

# Representation Learning with Adaptive Superpixel Coding

Mahmoud Khalil<sup>1</sup>, Ahmad Khalil<sup>1</sup>, Alioune Ngom<sup>1</sup>

<sup>1</sup>School of Computer Science, University of Windsor, Windsor, ON, Canada  
{khalil12, khalil71, angom}@uwindsor.ca

## Abstract

Deep learning vision models are typically tailored for specific modalities and often rely on domain-specific assumptions, such as the grid structures used by nearly all existing vision models. In this work, we propose a self-supervised model based on Transformers, which we call Adaptive Superpixel Coding (ASC). The key insight of our model is to overcome the limitations of traditional Vision Transformers, which depend on fixed-size and non-adaptive patch partitioning. Instead, ASC employs adaptive superpixel layers that dynamically adjust to the underlying image content.

We analyze key properties of the approach that make it effective, and find that our method outperforms widely-used alternatives on standard image downstream task benchmarks.

## 1 Introduction

Deep learning has revolutionized computer vision by enabling models to learn meaningful representations directly from unstructured sensory data, yielding significant performance gains across a wide range of tasks [Bommasani *et al.*, 2021]. In particular, self-supervised learning (SSL) has emerged as a powerful paradigm for learning visual representations without human annotations, achieving remarkable results in image classification, segmentation, and detection.

However, most SSL methods rely on fixed grid-based representations to encode image structure. Convolutional Neural Networks (CNNs) [He *et al.*, 2016; Sandler *et al.*, 2018] treat the image as a regular grid and extract local features using sliding windows. Vision Transformers (ViTs) [Dosovitskiy, 2020], though more flexible in modeling long-range dependencies, It begins by decomposing an image into non-overlapping patches arranged in a uniform grid. These grid structures lead to an inherent entanglement between the representation structure and image structure, constraining the model’s ability to adapt to variation in object shape, scale, or layout.

This entanglement between the representation structure and the image structure is especially limiting in downstream scene understanding tasks, where object-level reasoning, deformation invariance, and flexible spatial grouping are often essential. Many of these tasks require the model to localize,

compare, and track object parts across varying spatial arrangements—capabilities poorly supported by rigid, grid-aligned tokens.

Superpixels offer a natural alternative. Originating as a form of image over-segmentation [Ren and Malik, 2003], superpixels group pixels into visually and semantically coherent regions. These regions serve as low- to mid-level primitives that align well with object boundaries and exhibit strong inductive priors for efficient computation. They have been widely used in vision tasks including object detection [Shu *et al.*, 2013; Yan *et al.*, 2015], semantic segmentation [Gould *et al.*, 2008; Sharma *et al.*, 2014], saliency estimation [He *et al.*, 2015], and optical flow [Hu *et al.*, 2016], among others.

Yet, despite their utility, superpixels remain underutilized in transformer-based models, with only a few exceptions [Mei *et al.*, 2024; Zhu *et al.*, 2023]. One key reason is that transformers assume a fixed token layout, typically regular patches. However, superpixels are inherently irregular and content-adaptive, which clashes with the design of existing architectures that assume spatial consistency across layers.

In this work, we revisit the idea of superpixels and propose a token-level grouping mechanism for transformers. Rather than enforcing grid-based token structures, we introduce a novel transformer-compatible layer, *Adaptive Superpixel Coding (ASC)*, that adaptively merges tokens into semantically coherent regions. This layer computes pairwise token similarities, forms a weighted graph, and extracts connected components via a differentiable thresholding and grouping operation. The resulting components serve as *superpixels* that replace fixed patches in subsequent transformer layers.

Importantly, we demonstrate that this representation is especially effective in a self-supervised setting. Using contrastive training based on frame-level similarity across videos. The resulting representations are more aligned with object structure and transfer well to diverse downstream tasks.

### Our contributions are summarized as follows:

- We introduce *Adaptive Superpixel Coding (ASC)*, a novel transformer-compatible layer that integrates an adaptive superpixel mechanism, enabling the decoupling of the image’s *grid-based* structure from its representation structure.
- We demonstrate that the learned representation transfers effectively to a wide range of downstream vision tasks,

including image classification, object detection, and semantic segmentation. Our method outperforms strong self-supervised baselines such as BYOL [Grill *et al.*, 2020] and DINO [Caron *et al.*, 2021].

## 2 Related Work

**Grid-based Representation.** Deep learning methods for visual tasks [Krizhevsky *et al.*, 2012; Koonce and Koonce, 2021; Chen *et al.*, 2020a; Dosovitskiy, 2020] have garnered significant attention in the machine learning community in recent years. Early work, such as the use of CNN feature maps [Krizhevsky *et al.*, 2012], laid the foundation, followed by several efforts to integrate CNN-like architectures with self-attention mechanisms [Wang *et al.*, 2018; Carion *et al.*, 2020], and even replace convolutions altogether [Ramachandran *et al.*, 2019; Wang *et al.*, 2020]. Later innovations, such as the Vision Transformer (ViT) [Dosovitskiy, 2020], represent a natural evolution of this trend. Most of these approaches operate on a grid of image features ranging from patches or CNN feature maps to the final output layer. Although this grid-based representation has proven highly effective across a broad range of tasks, it inherently ties the learned features to fixed spatial 2D locations in image space.

**Off-the-Grid Representation.** A key method for decoupling model representations from the image grid is through the use of *cross-attention*, where one set of tokens is updated based on the values of another set. Related approaches, such as [Locatello *et al.*, 2020; Kipf *et al.*, 2021], extend this idea by leveraging slot-based attention mechanisms to model object-centric representations. The cross-attention approach has been particularly effective in object-centric tasks such as detection [Carion *et al.*, 2020; Zhu *et al.*, 2020], tracking [Kipf *et al.*, 2021; Meinhardt *et al.*, 2022], and instance segmentation [Cheng *et al.*, 2022; Kirillov *et al.*, 2023]. GroupViT [Xu *et al.*, 2022] is closely related to our approach, as it employs a large number of group tokens that are updated via cross-attention from visual tokens. In this framework, each group token serves as a compact and learnable representation that captures both local and global contextual information from the image. These tokens interact with image tokens through a series of cross-attention operations.

Inspired by GroupViT, our model introduces a novel layer into the Vision Transformer (ViT) architecture. However, in contrast to GroupViT, our method does not rely on explicit object queries or grouping tokens. Instead, we perform adaptive thresholding on similarity scores to modulate the resulting representations, allowing for more flexible and context-sensitive feature aggregation.

**Combining Tokens.** Our work is closely related to recent efforts aimed at improving the efficiency of Vision Transformers by reducing the number of tokens processed per layer. Notably, Token Merging (ToMe) [Bolya *et al.*, 2022] proposes a deterministic token coalescing mechanism based on token similarity, enabling faster inference without retraining. Other methods, such as AdaViT [Meng *et al.*, 2022] and DynamicViT [Rao *et al.*, 2021], introduce adaptive token pruning or merging strategies based on confidence scores or learned importance measures.

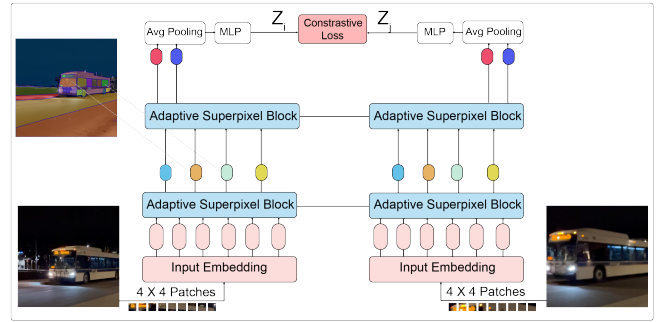


Figure 1: **The Architecture and Training Pipeline of ASC.** ASC is structured as a hierarchy of Transformer layers organized into stages, each processing increasingly larger visual segments. The first stage transforms raw pixels into superpixels representing object features while simultaneously eliminating redundancies. The bus image on the upper left shows the results of merging similar tokens and discarding noise. The dashed line illustrates the relationship between the superpixel representation and elements within the image.

While inspired by this broader class of approaches, our method differs in both *motivation* and *mechanism*. Instead of focusing on token efficiency, our goal is to induce *object-centric representations* by constructing a token affinity graph and identifying *connected components* to form semantically coherent groups.

In contrast to ToMe, which merges token pairs based on a greedy bipartite matching algorithm, our method constructs a *soft adjacency matrix* using a learnable gating mechanism and discovers groups via graph traversal (DFS). This enables us to form compositional groupings of similar tokens, rather than pairwise merges. Our formulation introduces two key distinctions: (1) a *differentiable connectivity threshold* jointly learned with the model, and (2) a *graph-theoretic* interpretation of token affinity, which structures the merging process around spatial and semantic coherence.

**Superpixels.** Recent methods such as SpFormer [Mei *et al.*, 2024] and Superpixel Transformer [Zhu *et al.*, 2023] incorporate superpixel-inspired structures into Transformer architectures to promote inductive biases that align with object boundaries. These methods *learn* to group pixels based on high-level semantic features through task-driven supervision (classification or segmentation objectives), rather than optimizing for low-level features of the pixel space. Although they incorporate a locality prior that biases adjacent pixels to cluster together, superpixel formation is ultimately guided by discriminative task performance. In contrast, our approach explicitly constructs a token affinity graph based on pairwise token similarity and performs grouping via a soft gating mechanism followed by graph traversal.

