arXiv:2603.11246v1 [cs.CV] 11 Mar 2026

When Slots Compete: Slot Merging in Object-Centric Learning

Christos Chatzisavvas[1,2], Panagiotis Rigas[3,4], George Ioannakis[2], Vassilis Katsouros[2], and Nikolaos Mitianoudis[1,2]

[1] Department of Electrical and Computer Engineering, Democritus University of Thrace, Greece
[2] Institute for Language and Speech Processing, Athena Research Center, Greece
[3] Department of Informatics & Telecommunications, National and Kapodistrian University of Athens, Greece
[4] Archimedes, Athena Research Center, Greece

Fig. 1: We introduce a merge operator over the slot set that adaptively refines factorization, producing coherent object-level representations.

Abstract. Slot-based object-centric learning represents an image as a set of latent slots with a decoder that combines them into an image or features. The decoder specifies how slots are combined into an output, but the slot set is typically fixed: the number of slots is chosen upfront and slots are only refined. This can lead to multiple slots *competing* for overlapping regions of the same entity rather than focusing on distinct regions. We introduce *slot merging*: a drop-in, lightweight operation on the slot set that *merges* overlapping slots during training. We quantify overlap with a *Soft-IoU* score between slot-attention maps and combine selected pairs via a barycentric update that preserves gradient flow. Merging follows a fixed policy, with the decision threshold inferred from overlap statistics, requiring no additional learnable modules. Integrated into the established feature-reconstruction pipeline of DINOSAUR, the proposed method improves object factorization and mask quality, surpassing other adaptive methods in object discovery and segmentation benchmarks.

Keywords: Object Centric Learning · Slot Attention · Adaptive Refinement · Scene Decomposition

1 Introduction

Human perception is fundamentally linked to the ability to decompose visual environments into discrete, meaningful entities. Object-Centric Learning (OCL) translates this cognitive principle into artificial systems by extracting structured, object-level representations from raw perceptual data. To bypass the supervision bottleneck inherent in manual annotation, recent advancements treat scene decomposition as an unsupervised reconstruction task. A prominent framework is Slot Attention [24], which serves as the architectural bottleneck enabling this decomposition. It operates on a set of K latent variables, referred to as *slots*, which are learnable vectors intended to encode individual objects or coherent scene components. Through an iterative, competitive attention mechanism, these slots selectively aggregate information from spatial feature embeddings to explain the distinct regions of a scene. The resulting representations can subsequently be decoded and leveraged in a range of object-centric downstream tasks, including reasoning [5], classification [30], and manipulation [4].

Slot attention has an innate limitation: it requires a fixed, apriori number of slots K, limiting applicability to real-world scenes. When K is larger than necessary, several slots may attend to the same object. We refer to this phenomenon as *slot competition*, where a single object becomes fragmented across multiple slots. While this is often viewed through the lens of variable object counts, we argue that it is fundamentally a representation issue: multiple slots compete to explain the same objects. Prior works attempt to address the fixed-K constraint by adapting slot cardinality. Adaslot [11] *prunes* redundant slots through a learned mechanism, while MetaSlot [23] *clusters* slots globally to infer a scene-dependent count. These approaches address *over-segmentation*, that is, the presence of more slots than necessary to explain the scene, through selection or suppression. In contrast, we interpret overlapping slots as fragmented representations and merge them into a single coherent slot. This perspective motivates a simple question: if slots compete over the same object, can their representations be merged rather than discarded?

Motivated by this formulation, we introduce a slot-merging algorithm. The method merges similar slots that compete to represent the same object. The merging itself is performed by a differentiable merge operator that combines slot representations through a mass-weighted convex interpolation, producing a single representation. Slot competition is quantified through the Intersection-over-Union (IoU) of slot attention maps. A fixed merge policy selects candidate slot pairs for merging based on their overlap structure, providing explicit control over when and how merging is applied. While the merge decision follows this fixed rule, the merge operator remains differentiable, allowing gradients to propagate through Slot Attention during training. The merging process is further regulated by an adaptive stopping rule derived from the empirical distribution of pairwise overlaps. The resulting mechanism acts as a refinement of Slot Attention. We evaluate the proposed method within the DINOSAUR [31] framework on four public benchmarks against representative slot-based baselines. The over-

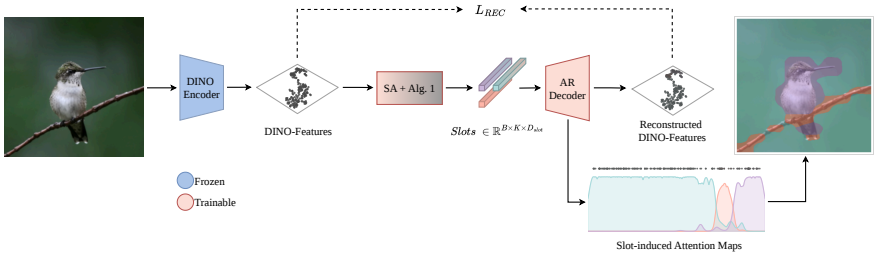

Fig. 2: Illustration of the proposed pipeline.

all pipeline is illustrated in Fig. 2, and segmentation results are shown in Fig. 1. To summarize, our contributions are:

- **Spatial slot competition score.** A probabilistic Intersection-over-Union (Soft-IoU) score over slot attention maps that quantifies spatial overlap and identifies slots competing to represent the same object. (Sec. 3.1)
- **Differentiable slot merge operator.** A differentiable merge operator that merges competing slots through a mass-weighted convex interpolation of their representations, producing a single consolidated slot during training. (Sec. 3.2)
- **Fixed merge policy.** A fixed merge policy that selects candidate slot pairs based on their overlap structure and regulates the merging process through a data-driven threshold derived from the distribution of pairwise overlaps. (Sec. 3.3)
- **Empirical evaluation.** Evaluation of the proposed method within the DINOSAUR framework on four public object-centric learning benchmarks, demonstrating consistent improvements over representative slot-based baselines. (Sec. 4)

2 Related Work

Unsupervised Object-Centric Learning. Object-Centric Learning aims to create representations for complex scenes that can encode each object separately [2]. Early approaches relied on sequential generative inference mechanisms. For instance, AIR [9] and SQAIR [21] discover objects by iteratively attending to image regions and representing each object through bounding box-based latent variables. Subsequent works shifted toward parallel mixture-based formulations. Autoencoder-based frameworks such as MONET [2], IODINE [13], and GENESIS [7] decompose a scene into a fixed set of latent object components and infer soft pixel-wise assignments using spatial mixture models trained with variational objectives, avoiding explicit box parametrization. Additional studies further explore auto-encoder based architectures for object discovery [19, 25, 26]. Building on this line of work, Slot Attention [24] replaces iterative variational refinement

with a single encoder forward pass, followed by iterative attention-based slot refinement, where slots compete to explain input features. This formulation enables efficient and permutation-equivariant learning of object-centric representations.

