MOCHA: Multi-modal Objects-aware Cross-arcHitecture Alignment

Elena Camuffo[1,2*], **Francesco Barbato**[1,2*], **Mete Ozay**[1], **Simone Milani**[2], **Umberto Michieli**[1]

[1]Samsung R&D Institute UK, United Kingdom
[2]University of Padova, Italy

Abstract

We introduce MOCHA (Multi-modal Objects-aware Cross-arcHitecture Alignment), a knowledge distillation approach that transfers region-level multimodal semantics from a large vision-language teacher (*e.g.*, LLaVa) into a lightweight vision-only object detector student (*e.g.*, YOLO). A translation module maps student features into a joint space, where the training of the student and translator is guided by a dual-objective loss that enforces both local alignment and global relational consistency. Unlike prior approaches focused on dense or global alignment, MOCHA operates at the object level, enabling efficient transfer of semantics without modifying the teacher or requiring textual input at inference. We validate our method across four personalized detection benchmarks under few-shot regimes. Results show consistent gains over baselines, with a +10.1 average score improvement. Despite its compact architecture, MOCHA reaches performance on par with larger multimodal models, proving its suitability for real-world deployment.

1 Introduction

Recent advances in vision-language models (VLMs) such as CLIP (Radford et al. 2021), Flamingo (Alayrac et al. 2022), and LLaVa (Liu et al. 2023a) have demonstrated remarkable zero-shot and open-vocabulary capabilities, owing to their rich representations and scale. These models, however, are typically large and computationally intensive, limiting their applicability in real-time or resource-constrained scenarios such as mobile devices (*e.g.*, smartphones or robots). In contrast, lightweight architectures such as YOLO (Jocher, Chaurasia, and Qiu 2023) offer fast and memory-efficient object detection, at the cost of degraded performance in low-data regimes and vulnerability to neural collapse (Kothapalli 2022; Papyan, Han, and Donoho 2020), where features become indistinguishably aligned, losing their discriminative power. In this work, we aim to bridge this gap by proposing MOCHA, a knowledge distillation approach that transfers object-centric multimodal embeddings from a vision-language teacher (*e.g.*, LLaVa) into a compact vision-only student detector. MOCHA exploits the property that semantically related concepts exhibit similar structures in the embedding space across different modalities under a well-generalized backbone (Huh et al. 2024; Girdhar et al. 2023; Camuffo, Michieli, and Milani 2023). By aligning and regularizing the student's feature space, it shows improved generalization in few-shot detection settings (Fig. 1). Our approach is built on three stages that include (1) pretraining of a student model, (2) knowledge distillation on rich joint visual and textual features from a frozen teacher, and (3) few-shot personalization with frozen student model and prototype-based classifier (Snell, Swersky, and Zemel 2017). Our contribution lies primarily in the knowledge distillation stage, where we introduce: (i) extraction of rich joint visual-language representations conditioned on image regions and class labels from a frozen teacher model; (ii) a translation module that maps student features into the teacher's multimodal space; and (iii) a distillation objective combining local alignment with relational embedding regularization across regions. We validate MOCHA on personal object detection datasets, where the goal is to adapt a general detector to user-specific personal classes (*e.g.*, distinguishing a generic *dog* from a particular *user's dog*). MOCHA achieves significant improvements

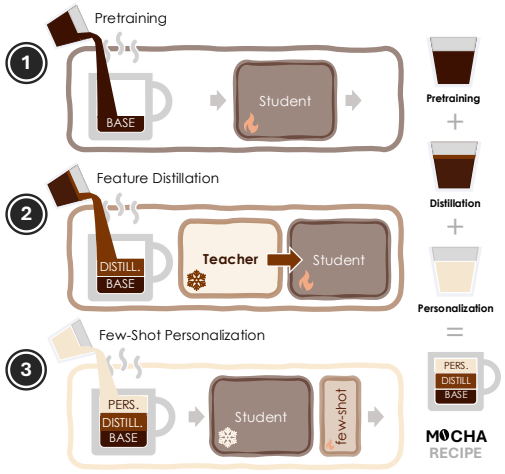

Figure 1: MOCHA recipe. (1) Pretraining student model. (2) Knowledge distillation on rich joint visual and textual features from a frozen teacher. (3) Few-shot personalization with frozen student and prototypical learner.

*Work partially done during an internship at Samsung R&D Institute UK.

over state-of-the-art methods while retaining minimal inference cost, making it practical for resource-constrained deployments.

2 Related Work

Cross-Architecture Knowledge Distillation. Traditional knowledge distillation methods (Romero et al. 2015; Zhang et al. 2020; Passalis, Tzelepi, and Tefas 2020; Luo et al. 2019; Gou et al. 2024; Hao et al. 2022; Li et al. 2024; Tu et al. 2025) typically assume that teacher and student networks share similar architectures. The adoption of large transformer-based models (Vaswani et al. 2017; Dosovitskiy 2020; Devlin et al. 2019; Liu et al. 2021) has raised new challenges in transferring knowledge into efficient CNNs, which offer better trade-offs for deployment. Recent works (Liu et al. 2022; Hao et al. 2023) address this by projecting multi-scale features from teacher and student into a shared space for supervision, echoing shortcut-based strategies (Szegedy et al. 2014; Lee et al. 2014) to improve convergence stability.

These methods primarily aim to improve performance on the same task shared by teacher and student. In contrast, our approach aligns more closely with AuXFT (Barbato et al. 2024b), which leverages teacher features to guide a student on a related auxiliary task without degrading performance on the original one. Specifically, AuXFT addresses few-shot instance-level personalized object detection by injecting high-level teacher features into a lightweight detector to support prototype-based classification. Our method differs in three ways: (i) we leverage multimodal supervision from both visual and textual signals, unlike purely visual clues used in prior art; (ii) we distill compact region-level embeddings rather than dense intermediate activations, improving efficiency; and (iii) we explicitly regularize the embedding space to encourage geometric and relational alignment between individual instances.

Cross-Modal Adaptation. Several cross-modal learning strategies have been proposed to exploit complementary information across modalities (Xia et al. 2023; Barbato et al. 2024a) or to enforce consistency between predictions from different modalities (Jaritz et al. 2022, 2020). In our work, we transfer rich multimodal semantics from a vision-language model into a compact student detector to embed multimodal priors into a lightweight architecture.

Personalized Scene Understanding. Personalization was first explored in Natural Language Processing (NLP) (Liu et al. 2023a; Vaswani et al. 2017; Devlin et al. 2019). Later, it was extended to computer vision, with early applications in personalized semantic segmentation (Zhang et al. 2021) and object detection (Michieli et al. 2024). The advent of foundation models such as CLIP (Radford et al. 2021) and SAM (Kirillov et al. 2023) has shifted personalization strategies toward prompt-based control of generalist architectures (Meng, Xia, and Ma 2024). For example, ViLD (Gu et al. 2022) aligns vision and language embeddings to enable open-vocabulary detection, paving the way for instance- and user-aware predictions through textual prompts. Similarly, SwissDINO (Paramonov et al. 2024), PerSAM (Zhang et al. 2024), Matcher (Liu et al. 2023b), and SegGPT (Wang et al. 2023) exploit prompt tuning and strong visual representations to tailor predictions to user preferences. Our work instead distills rich semantic features into a compact detector during server-side training, avoiding reliance on prompts or online adaptation and enabling efficient few-shot personalization for resource-constrained deployment.

3 Methodology

MOCHA (Multi-modal Objects-aware Cross-arcHitecture Alignment) distills multimodal semantics from a vision-language foundation model into a lightweight student detector by aligning region-level embeddings and regularizing the student feature space. MOCHA enhances feature separability and generalization for few-shot personalized detection. Our method is architecture-agnostic and combines multimodal supervision, feature translation, and a dual loss enforcing both local and relational alignment.

3.1 Problem Setup

MOCHA consists of three stages, as described next.

Base Pretraining. We begin with a standard detection model $\mathrm{m}_S(\cdot) = \mathrm{l}_S \circ \mathrm{g}_S$, where $\mathrm{l}_S(\cdot)$ is the detection head, $\mathrm{g}_S(\cdot)$ is the student backbone, and $\circ$ is the composition operator. We train it on a pretraining detection dataset using standard objectives $\mathcal{L}_{det}$ (Jocher, Chaurasia, and Qiu 2023). This stage results in a strong generic detector that identifies broad object categories but lacks personalized semantics or fine-grained embedding separation and serves as initialization for the student's weights.

Feature Distillation. After initialization, we perform feature distillation over an object detection dataset $\mathcal{D}_c$, using a frozen large vision-language model as the teacher (Fig. 2a). Each image $X \in \mathcal{D}_c$ in the dataset is labeled by a set of tuples $\{(b_i, c_i)\}_{i=1}^{n} \subset \mathcal{B} \times \mathcal{C}_c$, where $c_i \in \mathcal{C}_c$ are class labels and $b_i \in \mathcal{B} \subset \mathbb{R}^4$ are box coordinates identifying n image regions X_i. The teacher is assumed to comprise a visual encoder $\mathrm{g}_T(\cdot)$, and a large language model $\mathrm{f}_T(\cdot)$, enabling the joint processing of images and text. Note that, in general, the student and teacher features may not be aligned. For this reason, we use a feature translation module $\mathrm{t}_S(\cdot) : \mathbb{R}^{d_s} \mapsto \mathbb{R}^{d_t}$ (where d_s and d_t are the dimensions of the student and teacher features, respectively), to allow knowledge distillation from the student's to the teacher's space.