## 3 Method

In this section, we introduce Adaptive Superpixel Coding (Fig. 2; Section 3.2) and demonstrate how it can be integrated into an architecture for self-supervised contrastive learning (Fig. 1; Section 3.1).




Figure 2: Adaptive Superpixel Coding.

### 3.1 Overall Architecture

An overview of the Adaptive Superpixel Coding Transformer architecture is presented in (Fig. 1).

Inspired by Video Frame-level Similarity (VFS) [Xu and Wang, 2021], we begin by sampling two random frames,  $f_i$  and  $f_j$ , from a video clip, where each frame is represented as an RGB image  $f \in \mathbb{R}^{H \times W \times 3}$ . Each frame is uniformly partitioned into  $N$  non-overlapping patches, which serve as basic visual tokens.

For each patch, the raw RGB values are flattened into a vector of dimension  $4 \times 4 \times 3 = 48$  (assuming a patch size of  $4 \times 4$ ), and subsequently mapped into a latent embedding space of dimension  $C$  via a linear projection layer. This process yields token sets  $\{\mathbf{z}_t^{(i)}\}_{t=1}^T$  and  $\{\mathbf{z}_t^{(j)}\}_{t=1}^T$  for frames  $f_i$  and  $f_j$ , respectively, where superscripts indicate the source frame and subscripts index the patch.

The resulting patch embeddings are then processed by a series of Transformer blocks, equipped with an Adaptive Superpixel Layer 3.2. These token sequences, derived independently from frames  $f_i$  and  $f_j$ , are passed through two branches of a Siamese architecture: a predictor encoder  $P$  and a target encoder  $Q$ . Each encoder operates on the full token sequence, producing contextualized embeddings that are subsequently normalized using the L2 norm:

$$\mathbf{Z}_i = \frac{P(\{\mathbf{z}_t^{(i)}\}_{t=1}^T)}{\|P(\{\mathbf{z}_t^{(i)}\}_{t=1}^T)\|_2}, \quad \mathbf{Z}_j = \frac{Q(\{\mathbf{z}_t^{(j)}\}_{t=1}^T)}{\|Q(\{\mathbf{z}_t^{(j)}\}_{t=1}^T)\|_2} \quad (1)$$

where  $\mathbf{Z}_i$  and  $\mathbf{Z}_j$  denote the normalized embeddings of frames  $f_i$  and  $f_j$ , respectively, produced by the predictor encoder  $P$  and the target encoder  $Q$ . If these embeddings capture similar content, contrastive learning encourages their alignment while pushing apart representations from unrelated samples [Wu *et al.*, 2018].

### 3.2 Adaptive Superpixel Block

As illustrated in Fig. 2, the Adaptive Superpixel Block consists of a standard self-attention layer followed by Adaptive Superpixel Layer.

Given contextualized token embeddings produced by self-attention, we compute a superpixel-like representation by interpreting pairwise similarities as edges in a graph. Rather than

using a full attention distribution, we adopt the dot product to compute affinities:

$$S = \mathbf{Z}\mathbf{Z}^\top, \quad \text{where } \mathbf{Z} \in \mathbb{R}^{N \times d} \quad (2)$$

Each entry  $S_{ij}$  quantifies the similarity between token embeddings  $z_i$  and  $z_j$ . Since  $S$  is symmetric, it can be interpreted as a weighted adjacency matrix of an undirected graph  $G = (V, E)$ , where each token is a node, and the edge weights reflect the semantic affinity.

#### Theoretical Foundations and Assumptions.

- **Undirected Graph:** Since  $S^\top = S$ , the graph  $G$  is undirected. That is, if  $(i, j) \in E$ , then  $(j, i) \in E$  and  $s_{ij} = s_{ji}$ .
- **Graph Model:** Let  $V$  denote the set of tokens, and let an edge  $(i, j) \in E$  exist if  $A_{ij} > 0$ , where

$$A = \sigma(S - \theta),$$

$\sigma$  is the sigmoid function and  $\theta$  is a learnable threshold. The matrix  $A \in [0, 1]^{N \times N}$  encodes soft connectivity and controls the sparsity of the graph.

- **Connectedness and Objecthood:** Let  $\sim$  be the binary relation over  $V$  such that  $i \sim j$  if and only if there exists a path in  $G$  connecting  $z_i$  and  $z_j$ . The relation  $\sim$  is an equivalence relation (reflexive, symmetric, transitive), and partitions  $V$  into disjoint connected components. We formally define an **object** as such a connected component, consistent with Proposition 1 and 2.

**Adaptive Superpixel Layer Mechanism.** The goal of Adaptive Superpixel Layer is to construct object-centric embeddings by identifying and merging connected components in the similarity graph.

Let  $K = [k_1, \dots, k_N] \in \mathbb{R}^{N \times d_z}$  denote the key representations associated with tokens  $Z = [z_1, \dots, z_N]$ . Define the similarity matrix as:

$$S_{ij} = k_i^\top k_j, \quad A = \sigma(S - \theta), \quad (3)$$

A graph  $G = (V, E)$  is constructed where  $z_i, z_j \in V$  are connected by an undirected edge  $(i, j) \in E$  if  $A_{ij} > 0$ . By construction,  $A$  is symmetric and thus defines an undirected graph.

**Proposition 1 (Object Membership).** Two tokens  $z_i$  and  $z_j$  belong to the same object if and only if there exists a path in  $G$  connecting them.

**Proposition 2 (Transitivity of Connectivity).** If  $z_i \sim z_j$  and  $z_j \sim z_k$ , then  $z_i \sim z_k$ . Hence, all three tokens belong to the same connected component, forming a single object.

*The proof is in the supplementary material.*

**Component Aggregation.** The algorithmic implementation uses depth-first search (DFS) to enumerate connected components in  $G$ , as outlined in Algorithm 1. Each connected component  $C_i \subseteq V$  defines a group of semantically coherent tokens.

---

**Algorithm 1** Depth-First Search (DFS) for Connected Components

---

```
1: function DFS(node, A, vis, curr)
2:   vis[node]  $\leftarrow$  True
3:   append node to curr
4:   for nbr = 0 to  $N-1$  do
5:     if A[node, nbr] > 0 and  $\neg$  vis[nbr] then
6:       DFS(nbr, A, vis, curr)
7:   end if
8: end for
9: end function
10: function FINDCOMPONENTS(A)
11:   vis  $\leftarrow$  array of size  $N$  initialized to False
12:   comps  $\leftarrow$  empty list
13:   for node = 0 to  $N-1$  do
14:     if  $\neg$  vis[node] then
15:       curr  $\leftarrow$  empty list
16:       DFS(node, A, vis, curr)
17:       append curr to comps
18:     end if
19:   end for
20:   return comps
21: end function
```

---

**Feature Merging.** Once connected components are identified, the embeddings within each component  $C_i \subseteq \{1, \dots, N\}$  are merged via mean pooling:

$$\mathbf{z}_i^{l+1} = \frac{1}{|C_i|} \sum_{j \in C_i} \mathbf{z}_j^l \quad (4)$$

This yields a compact set of embeddings  $\tilde{Z} \in \mathbb{R}^{\tilde{N} \times d_z}$ , where  $\tilde{N} \leq N$ , and each row corresponds to an object-level representation. These object-centric tokens serve as input to the next layer in the Vision Transformer.

**Computational Complexity.** The Adaptive Superpixel Coding adds an  $\mathcal{O}(N^2d)$  overhead due to pairwise similarity computation and graph-based grouping. However, this cost is offset by a reduction in the number of tokens  $\tilde{N} < N$  passed to subsequent self-attention layers. The overall impact on runtime depends on the balance between this upfront cost and the savings from reduced token count, which can be quantified through ablation studies.

## 4 Experiments

We rigorously evaluate the learned representations from ASC on standard benchmarks, comparing with state-of-the-art self-supervised methods including DINO [Caron *et al.*, 2021], BYOL [Grill *et al.*, 2020], MoCo-v3 [Chen *et al.*, 2020b], and CLIP [Radford *et al.*, 2021]. We assess our model on three downstream tasks: image classification, object detection, and semantic segmentation.

This section first presents our detailed pre-training setup, evaluation protocols, and metrics. We then compare ASC with state-of-the-art methods across multiple benchmarks. Finally, we conduct ablation studies to analyze the contribution of each design choice in ASC.

### 4.1 Self-Supervised Pre-Training

**Datasets.** We use three video datasets for self-supervised pre-training:

- **Moments in Time** [Monfort *et al.*, 2019]: 1 million 3-second video clips spanning 339 action classes. We extract frames at 16 FPS, resulting in approximately 48 million frames at  $224 \times 224$  resolution.
- **Kinetics-700** [Carreira *et al.*, 2019]: 650,000 video clips of human actions. We sample at 15 FPS, yielding approximately 240 million frames at  $224 \times 224$  resolution.
- **Ego4D** [Grauman *et al.*, 2022]: 3,670 hours of first-person video. We sample frames at 5 FPS, resulting in approximately 66 million frames at  $224 \times 224$  resolution.

For all datasets, we maintain a strict separation between pre-training and downstream evaluation data to ensure fair assessment. When evaluating on standard benchmarks (e.g., ImageNet), we verify no overlap exists between pre-training videos and evaluation images.

**Architectures.** We follow the standard Vision Transformer (ViT) architecture [Dosovitskiy, 2020] with the following modifications. We reduce the patch size from  $16 \times 16$  to  $4 \times 4$ , using an input resolution of  $224 \times 224$ . We insert an Adaptive Superpixel Coding layer immediately after the self-attention module in each transformer block. This layer reduces the token sequence length through spatial aggregation. We use pre-normalization throughout the network.