Slot-Based Architectures. Since its introduction, Slot Attention [24] has become a core building block for modern object-centric architectures. Subsequent work extends it along several directions, including reconstruction targets, decoder design, training strategies, and input modalities. In terms of reconstruction targets, DINOSAUR [31] replaces pixel-space reconstruction with pretrained feature reconstruction, significantly improving object discovery in complex real-world scenes. Regarding decoder design, SLATE [32] employs an autoregressive (AR) transformer decoder to model compositional structure, whereas Latent Slot Diffusion (LSD) [16] integrates a diffusion-based decoder for higher-fidelity generation. From a training perspective, SPOT [17] improves object binding through a teacher–student training framework with patch-order permutations, while other extensions introduce auxiliary training objectives for Slot Attention-based models [34]. Finally, Slot Attention has been extended to the video domain through methods such as SAVi [20], SOLV [1], and VideoSAUR [37], which introduce recurrent slot updates for consistent tracking, agglomerative clustering of slot embeddings across frames to mitigate over-clustering, and feature similarity prediction to enforce temporal object persistence.

Adaptive Slot Refinement. A shared limitation of most Slot Attention-based approaches is the use of a fixed number of slots K as a static hyperparameter. While sufficient for controlled benchmarks, real-world scenes exhibit variable object counts, leading to an inherent trade-off between over-segmentation and under-segmentation. Recent works address this limitation primarily through *slot selection or suppression*. Adaslot [11] initializes with an inflated set of slots and learns a lightweight scoring module to predict which slots are informative, discarding the rest during decoding via masked reconstruction. Similarly, Metaslot [23] introduces a vector-quantization codebook of object prototypes and removes duplicate slots by mapping them to shared discrete indices, effectively suppressing redundant slots before refinement. Beyond Slot Attention, GENESIS-V2 [8] employs a stick-breaking process to cluster pixel features, allowing the model to infer a variable number of object components. Despite their differences, these methods address slot redundancy primarily through slot selection or suppression. However, slot redundancy may also reflect *representation fragmentation*, where multiple slots compete to attend to the same object. Rather than discarding such slots, we *explicitly merge* fragmented representations of the same entity into a single coherent slot. To this end, we formulate a differentiable slot-merging mechanism that operates directly on the slot set. Spatially overlapping slots are detected via Attention-IoU and merged through barycentric aggregation. This perspective reframes variable-slot learning from discrete slot selection to continuous representation consolidation.

3 Method

Preliminaries. Slot Attention [24] maps a set of N input features $X \in \mathbb{R}^{N \times d}$ to K object-centric embeddings (slots) $S \in \mathbb{R}^{K \times d_{slots}}$ through an iterative cross-attention mechanism. Conceptually, Slot Attention can be interpreted as a learned variant of soft k-means clustering, where slots act as cluster centroids that iteratively group input features into object-level representations. Slots are initialized from a learnable Gaussian distribution $S^{(0)} \sim \mathcal{N}(\mu, \mathrm{diag}(\sigma^2))$, which allows each slot to serve as a flexible prototype that can bind to arbitrary objects. At each iteration $t = 0, 1, \ldots, T-1$, slots compete to explain the input features via attention. The attention weights between inputs and slots are computed as:

$$A = \mathrm{softmax}\left(\frac{K_x Q_s^\top}{\sqrt{d}}\right) \in \mathbb{R}^{N \times K}, \quad \text{where } Q_s = q(S^{(t)}),\ K_x = k(X),\ V_x = v(X).$$

Here, $q(\cdot)$, $k(\cdot)$, and $v(\cdot)$ are learned linear maps. The softmax is applied across the slot dimension, ensuring that attention is normalized across slots for each patch: $\sum_{k=1}^{K} A_{n,k} = 1,\ \forall n$. Such normalization induces competition between slots, encouraging specialization across distinct regions or objects. For notational simplicity, the batch dimension is omitted in the Slot Attention formulation above and introduced explicitly in Alg. 1.

Using these attention weights, each slot gathers a weighted summary of the input features:

$$\hat{S} = \left(\frac{A_{i,j}}{\sum_{l=1}^{N} A_{l,j}}\right)^\top V.$$

This slot-wise summary $\hat{S}$ is then combined with the previous slot state $S^{(t)}$ using a Gated Recurrent Unit (GRU), followed by layer normalization and a lightweight MLP to produce the refined slots:

$$S^{(t+1)} = \mathrm{MLP}\Big(\mathrm{LayerNorm}\left(\mathrm{GRU}(S^{(t)}, \hat{S})\right)\Big).$$

After T refinement iterations, the final slots $S^{(T)}$ serve as object-centric representations for downstream modules. Our proposed merging mechanism operates on these refined slots, further optimizing them during training.

3.1 Quantifying Spatial Slot Competition

Effective slot merging requires identifying slots that correspond to the same underlying object. While slot vectors encode compact semantic representations, the spatial structure of their attention maps provides a more direct signal of object identity. When two slots attend to spatial locations with strong overlap, they compete to explain the same underlying structure, indicating an inefficient allocation of representational capacity.

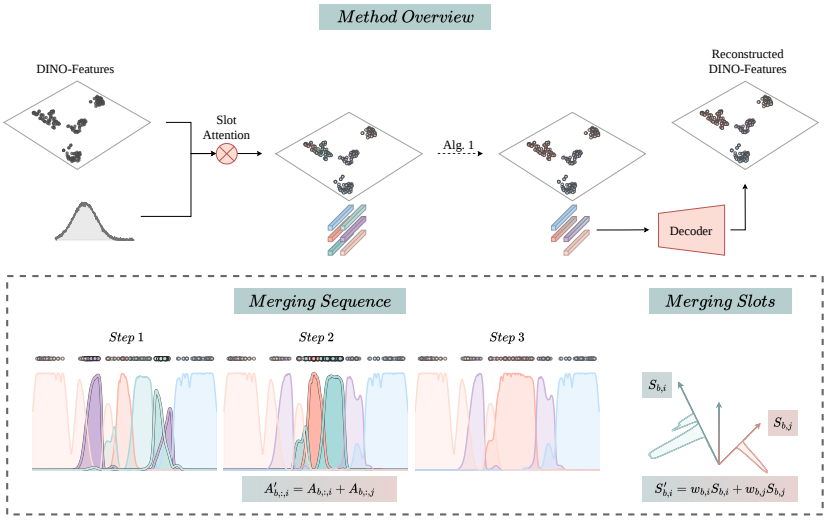

Fig. 3: Overview of the proposed training pipeline and merging mechanism. After Slot Attention produces refined slots and attention maps, pairwise Attention-IoU scores are computed to identify the most overlapping slot pair. The selected slots are merged via barycentric aggregation and their attention maps are summed. The procedure is repeated iteratively until no overlap exceeds the threshold, yielding a refined slot set passed to the decoder.

Let $A \in [0,1]^{N \times K}$ denote the normalized attention matrix produced by Slot Attention, where N is the number of spatial locations and K the number of slots. The attention maps associated with slots i and j are $p = A_{:,i}$ and $q = A_{:,j}$ with $p, q \in [0,1]^N$. Interpreting these vectors as soft spatial masks, we quantify their overlap using a probabilistic variant of the Intersection-over-Union (IoU) [29]. The resulting similarity measure is defined as

$$\mathrm{IoU}(p,q) = \frac{\sum_n p_n q_n}{\sum_n (p_n + q_n - p_n q_n)}. \tag{1}$$

Let $I_{ij} = \mathrm{IoU}(A_{:,i}, A_{:,j})$ denote the pairwise overlap score between slots i and j. The numerator represents the soft intersection, computed as the element-wise product of the two attention maps and summed over spatial locations. The denominator corresponds to the soft union, aggregating the spatial extent covered by each slot while subtracting the intersection term to avoid double counting. High IoU values therefore indicate substantial spatial agreement between slots, suggesting that they capture the same underlying entity. The resulting overlap scores are used to identify pairs of slots that compete for the same object and are therefore candidates for merging.