Few-Shot Personalization. After feature distillation, the student backbone is frozen and used to extract multiscale region-level features $\{F_j\}_j = \mathrm{g}_S(X)$ from a personal detection dataset $\mathcal{D}_f$, labeled with bounding boxes $\mathcal{B}$ and class labels $\mathcal{C}_f$. The features are translated into the shared embedding space via the frozen translation module (Fig. 2b) and fed to a prototype-based few-shot learner $\mathrm{p}(\cdot)$, like nearest class mean (Snell, Swersky, and Zemel 2017), which is trained and tested as in (Barbato et al. 2024b). More in detail, during personalization, a user provides a few labeled samples of personal instances of objects to the system (*e.g.*,

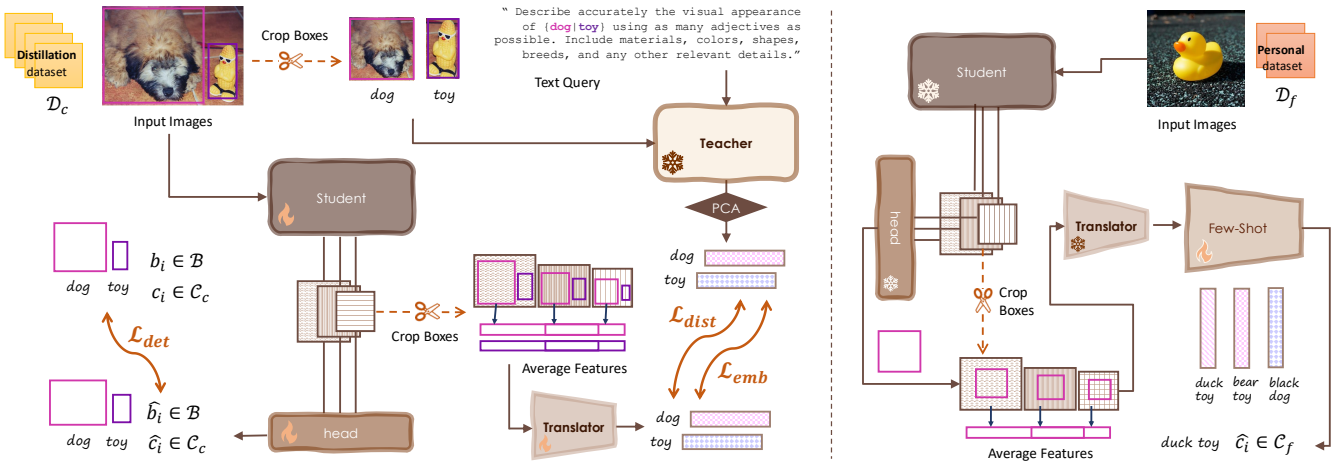

(a) Feature distillation stage of the student network from the frozen teacher.

(b) Personalization stage.

Figure 2: MOCHA system. **(a) Feature distillation:** Student detector is trained on dataset $\mathcal{D}_c$ aligning multiscale region-level features to PCA-pruned multimodal embeddings from a frozen vision-language teacher via the translation module $\mathrm{t}_S(\cdot)$. **(b) Personalization:** Student backbone and Translator are frozen and used to compute semantic prototypical features from a personal dataset $\mathcal{D}_f$. These are then used to train a prototype-based few-shot learner $\mathrm{p}(\cdot)$ for user-specific object classification.

1-5 samples). Later, during inference, the system automatically overrides the coarse predictions with the personal labels. This strategy enables the student to retain its compact architecture while adapting to novel visual concepts. Since no teacher inference or prompt engineering is required at personalization time, the method is efficient and suitable for deployment in real-world low-resource settings.

3.2 Multimodal Supervision Extraction

Without loss of generality, in our experiments, we instantiate the teacher with a LLaVa (Liu et al. 2023a) model, which includes a shared projection matrix W_T, in addition to $\mathrm{g}_T(\cdot)$ and $\mathrm{f}_T(\cdot)$. Each cropped image region X_i is processed by the teacher's visual encoder to obtain a visual embedding $Z_{V,i}$, projected into a joint text-vision representation space as $H_{V,i} = W_T Z_{V,i}$. Simultaneously, the class label c_i is embedded as a textual query $Z_{Q,i}$ and projected via $H_{Q,i} = W_T Z_{Q,i}$. The pair $(H_{V,i}, H_{Q,i})$ is passed to the frozen language model f_T, which computes a sequence of tokens describing the interaction between the region and the class. We average the output along the temporal dimension to obtain a single semantic embedding $h_i \in \mathbb{R}^{d_h}$ representing the fused visual-linguistic description of the region X_i.

Dimensionality Reduction. While the teacher's multimodal feature h_i captures rich semantics, it does not necessarily retain appearance-level cues. To combine both semantic and visual signals, we concatenate the multimodal representation with the d_z-dimensional class token $z_{V,i} \in Z_{V,i} \subset \mathbb{R}^{d_z}$ of its corresponding visual embedding:

$$u_i = \mathrm{concat}(\gamma\, z_{V,i}, h_i) \in \mathbb{R}^d, \quad d = d_z + d_h. \tag{1}$$

with $\gamma = ||h_i||$ as a rescaling factor that normalizes visual features to textual ones. This combined vector represents both the raw visual content of the object and its semantic interpretation conditioned on the label. However, the resulting representation is high-dimensional and could include redundant information. To reduce dimensionality and improve efficiency, we apply Principal Component Analysis (PCA) to u_i, offline on the same dataset used for distillation, and we take the first d_t elements:

$$\hat{u}'_i = \mathrm{PCA}(u_i) \in \mathbb{R}^{d_t}\ . \tag{2}$$

Finally, to equalize the contribution of each channel in the PCA representation, we estimate the standard deviation of each channel activation σ_c in the distillation dataset, and use it to rescale the channels as:

$$u'_i[c] = \hat{u}'_i[c]/\sigma_c \quad \text{for } c = 1, \ldots, d. \tag{3}$$

Remarkably, the deviations closely follow a hyperbolic shape, allowing us to estimate their value based only on the channel index (detailed discussion in Appendix).

3.3 Student Detector and Feature Aggregation

To perform object-level supervision, we apply the same cropping strategy used for the teacher: for each ground-truth bounding box b_i, we first resize the feature maps to a common size to account for resolution differences, ensuring consistent spatial alignment. Then, we extract the corresponding region-aligned features $F_{A,j,i}$ from each resolution level j, and apply spatial average pooling to obtain fixed-size descriptors $f_{A,j,i} \in \mathbb{R}^{d_s}$.

The final aggregated feature for region i is formed by concatenating the pooled representations at each level:

$$f_{A,i} = \mathrm{concat}_j(f_{A,j,i}) \in \mathbb{R}^{d_s}. \tag{4}$$

We denote with F_A the dense equivalent of these features, that is, the concatenation without region pooling.

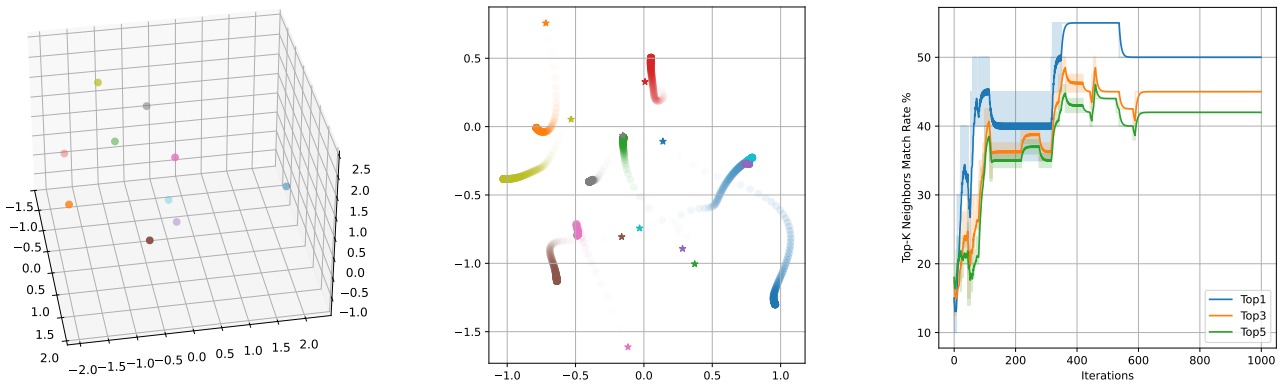

(a) 3D reference points (proxy for u'_i). (b) 2D points being optimized (proxy for $f'_{A,i}$). (c) k-NN match rate.

Figure 3: Toy example. Effect of $\mathcal{L}_{\text{emb}}$ on a set of ten 2D points optimized with respect to 3D ones. **(a)** 3D reference points, proxy for the teacher embeddings u'_i. **(b)** Evolution of the 2D points (proxy for student embeddings $f'_{A,i}$) updated via $\mathcal{L}_{\text{emb}}$ from the 3D reference points u'_i. $\star$ marks the original location, timesteps increase with color saturation. **(c)** Percent rate of 2D top-k nearest neighbors (k-NN) that match those of the reference 3D distribution.

Without loss of generalization, we adopt YOLOv8 (Jocher, Chaurasia, and Qiu 2023) as the student model pretrained on the COCO dataset (Lin et al. 2014) and then fine-tuned in MOCHA.

Feature Translation Module. The translation module t_S is modelled after a transformer encoder block and consists of a channel-wise multi-head self-attention block followed by a lightweight multi-layer perceptron (MLP). The attention mechanism enhances expressiveness by modelling inter-channel dependencies, while the MLP adapts the representation to the required dimensionality:

$$f'_{A,i} = \mathrm{t}_S(f_{A,i}), \tag{5}$$

where $f'_{A,i} \in \mathbb{R}^{d_t}$ is the translated student feature with the same dimension as the PCA-compressed teacher target u'_i.

In our implementation (Fig. 2a), the translation module is trained jointly with the student detector, allowing the feature alignment to evolve progressively during distillation.

3.4 Training Objectives and Strategy

To align the translated student features with the teacher targets, we define two complementary objectives: a pointwise distillation loss and a relational embedding loss.