The model comprises a *predictor encoder* and a *target encoder*, which form an asymmetric Siamese architecture. Both encoders share the same modified ViT backbone and a 3-layer MLP projector, whose parameters are tied. However, only the predictor encoder includes an additional 2-layer MLP predictor, which transforms the projected features to match the output of the target encoder. This asymmetry is critical: the target encoder provides a stable representation (updated via stop-gradient or exponential moving average), while the predictor encoder is trained to align with it. All batch normalization layers in the backbone, projector, and predictor use synchronized batch normalization (SyncBN) across devices, as in [Chen *et al.*, 2020a; Grill *et al.*, 2020].

**Learning Objectives.** Following a strategy similar to the [Xu and Wang, 2021] method, we begin with a video consisting of  $L$  frames  $\{f_1, f_2, \dots, f_L\}$ . We divide the video into 4 equal temporal intervals and randomly select one frame from each segment. These selected frames are then augmented. Two of the augmented frames are fed into the predictor encoder  $P$  and the target encoder  $T$ , respectively, producing normalized feature embeddings:  $Z_i = \frac{P(f_i)}{\|P(f_i)\|_2}$  and  $Z_j = \frac{Q(f_j)}{\|Q(f_j)\|_2}$ .

**Data augmentations** Besides using temporal signals to provide different views of training data, ASC uses the same set of augmentations as in [Wu *et al.*, 2018; Chen *et al.*, 2020b; Chen *et al.*, 2020a]. We apply spatial augmentations (random cropping to  $224 \times 224$  pixels with scale ranging from 0.2 to 1.0, horizontal flipping with probability 0.5) and color augmentations (color jittering with brightness, contrast, saturation, and hue factors of 0.4, 0.4, 0.4, and 0.1 respectively, random grayscale conversion with probability 0.2, and Gaussian blur with kernel size of  $23 \times 23$  and sigma 0.5).

| Method                              | Top-1       | Top-5       |
|-------------------------------------|-------------|-------------|
| MAE [He <i>et al.</i> , 2022]       | 72.6        | 93.8        |
| CLIP [Radford <i>et al.</i> , 2021] | 76.2        | 93.2        |
| DINO [Caron <i>et al.</i> , 2021]   | 78.2        | 94.3        |
| ToME [Bolya <i>et al.</i> , 2022]   | 77.7        | 92.8        |
| PiToME [Tran <i>et al.</i> , 2024]  | 79.1        | 95.4        |
| DINOv2 [Oquab <i>et al.</i> , 2023] | 83.2        | <i>NA</i>   |
| ASC (ours)                          | <b>82.1</b> | <b>96.4</b> |

Table 1: Top-1 and top-5 accuracy with ViT-B encoders on ImageNet.

**Pre-training Details** We trained ASC on our combined video dataset using only the self-supervised contrastive loss in (5). The model is trained on sub-sequences of 4 frames sampled at uniform intervals from each video clip. Our batch size is 512 distributed across 8 NVIDIA A100 GPUs with a learning rate initialized to 0.0016 and decayed via the cosine schedule [Loshchilov and Hutter, 2016]. We use the Adam [Diederik, 2014] optimizer with a weight decay of 0.05. We train for 200 epochs for ViT-B models.

**Hyperparameter Selection.** We performed hyperparameter tuning using a validation set comprising 5

- Learning rate: [0.0005, 0.001, 0.0016, 0.003, 0.005]
- Weight decay: [0.01, 0.05, 0.1, 0.2]
- Temperature for similarity threshold  $\theta$ : [0.05, 0.1, 0.2, 0.5]

Final values were selected based on validation loss and transferred to downstream tasks.

$$\mathcal{L}_{\text{pos}}(p_i, z_j) = \|p_i - z_j\|_2^2 = 2 - 2 \cdot \langle p_i, z_j \rangle \quad (5)$$

where:

- $\mathcal{L}_{\text{pos}}(p_i, z_j)$  is the contrastive loss between a positive pair of representations.
- $p_i$  is the output of the predictor network for the  $i$ -th view of a sample.
- $z_j$  is the output of the projector network (from the target encoder) for the  $j$ -th view of the same sample. This branch is typically stop-gradient.
- $\|p_i - z_j\|_2^2$  denotes the squared  $\ell_2$  distance between  $p_i$  and  $z_j$ .
- $\langle p_i, z_j \rangle$  is the dot product between the two vectors. When  $p_i$  and  $z_j$  are  $\ell_2$ -normalized, this corresponds to cosine similarity.

## 4.2 Linear evaluation on ImageNet

We assess the quality of the learned representations via linear probing on the ImageNet-1K dataset [Russakovsky *et al.*, 2015], adhering to standard evaluation protocols [Kolesnikov *et al.*, 2019; Kornblith *et al.*, 2019; Chen *et al.*, 2020a]. Table 1, 2 reports both top-1 and top-5 classification accuracy on the held-out test set, comparing our method with leading self-supervised approaches. Using a ViT-B backbone, ASC attains

| Method                              | Architecture             | Top-1       | Top-5       |
|-------------------------------------|--------------------------|-------------|-------------|
| SIMCLR [Chen <i>et al.</i> , 2020a] | ResNet-50 ( $4\times$ )  | 76.5        | 93.2        |
| BYOL [Grill <i>et al.</i> , 2020]   | ResNet-200 ( $2\times$ ) | 79.6        | 94.8        |
| DINO [Caron <i>et al.</i> , 2021]   | ResNet-50                | 75.3        | 95.0        |
| ASC (ours)                          | ViT-B                    | <b>82.1</b> | <b>96.4</b> |

Table 2: Top-1 and top-5 accuracy with various ResNet encoders on ImageNet.

82.1% top-1 and 96.4% top-5 accuracy, which is comparable to DINOv2 [Tran *et al.*, 2024] (83.2% top-1) and outperforms BYOL [Grill *et al.*, 2020] (78.6% top-1) and DINO [Caron *et al.*, 2021] (78.2% top-1) when using the same architecture.

It is worth noting that DINOv2 benefits from training on a carefully curated and filtered dataset. In contrast, our model is trained solely on raw, uncurated video frames without any manual filtering or supervision, yet achieves competitive performance.

## 4.3 Transfer to Diverse Classification Benchmarks

To evaluate the generality of the learned representations, we assess their performance across a suite of image classification tasks beyond ImageNet. This allows us to determine whether the features captured by our model are domain-agnostic or exhibit dataset-specific bias.

We conduct both linear probing and full fine-tuning on the following standard benchmarks:

- **CIFAR-10/100** [Krizhevsky *et al.*, 2009]: 10/100 classes of natural images (60,000  $32\times 32$  color images)
- **SUN397** [Xiao *et al.*, 2010a]: Scene understanding dataset with 397 categories (108,754 images)
- **VOC2007** [Everingham *et al.*, 2010a]: Object recognition with 20 classes (9,963 images)
- **DTD** [Cimpoi *et al.*, 2014a]: Describable Textures Dataset with 47 texture categories (5,640 images)
- **Flowers-102** [Nilsback and Zisserman, 2008]: 102 flower categories (8,189 images)

For all transfer learning experiments, we follow the protocols established in [Chen *et al.*, 2020a; Kornblith *et al.*, 2019] to ensure fair comparison. For linear probing, we train a linear classifier on frozen features extracted from our pre-trained model. For fine-tuning, we train the entire network end-to-end, initialized with our pre-trained weights. Evaluation is carried out using conventional metrics appropriate to each dataset, with performance reported on the held-out test set after model selection based on validation performance.

As summarized in Table 3, our method consistently outperforms BYOL across all evaluated tasks under both evaluation regimes. Notably, the learned representations transfer effectively to diverse visual domains, including low-resolution datasets such as CIFAR [Krizhevsky *et al.*, 2009], scene-centric datasets like SUN397 [Xiao *et al.*, 2010a] and VOC2007 [Everingham *et al.*, 2010a], and texture datasets such as DTD [Cimpoi *et al.*, 2014a].



| Method                              | Food101     | CIFAR10     | CIFAR100    | Birdsnap    | SUN397      | Cars        | Aircraft    | VOC2007     | DTD         | Pets        | Caltech-101 | Flowers     |
|-------------------------------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|
| <i>Linear evaluation:</i>           |             |             |             |             |             |             |             |             |             |             |             |             |
| ASC                                 | <b>76.2</b> | <b>91.9</b> | <b>79.1</b> | <b>58.5</b> | <b>63.0</b> | <b>68.6</b> | <b>61.5</b> | <b>83.4</b> | <b>76.3</b> | <b>91.1</b> | <b>94.8</b> | <b>96.6</b> |
| BYOL [Grill <i>et al.</i> , 2020]   | 75.3        | 91.3        | 78.4        | 57.2        | 62.2        | 67.8        | 60.6        | 82.5        | 75.5        | 90.4        | 94.2        | 96.1        |
| SIMCLR [Chen <i>et al.</i> , 2020a] | 68.4        | 90.6        | 71.6        | 37.4        | 58.8        | 50.3        | 50.3        | 80.5        | 74.5        | 83.6        | 90.3        | 91.2        |
| <i>Fine-tuned:</i>                  |             |             |             |             |             |             |             |             |             |             |             |             |
| ASC                                 | <b>89.4</b> | <b>98.2</b> | <b>86.9</b> | <b>77.1</b> | <b>64.5</b> | <b>92.3</b> | <b>88.7</b> | <b>86.1</b> | <b>76.9</b> | <b>92.3</b> | <b>94.4</b> | <b>97.5</b> |
| BYOL                                | 88.5        | 97.8        | 86.1        | 76.3        | 63.7        | 91.6        | 88.1        | 85.1        | 76.2        | 91.7        | 93.8        | 97.0        |
| SIMCLR [Chen <i>et al.</i> , 2020a] | 88.2        | 97.7        | 85.9        | 75.9        | 63.5        | 91.3        | 88.1        | 84.1        | 73.2        | 89.2        | 92.1        | 97.0        |