3.2 Slot Merge Operator

Given two competing slots i and j, identified through the overlap score defined in Sec. 3.1, we merge their representations using a mass-weighted operator. The operator combines both the slot vectors and their associated attention maps while preserving the total mass. Let $A \in [0,1]^{N \times K}$ denote the attention matrix and $S_i, S_j \in \mathbb{R}^D$ the corresponding slot representations. The attention mass of slot i is defined as the sum of its attention weights over all spatial locations,

$$\alpha_i = \sum_{n=1}^{N} A_{n,i}. \tag{2}$$

The relative contributions of the two slots are determined by normalized weights

$$w_i = \frac{\alpha_i}{\alpha_i + \alpha_j}, \qquad w_j = \frac{\alpha_j}{\alpha_i + \alpha_j}, \tag{3}$$

which satisfy $w_i + w_j = 1$. These coefficients ensure that each slot contributes proportionally to its spatial support. The merged slot representation is obtained through a barycentric update,

$$S_i \leftarrow w_i S_i + w_j S_j, \tag{4}$$

which replaces slot i with the convex combination of the two slot vectors.

The attention maps are updated accordingly by aggregation,

$$A_{:,i} \leftarrow A_{:,i} + A_{:,j}. \tag{5}$$

After the update, slot j is removed from the slot set. The resulting slot therefore preserves the total attention mass while aggregating the representations of the two competing slots. The merging operation is also illustrated in Fig. 3

3.3 Fixed Merge Policy

The merge operator defined in Sec. 3.2 specifies how two slots are merged. A policy determines which slot pairs are merged and when the process terminates. Given the pairwise overlap scores defined in Sec. 3.1, candidate pairs are selected according to the maximum overlap criterion

$$(i^*, j^*) = \arg\max_{i \neq j} I_{ij}. \tag{6}$$

This rule selects the pair of slots that exhibits the strongest spatial competition. Starting from the current slot set, the merge operator is applied to the selected pair (i^*, j^*). After each merge, the attention matrix and slot representations are updated and the pairwise overlap scores are recomputed. The procedure is repeated while the maximum overlap among all slot pairs exceeds a predefined threshold. Since merges are performed sequentially, the procedure is not transitive: if slot j overlaps with both i and k, merging j into i does not necessarily imply that the resulting slot will also merge with k.

Let $I_{\max} = \max_{i \neq j} I_{ij}$ denote the maximum pairwise overlap in the current slot set. Merging terminates when $I_{\max} < \tau$, where τ is a fixed threshold controlling the merging procedure. The threshold is estimated from the empirical distribution of pairwise overlap scores observed during training. Specifically, overlap values are aggregated over batches to form a histogram, and the triangle thresholding method [36] is applied to determine a separation point between low-overlap and high-overlap slot pairs.

Training Schedule. To avoid interference with early slot formation, merging is activated only after the slot representations have stabilized (Fig. 4). In practice, the merge policy is enabled after an initial training phase and subsequently applied during the remaining optimization steps. Alg. 1 summarizes the resulting merge procedure. The algorithm maintains a binary mask M indicating active slots for each batch element, while $\mathcal{K}$ denotes the set of currently active slot indices.

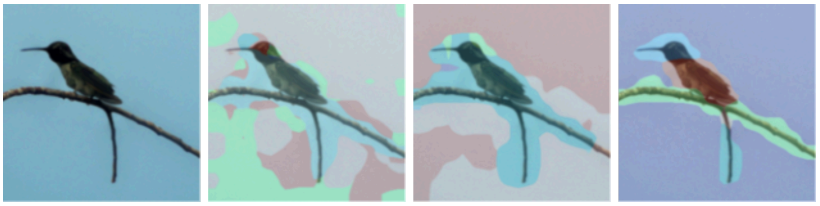

Fig. 4: Evolution of slot representations during training (left to right). Slot assignments transition from fragmented and unstable to progressively specialized and object-aligned. After stabilization (right), slot merging is enabled to resolve overlap between slots.

3.4 Computational Complexity

Slot Attention has time complexity $\mathcal{O}(TDNK)$ [24]. Let N denote the number of input patches and K the initial number of slots, and let $m \leq K-1$ denote the number of merges performed. A naive implementation recomputes all pairwise overlaps after each merge. At a stage with r active slots there are $\mathcal{O}(r^2)$ slot pairs and each overlap requires a linear pass over N patches, giving a cost $\mathcal{O}(Nr^2)$ per step. Summing across merge steps yields $\sum_{r=K-m+1}^{K} \mathcal{O}(Nr^2)$, which in the worst case $m = K-1$ becomes $\mathcal{O}(NK^3)$. In practice we implement merging with incremental updates. All pairwise soft intersections $\sum_n A_{n,i} A_{n,j}$ and slot masses $\sum_n A_{n,i}$ are precomputed once in $\mathcal{O}(NK^2)$. After contracting slot j into i via Eq. 5, only overlaps incident to slot i must be updated, which requires $\mathcal{O}(NK)$ work per merge. This yields an overall complexity of $\mathcal{O}(NK^2)$ per image. Consequently, the full pipeline scales as $\mathcal{O}(TDNK + NK^2)$ in practice. Since the number of slots is typically small, the additional merging cost remains minor compared to the Slot Attention updates.

Algorithm 1 Slot Merging

Require: slots $S^{(T)} \in \mathbb{R}^{B \times K \times D_{\text{slots}}}$, attention maps $A \in \mathbb{R}^{B \times N \times K}$, threshold τ
$S \leftarrow S^{(T)}$; $M \leftarrow \mathbf{1}^{B \times K}$
for $b = 1, \ldots, B$ **in parallel do**
  $\mathcal{K} \leftarrow \{i \in [K] : M_{b,i} = 1\}$ ▷ $\mathcal{K}$ denotes active slots
  while $|\mathcal{K}| > 1$ **do**
    compute $\{I_{ij}^{(b)}\}_{i<j,\, i,j \in \mathcal{K}}$ via Eq. 1
    $(i^*, j^*) \leftarrow \arg\max_{i<j,\, i,j \in \mathcal{K}} I_{ij}^{(b)}$
    if $I_{i^*,j^*}^{(b)} \leq \tau$ **then**
      break
    end if
    compute masses $\alpha_{b,i^*}, \alpha_{b,j^*}$ via Eq. 2
    compute weights w_{b,i^*}, w_{b,j^*} from $\alpha_{b,i^*}, \alpha_{b,j^*}$ via Eq. 3
    $S_{b,i^*} \leftarrow w_{b,i^*} S_{b,i^*} + w_{b,j^*} S_{b,j^*}$ via Eq. 4
    $A_{b,:,i^*} \leftarrow A_{b,:,i^*} + A_{b,:,j^*}$ via Eq. 5
    $M_{b,j^*} \leftarrow 0$
    $\mathcal{K} \leftarrow \mathcal{K} \setminus \{j^*\}$
  end while
end for
return S, M

4 Experiments

4.1 Setup

Datasets. Our method is evaluated on four widely used public benchmarks covering both real-world and synthetic data. For real-world object discovery, we use PASCAL VOC 2012 [10] and MS COCO 2017 [22]. PASCAL VOC provides diverse object categories in natural environments with moderate clutter. In contrast, MS COCO introduces substantially higher scene complexity, characterized by dense object layouts, heavy occlusion, and cluttered backgrounds. For synthetic evaluation, we adopt MOVi-C and MOVi-E [12], generated with approximately 1,000 scanned 3D assets. Although both are video datasets, we follow the protocol of [31] and sample frames independently, treating them as static images to ensure comparability with prior object-centric learning work. These datasets enable controlled evaluation across different object-density ranges. MOVi-C contains relatively sparse scenes (3–10 objects/regions per image), while MOVi-E presents significantly denser compositions (11–23 objects/regions per image).