Distillation Loss. The primary objective enforces direct alignment between translated student features $f'_{A,i}$ and their corresponding teacher vector u'_i. We define the distillation loss as the average of ℓ_1 and ℓ_2 distances, as in (Barbato et al. 2024b):

$$\mathcal{L}_{\text{dist}} = \frac{1}{n}\sum_{i=1}^{n}\left(\|f'_{A,i} - u'_i\|_1 + \|f'_{A,i} - u'_i\|_2\right). \tag{6}$$

This formulation captures both robustness to outliers (via the ℓ_1 term) and fine-grained magnitude alignment (via the ℓ_2 term), promoting stable and discriminative alignment.

Embedding Loss. Beyond individual alignment, we encourage the student to preserve the global geometry of the teacher's embedding space. To this end, we define a relational embedding loss based on pairwise Euclidean distances between all student and teacher features.

Given the translated student features $\{f'_{A,i}\}_i$ and the corresponding teacher targets $\{u'_i\}_i$, we compute two pairwise distance matrices D_{ff} and D_{uu}. We convert them into normalized distributions using temperature-scaled softmax:

$$P_{ff} = \mathrm{softmax}\left(-\frac{D_{ff}}{\tau}\right), \quad P_{uu} = \mathrm{softmax}\left(-\frac{D_{uu}}{\tau}\right), \tag{7}$$

where τ is a hyperparameter controlling distribution sharpness. The embedding loss is then defined as the cross-entropy between the two distributions:

$$\mathcal{L}_{\text{emb}} = -\frac{1}{n}\sum_{i=1}^{n}\sum_{j\neq i} P_{uu}[i,j]\log P_{ff}[i,j]. \tag{8}$$

This term ensures that the student preserves not only semantic content but also the relative structure of the embedding space, reflecting inter-class and intra-class relationships encoded by the teacher.

Final Objective. The final objective combines the detection loss with the distillation and embedding terms:

$$\mathcal{L} = \mathcal{L}_{\text{det}} + \lambda_{\text{dist}}\mathcal{L}_{\text{dist}} + \lambda_{\text{emb}}\mathcal{L}_{\text{emb}}, \tag{9}$$

where the λ coefficients regulate the contribution of each auxiliary term. By integrating multimodal feature supervision with structural regularization, MOCHA enables robust semantic transfer into compact detectors, improving performance in low-resource and personalized scenarios.

4 Experiments

In this section, we introduce the experimental setup to test MOCHA's ability to effectively distill knowledge from a

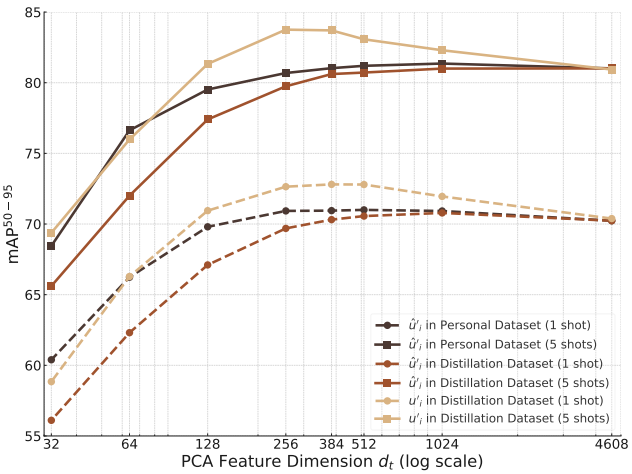

Figure 4: PCA study. Average mAP$^{50-95}$ across all evaluation datasets varying the feature dimension d_t.

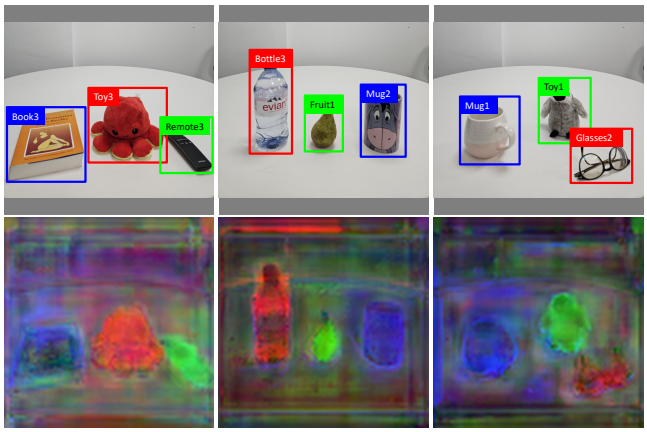

Figure 5: Feature similarity between the F_A embeddings and the teacher target u'_i is encoded in the R/G/B channels, one for each object in the input scene from POD dataset.

multimodal teacher into a lightweight student model, with the goal of achieving robust performance in few-shot personalized object detection across four personal datasets. We begin by discussing general analyses, followed by comparisons with state-of-the-art approaches, and conclude by presenting ablation studies.

4.1 Experimental Setup

Datasets. For pretraining, we adopt COCO (Lin et al. 2014) and OpenImagesV7 (Kuznetsova et al. 2020). Also, we consider starting from AuXFT pretrained weights (which start from COCO and fine-tune on OpenImages). Distillation is performed on OpenImages. Considered personalization datasets are: PerSeg (Zhang et al. 2024), POD (Barbato et al. 2024b), CORe50 (Lomonaco and Maltoni 2017), and iCubWorld (Fanello et al. 2013).

Models. Our approach uses YOLOv8n as the student model and LLaVa-1.5-7B as the teacher, which includes CLIP with ViT-B/32 as the visual encoder and LLaMA 7B as the text encoder. For comparison, we also include: (i) the text encoder from LLaVa, *i.e.*, LLaMA 7B (ii) the visual encoder from LLaVa, *i.e.*, CLIP, and (iii) DINOv2 (Oquab et al. 2023). Moreover, we compare against other distillation-based approaches: OFA (Hao et al. 2023), knowledge distillation at the output space via KL divergence (KL) (Touvron et al. 2021) and its combination with embedding space distillation via MSE (MSE + KL) (Romero et al. 2015), open-vocabulary detectors (VILD) (Gu et al. 2022), and the most relevant personalization method to date (AuXFT).

Metrics. For consistency and fair comparison, we follow the same implementation details as in AuXFT, except when distilling from AuXFT-pretrained model. In that case, the number of distillation epochs is reduced from 50 to 20, as the convergence time decreases. We also adopt the same evaluation metrics: mAP$^{50\text{-}95}$ for PerSeg and POD, and retrieval accuracy for CORe50 and iCubWorld. Few-shot personalized detection is led under 1- and 5-shot regimes.

4.2 Analyses

Embedding Loss. Fig. 3 illustrates the impact of our relational embedding loss $\mathcal{L}_{\text{emb}}$ on preserving the geometric structure of the teacher features onto student ones. We consider a toy setup, where we optimize the coordinates of a set of 2D points (proxies for student embeddings $f'_{A,i}$) to match the pairwise distance distribution of a fixed 3D reference configuration (proxy for teacher embeddings u'_i). Fig. 3a shows the 3D reference points sampled uniformly in space. Fig. 3b visualizes the optimization trajectory in the 2D space: each path starts from an initial position (marked by $\star$) and progressively adapts under the guidance of $\mathcal{L}_{\text{emb}}$, preserving local neighborhood relationships while globally rearranging to approximate the original distance geometry (color saturation indicates temporal evolution). Fig. 3c quantitatively tracks the alignment between neighborhood structures. We report the percentage rate of 2D top-k nearest neighbors that match those of the reference 3D distribution. We observe a steady improvement across all k with a rapid convergence. This analysis confirms the role of relational supervision in regularizing the student's embedding space beyond pointwise alignment, ultimately promoting feature diversity and semantic separability.

Impact of Teacher Feature Dimensionality. Fig. 4 evaluates the role of teacher supervision dimensionality (d_t) by comparing features: (i) of personal datasets $\mathcal{D}_f$ (as an upper bound), (ii) of the distillation dataset $\mathcal{D}_c$ unnormalized and (iii) PCA normalized (additional details and validation in Appendix). Overall, performance increases steadily up to $d_t = 512$, which strikes a good balance between compactness and expressivity. After that, it plateaus or slightly degrades up to the maximum dimension $d = 4608$. This indicates that most of the relevant semantics are preserved in a compact subspace, and that further dimensions may introduce redundancy or noise. Normalized features u'_i consistently outperform their unnormalized counterparts $\hat{u}'_i$ (especially in the 1-shot setting), obtaining better performance even if compared to personal datasets' features.

Teacher	Student	PerSeg	POD		CORe50		iCubWorld		Avg
		1 SHOT	1 SHOT	5 SHOT	1 SHOT	5 SHOT	1 SHOT	5 SHOT	
DINO	—	90.5 ±2.7	55.1 ±4.0	66.8 ±0.0	34.8 ±7.2	47.9 ±9.8	50.6 ±2.9	73.5 ±2.2	59.9
CLIP ($z_{V,i}$)	—	95.7 ±1.4	64.5 ±3.5	77.8 ±0.0	58.1 ±4.6	79.3 ±4.4	55.5 ±3.2	77.9 ±2.5	72.7
LLaVa (h_i)	—	82.2 ±3.3	59.0 ±4.1	72.7 ±0.5	44.1 ±3.4	63.1 ±4.2	46.7 ±3.1	67.6 ±2.7	62.2
LLaVa (h_i) + CLIP ($z_{V,i}, \gamma = 1$)	—	93.9 ±1.8	66.4 ±3.8	80.7 ±0.2	59.4 ±4.6	80.4 ±4.3	58.0 ±3.2	80.6 ±2.4	74.2
LLaVa (h_i) + CLIP ($z_{V,i}$)	—	95.1 ±1.7	70.0 ±3.4	80.6 ±0.0	57.0 ±4.0	79.7 ±3.7	56.4 ±3.6	81.2 ±2.5	**74.3**
—	YOLO	69.6 ±3.4	35.4 ±3.4	41.3 ±0.0	31.5 ±5.5	42.8 ±7.8	33.3 ±2.7	48.6 ±2.4	43.2
CLIP ($z_{V,i}$)	YOLO	75.5 ±3.5	40.9 ±3.2	51.7 ±0.0	35.0 ±3.9	47.3 ±5.7	40.6 ±2.9	60.0 ±2.3	50.7
LLaVa (h_i)	YOLO	73.0 ±3.6	37.1 ±4.2	47.6 ±0.0	33.9 ±3.7	45.9 ±5.4	37.4 ±3.1	55.8 ±2.6	48.0
LLaVa (h_i) + CLIP ($z_{V,i}$)	YOLO	82.0 ±3.4	45.7 ±3.5	58.1 ±0.0	34.5 ±4.1	47.8 ±5.9	42.9 ±3.1	62.7 ±2.2	**54.0**

Table 1: Oracle results using ground truth bounding boxes on personal datasets. Top block: performance when using frozen teacher embeddings directly in the personalization stage. Bottom block: performance of YOLO distilled with the teacher signals.