Table 3: Transfer learning results from ImageNet (IN) with the standard ResNet-50 architecture.

| Method                               | ImageNet    |             | CIFAR-10    | CIFAR-100   | VOC07       | SUN397      |
|--------------------------------------|-------------|-------------|-------------|-------------|-------------|-------------|
|                                      | Top-1       | Top-5       |             |             |             |             |
| MoCo-v3 [Chen <i>et al.</i> , 2020b] | 76.5        | 93.2        | 93.6        | 78.4        | 86.2        | 65.3        |
| BYOL [Grill <i>et al.</i> , 2020]    | 79.6        | 94.8        | 94.2        | 79.6        | 87.5        | 66.8        |
| DINO [Caron <i>et al.</i> , 2021]    | 78.2        | 94.3        | 95.1        | 81.2        | 89.3        | 69.7        |
| CLIP [Radford <i>et al.</i> , 2021]  | 76.2        | 93.2        | 95.7        | 82.9        | 91.2        | 72.8        |
| ASC (Ours)                           | <b>82.1</b> | <b>96.4</b> | <b>95.3</b> | <b>81.5</b> | <b>89.7</b> | <b>70.2</b> |

Table 4: Comparison with state-of-the-art self-supervised methods. All methods use a ViT-B backbone except CLIP, which uses ViT-B/32. Values represent accuracy (%) on linear evaluation.

#### 4.4 Comparison with State-of-the-Art Methods

Table 4 presents a comprehensive comparison between our approach and leading self-supervised methods across multiple benchmarks. For fairness, we compare methods using the same backbone architecture (ViT-B) whenever possible. Our method achieves competitive or superior performance compared to DINO [Caron *et al.*, 2021], BYOL [Grill *et al.*, 2020], and MoCo-v3 [Chen *et al.*, 2020b] across most evaluated tasks.

#### 4.5 Transfer to Other Vision Tasks

We assess the generality of the learned representation by transferring it to a diverse set of downstream tasks commonly encountered in computer vision, including semantic segmentation, object detection, and monocular depth estimation. These evaluations aim to determine whether ASC captures transferable features beyond image-level classification.

**Semantic Segmentation.** We evaluate ASC on the PASCAL VOC 2012 semantic segmentation benchmark, following the protocol detailed in Appendix F.1. The task involves assigning a semantic label to each pixel. To adapt our Vision Transformer backbone for this dense prediction task, we employ a transformer-compatible segmentation approach inspired by SegFormer [Xie *et al.*, 2021]. Our decoder leverages hierarchical multi-level features from transformer layers 3, 6, 9, and 12, preserving the global contextual information that is characteristic of transformer models.

For each feature level, we reshape the token sequence into 2D feature maps while maintaining their positional relationships. Rather than relying solely on convolutional operations, we employ a lightweight MLP decoder that preserves the contextual reasoning capabilities of the transformer. This decoder consists of layer-specific projection layers that unify channel dimensions to 256, followed by upsampling operations with skip connections. The final prediction head combines these multi-scale features through a fusion module before producing per-pixel class predictions. We fine-tune the model end-to-end using a combination of cross-entropy and Lovász-Softmax loss functions. As reported in Table 5, ASC achieves a sub-

stantial improvement over prior baselines, outperforming the Supervised-IN baseline by +1.9 mIoU and SIMCLR by +1.1 mIoU.

**Object Detection.** To evaluate object-level representation quality, we follow a transformer-compatible detection protocol on PASCAL VOC 2007. Rather than forcing our transformer features into CNN-based detection frameworks, we adapt the DETR [Carion *et al.*, 2020] approach which was specifically designed for transformer architectures. Our implementation maintains the query-based detection paradigm where a set of object queries interact with transformer features through cross-attention mechanisms.

Specifically, we extract hierarchical features from transformer layers 3, 6, 9, and 12, and process them with a transformer decoder consisting of alternating self-attention and cross-attention layers. The self-attention allows object queries to differentiate from each other, while cross-attention enables each query to attend to relevant image regions. This approach naturally leverages the global receptive field of transformers without requiring the spatial reshaping that CNN detectors need.

For fair comparison with established benchmarks, we report results using the standard protocol on PASCAL VOC 2007 trainval/test splits, measuring average precision at IoU threshold 0.5 (AP<sub>50</sub>). As shown in Table 5, ASC improves upon the Supervised-IN by +3.1 AP<sub>50</sub> and surpasses SIMCLR by +2.3 AP<sub>50</sub>, demonstrating that our self-supervised features better capture object-level semantics even when evaluated in a detection framework optimized for transformers.

**Monocular Depth Estimation** We evaluate monocular depth prediction on NYU Depth v2, using a transformer-specific protocol inspired by Dense Prediction Transformer (DPT) [Ranftl *et al.*, 2021] while maintaining fair comparison with baselines [Laina *et al.*, 2016].

Rather than directly adapting CNN-based decoders, we leverage the transformer’s inherent capabilities through a specialized depth prediction architecture. Our approach maintains the global context awareness that distinguishes transformers while efficiently generating dense spatial outputs:

- We extract multi-level features from transformer layers (3, 6, 9, 12), preserving both local patterns and global semantic information
- Instead of simple reshaping and convolution, we employ a cross-attention mechanism between transformer features and learnable query embeddings at each resolution level
- These query embeddings are arranged in a 2D grid matching target output resolutions and serve as a spatial structure for upsampling
- The cross-attention operation projects the transformer’s global features onto this spatially structured grid
- Multi-scale features are gradually merged through residual blocks that combine convolutions with self-attention mechanisms

This approach preserves the transformer’s ability to model long-range dependencies while effectively generating the spatially dense predictions required for depth estimation. We train

| Method                                  | AP <sub>50</sub> | mIoU        |
|---|------------------|-------------|
| Supervised-IN [He <i>et al.</i> , 2020] | 74.4             | 74.4        |
| MoCo [He <i>et al.</i> , 2020]          | 74.9             | 72.5        |
| SIMCLR (repro)                          | 75.2             | 75.2        |
| ASC                                     | <b>77.5</b>      | <b>76.3</b> |

Table 5: Transfer results on semantic segmentation and object detection.

| Method                                     | pct.<1.25   | Higher better<br>pct.<1.25 <sup>2</sup> | pct.<1.25 <sup>3</sup> | Lower better<br>rms | Lower better<br>rel |
|--|-------------|---|------------------------|---------------------|---------------------|
| Supervised-IN [Laina <i>et al.</i> , 2016] | 81.1        | 95.3                                    | 98.8                   | 0.573               | <b>0.127</b>        |
| SIMCLR (repro)                             | 83.3        | 96.5                                    | 99.1                   | 0.557               | 0.134               |
| BYOL                                       | 84.6        | 96.7                                    | 99.1                   | 0.541               | 0.129               |
| ASC  | <b>85.1</b> | <b>96.9</b>                             | <b>99.3</b>            | <b>0.533</b>        | 0.128               |

Table 6: Transfer results on NYU v2 depth estimation.

on the official NYU Depth v2 training set (approximately 24K RGB-D image pairs) and evaluate on the standard test set (654 images). The model is optimized using a combination of L1 depth loss and scale-invariant loss with weights 0.5 and 1.0, respectively.

As shown in Table 6, ASC achieves superior performance compared to self-supervised baselines across most metrics, with particularly notable improvements in accuracy thresholds and RMS error. The results demonstrate that our transformer-based architecture effectively captures both fine-grained details and global scene structure, essential for accurate depth estimation.

## 5 Model Ablations

To better understand the contribution of each component in the Adaptive Superpixel Coding (ASC) framework, we perform a series of ablation studies. Our goal is to quantify the impact of the Adaptive Superpixel Layer, the learnable threshold parameter, and the graph-based grouping mechanism on the learned representations and their transferability.

**Effect of Adaptive Superpixel Layer.** We first assess the importance of the Adaptive Superpixel Layer by removing it from the architecture and reverting to a standard Vision Transformer (ViT) with fixed-size patch embeddings. The model is trained under identical contrastive learning settings. As shown in Table 7, removing the superpixel mechanism leads to a consistent drop in accuracy across all downstream tasks, confirming that content-adaptive grouping contributes significantly to learning object-centric representations.

| Model Variant        | ImageNet Top-1 | VOC07 AP50  | NYU Depth pct.<1.25 |
|----------------------|----------------|-------------|---------------------|
| ASC (full)           | <b>82.1</b>    | <b>77.5</b> | <b>85.1</b>         |
| w/o Superpixel Layer | 78.6           | 74.8        | 82.4                |

Table 7: Effect of removing the Adaptive Superpixel Layer.