Evaluation Metrics. Following standard practice in object-centric learning, segmentation quality is evaluated using Mean Best Overlap (mBO) and Mean Intersection over Union (mIoU). Mean Best Overlap [28] measures how well each ground-truth mask is recovered by selecting the predicted mask with the highest overlap. Both instance-level (mBO^i) and class-level (mBO^c) variants are reported. mIoU computes one-to-one assignments between predicted and ground-truth masks using the Hungarian matching algorithm, and averages the IoU over the matched pairs. Consistent with recent literature [17,34], Foreground Adjusted Rand Index (FG-ARI) is not reported, as it is known to favor under-segmentation and can therefore yield misleading conclusions in dense scenes.

Implementation Details. The proposed slot merging mechanism is integrated into the DINOSAUR architecture [31] using a Vision Transformer B/16 [6] en-

Fig. 5: Qualitative results on COCO and PASCAL VOC.

coder initialized with DINO [3]. The model is optimized with Adam using a batch size of 64. Training is conducted for 100 epochs on VOC and COCO, and for 65 epochs on MOVi-C and MOVi-E. The model is initialized with 7, 8, 12, and 25 slots for VOC, COCO, MOVi-C, and MOVi-E, respectively. Slot merging is enabled from epoch 70 for VOC and COCO, and from epoch 30 for MOVi-C and MOVi-E. Regarding the merging threshold, we use dataset-specific values of $\tau = 0.036, 0.04, 0.035$, and 0.019 for VOC, COCO, MOVi-C, and MOVi-E, respectively, derived from overlap statistics. The threshold remains fixed during training. Additional experimental details are provided in the supplementary material. All models are trained on a single NVIDIA RTX A6000 GPU with 48Gbytes.

4.2 Analysis

In this section, we analyze the proposed method, focusing on slot merging dynamics during training, dataset-level merging statistics, and the impact of key implementation choices. The primary experiments are conducted on the VOC dataset, whose moderate object complexity makes it particularly suitable for assessing slot merging within the Slot Attention framework. Additional evaluations are conducted on MOVi-C.

Effect of Merging During Training. Tab. 1 compares three configurations: the vanilla baseline (DINOSAUR), a variant with merging applied only at inference, and the full model with merging incorporated during training. Applying merging solely at inference already improves performance over the baseline. However, the gain remains limited because the slot representations were optimized under a fixed-slot constraint. In contrast, incorporating merging into

Table 1: Comparison between inference-only and training-integrated slot merging on VOC and MOVi-C, evaluated with mBO (instance and class). Results are averaged over 3 seeds.

	VOC			MOVi-C	
	mBO^i	mBO^c	mIoU	mBO^i	mIoU
Baseline	$44.0_{\pm 1.9}$	$51.2_{\pm 1.9}$	–	42.4	–
Merging during Inference	$46.01_{\pm 0.7}$	$52.68_{\pm 0.65}$	$44.51_{\pm 0.55}$	$44.53_{\pm 1.7}$	$43.81_{\pm 1.6}$
Merging during Training	$\mathbf{47.83}_{\pm 0.17}$	$\mathbf{55.34}_{\pm 0.27}$	$\mathbf{46.01}_{\pm 0.18}$	$\mathbf{48.35}_{\pm 1.32}$	$\mathbf{47.43}_{\pm 1.26}$

Table 2: Ablation on gradient detachment in the barycentric slot merging operator during training on VOC and MOVi-C. Results are the mean and standard deviation of the decoder's metrics over 3 seeds.

	VOC			MOVi-C	
	mBO^i	mBO^c	mIoU	mBO^i	mIoU
With Detach	$45.75_{\pm 0.28}$	$53.13_{\pm 0.42}$	$43.97_{\pm 0.22}$	$45.39_{\pm 2.95}$	$44.46_{\pm 2.9}$
Without Detach	$\mathbf{47.83}_{\pm 0.17}$	$\mathbf{55.34}_{\pm 0.27}$	$\mathbf{46.01}_{\pm 0.18}$	$\mathbf{48.35}_{\pm 1.32}$	$\mathbf{47.43}_{\pm 1.26}$

the training process yields consistent gains across all metrics and datasets. On VOC, our method reaches 55.34 mBO^c, clearly outperforming both the baseline and the inference-only variant. Similar improvements are observed on MOVi-C. These results demonstrate that the benefit of merging extends beyond a post-hoc adjustment, since integrating it during training encourages more coherent representations and leads to improved segmentation quality. Fig. 5 visualizes the object-segmentation results of our method in comparison with DINOSAUR.

Effect of Differentiable Merging. To verify that the observed gains arise from gradient propagation rather than from the merging rule itself, we conduct an ablation in which gradients are detached during the merging step (Tab. 2). In this variant, merged slots weights are removed from the computational graph, preventing gradients from flowing back to the respective slots. Detaching gradients consistently degrades performance, which indicates that the improvement cannot be attributed to the forward merging operation alone. Allowing gradients to propagate through the merging layer ensures that slot representations are optimized jointly with the merging operation during training.

Effect of Attention Map Update. We also examine whether updating the attention maps after merging contributes to performance. Specifically, attention map updates using Eq. 5 are compared against retaining the original attention map of a single slot after merging (Tab. 3). Removing attention addition leads to consistent degradation, suggesting that attention map fusion supports effective slot merging.

Merging Statistics. To further analyze the behavior of the merging mechanism, Tab. 4 reports aggregated merging statistics. The merge frequency per image remains consistent between training and inference, indicating stable operation. On COCO, which contains dense scenes with approximately 7.7 objects

Table 3: Effect of attention map aggregation during slot merging. Results are averaged over 3 seeds.

	mBO^i	mBO^c	mIoU
Without Aggregation	$47.38_{\pm 0.52}$	$54.83_{\pm 0.58}$	$45.53_{\pm 0.49}$
With Aggregation	$\mathbf{47.83}_{\pm 0.17}$	$\mathbf{55.34}_{\pm 0.27}$	$\mathbf{46.01}_{\pm 0.18}$

per image, the mean number of merges is lowest (0.69) despite the initialization $K = 8$, suggesting that most slots are already specialized with limited overlap. In contrast, MOVi-E starts from $K = 25$ slots, exceeding the typical object count, and exhibits substantially more merging (7.96 on average). MOVi-C shows a similar pattern: with $K = 12$, the model performs about 2.2 merges per image. Overall, the number of merges correlates with the degree of slot overlap induced by the initial slot allocation and scene complexity.

Table 4: Merging statistics during training and inference across all four datasets, reporting the mean and standard deviation of the number of merged slots per image.