Teacher	Supervision	Student	PerSeg	POD		CORe50		iCubWorld		Avg
			1 SHOT	1 SHOT	5 SHOT	1 SHOT	5 SHOT	1 SHOT	5 SHOT	
—	—	YOLO	41.2 ±2.6	23.6 ±2.4	30.4 ±0.0	57.8 ±6.3	67.3 ±7.7	51.2 ±3.0	68.4 ±2.4	48.6
DINO	AuXFT	YOLO	48.8 ±3.4	31.5 ±2.7	38.8 ±0.0	58.8 ±6.4	69.3 ±6.0	55.0 ±3.0	74.5 ±2.5	53.8
CLIP ($z_{V,i}$)	ViLD	YOLO	44.4 ±3.2	24.8 ±2.8	33.3 ±0.0	54.6 ±5.2	64.9 ±4.9	57.7 ±3.0	70.4 ±2.4	50.9
LLaVa (u_i)	KL div.	YOLO	43.7 ±2.9	23.4 ±2.5	27.8 ±0.0	54.4 ±5.0	65.2 ±5.0	54.7 ±3.0	69.5 ±2.3	49.0
LLaVa (u_i)	KL div. + MSE	YOLO	50.1 ±3.0	27.6 ±2.4	30.5 ±0.0	53.3 ±4.5	65.7 ±4.6	61.1 ±2.9	**78.5** ±1.9	52.1
LLaVa (u_i)	OFA	YOLO	47.2 ±3.4	23.3 ±2.6	28.0 ±0.0	55.3 ±4.3	63.6 ±4.3	60.4 ±2.5	72.5 ±2.4	50.1
CLIP ($z_{V,i}$)	MOCHA (visual)	YOLO	53.0 ±2.8	26.7 ±3.1	37.5 ±0.0	57.2 ±4.7	67.2 ±4.8	62.6 ±2.8	76.1 ±2.3	54.6
LLaVa (h_i)	MOCHA (text)	YOLO	52.3 ±2.8	28.8 ±2.6	33.7 ±0.0	56.2 ±4.7	67.1 ±4.7	**62.8** ±2.8	76.8 ±2.0	54.0
LLaVa (u_i)	**MOCHA (COCO)**	YOLO	55.4 ±3.3	33.9 ±2.8	38.9 ±0.0	55.8 ±4.7	65.4 ±4.7	60.6 ±3.3	77.6 ±2.5	56.0
LLaVa (u_i)	**MOCHA (AuXFT)**	YOLO	**59.1** ±3.3	**36.3** ±3.3	**45.9** ±0.0	**60.9** ±4.4	**70.6** ±4.5	61.4 ±3.1	77.0 ±2.3	**58.7**

Table 2: Few-shot personalized detection results on personal datasets. First row: YOLO baseline with no distillation. First block: MOCHA compared against prior state-of-the-art supervision methods. Second block: evaluation of different MOCHA variants.

Feature Similarity. Fig. 5 shows the similarity between the teacher vision-language descriptors u'_i of the three objects in each scene and the features output by the student, F_A. In each bottom image, the three similarity matrices are visualized in R/G/B channels, showing that MOCHA successfully maps sparse VLM object-level features (inside object boxes) into the dense features of an efficient detector.

4.3 Main Results

We begin our evaluation in the oracle setup, where models are evaluated using ground-truth bounding boxes during the few-shot personalization stage. This setting removes the detection head and focuses purely on the classification capabilities of each approach, allowing us to compare different teacher models and types of supervision embeddings (linguistic, visual, and multimodal). Tab. 1 is structured in two blocks. The top block reports results from using teacher embeddings directly in the personalization stage, while the bottom block shows YOLO models distilled using those same supervision signals. In the first block, visual teachers show strong performance. Notably, CLIP outperforms DINO (used in AuXFT). On the other hand, LLaVa's features (h_i) perform worse, but combining them with CLIP features ($z_{V,i}$) — *i.e.*, using u_i for supervision — yields notable gains at no added cost (as CLIP is already part of LLaVa), demonstrating the complementarity of visual and textual cues. Normalizing the visual features also leads to a slight improvement over their unnormalized version ($\gamma = 1$). In the second block, we evaluate distillation using either visual supervision (CLIP, $z_{V,i}$), textual supervision (LLaVa, h_i), or their combination (LLaVa, u_i). All models are distilled through MOCHA supervision, showing consistent gains of our approach over the non-distilled baseline (YOLO only). Still, using a single modality alone proves less effective than leveraging full multimodal supervision.

Comparisons. We then evaluate the complete approach, including both backbone and detection head, in the standard few-shot personalization setting (without ground-truth bounding boxes). Tab. 2 shows a comparison between MOCHA and recent approaches for personalized detection. Among distillation-based methods, KL divergence and OFA yield only marginal improvements over the YOLO baseline, while combining KL with MSE provides slightly better performance. ViLD struggles in highly personalized scenarios, despite its open-vocabulary capabilities. AuXFT achieves the strongest average performance among prior methods, but still underperforms in more challenging domains (*e.g.*, PerSeg and POD), where fine-grained semantic alignment is critical. In contrast, MOCHA surpasses all other approaches showing consistently strong results across all datasets and settings, obtaining a +10.1 average score improvement over

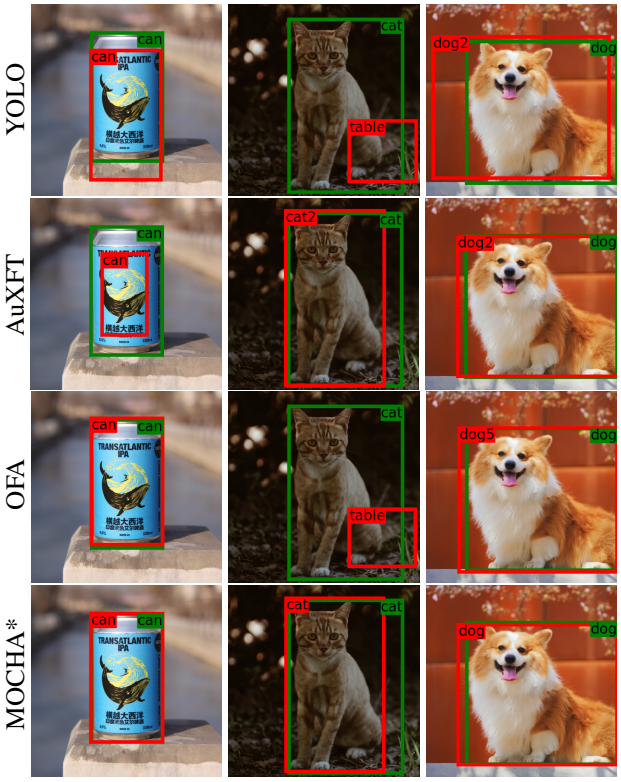

Figure 6: Qualitative results on PerSeg. Ground truth in green, prediction with the highest confidence in red (class names shown refer to personal class labels in PerSeg). YOLO refers to the baseline. *: refers to MOCHA (AuXFT).

the YOLO baseline and +4.9 on the best competitor, AuXFT. When distilled from COCO-pretrained weights, MOCHA already exceeds the performance of state-of-the-art methods. Starting from AuXFT-pretrained weights yields further gains, confirming the compatibility and effectiveness of the initialization. Overall, the results validate the effectiveness of MOCHA in transferring high-level semantics from frozen vision-language models into compact detectors, enabling robust personalization under few-shot constraints.

Qualitative Results. Fig. 6 presents a qualitative comparison on the PerSeg dataset, featuring three sample images: a *can*, a *cat*, and a *dog*. Green bounding boxes indicate ground truth annotations, while red boxes show the highest-confidence predictions. MOCHA (AuXFT) consistently demonstrates superior performance across all samples, with correct class predictions and precise bounding boxes. In contrast, competing methods exhibit occasional misclassifications and less accurate detections. The results emphasize the effectiveness of MOCHA's feature translation and relational alignment strategies in enhancing spatial precision and adaptability to novel classes.