**Learnable vs. Fixed Threshold.** We next evaluate the role of the learnable threshold  $\theta$  used in forming the affinity graph for superpixel grouping. We compare ASC with a fixed threshold (e.g.,  $\theta = 0.2$ ) versus a learnable threshold that is optimized end-to-end. The learnable threshold consistently yields better performance.

| Threshold Type                | ImageNet Top-1 | CIFAR100    | VOC07 mIoU  |
|-------------------------------|----------------|-------------|-------------|
| Fixed ( $\theta = 0.2$ )      | 80.3           | 78.4        | 74.8        |
| Learnable ( $\theta$ trained) | <b>82.1</b>    | <b>79.1</b> | <b>76.3</b> |

Table 8: Impact of fixed vs. learnable threshold.

**Graph Traversal Strategy.** To test the effectiveness of our connected-component-based grouping via DFS, we replace it with a simpler token-merging strategy similar to ToMe [Bolya *et al.*, 2022]. The results in Table 9 show that the connected component strategy performs better.

| Grouping Mechanism         | PASCAL mIoU | NYU Depth rel |
|----------------------------|-------------|---------------|
| Connected Components (DFS) | <b>76.3</b> | <b>0.128</b>  |
| ToMe-style Merging         | 74.1        | 0.137         |

Table 9: Comparison of token grouping strategies.

**Impact of Component Aggregation Strategy.** We also compare mean pooling with alternative strategies such as max pooling and attention-based fusion. As shown in Table 10, mean pooling performs best overall.

| Aggregation Method | ImageNet Top-1 | Flowers Accuracy |
|--------------------|----------------|------------------|
| Mean Pooling       | <b>82.1</b>    | <b>96.6</b>      |
| Max Pooling        | 80.8           | 95.2             |
| Attention Fusion   | 81.4           | 96.1             |

Table 10: Comparison of token aggregation strategies.

## 6 Conclusion

We introduced ASC, a novel architectural component for self-supervised visual representation learning that adaptively groups tokens into object-level embeddings based on similarity across video frames. The method operates without negative pairs and integrates seamlessly with Vision Transformers. Under the standard linear evaluation protocol on ImageNet with a ViT-B backbone, ASC achieves a top-1 accuracy of 82.1%.

While effective, ASC has several limitations: (i) sensitivity to the gating threshold and merge strategy; (ii) lack of explicit mechanisms for ensuring object-level invariance to pose, occlusion, or viewpoint changes; and (iii) additional computational cost introduced by graph construction and traversal.

## References

- [Berg and Belhumeur, 2014] Thomas Berg and Peter N Belhumeur. Birdsnap: Large-scale fine-grained visual categorization of birds. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2019–2026, 2014.
- [Bolya *et al.*, 2022] Daniel Bolya, Cheng-Yang Fu, Xiaoliang Dai, Peizhao Zhang, Christoph Feichtenhofer, and Judy Hoffman. Token merging: Your vit but faster. *arXiv preprint arXiv:2210.09461*, 2022.
- [Bommasani *et al.*, 2021] Rishi Bommasani, Drew A Hudson, Ehsan Adeli, Russ Altman, Simran Arora, Sydney von Arx, Michael S Bernstein, Jeannette Bohg, Antoine Bosselut, Emma Brunskill, et al. On the opportunities and risks of foundation models. *arXiv preprint arXiv:2108.07258*, 2021.
- [Bossard *et al.*, 2014] Lukas Bossard, Matthieu Guillaumin, and Luc Van Gool. Food-101—mining discriminative components with random forests. In *European Conference on Computer Vision (ECCV)*, pages 446–461. Springer, 2014.
- [Carion *et al.*, 2020] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *European conference on computer vision*, pages 213–229. Springer, 2020.
- [Caron *et al.*, 2021] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 9650–9660, 2021.
- [Carreira *et al.*, 2019] Joao Carreira, Eric Noland, Chloe Hillier, and Andrew Zisserman. A short note on the kinetics-700 human action dataset. *arXiv preprint arXiv:1907.06987*, 2019.
- [Chen *et al.*, 2020a] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, pages 1597–1607. PMLR, 2020.
- [Chen *et al.*, 2020b] Xinlei Chen, Haoqi Fan, Ross Girshick, and Kaiming He. Improved baselines with momentum contrastive learning. *arXiv preprint arXiv:2003.04297*, 2020.
- [Cheng *et al.*, 2022] Bowen Cheng, Ishan Misra, Alexander G Schwing, Alexander Kirillov, and Rohit Girdhar. Masked-attention mask transformer for universal image segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 1290–1299, 2022.
- [Cimpoi *et al.*, 2014a] Mircea Cimpoi, Subhransu Maji, Iasonas Kokkinos, Sammy Mohamed, and Andrea Vedaldi. Describing textures in the wild. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3606–3613, 2014.
- [Cimpoi *et al.*, 2014b] Mircea Cimpoi, Subhransu Maji, Iasonas Kokkinos, Sammy Mohamed, and Andrea Vedaldi. Describing textures in the wild. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3606–3613, 2014.
- [Diederik, 2014] P Kingma Diederik. Adam: A method for stochastic optimization. (*No Title*), 2014.
- [Dosovitskiy, 2020] Alexey Dosovitskiy. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020.
- [Everingham *et al.*, 2010a] Mark Everingham, Luc Van Gool, Christopher KI Williams, John Winn, and Andrew Zisserman. The pascal visual object classes (voc) challenge. *International journal of computer vision*, 88:303–338, 2010.
- [Everingham *et al.*, 2010b] Mark Everingham, Luc Van Gool, Christopher KI Williams, John Winn, and Andrew Zisserman. The pascal visual object classes (voc) challenge. *International Journal of Computer Vision*, 88(2):303–338, 2010.
- [Fei-Fei *et al.*, 2007] Li Fei-Fei, Rob Fergus, and Pietro Perona. Learning generative visual models from few training examples: An incremental bayesian approach tested on 101 object categories. *Computer Vision and Image Understanding*, 106(1):59–70, 2007.
- [Gould *et al.*, 2008] Stephen Gould, Jim Rodgers, David Cohen, Gal Elidan, and Daphne Koller. Multi-class segmentation with relative location prior. *International journal of computer vision*, 80:300–316, 2008.
- [Grauman *et al.*, 2022] Kristen Grauman, Andrew Westbury, Eugene Byrne, Zachary Chavis, Antonino Furnari, Rohit Girdhar, Jackson Hamburger, Hao Jiang, Miao Liu, Xingyu Liu, et al. Ego4d: Around the world in 3,000 hours of egocentric video. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18995–19012, 2022.
- [Grill *et al.*, 2020] Jean-Bastien Grill, Florian Strub, Florent Altché, Corentin Tallec, Pierre Richemond, Elena Buchatskaya, Carl Doersch, Bernardo Avila Pires, Zhao-han Guo, Mohammad Gheshlaghi Azar, et al. Bootstrap your own latent—a new approach to self-supervised learning. *Advances in neural information processing systems*, 33:21271–21284, 2020.
- [He *et al.*, 2015] Shengfeng He, Rynson WH Lau, Wenxi Liu, Zhe Huang, and Qingxiong Yang. Supercnn: A superpixelwise convolutional neural network for salient object detection. *International journal of computer vision*, 115:330–344, 2015.
- [He *et al.*, 2016] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [He *et al.*, 2020] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9729–9738, 2020.