	Training / Inference			
	VOC	COCO	MOVi-C	MOVi-E
Mean	1.67 / 1.73	0.69 / 0.82	2.2 / 2.2	7.96 / 7.91
Std	1.0 / 0.98	0.76 / 0.81	1.32 / 1.32	2.13 / 2.08

Effect of Initialization Epoch. Tab. 5 analyzes the effect of the epoch at which merging is introduced during training. On VOC, earlier activation yields the strongest performance, while delaying merging leads to consistent degradation. A similar trend is observed on MOVi-C, where early activation outperforms later introduction. The performance gap confirms that sufficient optimization time after enabling merging is critical. Earlier activation allows the model to adapt slot representations to the merging dynamics, whereas late activation restricts this adaptation window and limits the attainable gains. Additional ablation studies and results are provided in the supplementary material.

Table 5: Influence of the training epoch at which slot merging is activated, evaluated on VOC and MOVi-C.

	VOC				MOVi-C	
Epoch	mBO^i	mBO^c	mIoU	Epoch	mBO^i	mIoU
90	46.15	52.91	44.65	50	47.05	46.19
80	46.96	54.10	45.25	40	46.83	46.00
70	**47.83**	**55.33**	**46.01**	30	**48.35**	**47.43**

4.3 Comparison with Object-Centric Methods

Tabs. 6 and 7 summarize comparisons with prior object-centric learning approaches. Our method consistently improves over its backbone baseline, DINOSAUR, across all four datasets. On VOC, it achieves gains of more than 4 and 3.8 points in mBO^c and mBO^i, respectively, compared to DINOSAUR. On the COCO dataset, characterized by dense and cluttered scenes, our method surpasses strong extensions such as DINOSAUR+FS+RC [34]. Improvements are also observed on the synthetic benchmarks MOVi-C and MOVi-E. Overall, the empirical findings show that the proposed merging mechanism consistently enhances object representations across both real-world and synthetic environments.

Table 6: Comparison with object-centric methods on the VOC and COCO datasets. Results for our approach are reported as mean ± standard deviation over 3 seeds, using a ViT-B/16 backbone and 224 × 224 input resolution. † indicates results taken from [17]. Otherwise results are from the original papers.

	VOC		COCO	
Method	mBO^i	mBO^c	mBO^i	mBO^c
SA [24]† NeurIPS'20	24.6	24.9	17.2	19.2
SLATE [32]† ICLR'22	35.9	41.5	29.1	33.6
CAE [26]† TMLR	$32.9_{\pm 0.9}$	$37.4_{\pm 1.0}$	–	–
Rotating Features [25]† NeurIPS'23	$40.7_{\pm 0.1}$	$46.0_{\pm 0.1}$	–	–
DINOSAUR-MLP [31]† ICLR'23	$39.5_{\pm 0.1}$	$40.9_{\pm 0.1}$	$27.7_{\pm 0.2}$	$30.9_{\pm 0.2}$
DINOSAUR [31]† ICLR'23	$44.0_{\pm 1.9}$	$51.2_{\pm 1.9}$	$32.3_{\pm 0.4}$	$38.8_{\pm 0.4}$
DINOSAUR+FS+RC [34] CVPR'25	<u>45.2</u>	<u>52.9</u>	31.2	<u>40.1</u>
$DIAS_{image}$ [38] ACM MM'25	$44.8_{\pm 9.8}$	–	<u>$32.8_{\pm 0.1}$</u>	–
AdaSlot [11] CVPR'24	–	–	27.4	–
MetaSlot-MLP [23] NeurIPS'25	$43.9_{\pm 0.3}$	–	$29.5_{\pm 0.2}$	–
MetaSlot-Tfd [23] NeurIPS'25	$39.0_{\pm 0.3}$	–	$28.2_{\pm 0.7}$	–
Ours	$\mathbf{47.83}_{\pm 0.17}$	$\mathbf{55.34}_{\pm 0.27}$	$\mathbf{32.88}_{\pm 0.40}$	$\mathbf{42.78}_{\pm 0.22}$

We further compare our method with recent approaches for variable slot cardinality, namely Adaslot [11] and Metaslot [23]. Unlike these methods, which rely on discrete pruning or clustering to adjust the number of active slots, our approach models slot interaction as a differentiable optimization process. Instead of discarding slots, overlapping representations are fused through barycentric aggregation while preserving gradient flow. Our method surpasses both approaches across all compared datasets, indicating that differentiable aggregation offers a more effective alternative to hard slot selection strategies.

4.4 Limitations

The proposed merging mechanism is intentionally simple and relies on pairwise overlap statistics to regulate slot consolidation. While this design keeps the pro-

Table 7: Comparison with object-centric methods on MOVi-C and MOVi-E. Results for our approach are reported as mean ± standard deviation over 3 seeds. † indicates results taken from [17]. Otherwise results are from the original papers.

	MOVi-C		MOVi-E	
Method	mBO^i	mIoU	mBO^i	mIoU
SA [24]† NeurIPS'20	$26.2_{\pm 1.0}$	–	$24.0_{\pm 1.2}$	–
SLASH [19]† CVPR'23	–	$27.7_{\pm 5.9}$	–	–
SLATE [32]† ICLR'22	$39.4_{\pm 0.8}$	$37.8_{\pm 0.7}$	$30.2_{\pm 1.7}$	$\underline{28.6}_{\pm 1.7}$
DINOSAUR-MLP [31]† ICLR'23	$39.1_{\pm 0.2}$	–	$\underline{35.5}_{\pm 0.2}$	–
DINOSAUR [31]† ICLR'23	42.4	–	–	–
DINOSAUR+FS+RC [34] CVPR'25	<u>46.0</u>	<u>44.8</u>	–	–
AdaSlot [11] CVPR'24	35.6	–	29.8	–
MetaSlot-MLP [23] NeurIPS'25	35.0	–	–	–
Ours	$\mathbf{48.35}_{\pm 1.32}$	$\mathbf{47.43}_{\pm 1.26}$	$\mathbf{39.56}_{\pm 0.55}$	$\mathbf{37.75}_{\pm 0.55}$

cedure efficient, it does not explicitly model higher-order relations among multiple slots. Extending the policy to incorporate group-level similarity or set-based distance measures could provide additional context when resolving complex interactions between objects. More broadly, the current stopping rule is derived from the empirical distribution of overlap scores and remains fixed during training. Developing more interpretable or adaptive stopping criteria that account for scene composition or global structure represents an interesting direction for future work.

5 Conclusion

Slot Attention specifies how slots compose an image, but it does not define operations on the slot set itself. As a result, representational fragmentation may happen, where multiple slots bind to the same entity and spend capacity resolving mutual competition. To address this, we introduce a merge operator that acts directly on the slot set. The operator identifies overlapping slots and replaces them with a single aggregated representation, converting competitive redundancy into cooperative refinement at the representation level. The proposed mechanism is simple, local, and leaves the underlying Slot Attention dynamics unchanged. Empirically, merging stabilizes slot factorization and improves object-level grouping across datasets. More broadly, our results highlight the importance of explicit structural operations on the slot set, complementing existing efforts that focus on improving feature binding or architectural design.

Acknowledgments

This work has been partially supported by project MIS 5154714 of the National Recovery and Resilience Plan Greece 2.0 funded by the European Union un-

der the NextGenerationEU Program, and by the European Union's Horizon Europe research and innovation programme under Grant Agreement No. 101158328 (TEXTaiLES).

