y., Yu, W., Wang, L., Zhang, W.: Products-10K: A large-scale product recognition dataset. *arXiv:2008.10545 [cs.CV]* (2020)
4. Barbarani, G., Mostafa, M., Bayramov, H., Trivigno, G., Berton, G., Masone, C., Caputo, B.: Are local features all you need for cross-domain visual place recognition? In: CVPR. pp. 6155–6165 (2023)
5. Barsellotti, L., Bigazzi, R., Cornia, M., Baraldi, L., Cucchiara, R.: Personalized instance-based navigation toward user-specific objects in realistic environments. In: NeurIPS. pp. 11228–11250 (2024)
6. Berton, G., Goletto, G., Trivigno, G., Stoken, A., Caputo, B., Masone, C.: Earth-Match: Iterative coregistration for fine-grained localization of astronaut photography. In: CVPRW. pp. 4264–4274 (2024)
7. Berton, G., Mereu, R., Trivigno, G., Masone, C., Csurka, G., Sattler, T., Caputo, B.: Deep visual geo-localization benchmark. In: CVPR. pp. 5396–5407 (2022)
8. Berton, G., Trivigno, G., Caputo, B., Masone, C.: JIST: Joint image and sequence training for sequential visual place recognition. *IEEE Robotics Autom. Lett.* **9**(2), 1310–1317 (2023)
9. Cao, B., Araujo, A., Sim, J.: Unifying deep local and global features for image search. In: ECCV. vol. 12365, pp. 726–743 (2020)
10. Chen, S., Ge, C., Tong, Z., Wang, J., Song, Y., Wang, J., Luo, P.: AdaptFormer: Adapting vision transformers for scalable visual recognition. In: NeurIPS. vol. 35, pp. 16664–16678 (2022)
11. Cohen, N., Gal, R., Meirom, E.A., Chechik, G., Atzmon, Y.: “This is my unicorn, Fluffy”: Personalizing frozen vision-language representations. In: ECCV. vol. 13680, pp. 558–577 (2022)
12. Cuttano, C., Trivigno, G., Averta, G., Masone, C.: SANSa: Unleashing the hidden semantics in SAM2 for few-shot segmentation. In: NeurIPS (2025)
13. Cuttano, C., Trivigno, G., Rosi, G., Masone, C., Averta, G.: Samwise: Infusing wisdom in SAM2 for Text-Driven Video Segmentation. In: CVPR. pp. 3395–3405 (2025)
14. Deng, K., Yang, Y., Sun, J., Liu, X., Liu, Y., Liang, D., Cao, Y.P.: GeoSAM2: Unleashing the power of SAM2 for 3D part segmentation. *arXiv:2508.14036 [cs.CV]* (2025)
15. Dutto, M., Berton, G., Caldarola, D., Fani, E., Trivigno, G., Masone, C.: Collaborative visual place recognition through federated learning. In: CVPRW. pp. 4215–4225 (2024)
16. Edstedt, J., Nordström, D., Zhang, Y., Bökman, G., Astermark, J., Larsson, V., Heyden, A., Kahl, F., Wadenbäck, M., Felsberg, M.: RoMa v2: Harder better faster denser feature matching. *arXiv:2511.15706 [cs.CV]* (2025)
17. Edstedt, J., Sun, Q., Bökman, G., Wadenbäck, M., Felsberg, M.: RoMa: Robust dense feature matching. In: CVPR. pp. 19790–19800 (2024)
18. El-Nouby, A., Neverova, N., Laptev, I., Jégou, H.: Training vision transformers for image retrieval. *arXiv:2102.05644 [cs.CV]* (2021)

19. Forni, B., Lombardi, G., Pozzi, F., Planamente, M.: FS-SAM2: Adapting Segment Anything model 2 for few-shot semantic segmentation via low-rank adaptation. In: ICIAP. vol. 16167, pp. 67–78 (2025)
20. Ge, Y., Zhang, R., Wang, X., Tang, X., Luo, P.: DeepFashion2: A versatile benchmark for detection, pose estimation, segmentation and re-identification of clothing images. In: CVPR. pp. 5337–5345 (2019)
21. He, X., Yu, H., Peng, S., Tan, D., Shen, Z., Bao, H., Zhou, X.: MatchAnything: Universal cross-modality image matching with large-scale pre-training. arXiv:2501.07556 [cs.CV] (2025)
