

SILC: Improving Vision Language Pretraining with Self-Distillation

Muhammad Ferjad Naeem^{1*} Yongqin Xian^{2*} Xiaohua Zhai^{3*}
Lukas Hoyer^{1,2^o} Luc Van Gool¹ Federico Tombari^{2,4}

¹ETH Zurich ²Google ³Google Deepmind ⁴TU Munich

Abstract. Image-Text pretraining on web-scale image caption datasets has become the default recipe for open vocabulary classification and retrieval models thanks to the success of CLIP and its variants. Several works have also used CLIP features for dense prediction tasks and have shown the emergence of open-set abilities. However, the contrastive objective used by these models only focuses on image-text alignment and does not incentivise image feature learning for dense prediction tasks. In this work, we introduce SILC, a novel framework for vision language pretraining. SILC improves image-text contrastive learning with the simple addition of local-to-global correspondence learning by self-distillation. We show that distilling local image features from an exponential moving average (EMA) teacher model significantly improves model performance on dense prediction tasks like detection and segmentation, while also providing improvements on image-level tasks such as classification and retrieval. SILC models sets a new state of the art for zero-shot classification, few shot classification, image and text retrieval, zero-shot segmentation, and open vocabulary segmentation. We further show that SILC features greatly benefit open vocabulary detection, captioning and visual question answering.

Keywords: Vision Language Models · Self-supervised Learning · Open-Vocabulary

1 Introduction.

Recent advancements in self-supervised learning [9, 11, 23, 45] and weakly supervised learning on web data [26, 47, 64] has spearheaded the development of foundational language [14, 48] and vision-language models [26, 47, 64]. These methods get around the long term challenge of obtaining large labelled dataset by developing self-supervision objectives. Developing open vocabulary computer vision models that can reason beyond a pre-determined set of classes has been

★ Research Consultant with Google, * Equal advising, ○ Intern at Google during the project.

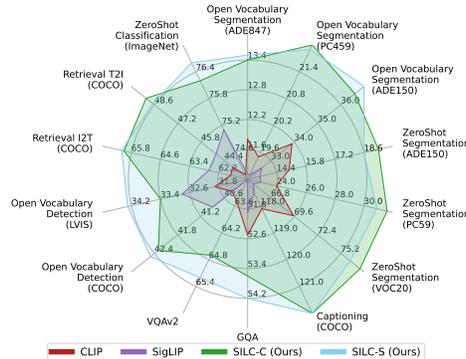


Fig. 1: SILC improves image-text contrastive learning with the addition of local-to-global correspondence learning by self-distillation. As a result, SILC models learn more locally aware visual features that are also grounded in language. SILC models offer significant improvements over CLIP (WebLI) and SigLIP over a wide variety of computer vision tasks including classification, segmentation, detection, captioning, VQA and retrieval.

a long-term challenge. The introduction of web image-text datasets and the progress in compute have enabled significant advances in this field. Popularized by CLIP [47], contrastive pretraining utilizes large datasets with paired image and text from the web and trains a vision-language model (VLM) to embed them to a shared latent space. Since these models are trained on a wide set of concepts, the learned VLM allows for open vocabulary inference [47]. However, developing open vocabulary dense prediction models for segmentation and detection is still an open challenge, since internet-scale datasets do not have dense pixel-level labels. Several works have found that incorporating VLMs in segmentation and detection models can unlock some open vocabulary abilities [13, 16, 27, 57, 62]. Since CLIP is not trained for these tasks, these methods get around its limitations by tuning the learned model with some dense prediction labelled dataset. One set of methods utilizes a normal segmentation / detection model for class agnostic inference and then predict the class logits with CLIP [13, 34]. Another family of methods aims to distill VLMs directly into a dense prediction model and utilize the text transformer to generate the class weights to predict logits [21, 28]. These works have been highly impactful towards expanding open vocabulary abilities of dense prediction models. However, since the contrastive pretraining objective does not explicitly encourage learning good local features for dense prediction tasks, these methods are limited by the VLM’s intrinsic performance [45] as we also show later in our experiments.

In the self-supervised literature, enforcing local-to-global consistency by self-distillation has emerged as a powerful pretraining objective [9, 45, 73] to learn vision backbones that are competitive on classification as well as dense prediction tasks, e.g. segmentation and detection. However, these backbones can not directly be used for zero-shot or open vocabulary inference as they do not contain any notion of class or language in the model. In this work, we propose SILC, which combines the advantages of these two branches and unifies image-text

contrastive pretraining and local-to-global consistency learning. SILC utilises a web image-text dataset to learn one model that improves VLM performance on existing classification and retrieval tasks while especially improving performance on zero-shot and open vocabulary segmentation, open vocabulary detection, captioning and Visual Question Answering (VQA).