4.4 Ablation Studies

Tab. 3 reports an extensive ablation study, isolating the contributions of key design and training components of

Pretr.	d_t	Case	PerSeg 1 SHOT	iCubWorld 1 SHOT	iCubWorld 5 SHOT
OI	—	(YOLO)	41.2 ±2.6	51.2 ±3.0	68.4 ±2.4
AuXFT	256	—	56.4 ±3.3	61.2 ±2.6	75.7 ±2.1
AuXFT	384	—	56.2 ±3.5	59.6 ±3.0	74.7 ±2.0
AuXFT	1024	—	53.8 ±3.5	58.4 ±3.1	77.0 ±2.3
AuXFT	512	$\mathcal{L}_{emb}$	56.6 ±3.3	59.8 ±2.9	74.9 ±2.0
AuXFT	512	t_S no attn.	56.8 ±3.3	57.5 ±3.1	73.2 ±2.1
AuXFT	512	t_S no ffn	55.7 ±3.4	**61.4** ±3.0	76.6 ±2.3
AuXFT	512	$\mathcal{D}_c$ augm.	53.3 ±3.9	58.4 ±2.9	73.4 ±2.1
AuXFT	512	2×lr (enc)	53.6 ±2.9	59.0 ±3.0	75.8 ±2.5
AuXFT	512	2×lr (dec)	53.9 ±2.9	59.7 ±3.1	75.9 ±2.5
OI	512	(**MOCHA**)	53.8 ±3.2	56.8 ±2.7	70.4 ±2.3
COCO	512	(**MOCHA**)	55.4 ±3.3	60.6 ±3.3	**77.6** ±2.5
AuXFT	512	(**MOCHA**)	**59.1** ±3.3	**61.4** ±3.1	77.0 ±2.3

Table 3: Ablation studies on MOCHA. First block ablates on PCA dimension d_t. The second block changes different MOCHA components. The third block compares different pretraining: OpenImages (OI), COCO, and AuXFT.

MOCHA across the PerSeg and iCubWorld datasets. We first assess the impact of the PCA dimensionality d_t used to compress teacher embeddings. As previously highlighted in Fig. 4, compressing the teacher's embedding improves performance compared to using full-resolution features. A PCA dimension of $d_t = 512$ strikes a good balance between compactness and expressiveness. More aggressive compression (*e.g.*, 256 or 384) leads to reduced performance, indicating a loss of relevant semantic content. Next, we examine the role of core MOCHA components. Removing the embedding alignment loss ($\mathcal{L}_{\text{emb}}$) degrades accuracy on both datasets, underscoring its importance for preserving global teacher structure. Disabling attention mechanisms in the translator (t_S no attn.) also leads to a consistent drop, suggesting that attention layers play a key role in adapting multimodal features. Likewise, removing the translator's final layer (t_S no ffn) reduces performance. Enabling data augmentation (standard YOLO augmentations only varying color, $\mathcal{D}_c$ augm.) lowers robustness, while manipulating learning rates — doubling them for encoder or decoder separately — slightly harms performance. Finally, we compare different pretraining strategies. When pretraining is performed on OpenImages only, performance remains below that of MOCHA configurations. COCO-pretrained weights yield a substantial boost, but AuXFT-pretrained weights lead to the best results overall, thanks to their strong initial performance and architectural compatibility with our MOCHA.

5 Conclusion

In this work, we introduced MOCHA, a novel approach for multimodal, cross-architecture distillation tailored to few-shot personalized object detection. MOCHA aligns compact student features to region-level embeddings derived from a large-scale frozen vision-language model, removing redundancy via PCA. Our approach leverages both intra-class and inter-class alignment, through a distillation loss and an em-

bedding loss, respectively, achieving significant gains over the state of the art. MOCHA requires minimal overhead and operates without the teacher model at test time, making it compatible with lightweight detectors and suitable for resource-constrained applications such as mobile robotics.

References

Alayrac, J.-B.; Donahue, J.; Luc, P.; Miech, A.; Barr, I.; Hasson, Y.; Lenc, K.; Mensch, A.; Millicah, K.; Reynolds, M.; Ring, R.; Rutherford, E.; Cabi, S.; Han, T.; Gong, Z.; Samangooei, S.; Monteiro, M.; Menick, J.; Borgeaud, S.; Brock, A.; Nematzadeh, A.; Sharifzadeh, S.; Binkowski, M.; Barreira, R.; Vinyals, O.; Zisserman, A.; and Simonyan, K. 2022. Flamingo: a visual language model for few-shot learning. In *International Conference on Neural Information Processing Systems (NeurIPS)*. Red Hook, NY, USA: Curran Associates Inc. ISBN 9781713871088.

Barbato, F.; Camuffo, E.; Milani, S.; and Zanuttigh, P. 2024a. Continual road-scene semantic segmentation via feature-aligned symmetric multi-modal network. In *IEEE International Conference on Image Processing (ICIP)*, 722–728. IEEE.

Barbato, F.; Michieli, U.; Moon, J.; Zanuttigh, P.; and Ozay, M. 2024b. Cross-Architecture Auxiliary Feature Space Translation for Efficient Few-Shot Personalized Object Detection. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE.

Camuffo, E.; Michieli, U.; and Milani, S. 2023. Learning from mistakes: Self-regularizing hierarchical representations in point cloud semantic segmentation. *IEEE Transactions on Multimedia*.

Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *North American Chapter of the Association for Computational Linguistics*.

Dosovitskiy, A. 2020. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*.

Fanello, S. R.; Ciliberto, C.; Santoro, M.; Natale, L.; Metta, G.; Rosasco, L.; and Odone, F. 2013. iCub World: Friendly Robots Help Building Good Vision Data-Sets. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, 700–705.

Girdhar, R.; El-Nouby, A.; Liu, Z.; Singh, M.; Alwala, K. V.; Joulin, A.; and Misra, I. 2023. Imagebind: One embedding space to bind them all. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 15180–15190.

Gou, J.; Chen, Y.; Yu, B.; Liu, J.; Du, L.; Wan, S.; and Yi, Z. 2024. Reciprocal Teacher-Student Learning via Forward and Feedback Knowledge Distillation. *IEEE Transactions on Multimedia*, 26: 7901–7916.

Gu, J.; Ghiasi, G.; Cui, Y.; Wang, Z.; et al. 2022. ViLD: Open-Vocabulary Object Detection via Vision and Language Knowledge Distillation. *International Conference on learning Representations (ICLR)*.

Hao, Z.; Guo, J.; Han, K.; Tang, Y.; Hu, H.; Wang, Y.; and Xu, C. 2023. One-for-All: Bridge the Gap Between Heterogeneous Architectures in Knowledge Distillation. In *IEEE International Conference on Neural Information Processing Systems (NeurIPS)*.

Hao, Z.; Luo, Y.; Wang, Z.; Hu, H.; and An, J. 2022. CDFKD-MFS: Collaborative Data-Free Knowledge Distillation via Multi-Level Feature Sharing. *IEEE Transactions on Multimedia*, 24: 4262–4274.

Huh, M.; Cheung, B.; Wang, T.; and Isola, P. 2024. The Platonic Representation Hypothesis. *Proceedings of Machine Learning Research (PMLR)*.

Jaritz, M.; Vu, T.-H.; Charette, R. d.; Wirbel, E.; and Pérez, P. 2020. xmuda: Cross-modal unsupervised domain adaptation for 3d semantic segmentation. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 12605–12614.

Jaritz, M.; Vu, T.-H.; De Charette, R.; Wirbel, É.; and Pérez, P. 2022. Cross-modal learning for domain adaptation in 3d semantic segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, 45(2): 1533–1544.

Jocher, G.; Chaurasia, A.; and Qiu, J. 2023. Ultralytics YOLOv8 [Computer Software]. https://github.com/ultralytics/ultralytics. Accessed July 2024.

Kirillov, A.; Mintun, E.; Ravi, N.; Mao, H.; Rolland, C.; Gustafson, L.; Xiao, T.; Whitehead, S.; Berg, A. C.; Lo, W.-Y.; et al. 2023. Segment anything. In *IEEE/CVF International Conference on Computer Vision (ICCV)*, 4015–4026. IEEE.

Kothapalli, V. 2022. Neural Collapse: A Review on Modelling Principles and Generalization. *Transactions on Machine Learning Research*.

Kuznetsova, A.; Rom, H.; Alldrin, N.; Uijlings, J.; Krasin, I.; Pont-Tuset, J.; Kamali, S.; Popov, S.; Malloci, M.; Kolesnikov, A.; et al. 2020. The open images dataset v4: Unified image classification, object detection, and visual relationship detection at scale. *International Journal of Computer Vision (IJCV)*, 128(7): 1956–1981.

Lee, C.-Y.; Xie, S.; Gallagher, P.; Zhang, Z.; and Tu, Z. 2014. Deeply-Supervised Nets. *Proceedings of Machine Learning Research (PMLR)*.

Li, M.; Zhang, L.; Zhu, M.; Huang, Z.; Yu, G.; Fan, J.; and Chen, T. 2024. Lightweight Model Pre-Training via Language Guided Knowledge Distillation. *IEEE Transactions on Multimedia*, 26: 10720–10730.

Lin, T.-Y.; Maire, M.; Belongie, S.; Hays, J.; Perona, P.; Ramanan, D.; Dollár, P.; and Zitnick, C. L. 2014. Microsoft coco: Common objects in context. In *European Conference on Computer Vision (ECCV)*, 740–755. Springer.

Liu, H.; Li, C.; Wu, Q.; and Lee, Y. J. 2023a. Visual instruction tuning. *Advances in Neural Information Processing Systems*, 36: 34892–34916.

Liu, Y.; Cao, J.; Li, B.; Hu, W.; Ding, J.; and Li, L. 2022. Cross-architecture knowledge distillation. In *IEEE/CVF Proceedings of the Asian Conference on Computer Vision (ACCV)*, 3396–3411.

Liu, Y.; Zhu, M.; Li, H.; Chen, H.; Wang, X.; and Shen, C. 2023b. Matcher: Segment anything with one shot using all-purpose feature matching. *International Conference on Learning Representations (ICLR)*.

Liu, Z.; Lin, Y.; Cao, Y.; Hu, H.; Wei, Y.; Zhang, Z.; Lin, S.; and Guo, B. 2021. Swin transformer: Hierarchical vision transformer using shifted windows. In *IEEE/CVF International Conference on Computer Vision (ICCV)*, 10012–10022. IEEE.

Lomonaco, V.; and Maltoni, D. 2017. CORe50: a New Dataset and Benchmark for Continuous Object Recognition. arXiv:1705.03550.