- [He *et al.*, 2022] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 16000–16009, 2022.
- [Hu *et al.*, 2016] Yinlin Hu, Rui Song, Yunsong Li, Peng Rao, and Yangli Wang. Highly accurate optical flow estimation on superpixel tree. *Image and Vision Computing*, 52:167–177, 2016.
- [Kipf *et al.*, 2021] Thomas Kipf, Gamaleldin F Elsayed, Aravindh Mahendran, Austin Stone, Sara Sabour, Georg Heigold, Rico Jonschkowski, Alexey Dosovitskiy, and Klaus Greff. Conditional object-centric learning from video. *arXiv preprint arXiv:2111.12594*, 2021.
- [Kirillov *et al.*, 2023] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C Berg, Wan-Yen Lo, et al. Segment anything. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4015–4026, 2023.
- [Kolesnikov *et al.*, 2019] Alexander Kolesnikov, Xiaohua Zhai, and Lucas Beyer. Revisiting self-supervised visual representation learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 1920–1929, 2019.
- [Koonce and Koonce, 2021] Brett Koonce and Brett Koonce. Resnet 50. *Convolutional neural networks with swift for tensorflow: image recognition and dataset categorization*, pages 63–72, 2021.
- [Kornblith *et al.*, 2019] Simon Kornblith, Jonathon Shlens, and Quoc V Le. Do better imagenet models transfer better? In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2661–2671, 2019.
- [Krause *et al.*, 2013] Jonathan Krause, Michael Stark, Jia Deng, and Li Fei-Fei. 3d object representations for fine-grained categorization. In *IEEE International Conference on Computer Vision Workshops (ICCVW)*, pages 554–561, 2013.
- [Krizhevsky *et al.*, 2009] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009.
- [Krizhevsky *et al.*, 2012] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. *Advances in neural information processing systems*, 25, 2012.
- [Krizhevsky, 2009] Alex Krizhevsky. Learning multiple layers of features from tiny images. Technical Report TR-2009, University of Toronto, 2009.
- [Laina *et al.*, 2016] Iro Laina, Christian Rupprecht, Vasileios Belagiannis, Federico Tombari, and Nassir Navab. Deeper depth prediction with fully convolutional residual networks. In *Proceedings of the International Conference on 3D Vision (3DV)*, pages 239–248. IEEE, 2016.
- [Locatello *et al.*, 2020] Francesco Locatello, Dirk Weisenborn, Thomas Unterthiner, Aravindh Mahendran, Georg Heigold, Jakob Uszkoreit, Alexey Dosovitskiy, and Thomas Kipf. Object-centric learning with slot attention. *Advances in neural information processing systems*, 33:11525–11538, 2020.
- [Loshchilov and Hutter, 2016] Ilya Loshchilov and Frank Hutter. Sgdr: Stochastic gradient descent with warm restarts. *arXiv preprint arXiv:1608.03983*, 2016.
- [Maji *et al.*, 2013] Subhansu Maji, Esa Rahtu, Juho Kannala, Matthew Blaschko, and Andrea Vedaldi. Fine-grained visual classification of aircraft. In *arXiv preprint arXiv:1306.5151*, 2013.
- [Mei *et al.*, 2024] Jieru Mei, Liang-Chieh Chen, Alan Yuille, and Cihang Xie. Spformer: Enhancing vision transformer with superpixel representation. *arXiv preprint arXiv:2401.02931*, 2024.
- [Meinhardt *et al.*, 2022] Tim Meinhardt, Alexander Kirillov, Laura Leal-Taixe, and Christoph Feichtenhofer. Trackformer: Multi-object tracking with transformers. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8844–8854, 2022.
- [Meng *et al.*, 2022] Lingchen Meng, Hengduo Li, Bor-Chun Chen, Shiyi Lan, Zuxuan Wu, Yu-Gang Jiang, and Ser-Nam Lim. Adavit: Adaptive vision transformers for efficient image recognition. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 12309–12318, 2022.
- [Monfort *et al.*, 2019] Mathew Monfort, Alex Andonian, Bolei Zhou, Kandan Ramakrishnan, Sarah Adel Bargal, Tom Yan, Lisa Brown, Quanfu Fan, Dan Gutfreund, Carl Vondrick, et al. Moments in time dataset: one million videos for event understanding. *IEEE transactions on pattern analysis and machine intelligence*, 42(2):502–508, 2019.
- [Nilsback and Zisserman, 2008] Maria-Elena Nilsback and Andrew Zisserman. Automated flower classification over a large number of classes. In *2008 Sixth Indian Conference on Computer Vision, Graphics & Image Processing*, pages 722–729. IEEE, 2008.
- [Oquab *et al.*, 2023] Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, et al. Dinov2: Learning robust visual features without supervision. *arXiv preprint arXiv:2304.07193*, 2023.
- [Parkhi *et al.*, 2012] Omkar M Parkhi, Andrea Vedaldi, Andrew Zisserman, and CV Jawahar. Cats and dogs. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3498–3505, 2012.
- [Radford *et al.*, 2021] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PmLR, 2021.

- [Ramachandran *et al.*, 2019] Prajit Ramachandran, Niki Parmar, Ashish Vaswani, Irwan Bello, Anselm Levskaya, and Jon Shlens. Stand-alone self-attention in vision models. *Advances in neural information processing systems*, 32, 2019.
- [Ranftl *et al.*, 2021] René Ranftl, Alexey Bochkovskiy, and Vladlen Koltun. Vision transformers for dense prediction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 12179–12188, 2021.
- [Rao *et al.*, 2021] Yongming Rao, Wenliang Zhao, Benlin Liu, Jiwen Lu, Jie Zhou, and Cho-Jui Hsieh. Dynamicvit: Efficient vision transformers with dynamic token sparsification. *Advances in neural information processing systems*, 34:13937–13949, 2021.
- [Ren and Malik, 2003] Ren and Malik. Learning a classification model for segmentation. In *Proceedings ninth IEEE international conference on computer vision*, pages 10–17. IEEE, 2003.
- [Russakovsky *et al.*, 2015] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C Berg, and Li Fei-Fei. Imagenet large scale visual recognition challenge. *International Journal of Computer Vision*, 115(3):211–252, 2015.
- [Sandler *et al.*, 2018] Mark Sandler, Andrew Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. Mobilenetv2: Inverted residuals and linear bottlenecks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4510–4520, 2018.
- [Sharma *et al.*, 2014] Abhishek Sharma, Oncel Tuzel, and Ming-Yu Liu. Recursive context propagation network for semantic scene labeling. *Advances in Neural Information Processing Systems*, 27, 2014.
- [Shu *et al.*, 2013] Guang Shu, Afshin Dehghan, and Mubarak Shah. Improving an object detector and extracting regions using superpixels. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3721–3727, 2013.
- [Tran *et al.*, 2024] Chau Tran, Duy MH Nguyen, Manh-Duy Nguyen, TrungTin Nguyen, Ngan Le, Pengtao Xie, Daniel Sonntag, James Y Zou, Binh Nguyen, and Mathias Niepert. Accelerating transformers with spectrum-preserving token merging. *Advances in Neural Information Processing Systems*, 37:30772–30810, 2024.
- [Wang *et al.*, 2018] Xiaolong Wang, Ross Girshick, Abhinav Gupta, and Kaiming He. Non-local neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7794–7803, 2018.
- [Wang *et al.*, 2020] Huiyu Wang, Yukun Zhu, Bradley Green, Hartwig Adam, Alan Yuille, and Liang-Chieh Chen. Axial-deeplab: Stand-alone axial-attention for panoptic segmentation. In *European conference on computer vision*, pages 108–126. Springer, 2020.
- [Wu *et al.*, 2018] Zhirong Wu, Yuanjun Xiong, Stella X Yu, and Dahua Lin. Unsupervised feature learning via non-parametric instance discrimination. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3733–3742, 2018.
- [Xiao *et al.*, 2010a] Jianxiong Xiao, James Hays, Krista A Ehinger, Aude Oliva, and Antonio Torralba. Sun database: Large-scale scene recognition from abbey to zoo. In *2010 IEEE computer society conference on computer vision and pattern recognition*, pages 3485–3492. IEEE, 2010.
- [Xiao *et al.*, 2010b] Jianxiong Xiao, James Hays, Krista A Ehinger, Aude Oliva, and Antonio Torralba. Sun database: Large-scale scene recognition from abbey to zoo. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3485–3492, 2010.
- [Xie *et al.*, 2021] Enze Xie, Wenhai Wang, Zhiding Yu, Anima Anandkumar, Jose M. Alvarez, and Ping Luo. Segformer: Simple and efficient design for semantic segmentation with transformers. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2021.
- [Xu and Wang, 2021] Jiarui Xu and Xiaolong Wang. Re-thinking self-supervised correspondence learning: A video frame-level similarity perspective. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10075–10085, 2021.
- [Xu *et al.*, 2022] Jiarui Xu, Shalini De Mello, Sifei Liu, Wonmin Byeon, Thomas Breuel, Jan Kautz, and Xiaolong Wang. Groupvit: Semantic segmentation emerges from text supervision. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18134–18144, 2022.
- [Yan *et al.*, 2015] Junjie Yan, Yinan Yu, Xiangyu Zhu, Zhen Lei, and Stan Z Li. Object detection by labeling superpixels. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5107–5116, 2015.
- [Zhu *et al.*, 2020] X Zhu, W Su, L Lu, B Li, X Wang, J Dai, and Deformable DETR. deformable transformers for end-to-end object detection. URL: <https://arxiv.org/abs/2010.04159> (Accessed 29.11. 2023), 2020.
- [Zhu *et al.*, 2023] Alex Zihao Zhu, Jieru Mei, Siyuan Qiao, Hang Yan, Yukun Zhu, Liang-Chieh Chen, and Henrik Kretzschmar. Superpixel transformers for efficient semantic segmentation. In *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 7651–7658. IEEE, 2023.

# Supplementary Material for Representation Learning with Adaptive Superpixel Coding

## A Proofs

Let  $G = (V, E, \Psi)$  be an undirected graph, where each node  $v \in V$  corresponds to a feature embedding in an image, and each edge  $e \in E$  connects two embeddings such that  $\Psi(e) = \{u, v\}$  for some  $u, v \in V$ .

**Definition:**

- A **walk** in  $G$  is an alternating sequence of vertices and edges:

$$w_0, e_1, w_1, e_2, \dots, e_k, w_k$$

such that  $\Psi(e_i) = \{w_{i-1}, w_i\}$  for all  $1 \leq i \leq k$ . A walk may contain repeated edges and vertices.

- A **trail** is a walk in which no edge is repeated.
- A **path** is a trail in which no vertex is repeated.

**Definition of an Object.** An **object** in an image is defined as a maximal set of feature embeddings (vertices in the graph) such that for all  $x, y$  in this set, there exists a path in  $G$  from  $x$  to  $y$ , and for any  $z \notin \text{object}$ , there is no path from  $x$  to  $z$ . This definition is equivalent to a connected component in graph theory.

**Proposition 1.** Let  $x, y \in V$ . Then  $x$  and  $y$  belong to the same object if and only if there exists a path in  $G$  connecting  $x$  and  $y$ .

**Proof of Proposition 1.** Let  $P(x, y)$  be the proposition: “There exists a path in  $G$  connecting  $x$  and  $y$ .” Let  $O(x, y)$  be the proposition: “ $x$  and  $y$  belong to the same object.”

We aim to prove:

$$P(x, y) \leftrightarrow O(x, y)$$

( $\Rightarrow$ ) Suppose  $P(x, y)$  holds, i.e., there exists a path from  $x$  to  $y$ . By the definition of an object as a connected subgraph (connected component), any two vertices connected by a path are in the same component. Hence,  $x$  and  $y$  must belong to the same object. Thus,  $P(x, y) \Rightarrow O(x, y)$ .