References

1. Aydemir, G., Xie, W., Guney, F.: Self-supervised object-centric learning for videos. Advances in Neural Information Processing Systems **36**, 32879–32899 (2023)
2. Burgess, C.P., Matthey, L., Watters, N., Kabra, R., Higgins, I., Botvinick, M., Lerchner, A.: Monet: Unsupervised scene decomposition and representation. arXiv preprint arXiv:1901.11390 (2019)
3. Caron, M., Touvron, H., Misra, I., Jégou, H., Mairal, J., Bojanowski, P., Joulin, A.: Emerging properties in self-supervised vision transformers. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 9650–9660 (2021)
4. Chapin, A., Dellandrea, E., Chen, L.: Is an object-centric representation beneficial for robotic manipulation? In: International Conference on Robotics, Computer Vision and Intelligent Systems. pp. 510–523. Springer (2025)
5. Chi, D., Kim, H., Oh, Y., Kim, Y., Lee, D., Jo, D., Kim, J., Baek, J., Ahn, S., Kim, S.: Slot-mllm: Object-centric visual tokenization for multimodal llm. arXiv preprint arXiv:2505.17726 (2025)
6. Dosovitskiy, A.: An image is worth 16x16 words: Transformers for image recognition at scale. arXiv preprint arXiv:2010.11929 (2020)
7. Engelcke, M., Kosiorek, A.R., Parker Jones, O., Posner, I.: GENESIS: Generative Scene Inference and Sampling with Object-Centric Latent Representations. In: International Conference on Learning Representations (ICLR) (2020)
8. Engelcke, M., Parker Jones, O., Posner, I.: Genesis-v2: Inferring unordered object representations without iterative refinement. Advances in Neural Information Processing Systems **34**, 8085–8094 (2021)
9. Eslami, S., Heess, N., Weber, T., Tassa, Y., Szepesvari, D., Hinton, G.E., et al.: Attend, infer, repeat: Fast scene understanding with generative models. Advances in neural information processing systems **29** (2016)
10. Everingham, M., Van Gool, L., Williams, C.K., Winn, J., Zisserman, A.: The pascal visual object classes challenge 2012 (voc2012). In: Results (2012)
11. Fan, K., Bai, Z., Xiao, T., He, T., Horn, M., Fu, Y., Locatello, F., Zhang, Z.: Adaptive slot attention: Object discovery with dynamic slot number. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 23062–23071 (2024)
12. Greff, K., Belletti, F., Beyer, L., Doersch, C., Du, Y., Duckworth, D., Fleet, D.J., Gnanapragasam, D., Golemo, F., Herrmann, C., et al.: Kubric: A scalable dataset generator. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 3749–3761 (2022)
13. Greff, K., Kaufman, R.L., Kabra, R., Watters, N., Burgess, C., Zoran, D., Matthey, L., Botvinick, M., Lerchner, A.: Multi-object representation learning with iterative variational inference. In: International conference on machine learning. pp. 2424–2433. PMLR (2019)
14. Hariharan, B., Arbeláez, P., Bourdev, L., Maji, S., Malik, J.: Semantic contours from inverse detectors. In: 2011 international conference on computer vision. pp. 991–998. IEEE (2011)

15. He, K., Chen, X., Xie, S., Li, Y., Dollár, P., Girshick, R.: Masked autoencoders are scalable vision learners. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 16000–16009 (2022)
16. Jiang, J., Deng, F., Singh, G., Ahn, S.: Object-centric slot diffusion. In: Advances in Neural Information Processing Systems. vol. 36, pp. 8563–8601 (2023), https://proceedings.neurips.cc/paper_files/paper/2023/file/1b3ceb8a495a63ced4a48f8429ccdcd8-Paper-Conference.pdf
17. Kakogeorgiou, I., Gidaris, S., Karantzalos, K., Komodakis, N.: Spot: Self-training with patch-order permutation for object-centric learning with autoregressive transformers. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 22776–22786 (June 2024)
18. Kim, D., Kim, S., Kwak, S.: Bootstrapping top-down information for self-modulating slot attention. Advances in Neural Information Processing Systems **37**, 103751–103773 (2024)
19. Kim, J., Choi, J., Choi, H.J., Kim, S.J.: Shepherding slots to objects: Towards stable and robust object-centric learning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 19198–19207 (2023)
20. Kipf, T., Elsayed, G.F., Mahendran, A., Stone, A., Sabour, S., Heigold, G., Jonschkowski, R., Dosovitskiy, A., Greff, K.: Conditional Object-Centric Learning from Video. In: International Conference on Learning Representations (ICLR) (2022)
21. Kosiorek, A., Kim, H., Teh, Y.W., Posner, I.: Sequential attend, infer, repeat: Generative modelling of moving objects. Advances in Neural Information Processing Systems **31** (2018)
22. Lin, T.Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft coco: Common objects in context. In: European conference on computer vision. pp. 740–755. Springer (2014)
23. Liu, H., Zhao, R., Chen, H., Pajarinen, J.: Metaslot: Break through the fixed number of slots in object-centric learning. In: The Thirty-ninth Annual Conference on Neural Information Processing Systems (2025)
24. Locatello, F., Weissenborn, D., Unterthiner, T., Mahendran, A., Heigold, G., Uszkoreit, J., Dosovitskiy, A., Kipf, T.: Object-centric learning with slot attention. Advances in neural information processing systems **33**, 11525–11538 (2020)
25. Löwe, S., Lippe, P., Locatello, F., Welling, M.: Rotating features for object discovery. Advances in Neural Information Processing Systems **36**, 59606–59635 (2023)
26. Löwe, S., Lippe, P., Rudolph, M., Welling, M.: Complex-valued autoencoders for object discovery. Transactions on Machine Learning Research (2022)
27. 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. arXiv preprint arXiv:2304.07193 (2023)
28. Pont-Tuset, J., Arbelaez, P., Barron, J.T., Marques, F., Malik, J.: Multiscale combinatorial grouping for image segmentation and object proposal generation. IEEE transactions on pattern analysis and machine intelligence **39**(1), 128–140 (2016)
29. Rahman, M.A., Wang, Y.: Optimizing intersection-over-union in deep neural networks for image segmentation. In: International symposium on visual computing. pp. 234–244. Springer (2016)
30. Rubinstein, A., Prabhu, A., Bethge, M., Oh, S.J.: Are we done with object-centric learning? arXiv preprint arXiv:2504.07092 (2025)
31. Seitzer, M., Horn, M., Zadaianchuk, A., Zietlow, D., Xiao, T., Simon-Gabriel, C., He, T., Zhang, Z., Schölkopf, B., Brox, T., et al.: Bridging the gap to real-world object-centric learning. In: International Conference on Learning Representations (ICLR) (2023)

32. Singh, G., Deng, F., Ahn, S.: Illiterate dall-e learns to compose. In: International Conference on Learning Representations (2022)
33. Singh, K., Schaub-Meyer, S., Roth, S.: Glass: Guided latent slot diffusion for object-centric learning. In: Proceedings of the Computer Vision and Pattern Recognition Conference. pp. 28673–28683 (2025)
34. Tian, P., Yang, S., Yu, H., Kot, A.: Pay attention to the foreground in object-centric learning. In: Proceedings of the Computer Vision and Pattern Recognition Conference. pp. 30281–30290 (2025)
35. Wu, Z., Hu, J., Lu, W., Gilitschenski, I., Garg, A.: Slotdiffusion: Object-centric generative modeling with diffusion models. Advances in Neural Information Processing Systems **36**, 50932–50958 (2023)
36. Zack, G.W., Rogers, W.E., Latt, S.A.: Automatic measurement of sister chromatid exchange frequency. Journal of Histochemistry & Cytochemistry **25**(7), 741–753 (1977)
37. Zadaianchuk, A., Seitzer, M., Martius, G.: Object-centric learning for real-world videos by predicting temporal feature similarities. Advances in Neural Information Processing Systems **36**, 61514–61545 (2023)
38. Zhao, R., Zhao, Y., Kannala, J., Pajarinen, J.: Slot attention with re-initialization and self-distillation. In: Proceedings of the 33rd ACM International Conference on Multimedia. pp. 4185–4192 (2025)

A Supplementary Material

A.1 Additional Qualitative Results

Fig. 6 demonstrates the merging steps for specific images starting from the representation that Slot Attention produces until the final step where the corresponding slots are passed to the decoder for further refinement. Each step merges the slots with the highest Soft-IoU score.