22. Houdouby, N., Giurghi, A., Jastrzebski, S., Morrone, B., De Laroussilhe, Q., Gesmundo, A., Attariyan, M., Gelly, S.: Parameter-efficient transfer learning for NLP. In: ICML. vol. 97, pp. 2790–2799 (2019)
23. Hu, E.J., Shen, Y., Wallis, P., Allen-Zhu, Z., Li, Y., Wang, S., Wang, L., Chen, W., et al.: LoRA: Low-rank adaptation of large language models. In: ICLR (2022)
24. Izquierdo, S., Civera, J.: Close, but not there: Boosting geographic distance sensitivity in visual place recognition. In: ECCV. vol. 15131, pp. 240–257 (2024)
25. Jin Kim, H., Dunn, E., Frahm, J.M.: Learned contextual feature reweighting for image geo-localization. In: CVPR. pp. 2136–2145 (2017)
26. Kirillov, A., Mintun, E., Ravi, N., Mao, H., Rolland, C., Gustafson, L., Xiao, T., Whitehead, S., Berg, A.C., Lo, W.Y., Dollar, P., Girshick, R.: Segment Anything. In: ICCV. pp. 4015–4026 (2023)
27. Kordopatis-Zilos, G., Stojnić, V., Manko, A., Suma, P., Ypsilantis, N.A., Efthymiadis, N., Laskar, Z., Matas, J., Chum, O., Tolas, G.: ILIAS: Instance-level image retrieval at scale. In: CVPR. pp. 14777–14787 (2025)
28. Kwon, T., Choi, D., Kim, H., Kim, S., Moon, S., Kwak, B.w., Huang, K.H., Yeo, J.: Embodied agents meet personalization: Investigating challenges and solutions through the lens of memory utilization. In: ICLR (2026)
29. Leroy, V., Cabon, Y., Revaud, J.: Grounding image matching in 3D with MAST3R. In: ECCV. vol. 15130, pp. 71–91 (2024)
30. Lindenberger, P., Sarlin, P.E., Pollefeys, M.: LightGlue: Local feature matching at light speed. In: ICCV. pp. 17627–17638 (2023)
31. Liu, L., Liu, Z., Zhang, B., Li, J., Niu, L., Liu, Q., Zhang, L.: Opa: Object placement assessment dataset (2022), <https://arxiv.org/abs/2107.01889>
32. Liu, Y., Zhu, M., Li, H., Chen, H., Wang, X., Shen, C.: Matcher: Segment anything with one shot using all-purpose feature matching. In: ICLR (2024)
33. Meng, L., Lan, S., Li, H., Alvarez, J.M., Wu, Z., Jiang, Y.G.: SEGIC: Unleashing the emergent correspondence for in-context segmentation. In: ECCV. vol. 15096, pp. 203–220 (2024)
34. Mereu, R., Trivigno, G., Berton, G., Masone, C., Caputo, B.: Learning sequential descriptors for sequence-based visual place recognition. IEEE Robotics Autom. Lett. **7**(4), 10383–10390 (2022)
35. Musgrave, K., Belongie, S., Lim, S.N.: A metric learning reality check. In: ECCV. vol. 12370, pp. 681–699 (2020)
36. Nie, J., Xing, Y., An, W., Zhao, Q., Shao, J., Tan, Y.P., Kot, A.C., Lu, S., Li, X.: Boosting SAM for cross-domain few-shot segmentation via conditional point sparsification. arXiv:2602.05218 [cs.CV] (2026)
37. Nitzan, Y., Aberman, K., He, Q., Liba, O., Yarom, M., Gandelsman, Y., Mosseri, I., Pritch, Y., Cohen-Or, D.: MyStyle: A personalized generative prior. ACM Trans. Graph. **41**(6), 1–10 (2022)

38. Oquab, M., Darcet, T., Moutakanni, T., Vo, H., Szafraniec, M., Khalidov, V., Fernandez, P., Haziza, D., Massa, F., El-Nouby, A., et al.: DINOv2: Learning robust visual features without supervision. *Trans. Mach. Learn. Res.* (2024)
39. Panek, V., Kukulova, Z., Sattler, T.: MeshLoc: Mesh-based visual localization. In: *ECCV*. vol. 13682, pp. 589–609 (2022)