Luo, S.; Wang, X.; Fang, G.; Hu, Y.; Tao, D.; and Song, M. 2019. Knowledge amalgamation from heterogeneous networks by common feature learning. *International Joint Conference on Artificial Intelligence (IJCAI)*.

Meng, Z.; Xia, X.; and Ma, J. 2024. Toward Foundation Models for Inclusive Object Detection: Geometry- and Category-Aware Feature Extraction Across Road User Categories. *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, 54(11): 6570–6580.

Michieli, U.; Moon, J.; Kim, D.; and Ozay, M. 2024. Object-Conditioned Bag of Instances for Few-Shot Personalized Instance Recognition. In *IEEE/SPS International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 7885–7889. IEEE.

Oquab, M.; Darcet, T.; Moutakanni, T.; Vo, H. V.; Szafraniec, M.; Khalidov, V.; Fernandez, P.; Haziza, D.; Massa, F.; El-Nouby, A.; Howes, R.; Huang, P.-Y.; Xu, H.; Sharma, V.; Li, S.-W.; Galuba, W.; Rabbat, M.; Assran, M.; Ballas, N.; Synnaeve, G.; Misra, I.; Jegou, H.; Mairal, J.; Labatut, P.; Joulin, A.; and Bojanowski, P. 2023. DINOv2: Learning Robust Visual Features without Supervision. *arXiv preprint arXiv:2304.07193*.

Papyan, V.; Han, X. Y.; and Donoho, D. L. 2020. Prevalence of neural collapse during the terminal phase of deep learning training. *Proceedings of the National Academy of Sciences*, 117(40): 24652–24663.

Paramonov, K.; Zhong, J.-X.; Michieli, U.; Moon, J.; and Ozay, M. 2024. Swiss dino: Efficient and versatile vision framework for on-device personal object search. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2564–2571. IEEE.

Passalis, N.; Tzelepi, M.; and Tefas, A. 2020. Heterogeneous knowledge distillation using information flow modeling. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2339–2348. IEEE.

Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; Krueger, G.; and Sutskever, I. 2021. Learning Transferable Visual Models From Natural Language Supervision. *Proceedings of Machine Learning Research (PMLR)*.

Romero, A.; Ballas, N.; Kahou, S. E.; Chassang, A.; Gatta, C.; and Bengio, Y. 2015. FitNets: Hints for Thin Deep Nets. In *International Conference on Learned Representations (ICLR)*.

Snell, J.; Swersky, K.; and Zemel, R. 2017. Prototypical Networks for Few-shot Learning. In *Advances in Neural Information Processing Systems*, volume 30.

Szegedy, C.; Liu, W.; Jia, Y.; Sermanet, P.; Reed, S.; Anguelov, D.; Erhan, D.; Vanhoucke, V.; and Rabinovich, A. 2014. Going Deeper with Convolutions. *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.

Touvron, H.; Cord, M.; Douze, M.; Massa, F.; Sablayrolles, A.; and Jegou, H. 2021. Training data-efficient image transformers & distillation through attention. In *International Conference on Machine Learning (ICML)*, volume 139 of *Proceedings of Machine Learning Research (PMLR)*, 10347–10357.

Tu, Z.; Zhou, W.; Qian, X.; and Yan, W. 2025. Hybrid Knowledge Distillation Network for RGB-D Co-Salient Object Detection. *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, 1–12.

Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. *Advances in Neural Information Processing Systems*, 30.

Wang, X.; Zhang, X.; Cao, Y.; Wang, W.; Shen, C.; and Huang, T. 2023. Seggpt: Segmenting everything in context. *IEEE/CVF International Conference on Computer Vision (ICCV)*.

Xia, R.; Zhao, C.; Zheng, M.; Wu, Z.; Sun, Q.; and Tang, Y. 2023. Cmda: Cross-modality domain adaptation for nighttime semantic segmentation. In *IEEE/CVF International Conference on Computer Vision (ICCV)*, 21572–21581.

Zhang, L.; Shi, Y.; Shi, Z.; Ma, K.; and Bao, C. 2020. Task-oriented feature distillation. *Advances in Neural Information Processing Systems*, 33: 14759–14771.

Zhang, R.; Jiang, Z.; Guo, Z.; Yan, S.; Pan, J.; Dong, H.; Qiao, Y.; Gao, P.; and Li, H. 2024. Personalize Segment Anything Model with One Shot. In *International Conference on Learning Representations (ICLR)*.

Zhang, Y.; Zhang, C.-B.; Jiang, P.-T.; Cheng, M.-M.; and Mao, F. 2021. Personalized image semantic segmentation. In *IEEE/CVF International Conference on Computer Vision (ICCV)*, 10549–10559. IEEE.

Appendix

In this appendix, we provide additional experiments and analyses to validate our methodology further. We start by providing a statistical validation of MOCHA using the Wilcoxon signed-rank test to assess the significance of our findings. We continue by presenting additional results, including PCA curve fitting analyses and detailed performance breakdowns across individual datasets. We then outline the implementation details of our experiments and provide a pseudocode description of the MOCHA methodology. Finally, we conclude by presenting some qualitative results and discussing MOCHA's limitations.

A.1 Wilcoxon Experiment

To verify the statistical significance of the results presented in the main document, in Tab. A.4 we report the Wilcoxon signed rank metric computed on the mAP achieved by MOCHA and its competitors in the PerSeg and POD datasets. The analysis is performed on the same episodic experiments in Tab. 1 (in the main document), where average scores are reported. The analysis strongly confirms the statistical superiority of MOCHA, achieving, in most cases, an extremely low p-value $< 10^{-17}$.

Supervision	**PerSeg** 1 SHOT	**POD** 1 SHOT
—	1.9×10^{-18}	1.9×10^{-18}
AuXFT	4.5×10^{-18}	1.1×10^{-9}
ViLD	2.1×10^{-18}	1.9×10^{-18}
KL div.	1.9×10^{-18}	1.9×10^{-18}
MSE + KL div.	2.1×10^{-18}	1.9×10^{-18}
OFA	1.9×10^{-18}	1.9×10^{-18}
MOCHA (visual)	2.0×10^{-18}	1.9×10^{-18}
MOCHA (text)	1.9×10^{-18}	2.0×10^{-18}
MOCHA (COCO)	5.5×10^{-13}	7.7×10^{-9}

Table A.4: p-value of Wilcoxon signed rank metric on PerSeg/POD 1-shot between MOCHA (AuXFT) and competitors.

A.2 Dimensionality Reduction

Fig. A.9 shows the variance profile of teacher embeddings, with a fitted decay curve of the form $\sigma \simeq \frac{a}{(x+1)^b} + c$, confirming the typical power-law structure that motivates dimensionality reduction. The standard deviation of each channel activation in the distillation dataset is estimated and used to rescale the channel c as: $\sigma_c \simeq \frac{18}{(c+1)^{0.47}} - 0.26$ (Eq. 3 in the main document). Tables **??** report extensive results across all datasets and protocols for different PCA settings under three configurations:

- PCA on $\hat{u}'_i$ in Personal Dataset (Tab. A.5), where PCA is applied on teacher embeddings of each personal dataset (ideal case);
- PCA on $\hat{u}'_i$ in Distillation Dataset (Tab. A.6), where PCA is applied on teacher embeddings of distillation dataset;
- PCA on u'_i in Distillation Dataset (Tab. A.7), where the channel normalization of Fig. A.9 is applied prior to PCA.

Across all experiments, we observe that moderate compression (*e.g.*, retaining 128–256 dimensions) achieves competitive performance, with minimal loss relative to full-dimensional embeddings, while substantially reducing memory and compute costs. As expected, aggressive compression (*e.g.*, 16–32 dimensions) leads to consistent drops across datasets, particularly on POD and iCubWorld. This analysis confirms that PCA acts as an effective mechanism to control embedding size while preserving sufficient discriminative information for distillation. The results validate the robustness of our approach to dimensionality reduction, supporting the choice of PCA dimension $d_t = 512$.

A.3 Qualitative Results

Figures **??** present qualitative comparisons across three models—YOLO (baseline), AuXFT, and MOCHA (AuXFT)—in both the 1-shot and 5-shot personalization settings. Results are shown on three representative datasets: POD, CoRE50, and iCubWorld. Ground truth annotations are depicted in green, while the prediction with the highest confidence is shown in red. In the 1-shot setting (Figure A.7), MOCHA already improves over both YOLO and AuXFT, correctly localizing and classifying objects that other models miss or mislabel (*e.g.*, *remote2* in POD, *case0* in iCubWorld). With 5-shot supervision (Figure A.8), MOCHA further enhances precision, handling subtle distinctions (*e.g.*, *pen5* vs. *mug5* in CoRE50), and resolving multiple similar instances in cluttered scenes. These examples qualitatively validate MOCHA's effectiveness in aligning student features to the teacher's multimodal embedding space, enabling more reliable object personalization with minimal data.

A.4 Implementation Details

All experiments were run on a RHEL8 (RedHat) Unix-based machine with kernel version 6.12.9-1, equipped with 8 NVIDIA L40S GPUs (48GB of VRAM, CUDA 12.8, Driver 570.86.15), 2 AMD EPYC 9224 24-Core processors, 1.5TB of RAM, and Python version 3.11.9. Each pretraining experiment uses 4 GPUs, 24 CPU cores, and 64GB of RAM, while for few-shot training and inference, a single GPU and 6 CPU cores are sufficient. A complete MOCHA distillation (50 epochs) takes about 48 hours, which decreases to around 20 hours when distilling from the AuXFT checkpoints (20 epochs).

A.4.1 MOCHA Pseudocode

In this section, we report the pseudocode implementation of MOCHA's feature distillation and few-shot learning procedures. Algorithm A.1 reports the steps necessary to distill visual-language knowledge into an efficient vision-only detector, while Algorithms A.2 and A.3 explain how to train MOCHA's few-shot module and perform inference, respectively. Note that # marks a comment.