( $\Leftarrow$ ) Suppose  $O(x, y)$  holds, i.e.,  $x$  and  $y$  belong to the same object. By definition of object, this implies the existence of a path connecting  $x$  and  $y$ . Hence,  $O(x, y) \Rightarrow P(x, y)$ .

Therefore,  $P(x, y) \leftrightarrow O(x, y)$ . □

**Proposition 2.** Connectivity is transitive: For any  $a, b, c \in V$ , if there exists a path from  $a$  to  $b$  and a path from  $b$  to  $c$ , then there exists a path from  $a$  to  $c$ .

**Proof of Proposition 2.** Let  $P(u, v)$  be the proposition: “There exists a path in  $G$  connecting  $u$  and  $v$ .”

We are given:

$$P(a, b) \wedge P(b, c)$$

We want to prove:

$$P(a, c)$$

Let  $p_1 = (a = v_0, v_1, \dots, v_k = b)$  be a path from  $a$  to  $b$ , and  $p_2 = (b = u_0, u_1, \dots, u_\ell = c)$  a path from  $b$  to  $c$ .

Since paths are sequences of adjacent vertices with no repetition, we can concatenate  $p_1$  and  $p_2$  (excluding the repeated node  $b$ ) and remove any resulting cycles to construct a new path  $p_3$  from  $a$  to  $c$ . This process is guaranteed by the fact that any walk can be converted to a path by removing cycles (a well-known result in graph theory).

Hence,  $P(a, c)$  holds. Therefore,  $P(a, b) \wedge P(b, c) \Rightarrow P(a, c)$ , proving transitivity of connectivity. □

**Conclusion.** By Propositions 1 and 2, the relation “being connected by a path” is an equivalence relation on  $V$  (reflexive, symmetric, transitive). Thus, the graph is partitioned into equivalence classes of connected vertices, i.e., connected components. Each connected component corresponds to a unique object in the image. Therefore, finding objects reduces to computing connected components in  $G$ .

## B Datasets

In this work, we use multiple datasets for every experiment. We make use of publicly available datasets that are released under MIT License and that are open to all research work.

We use three large-scale video datasets for self-supervised pre-training

### B.1 Moments in Time

[Monfort *et al.*, 2019]: A diverse dataset of one million 3-second video clips covering dynamic scenes and actions. We use the official training split containing 802,264 videos across 339 action classes.

## B.2 Ego4D

[Grauman *et al.*, 2022]: A large-scale egocentric video dataset capturing real-world, first-person interactions across diverse scenarios. We use 3,670 hours of video from the official training set, excluding any sequences that overlap with benchmark evaluation domains.

## B.3 Kinetics-700

[Carreira *et al.*, 2019]: A benchmark dataset consisting of over 650K video clips annotated with human actions. We use the official training split containing 545,317 videos spanning 700 action classes.

We evaluate our method on a diverse collection of publicly available datasets for both linear evaluation and fine-tuning protocols. All datasets are released under permissive licenses (e.g., MIT, CC BY) and are widely used in self-supervised learning literature.

## B.4 ImageNet

[Russakovsky *et al.*, 2015] For the main benchmark, we follow the standard linear evaluation protocol on ImageNet-1K, which consists of 1.28 million training images and 50,000 validation images across 1,000 categories. All models are evaluated using a frozen backbone and a single-layer linear classifier trained on top of the final layer features. Images are resized and center-cropped to  $224 \times 224$  resolution, consistent with prior work.

**Downstream Classification Datasets** We further evaluate the generalization of learned representations across 12 diverse classification tasks. For each dataset, a linear classifier or full fine-tuning is applied on top of the pretrained encoder.

- **Food101** [Bossard *et al.*, 2014]: 101 food categories with 101,000 images.
- **CIFAR-10/100** [Krizhevsky, 2009]: 10 and 100-class image classification datasets with low-resolution  $32 \times 32$  images.
- **Birdsnap** [Berg and Belhumeur, 2014]: 500 bird species with over 49,000 high-resolution images.
- **SUN397** [Xiao *et al.*, 2010b]: Scene recognition dataset with 397 categories and over 100,000 images.
- **Stanford Cars** [Krause *et al.*, 2013]: Fine-grained car classification with 196 classes.
- **FGVC Aircraft** [Maji *et al.*, 2013]: Aircraft model recognition with 100 classes.
- **VOC2007** [Everingham *et al.*, 2010b]: 20-class object classification, evaluated using classification labels only.
- **DTD** [Cimpoi *et al.*, 2014b]: A dataset of texture images categorized by human-centric attributes.
- **Oxford Pets** [Parkhi *et al.*, 2012]: 37 pet categories with annotations for both class and breed.
- **Caltech-101** [Fei-Fei *et al.*, 2007]: 101 object categories and a background class.
- **Oxford Flowers-102** [Nilsback and Zisserman, 2008]: Flower classification with 102 categories.

All datasets are preprocessed following the protocols established in [Grill *et al.*, 2020; Chen *et al.*, 2020a], including standard resizing and normalization. Where applicable, training-validation splits and evaluation metrics follow the respective official implementations or prior literature for fair comparison.

## C Metrics

**Image Classification (e.g., ImageNet, CIFAR, SUN397).** We evaluate classification performance using **Top-1** and **Top-5 accuracy**:

$$\text{Top-1 Accuracy} = \frac{1}{N} \sum_{i=1}^N \mathbb{1}(y_i = \hat{y}_i) \quad (6)$$

$$\text{Top-5 Accuracy} = \frac{1}{N} \sum_{i=1}^N \mathbb{1}(y_i \in \text{Top-5}(\hat{p}_i)) \quad (7)$$

where  $y_i$  is the ground-truth label,  $\hat{y}_i$  is the top-1 predicted label, and  $\hat{p}_i$  is the model’s class probability vector.

**Object Detection (PASCAL VOC 2007).** We report **Average Precision at IoU = 0.5** (AP@50):

$$\text{AP}_{50} = \int_0^1 p(r) dr \quad (8)$$

where  $p(r)$  is the precision at recall  $r$ , and predictions are considered correct if:

$$\text{IoU} = \frac{|\mathbf{B}_{\text{pred}} \cap \mathbf{B}_{\text{gt}}|}{|\mathbf{B}_{\text{pred}} \cup \mathbf{B}_{\text{gt}}|} \geq 0.5 \quad (9)$$

**Semantic Segmentation (PASCAL VOC 2012).** We evaluate segmentation using **mean Intersection over Union (mIoU)**:

$$\text{mIoU} = \frac{1}{C} \sum_{c=1}^C \frac{TP_c}{TP_c + FP_c + FN_c} \quad (10)$$

where  $TP_c$ ,  $FP_c$ , and  $FN_c$  are the true positives, false positives, and false negatives for class  $c$ , and  $C$  is the total number of classes.

**Monocular Depth Estimation (NYU Depth v2).** We use both accuracy and error metrics:

- **Threshold accuracy** at  $\delta \in \{1.25, 1.25^2, 1.25^3\}$ :

$$\text{Accuracy}_\delta = \frac{1}{N} \sum_{i=1}^N \mathbb{I} \left( \max \left( \frac{d_i}{\hat{d}_i}, \frac{\hat{d}_i}{d_i} \right) < \delta \right) \quad (11)$$

- **Root Mean Square Error (RMSE)**:

$$\text{RMSE} = \sqrt{\frac{1}{N} \sum_{i=1}^N (d_i - \hat{d}_i)^2} \quad (12)$$

- **Relative Error (REL)**:

$$\text{REL} = \frac{1}{N} \sum_{i=1}^N \frac{|d_i - \hat{d}_i|}{\hat{d}_i} \quad (13)$$

Here,  $d_i$  and  $\hat{d}_i$  denote the predicted and ground-truth depth at pixel  $i$ , respectively.

## D Evaluation Protocols

**Linear Evaluation on ImageNet.** Our linear evaluation protocol strictly follows established practices in self-supervised learning research:

1. Pre-trained model weights are frozen (no backpropagation through backbone)
2. A single linear layer is trained on top of the frozen features
3. Standard ImageNet augmentation protocol is used (resize to 256px, random crop to 224px)
4. Training uses SGD with momentum (0.9) for 100 epochs
5. Learning rate starts at 0.1 and is decayed by a factor of 10 at epochs 60 and 80
6. Weight decay is set to 0.0001

The reported accuracy is calculated on the official ImageNet validation set of 50,000 images.

**Transfer Learning Protocol** For fine-tuning on downstream datasets, we use:

- **Optimizer:** SGD with momentum (0.9)
- **Learning rate:** 0.01 for linear probing, 0.0001 for full fine-tuning, with cosine decay
- **Weight decay:** 0.0001
- **Epochs:** 100 for all datasets
- **Early stopping:** Based on validation performance
- **Augmentations:** Random resize, crop, and flip consistent with standard practices

For all transfer learning experiments, we divide each dataset into standard train/val/test splits (or use official splits where available), and report results on the designated test splits.