40. Philbin, J., Chum, O., Isard, M., Sivic, J., Zisserman, A.: Lost in quantization: Improving particular object retrieval in large scale image databases. In: *CVPR*. pp. 1–8 (2008)
41. Potje, G., Cadar, F., Araujo, A., Martins, R., Nascimento, E.R.: XFeat: Accelerated features for lightweight image matching. In: *CVPR*. pp. 2682–2691 (2024)
42. Radenović, F., Iscen, A., Tolias, G., Avrithis, Y., Chum, O.: Revisiting Oxford and Paris: Large-scale image retrieval benchmarking. In: *CVPR*. pp. 5706–5715 (2018)
43. Ravi, N., Gabeur, V., Hu, Y.T., Hu, R., Ryali, C., Ma, T., Khedr, H., Rädle, R., Rolland, C., Gustafson, L., et al.: SAM 2: Segment anything in images and videos. In: *ICLR* (2025)
44. Revaud, J., Almazán, J., Rezende, R.S., Souza, C.R.d.: Learning with average precision: Training image retrieval with a listwise loss. In: *ICCV*. pp. 5107–5116 (2019)
45. Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models. In: *CVPR*. pp. 10684–10695 (2022)
46. Rong, F., Lan, M., Zhang, Q., Zhang, L.: MPG-SAM 2: Adapting SAM 2 with mask priors and global context for referring video object segmentation. In: *ICCV*. pp. 23979–23989 (2025)
47. Rong, F., Lan, M., Zhang, Q., Zhang, L.: Customized SAM 2 for referring remote sensing image segmentation. In: *AAAI* (2026)
48. Sakaguchi, T., Taniguchi, A., Hagiwara, Y., El Hafi, L., Hasegawa, S., Taniguchi, T.: Object instance retrieval in assistive robotics: Leveraging fine-tuned SimSiam with multi-view images based on 3D semantic map. In: *IROS*. pp. 7817–7824 (2024)
49. Samuel, D., Ben-Ari, R., Levy, M., Darshan, N., Chechik, G.: Where’s Waldo: Diffusion features for personalized segmentation and retrieval. In: *NeurIPS*. vol. 37, pp. 128160–128181 (2024)
50. Sanakoyeu, A., Tschernetzki, V., Buchler, U., Ommer, B.: Divide and conquer the embedding space for metric learning. In: *CVPR*. pp. 471–480 (2019)
51. Sferrazza, D., Berton, G., Trivigno, G., Masone, C.: To Match or Not to Match: Revisiting image matching for reliable visual place recognition. In: *CVPR*. pp. 2849–2860 (2025)
52. Sipiran, I., Lazo, P., Lopez, C., Jimenez, M., Bagewadi, N., Bustos, B., Dao, H., et al.: SHREC 2021: Retrieval of cultural heritage objects. *Comput. Graph.* **100**(C), 1–20 (2021)
53. Suma, P., Kordopatis-Zilos, G., Iscen, A., Tolias, G.: AMES: Asymmetric and memory-efficient similarity estimation for instance-level retrieval. In: *ECCV*. vol. 15117, pp. 307–325 (2024)
54. Sun, J., Shen, Z., Wang, Y., Bao, H., Zhou, X.: LoFTR: Detector-free local feature matching with transformers. In: *CVPR*. pp. 8922–8931 (2021)
55. Sun, Y.C., Sun, C., Lin, C.Y., Yang, F.E., Chen, M.H., Lin, Y.Y., Liu, Y.L.: 3AM: Segment Anything with geometric consistency in videos. *arXiv:2601.08831 [cs.CV]* (2026)
56. Tan, F., Yuan, J., Ordonez, V.: Instance-level image retrieval using reranking transformers. In: *ICCV*. pp. 12105–12115 (2021)

57. Tang, L., Jiang, P.T., Xiao, H., Li, B.: Towards training-free open-world segmentation via image prompt foundation models. *Int. J. Comput. Vis.* **133**(1), 1–15 (2025)
58. Thomee, B., Shamma, D.A., Friedland, G., Elizalde, B., Ni, K., Poland, D., Borth, D., Li, L.J.: YFCC100M: The new data in multimedia research. *Commun. ACM* **59**(2), 64–73 (2016)
59. Trivigno, G., Berton, G., Aragon, J., Caputo, B., Masone, C.: Divide&Classify: Fine-grained classification for city-wide visual geo-localization. In: *ICCV*. pp. 11142–11152 (2023)