Algorithm A.1: Pseudocode for MOCHA's feature distillation.

Require: Pretrained detection model $\mathrm{m}_S = \mathrm{l}_S \circ \mathrm{g}_S$, visual encoder g_T, language model f_T dataset $\mathcal{D}_c$, translation module t_S, batch size K, number of epochs E

for $e \leftarrow 1 \dots E$ **do**
  for $k \leftarrow 1 \dots \lfloor |\mathcal{D}_c|/K \rfloor$ **do**
    Sample $\mathcal{K} \sim \mathcal{D}_c$, with $|\mathcal{K}| = K$
    Initialize batch loss $l \leftarrow 0$
    for $(X, \mathcal{Y}) \in \mathcal{K}$ **do**
      $\{F_j\} \leftarrow \mathrm{g}_S(X)$ # Compute detection features
      $\hat{\mathcal{Y}} \leftarrow \mathrm{l}_S(\{F_j\})$ # Compute detection predictions
      $l \leftarrow l + \mathcal{L}_{\mathrm{det}}(\hat{\mathcal{Y}}, \mathcal{Y})$ # Accumulate loss
      for $(b_i, c_i) \in \mathcal{Y}$ **do**
        Use b_i to crop region X_i from X
        $z_{V,i} \leftarrow \mathrm{g}_T(X_i)$ # Vision branch of teacher
        $h_i \leftarrow \mathrm{f}_T(c_i)$ # Text branch of teacher
        Compute u_i from $z_{V,i}$ and h_i using Eq. 1
        Use b_i to crop region $F_{A,j,i}$ from F_j, $\forall j$
        $f_{A,i} \leftarrow \mathrm{concat}_j(\mathrm{AvgPool}(F_{A,j,i}))$
        $f'_{A,i} \leftarrow \mathrm{t}_S(f_{A,i})$ # Translate features
        $l \leftarrow l + \mathcal{L}_{\mathrm{dist}}(f'_{A,i}, u_i)/n$ # n = num of boxes
      end for
      Compute P_{ff}, P_{uu} from $\{f'_{A,i}\}, \{u_i\}$ via Eq. 7
      $l \leftarrow l + \mathcal{L}_{\mathrm{emb}}(P_{ff}, P_{uu})$
    end for
    $(\mathrm{m}_S, \mathrm{t}_S) \leftarrow \mathrm{AdamOptim}(\mathrm{m}_S, \mathrm{t}_S, l)$ # Perform gradient descent to update model and translator
  end for
end for

A.5 MOCHA Limitations

The improvements brought about by MOCHA are substantial, but we wish to highlight some limitations that we consider potential directions for future research. First, due to the nature of the MOCHA architecture, the final task is constrained to object detection, which may limit applicability in other domains. Second, while inference requires minimal overhead, the reliance on foundation models during training incurs significant computational cost, which is undesirable in some contexts. Finally, the prototypical network in our method is inherently inefficient at large scales, meaning that performance may degrade when a large number of personal object instances are involved.

Algorithm A.2: Pseudocode for MOCHA's FSL training.

Require: Frozen distilled backbone g_S, frozen translator t_S, prototype classifier $\mathrm{p}(\cdot)$, personal dataset $\mathcal{D}_f$ (train)

Initialize empty support set $\mathcal{S} \leftarrow \emptyset$
for $(X, \mathcal{Y}) \in \mathcal{D}_f$ **do** # $\mathcal{Y} \subset \mathcal{B} \times \mathcal{C}_c$
  $\{F_j\} \leftarrow \mathrm{g}_S(X)$ # Extract multiscale features
  for $(b_i, c_i) \in \mathcal{Y}$ **do**
    Use b_i to crop region $F_{A,j,i}$ from F_j, $\forall j$
    $f_{A,i} \leftarrow \mathrm{concat}_j(\mathrm{AvgPool}(F_{A,j,i}))$
    $f'_{A,i} \leftarrow \mathrm{t}_S(f_{A,i})$ # Translated feature
    Add pair $(f'_{A,i}, c_i)$ to support set $\mathcal{S}$
  end for
end for
Add $\mathcal{S}$ to $\mathrm{p}(\cdot)$'s prototypes store, internal to the local device # Prototype classifier training

Algorithm A.3: Pseudocode for MOCHA's FSL inference.

Input: Input image X
Require: Frozen distilled model $\mathrm{m}_S = \mathrm{l}_S \circ \mathrm{g}_S$, frozen translator t_S, trained prototype classifier $\mathrm{p}(\cdot)$

$\{F_j\} \leftarrow \mathrm{g}_S(X)$ # Extract multiscale features
$\hat{\mathcal{Y}} \leftarrow \mathrm{l}_S(\{F_j\})$ # Compute detection predictions for general classes
for $(\hat{b}_i, \hat{c}_i) \in \hat{\mathcal{Y}}$ **do** # $\hat{\mathcal{Y}} \subset \mathcal{B} \times \mathcal{C}_c$
  Use $\hat{b}_i$ to crop region $F_{A,j,i}$ from F_j $\forall j$
  $f_{A,i} \leftarrow \mathrm{concat}_j(\mathrm{AvgPool}(F_{A,j,i}))$
  $f'_{A,i} \leftarrow \mathrm{t}_S(f_{A,i})$ # Translate to shared space
  $\hat{c}_i \leftarrow \mathrm{p}(f'_{A,i})$ # Update class via prototype matching over fine-grained class set
end for
return Predictions $\hat{\mathcal{Y}}$ # Same boxes, personal classes

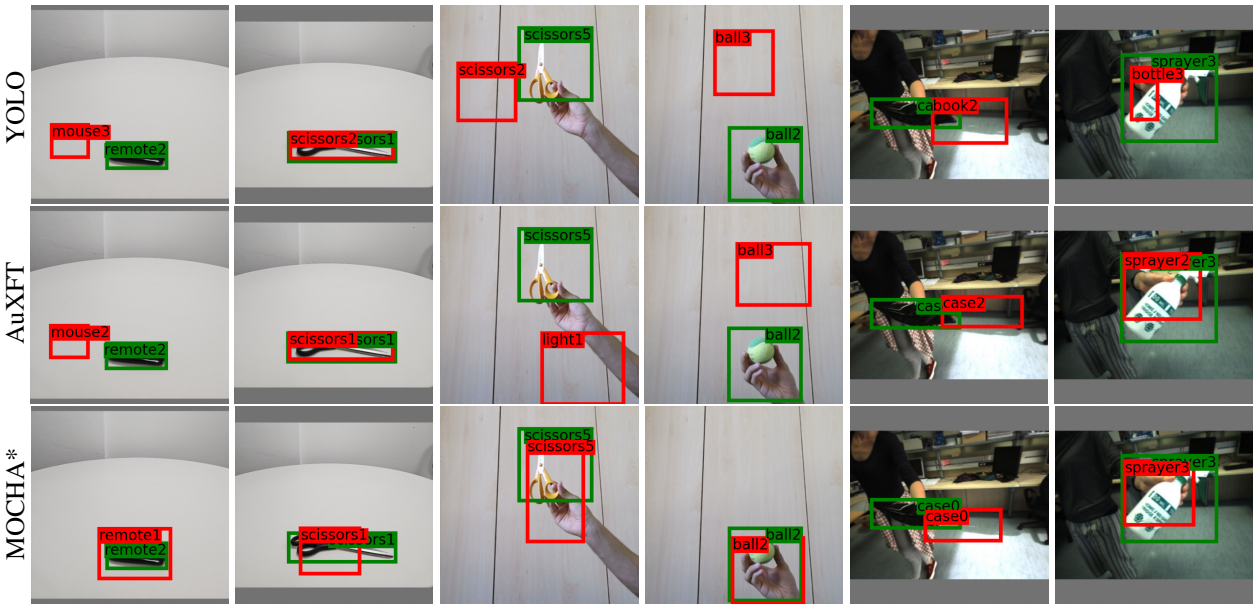

Figure A.7: Qualitative results in the 1-shot setting. Ground truth in green, prediction with the highest confidence in red (class names shown refer to personal class labels in each dataset). YOLO refers to the baseline. *: refers to MOCHA (AuXFT).

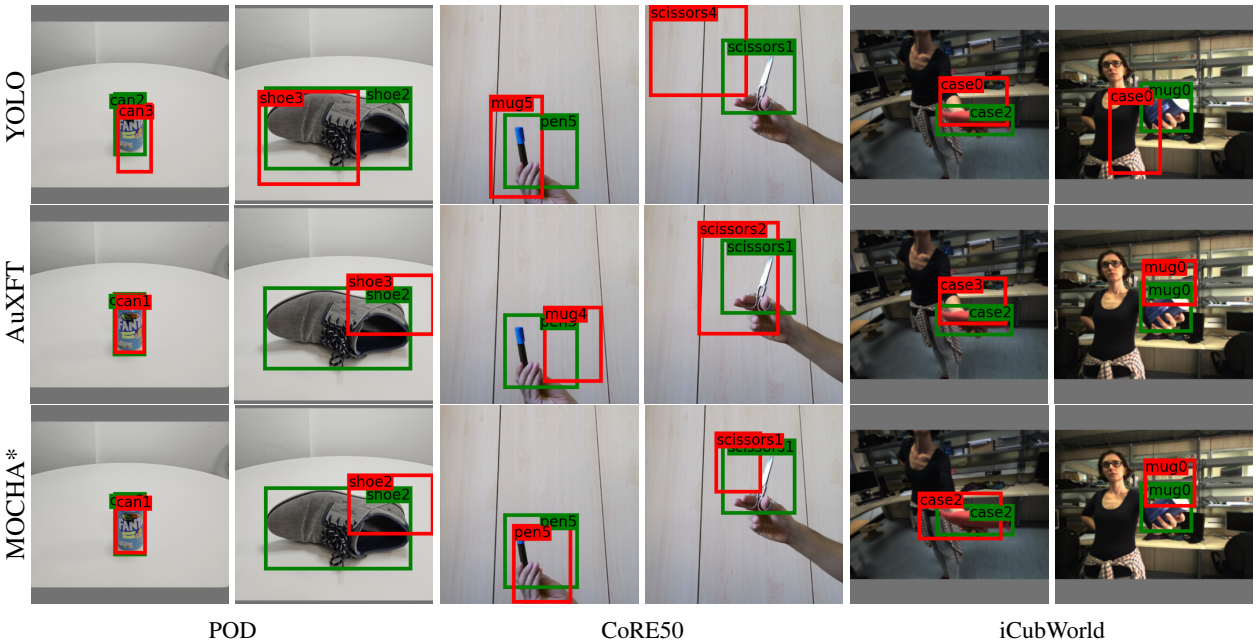

Figure A.8: Qualitative results in the 5-shot setting. Ground truth in green, prediction with the highest confidence in red (class names shown refer to personal class labels in each dataset). YOLO refers to the baseline. *: refers to MOCHA (AuXFT).