Fig. 6: Visualizations of the iterative slot merging process. Starting from slot representations produced by Slot Attention, the algorithm repeatedly merges the pair of slots with the highest Soft-IoU overlap. The final merged slot set is then passed to the decoder for reconstruction.

A.2 Implementation Details

Training Configuration. We train all models using the Adam optimizer with $\beta_1 = 0.9$, $\beta_2 = 0.99$, no weight decay, and a batch size of 64. The learning

rate follows a linear warm-up from 0 to its peak value during the first 10,000 training iterations and then decreases through cosine annealing. For experiments on VOC [10] and COCO [22] using DINO [3], the peak learning rate is 4×10^{-4} and the minimum value is 4×10^{-7}. For MOVi-C and MOVi-E [12], the peak and minimum values are 2×10^{-4} and 4×10^{-5}, respectively. When using DINOv2 [27] or MAE [15] on the VOC dataset, we use the same schedule with a peak learning rate of 2×10^{-4} and the minimum learning rate of 4×10^{-5}. Initialization happens through a shared gaussian and gradients are clipped at 0.3 infinity norm value.

Training is conducted for 100 epochs on VOC and COCO, and for 65 epochs on MOVi-C and MOVi-E. As the encoder, we use a ViT-B/16 [6] initialized with DINO. The AR transformer decoder consists of 4 transformer blocks with 6 attention heads each one. In all experiments, we use 3 iterations of the Slot Attention module [24] with a slot dimension of 256. Merging operation is applied to slots produced after the third Slot Attention iteration. The model is initialized with 7, 8, 12, and 25 slots for VOC, COCO, MOVi-C, and MOVi-E, respectively.

Data-driven Threshold Selection. To regulate the iterative merging process, an adaptive stopping criterion is introduced. Merging is activated only once the training dynamics stabilize, ensuring that slot representations are sufficiently discriminative. Merging threshold τ is determined automatically by analyzing the distribution of pairwise Soft IoU scores over 11 batches, one epoch before merging starts. For each batch b, the IoU values are discretized into 100 bins to form a histogram H_b. To mitigate the influence of extreme counts and skewed distributions, a logarithmic transformation is applied.

Following [36], a candidate threshold is extracted from histogram geometry. Specifically, a line is drawn between the histogram peak and farthest nonzero bin on the longer tail. The bin with the maximum perpendicular distance $d_\perp$ to this line separates high-overlap and low-overlap regions, and is selected as the batch-wise threshold candidate. To ensure robustness against batch-level variability, the candidate thresholds are aggregated using first-order statistics, such as mean or variance-adjusted mean. For VOC and MOVi-E, we select mean minus standard deviation among candidate thresholds, and mean value as threshold for COCO and MOVi-C. Selection is based on dataset statistics, leaving space for merging on datasets that can support it. Finally, the thresholds used for the experiments are: 0.036, 0.04, 0.035, and 0.019 for VOC, COCO, MOVi-C, and MOVi-E, respectively. Thresholding process is also depicted in Fig. 7a.

A.3 Additional Experimental Results

Comparison with Prior Object-Centric Methods. In Tab. 10, we report extended benchmark results including both mBO and mIoU across four datasets. On MOVi-C, our approach achieves the second-best performance among all compared methods. Furthermore, it surpasses adaptive slot-number approaches such as Adaslot [11] and Metaslot [23], indicating that the proposed merging strategy improves object discovery while preserving accurate segmentation. At the same

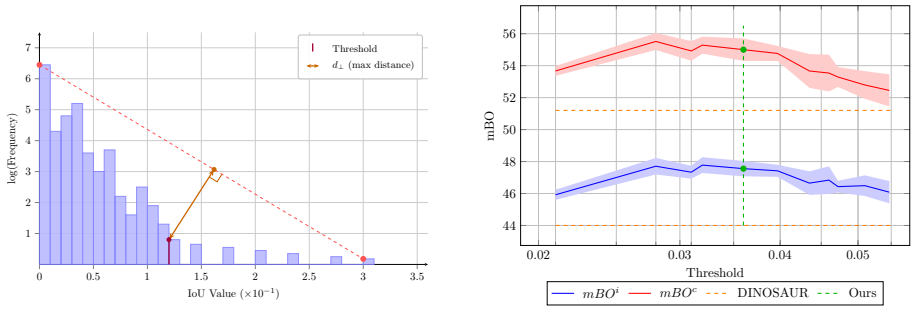

(a) **(b)**

Fig. 7: (a) Data-driven estimation of threshold from the histogram of pairwise Soft-IoU scores using the maximum perpendicular distance criterion. (b) Sensitivity analysis of the merging threshold on the VOC dataset measured using mBO^i and mBO^c.

Table 8: Different encoders comparison. DINO [3] and MAE [15] use ViT-B/16 architecture [6] while DINOv2 [27] employs ViT-B/14.

Encoder	Method	mBO^i	mBO^c	mIoU
DINO	DINOSAUR	44.0	51.2	–
	Ours	47.8	55.3	46.0
DINOv2	DINOSAUR	43.2	47.3	42.1
	Ours	44.4	50.5	43.2
MAE	DINOSAUR	37.8	42.7	36.5
	Ours	38.3	43.4	36.7

time, it remains competitive with more complex architectures such as SPOT [17], a teacher-student architecture, despite relying on a significantly simpler training procedure. Diffusion-based methods [16,33,35] achieve the highest absolute performance but involve more complex architectures and training pipelines, making direct metric comparison less intuitive.

Threshold Analysis. Fig. 7b evaluates the effect of merging threshold on VOC dataset. Performance remains stable across a wide range of values, showing robustness. Both mBO^i and mBO^c exhibit only minor variations as threshold changes. Across all evaluated thresholds, the proposed method consistently outperforms DINOSAUR baseline.

Comparison with other Pre-trained Feature Extractors. In our experiments, we use DINO featuresas training signal. We also evaluate DINOv2 [27] and MAE [15] as alternative feature extractors on VOC. As shown in Tab. 8, our method improves over DINOSAUR baseline across all encoders, with the largest gains observed when using DINO features.

Table 9: Training time comparison between DINOSAUR and our method.

Method	VOC	MOViC
DINOSAUR	$\sim$ 53 minutes	$\sim$ 4h 44 minutes
Ours	$\sim$ 53.5 minutes	$\sim$ 4h 50 minutes

Training Time. We measure training time for VOC and MOVi-C datasets. Training time is defined as the time spent within the training loop, including forward pass, backward pass, gradient updates, and data loading. Tab. 9 reports total training time on the two datasets. As shown, on VOC, the difference between DINOSAUR and our method is negligible, indicating that the proposed modification introduces only a minimal practical computational overhead on datasets with a relatively small number of slots K. On MOVi-C, the runtime increases slightly, but the overall overhead remains small.