60. Trivigno, G., Masone, C., Caputo, B., Sattler, T.: The unreasonable effectiveness of pre-trained features for camera pose refinement. In: *CVPR*. pp. 12786–12798 (2024)
61. Wang, W., Bao, H., Dong, L., Bjorck, J., Peng, Z., Liu, Q., Aggarwal, K., Mohammed, O.K., Singhal, S., Som, S., et al.: Image as a foreign language: BEiT pretraining for all vision and vision-language tasks. In: *CVPR*. pp. 19175–19186 (2023)
62. Wang, X., Zhang, X., Cao, Y., Wang, W., Shen, C., Huang, T.: SegGPT: Towards segmenting everything in context. In: *ICCV*. pp. 1130–1140 (2023)
63. Wang, X., Han, X., Huang, W., Dong, D., Scott, M.R.: Multi-Similarity loss with general pair weighting for deep metric learning. In: *CVPR*. pp. 5022–5030 (2019)
64. Wang, Y., He, X., Peng, S., Tan, D., Zhou, X.: Efficient LoFTR: Semi-dense local feature matching with sparse-like speed. In: *CVPR*. pp. 21666–21675 (2024)
65. Warburg, F., Hauberg, S., Lopez-Antequera, M., Gargallo, P., Kuang, Y., Civera, J.: Mapillary street-level sequences: A dataset for lifelong place recognition. In: *CVPR*. pp. 2626–2635 (2020)
66. Weyand, T., Araujo, A., Cao, B., Sim, J.: Google Landmarks Dataset v2—A large-scale benchmark for instance-level recognition and retrieval. In: *CVPR*. pp. 2575–2584 (2020)
67. Wu, C.Y., Manmatha, R., Smola, A.J., Krahenbuhl, P.: Sampling matters in deep embedding learning. In: *ICCV*. pp. 2840–2848 (2017)
68. Xiao, S., Kabra, R., Li, Y., Lee, D., Carreira, J., Panda, P.: OpenWorldSAM: Extending SAM2 for universal image segmentation with language prompts. In: *NeurIPS* (2025)
69. Xu, Q., Zhu, L., Liu, X., Lin, G., Long, C., Li, Z., Zhao, R.: Unlocking the power of SAM2 for few-shot segmentation. In: *ICML*. vol. 267, pp. 69950–69966 (2025)
70. Ypsilantis, N.A., Chen, K., Cao, B., Lipovskỳ, M., Dogan-Schönberger, P., Makosa, G., Bluntschli, B., Seyedhosseini, M., Chum, O., Araujo, A.: Towards universal image embeddings: A large-scale dataset and challenge for generic image representations. In: *ICCV*. pp. 11290–11301 (2023)
71. Ypsilantis, N.A., Garcia, N., Han, G., Ibrahimi, S., van Noord, N., Tolias, G.: The Met dataset: Instance-level recognition for artworks. In: *NeurIPS D&B*. vol. 1 (2021)
72. Zhai, X., Mustafa, B., Kolesnikov, A., Beyer, L.: Sigmoid loss for language image pre-training. In: *ICCV*. pp. 11975–11986 (2023)
73. Zhang, A., Gao, G., Jiao, J., Liu, C.H., Wei, Y.: Bridge the points: Graph-based few-shot segment anything semantically. *NeurIPS* **37**, 33232–33261 (2024)
74. Zhang, R., Jiang, Z., Guo, Z., Yan, S., Pan, J., Ma, X., Dong, H., Gao, P., Li, H.: Personalize Segment Anything model with one shot. In: *ICLR* (2024)
75. Zhu, M., Liu, Y., Luo, Z., Jing, C., Chen, H., Xu, G., Wang, X., Shen, C.: Unleashing the potential of the diffusion model in few-shot semantic segmentation. *NeurIPS* **37**, 42672–42695 (2024)

- 76. Zhu, S., Yang, L., Chen, C., Shah, M., Shen, X., Wang, H.: R2Former: Unified retrieval and reranking transformer for place recognition. In: CVPR. pp. 19370–19380 (2023)
- 77. Zou, X., Yang, J., Zhang, H., Li, F., Li, L., Wang, J., Wang, L., Gao, J., Lee, Y.J.: Segment everything everywhere all at once. In: NeurIPS. vol. 36, pp. 19769–19782 (2023)