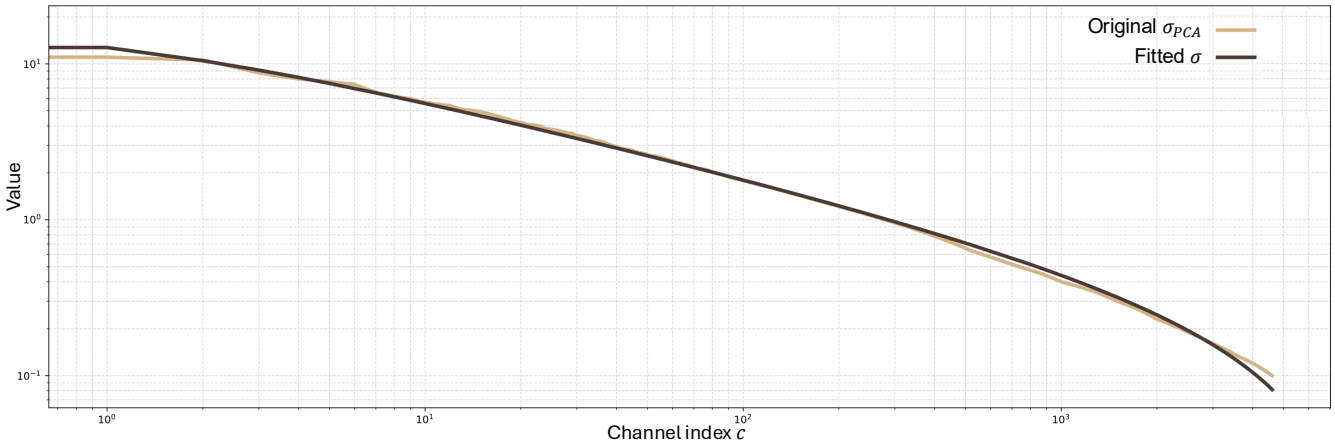

Figure A.9: Fitting the PCA curve $\sigma \simeq \frac{a}{(x+1)^b} + c$, $a = 18, b = 0.47, c = -0.26$.

Model	PerSeg	POD		CORe50		iCubWorld		Avg
	1 SHOT	1 SHOT	5 SHOT	1 SHOT	5 SHOT	1 SHOT	5 SHOT	
LLaVa (u_i) (4608 channels)	94.0 ± 1.8	66.4 ± 3.8	80.7 ± 0.0	59.4 ± 4.6	80.4 ± 4.3	58.0 ± 3.2	80.6 ± 2.4	74.2
PCA 1024 (4.5x less channels)	94.2 ± 1.7	67.9 ± 3.4	80.6 ± 0.0	60.1 ± 4.6	80.8 ± 4.3	59.3 ± 3.1	81.2 ± 2.3	**74.9**
PCA 512 (9x less channels)	94.3 ± 1.7	67.9 ± 3.5	80.4 ± 0.0	60.3 ± 4.6	80.6 ± 4.3	59.3 ± 3.1	80.9 ± 2.3	74.8
PCA 384 (12x less channels)	94.4 ± 1.7	67.5 ± 3.4	80.0 ± 0.0	60.4 ± 4.6	80.6 ± 4.3	59.3 ± 3.0	80.8 ± 2.2	74.7
PCA 256 (18x less channels)	94.2 ± 1.6	67.6 ± 3.3	79.6 ± 0.0	60.6 ± 4.6	80.5 ± 4.5	59.2 ± 3.1	80.2 ± 2.2	74.6
PCA 128 (36x less channels)	92.9 ± 1.5	66.4 ± 3.1	79.4 ± 0.0	60.5 ± 4.6	79.8 ± 4.3	57.5 ± 3.1	77.6 ± 2.4	73.4
PCA 64 (72x less channels)	88.0 ± 2.9	63.9 ± 3.4	79.9 ± 0.0	59.0 ± 4.5	77.1 ± 4.6	52.7 ± 3.3	71.2 ± 2.4	70.3
PCA 32 (144x less channels)	81.5 ± 2.8	60.0 ± 3.5	70.6 ± 0.0	53.5 ± 4.7	71.4 ± 5.0	45.9 ± 3.1	62.1 ± 2.7	63.6

Table A.5: PCA on $\hat{u}'_i$ in Personal Dataset (mAP/mAcc ± std). The last column (Avg) reports the average across the 7 columns. Overall best in bold.

Model	PerSeg	POD		CORe50		iCubWorld		Avg
	1 SHOT	1 SHOT	5 SHOT	1 SHOT	5 SHOT	1 SHOT	5 SHOT	
LLaVa (u_i) (4608 channels)	94.0 ± 1.8	66.4 ± 3.8	80.7 ± 0.0	59.4 ± 4.6	80.4 ± 4.3	58.0 ± 3.2	80.6 ± 2.4	74.2
PCA 1024 (4.5x less channels)	94.4 ± 1.7	67.9 ± 3.5	80.3 ± 0.0	59.7 ± 4.5	80.3 ± 4.3	59.2 ± 3.3	81.2 ± 2.2	**74.7**
PCA 512 (9x less channels)	94.2 ± 1.6	67.8 ± 3.5	80.1 ± 0.0	59.8 ± 4.5	80.1 ± 4.3	58.9 ± 3.3	80.8 ± 2.2	74.5
PCA 384 (12x less channels)	93.7 ± 1.7	67.7 ± 3.4	80.4 ± 0.0	59.8 ± 4.6	79.9 ± 4.4	58.6 ± 3.3	80.3 ± 2.3	74.3
PCA 256 (18x less channels)	93.0 ± 1.6	67.4 ± 3.5	79.2 ± 0.0	59.3 ± 4.6	79.6 ± 4.4	57.6 ± 3.3	79.3 ± 2.2	73.6
PCA 128 (36x less channels)	90.4 ± 2.0	65.9 ± 3.1	78.6 ± 0.0	56.7 ± 4.7	77.2 ± 4.5	54.3 ± 3.3	75.3 ± 2.3	71.2
PCA 64 (72x less channels)	85.8 ± 3.1	62.5 ± 3.3	74.8 ± 0.0	51.9 ± 4.6	72.2 ± 4.7	48.4 ± 3.2	67.9 ± 2.5	66.2
PCA 32 (144x less channels)	79.3 ± 3.2	57.0 ± 3.6	72.1 ± 0.0	44.8 ± 4.7	63.4 ± 5.1	43.7 ± 3.2	60.7 ± 2.6	60.1

Table A.6: PCA on $\hat{u}'_i$ in Distillation Dataset (mAP/mAcc ± std). The last column (Avg) reports the average across the 7 columns. Overall best in bold.

Model	PerSeg	POD		CORe50		iCubWorld		Avg
	1 SHOT	1 SHOT	5 SHOT	1 SHOT	5 SHOT	1 SHOT	5 SHOT	
LLaVa (u_i) (4608 channels)	95.1 ± 1.7	70.0 ± 3.4	80.6 ± 0.0	57.0 ± 4.0	79.7 ± 3.7	56.4 ± 3.6	81.2 ± 2.5	74.3
PCA 1024 (4.5x less channels)	96.6 ± 1.3	68.5 ± 3.7	81.7 ± 0.0	58.9 ± 4.3	80.6 ± 3.9	61.2 ± 3.6	83.4 ± 2.2	75.8
PCA 512 (9x less channels)	96.9 ± 1.1	68.9 ± 3.7	81.9 ± 0.0	59.6 ± 4.3	80.6 ± 4.0	63.7 ± 3.4	85.5 ± 2.1	76.7
PCA 384 (12x less channels)	96.4 ± 1.2	68.5 ± 3.7	83.3 ± 0.0	60.4 ± 4.4	80.6 ± 4.2	64.0 ± 3.2	86.0 ± 2.0	**77.0**
PCA 256 (18x less channels)	95.2 ± 1.3	68.6 ± 3.6	83.2 ± 0.0	60.9 ± 4.4	81.0 ± 4.2	64.1 ± 3.3	86.0 ± 1.9	**77.0**
PCA 128 (36x less channels)	92.2 ± 1.9	67.7 ± 3.7	79.3 ± 0.0	59.9 ± 4.4	79.6 ± 4.3	62.6 ± 3.2	84.2 ± 1.9	75.1
PCA 64 (72x less channels)	87.9 ± 2.2	65.7 ± 3.6	75.7 ± 0.0	55.3 ± 4.5	74.5 ± 4.7	55.0 ± 3.1	76.4 ± 2.4	70.1
PCA 32 (144x less channels)	80.2 ± 3.2	60.6 ± 3.8	74.6 ± 0.0	46.4 ± 4.6	64.7 ± 5.5	48.0 ± 3.2	68.0 ± 2.8	63.2

Table A.7: PCA on u'_i in Distillation Dataset (mAP/mAcc ± std). The last column (Avg) reports the average across the 7 columns. Overall best in bold.