A.4 Datasets and Evaluation

PASCAL VOC. We train on PASCAL VOC 2012 trainaug set (10,582 images), which combines 1,464 images from the segmentation train set and 9,118 images from the SBD dataset [14]. This split is used following prior work [17,31]. During training, the image minor axis is resized to 224 pixels, followed by random cropping to 224×224 and random horizontal flipping with probability 0.5. Evaluation is performed on the official segmentation validation set (1,449 images), ignoring unlabeled pixels. We resize the minor axis to 320 pixels, apply center cropping, and evaluate the masks at 320×320 resolution.

COCO. We train on the COCO 2017 dataset (118,287 images) and evaluate or method on the validation set (5,000 images). During training, the image minor axis is resized to 224×224, followed by a 224×224 center crop and random horizontal flipping with probability 0.5. For evaluation, the minor axis is resized to 320 pixels, center cropped, and masks are evaluated at a resolution of 320×320 as described in [17, 31].

MOVi-C/E. We convert the MOVi-C/E video datasets into image datasets by sampling nine random frames from each training clip, resulting in 87,633 images for MOVi-C and 87,741 for MOVi-E. During training, images are resized to 224×224. For evaluation, we use all frames from the 250 validation clips (6,000 images), following prior work [17, 31], and evaluate the masks at the full 128×128 resolution.

Table 10: Comparison with object-centric methods on VOC, COCO, MOVi-C, and MOVi-E datasets. Results for our approach are reported as mean ± standard deviation over 3 seeds. † indicates results taken from [17]. Otherwise results are from the original papers. The best results are shown in **bold** and the second-best results are underlined.

(a) VOC and COCO

Method	VOC			COCO		
	mBO^i	mBO^c	mIoU	mBO^i	mBO^c	mIoU
SA [24]† NeurIPS'20	24.6	24.9	–	17.2	19.2	–
SLATE [32]† ICLR'22	35.9	41.5	–	29.1	33.6	–
CAE [26]† TMLR	$32.9_{\pm 0.9}$	$37.4_{\pm 1.0}$	–	–	–	–
Top-Down Slot Attention [18] NeurIPS'24	$43.9_{\pm 2.6}$	$51.0_{\pm 2.5}$	$42.0_{\pm 2.8}$	$33.0_{\pm 0.3}$	$40.3_{\pm 0.2}$	$31.2_{\pm 0.3}$
Rotating Features [25]† NeurIPS'23	$40.7_{\pm 0.1}$	$46.0_{\pm 0.1}$	–	–	–	–
$DIAS_{image}$ [38] ACM MM'25	$44.8_{\pm 9.8}$	–	$42.8_{\pm 1.7}$	$32.8_{\pm 0.1}$	–	$30.1_{\pm 0.3}$
DINOSAUR-Based						
DINOSAUR-MLP [31]† ICLR'23	$39.5_{\pm 0.1}$	$40.9_{\pm 0.1}$	–	$27.7_{\pm 0.2}$	$30.9_{\pm 0.2}$	–
DINOSAUR [31]† ICLR'23	$44.0_{\pm 1.9}$	$51.2_{\pm 1.9}$	–	$32.3_{\pm 0.4}$	$38.8_{\pm 0.4}$	–
DINOSAUR+FS+RC [34] CVPR'25	45.2	52.9	–	31.2	40.1	–
Autoregressive Teacher-Student						
SPOT [31] ICLR'23	$48.3_{\pm 0.4}$	$55.6_{\pm 0.4}$	$\underline{46.8_{\pm 0.4}}$	$35.0_{\pm 0.1}$	$44.7_{\pm 0.3}$	$\underline{33.0_{\pm 0.1}}$
SPOT+FS+RC [34] CVPR'25	49.3	<u>56.5</u>	–	<u>35.7</u>	<u>45.3</u>	–
Diffusion						
SlotDiffusion [35]† NeurIPS'23	<u>50.4</u>	55.3	–	31.0	35.0	–
Stable-LSD [16]† NeurIPS'23	–	–	–	30.4	–	–
GLASS [33] CVPR'25	**58.9**	**62.2**	**58.1**	**40.8**	**48.7**	**39.0**
Adaptive Slot Attention						
AdaSlot [11] CVPR'24	–	–	–	27.4	–	–
MetaSlot-MLP [23] NeurIPS'25	$43.9_{\pm 0.3}$	–	$42.1_{\pm 0.2}$	$29.5_{\pm 0.2}$	–	$27.9_{\pm 0.2}$
MetaSlot-Tfd [23] NeurIPS'25	$39.0_{\pm 0.3}$	–	$37.8_{\pm 0.5}$	$28.2_{\pm 0.7}$	–	$26.7_{\pm 0.6}$
MetaSlot-Dfz [23] NeurIPS'25	$36.5_{\pm 0.1}$	–	$35.3_{\pm 0.1}$	$27.2_{\pm 0.2}$	–	$25.8_{\pm 0.1}$
Ours	$47.83_{\pm 0.17}$	$55.34_{\pm 0.27}$	$46.01_{\pm 0.18}$	$32.88_{\pm 0.40}$	$42.78_{\pm 0.22}$	$30.75_{\pm 0.55}$

(b) MOVi-C and MOVi-E

Method	MOVi-C		MOVi-E	
	mBO^i	mIoU	mBO^i	mIoU
SA [24]† NeurIPS'20	$26.2_{\pm 1.0}$	–	$24.0_{\pm 1.2}$	–
SLATE [32]† ICLR'22	$39.4_{\pm 0.8}$	$37.8_{\pm 0.7}$	$30.2_{\pm 1.7}$	$28.6_{\pm 1.7}$
SLASH [19]† CVPR'23	–	$27.7_{\pm 5.9}$	–	–
Top-Down Slot Attention [18] NeurIPS'24	$46.8_{\pm 2.4}$	$45.9_{\pm 2.5}$	$39.3_{\pm 1.8}$	$\underline{38.3_{\pm 1.9}}$
DINOSAUR-Based				
DINOSAUR-MLP [31]† ICLR'23	$39.1_{\pm 0.2}$	–	$35.5_{\pm 0.2}$	–
DINOSAUR [31]† ICLR'23	42.4	–	–	–
DINOSAUR+FS+RC [34] CVPR'25	46.0	44.8	–	–
Autoregressive Teacher-Student				
SPOT [31] ICLR'23	$47.3_{\pm 1.2}$	$46.7_{\pm 1.3}$	$\mathbf{40.1_{\pm 1.2}}$	$\mathbf{39.3_{\pm 1.2}}$
SPOT+FS+RC [34] CVPR'25	**49.0**	**47.8**	–	–
Diffusion				
SlotDiffusion [35]† NeurIPS'23	–	–	30.2	30.2
LSD [16]† NeurIPS'23	$45.6_{\pm 0.8}$	$44.2_{\pm 0.9}$	$39.0_{\pm 0.5}$	$37.6_{\pm 0.5}$
Adaptive Slot Attention				
AdaSlot [11] CVPR'24	35.6	–	29.8	–
MetaSlot-MLP [23] NeurIPS'25	35.0	–	–	–
Ours	$\underline{48.35_{\pm 1.32}}$	$\underline{47.43_{\pm 1.26}}$	$\underline{39.56_{\pm 0.55}}$	$37.75_{\pm 0.55}$