**Data Split Information** For each downstream evaluation benchmark, we use the following splits:

- **ImageNet:** 1.28M training images, 50K validation images (used as test set)
- **CIFAR-10/100:** 50K training images, 10K test images
- **VOC2007:** Train/val split (5,011 images), test split (4,952 images)
- **SUN397:** 19,850 training images, 19,850 validation images, 19,850 test images (official split)
- **DTD:** 3,760 training images, 1,880 test images
- **Flowers-102:** 2,040 training images, 6,149 test images



## E Experimental Details

**Dataset Processing Pipeline.** For all video datasets, we apply the following processing:

1. Decode videos at native resolution
2. Sample frames at specified FPS rates (16 FPS for Moments in Time, 15 FPS for Kinetics-700, 5 FPS for Ego4D)
3. Resize frames to maintain aspect ratio with shorter side = 256 pixels
4. Apply center crop to obtain  $224 \times 224$  pixel frames for pre-training

**Training Setup.** We train ASC using a self-supervised contrastive loss (Equation (5)) applied over 4-frame video subsequences. Each batch consists of 512 video clips. Training is conducted using:

- **Optimizer:** Adam [Diederik, 2014] with  $\beta_1 = 0.9$ ,  $\beta_2 = 0.999$ ,  $\epsilon = 10^{-8}$
- **Learning rate:** 0.0016, decayed via cosine schedule [Loshchilov and Hutter, 2016] with 5 epochs of warm-up
- **Weight decay:** 0.05
- **Batch size:** 512
- **Input:** Sub-sequences of 4 consecutive frames
- **Training epochs:** 200 for all models
- **Target encoder update:** Exponential moving average with momentum 0.996

**Hyperparameter Optimization.** We conduct a grid search over key hyperparameters using 5% of the training data as a validation set. The parameters and their search spaces include:

- Learning rate: [0.0005, 0.001, 0.0016, 0.003, 0.005]
- Weight decay: [0.01, 0.05, 0.1, 0.2]
- Similarity threshold  $\theta$ : [0.05, 0.1, 0.2, 0.5]
- Batch size: [1024, 2048, 4096]

The optimal values were selected based on validation loss convergence and downstream performance on a small subset of ImageNet classes.

**Compute Resources.** Experiments were conducted on 8× NVIDIA A100 GPUs (40GB memory each). Training on the combined dataset takes approximately 24 hours for ViT-Tiny and 2.5 days for ViT-B. The total compute for all reported experiments is estimated at approximately 200 GPU-hours, aligning with the computational budget of comparable prior work.

## F Downstream Task Architectures and Protocols

### F.1 Semantic Segmentation

For semantic segmentation, we evaluate on the PASCAL VOC 2012 dataset using the standard split provided by `augmented_voc`. The dataset contains 21 classes (including background).

**Architecture.** Our segmentation architecture adapts the ViT backbone to dense prediction through a carefully designed decoder structure. The decoder follows a Feature Pyramid Network (FPN) design that leverages multi-scale features from intermediate transformer layers:

- **Feature extraction:** We extract features from transformer layers 3, 6, 9, and 12, providing hierarchical representations at different semantic levels.
- **Feature normalization:** Each feature set is processed through a  $1 \times 1$  convolutional layer to unify channel dimensions to 256.
- **Upsampling pathway:** Starting from the deepest layer, features are progressively upsampled using bilinear interpolation and merged with corresponding shallower features through lateral connections.
- **Refinement:** Each merged feature map undergoes a  $3 \times 3$  convolutional refinement step with batch normalization and ReLU activation.
- **Prediction head:** The final feature map is processed by a  $1 \times 1$  convolution that outputs logits for the 21 semantic classes.

**Training Details.** Our ViT backbone is initialized with weights pretrained via ASC and fine-tuned end-to-end with the segmentation decoder. Input images are resized to  $512 \times 512$  pixels.

Training is conducted using the AdamW optimizer with an initial learning rate of  $5 \times 10^{-5}$ , cosine annealing scheduler, and a batch size of 16. We employ a combination of cross-entropy and Lovász-Softmax loss functions with weights of 1.0 and 0.5 respectively. We fine-tune for 20k iterations with a linear warmup for the first 1500 iterations. Data augmentation includes random scaling (0.5-2.0), random cropping, horizontal flipping, and color jittering. Evaluation is done using the mean Intersection-over-Union (mIoU) metric.

## F.2 Object Detection

We follow the object detection protocol on PASCAL VOC 2007 using Faster R-CNN as our detection framework.

**Architecture.** Since standard detection frameworks are designed for convolutional backbones, we integrate our pretrained ViT with Faster R-CNN through a specialized feature adapter:

- **Token reshaping:** The final layer patch tokens from the ViT ( $N$  tokens with dimension  $D$ ) are reshaped into a 2D spatial grid with dimensions  $H \times W \times D$ , where  $H \times W = N$ .
- **Feature adapter:** This grid is processed through a small convolutional network consisting of:
  - A  $1 \times 1$  convolution that maps  $D$  dimensions to 512 channels
  - Two residual blocks, each with two  $3 \times 3$  convolutions and skip connections
  - A final  $1 \times 1$  convolution that maps to the expected backbone output channels (256 for C4 Faster R-CNN)
- **Multi-scale features:** For improved detection of objects at different scales, we extract features from intermediate transformer layers (3, 6, 9) and process them through similar adapter modules before feeding them into the FPN structure of Faster R-CNN.
- **Detection head:** We use the standard Faster R-CNN detection head with RPN, RoI Pooling, and classification/bounding box regression heads.

**Training Details.** The detector is trained on the `trainval2007` split and evaluated on `test2007`. We use stochastic gradient descent with momentum (0.9), an initial learning rate of  $1 \times 10^{-3}$  with step decay at iterations 18k and 22k, weight decay of  $1 \times 10^{-4}$ , and train for 24k iterations with a batch size of 8.

We implement gradient clipping at a max norm of 10.0 to stabilize training. Standard detection augmentations include horizontal flipping, scale jittering, and color augmentation. Results are reported using  $AP_{50}$  following standard PASCAL VOC evaluation protocol.

## F.3 Monocular Depth Estimation

We evaluate monocular depth prediction on the NYU Depth v2 dataset, which consists of RGB-D indoor scenes.

**Architecture.** Our depth estimation approach is inspired by Dense Prediction Transformers (DPT) [Ranftl *et al.*, 2021], which is specifically designed to leverage transformer architectures for dense prediction tasks. The architecture consists of:

- **Feature extraction:** We extract token embeddings from transformer layers 3, 6, 9, and 12, providing a hierarchical representation with varying receptive fields.
- **Reassembly blocks:** For each feature level, we use:
  - A cross-attention mechanism where learnable query embeddings (arranged in a 2D grid) attend to the transformer tokens
  - This effectively projects global features back into a spatially structured representation
  - Each query embedding corresponds to a spatial location in the final prediction
- **Fusion blocks:** Features from different levels are progressively fused using:
  - Residual connections to preserve information flow
  - A combination of convolutions ( $3 \times 3$ ) for local context and self-attention for global context
  - Layer normalization between operations
- **Depth head:** The final fused features are processed by a lightweight MLP that outputs per-pixel depth values.

This architecture better leverages the transformer’s inherent global context modeling while maintaining spatial coherence through the structured query embeddings. Unlike conventional CNN-based upsampling decoders, it preserves the transformer’s ability to model long-range dependencies throughout the decoding process.

**Training Details.** We train on the official NYU Depth v2 dataset using an input resolution of  $640 \times 480$  pixels. The model is optimized using:

- Adam optimizer with learning rate  $1 \times 10^{-4}$
- Cosine learning rate schedule with 2000-step warmup
- Batch size of 8
- Training for 25 epochs
- Loss function combining L1 (weight 0.5), scale-invariant (weight 1.0), and edge-aware gradient loss (weight 0.5)
- Data augmentation including random horizontal flipping, color jittering (brightness, contrast, saturation), and random cropping to  $416 \times 352$  pixels

Our edge-aware gradient loss specifically helps preserve structural boundaries by penalizing depth discontinuities that don't align with RGB edges, which is particularly beneficial for transformer-based models that might otherwise produce overly smooth depth maps.

Evaluation metrics include relative error (rel), root mean squared error (rms), and the percentage of predicted depths within thresholds of  $1.25$ ,  $1.25^2$ , and  $1.25^3$  compared to ground truth, following standard protocols in depth estimation literature.

## G Limitations

We identify several limitations of the Adaptive Superpixel Coding that may guide future research:

**Hyperparameter sensitivity.** The grouping behavior is influenced by the learnable gating threshold  $\theta$ , which determines the sparsity of the similarity graph. While effective empirically, its interpretability and generalization across datasets remain under-explored.

**Merge policy and representation bias.** Connected components are merged using mean pooling, which assumes uniformity within each group. This may suppress distinctive or minority features within an object, especially near boundaries or in regions with semantic heterogeneity. More expressive aggregation mechanisms may yield richer object-level embeddings.

**Computational overhead.** The module introduces additional complexity per layer via  $\mathcal{O}(N^2d)$  similarity computation and DFS traversal. While the reduced token count in subsequent layers may compensate for this overhead, the overall trade-off remains unquantified and warrants empirical runtime analysis.

**Object granularity.** The notion of an object is inferred solely from local pairwise feature similarities, which can lead to over- or under-grouping in cluttered scenes or under occlusion. Without global context or top-down constraints, the grouping may not align with semantically meaningful object boundaries.

**Layer-wise inconsistency.** Grouping is performed independently at each layer, with no mechanism to enforce consistency or track object identity across layers. This can result in unstable representations, particularly in deep architectures. Incorporating recurrent or hierarchical grouping could improve temporal coherence.

**Symmetry and object independence.** The module does not explicitly enforce object symmetry or viewpoint invariance. As a result, the same object under different poses, orientations, or occlusions may be fragmented into separate groups. Achieving object-level invariance in grouping remains a key challenge for generalizing to complex real-world scenes.