Our contributions are as follows: 1. We propose a novel training framework for VLMs that pairs contrastive pretraining on image-text data with self-distillation on web images. 2. While conceptually very simple, we show that by learning stronger visual features with better local understanding, SILC models offer consistent improvements on multitude of computer vision tasks. These improvements are especially apparent on tasks that require better local understanding including zero-shot segmentation, open vocabulary segmentation, open vocabulary detection, captioning and Visual Question Answering (VQA). 3. We contribute a new foundation model that sets a new state of the art on zero-shot classification, few-shot classification, image-to-text and text-to-image retrieval, zero-shot semantic segmentation and open vocabulary semantic segmentation.

2 Related Works.

Image-Text Pretraining. Vision-language model (VLM) pretraining [12, 26, 30, 47] aims to learn generic multimodal representations that generalize to a wide range of downstream tasks. Substantial progress has been made in this field towards better pretraining objectives [26, 55] and better large-scale image-text dataset [12, 47]. One of the most popular objective functions is contrastive learning [26, 47] that pulls positive image and text pairs close and pushes negative ones apart in the joint embedding space. It is capable of scaling to a large-scale pretraining dataset and learning highly discriminative image and text features. Many works [20, 33, 42, 43, 52, 60, 64, 65] in this direction have demonstrated improvements across zero-shot image classification and retrieval benchmarks.

Another line of research focuses on generative learning via autoregressive text generation [53–55]. Compared to the contrastive learning, generative learning usually performs better on text generation tasks e.g., image captioning and VQA. Finally, there are hybrid methods [1, 30, 31, 35, 51, 61] that combine multiple objective functions including generative, contrastive and multi-task losses. While many VLMs [47, 55] mainly focus on learning global image-text alignment that benefit image-level downstream tasks, our work aims to develop a new VLM that benefits both image-level and pixel-level tasks. There have been a few attempts [17, 18, 36, 68] to improve VLMs for dense prediction tasks including object detection and semantic segmentation. However, they are either modeling the fine-grained patch-text interactions that are not scalable [18, 36] or rely on additional bounding box annotations [32, 68]. In this work we propose to pair image-text contrastive learning with self-distillation to learn a VLM.

Self-supervised Learning. Self-supervised learning is another popular pre-training paradigm where features are learned from image data itself. One branch of methods optimize the network to solve pretext tasks e.g., image coloring [67],

inpainting [46], transformation prediction [22], and patch ordering [38]. Another family of approaches adopt instance-level discriminative learning via contrastive learning [11, 25] and clustering [7, 8]. Recently, [24] shows that masked auto-encoder is also a scalable self-supervised learner. Several works have tried incorporating SSL in VLM pretraining with SimCLR loss [11, 40], distilling image captions [72] or distilling from masked images [17]. However, these works have not incorporated the most promising local-to-global consistency learning paradigm from this literature. We found that this requires careful design as outlined in our method section. Comparisons with the strongest baseline MaskCLIP [17] are available in Supplementary. Our work is inspired by DINO [9] which shows that segmentation emerges from learning local and global-views consistency. However, DINO cannot be directly used for zero-shot and open-vocabulary inference because it only learns image features. In contrast, our method is trained on image and text data jointly. We show that together with text data, the DINO objective allows the model to develop an understanding of local features and their semantic classes. Therefore our model can potentially directly benefit far more computer-vision applications.

Zero-shot Semantic Segmentation. Zero-shot semantic segmentation aims to segment arbitrary visual concepts in the wild without dense annotations [56]. Methods in this area rely on image-text pairs from a combination of image captioning and web image-text dataset. Since these datasets do not have dense labels, they utilize a self-supervised image region to text attention criterion. Group-VIT [56] proposes to introduce grouping tokens that cluster similar image patches under each group token. MaskCLIP [71] and CLIPpy [49] found that normal CLIP training results in zero-shot segmentation emerging. ReCo [50] proposes a refinement process on top of MaskCLIP by retrieval and co-segmentation. Finally, the current state-of-the-art TCL [10] learns a decoder to upsample the grounded patch embeddings and learns a region to text attention.

Open Vocabulary Segmentation and Detection. Open-vocabulary semantic segmentation methods aim to segment images according to a vocabulary of class categories provided at test-time containing additional unseen classes. In contrast to zero-shot segmentation, open-vocabulary semantic segmentation has access to a semantic segmentation dataset with a limited vocabulary for training. Recent methods transfer the open-vocabulary capabilities of CLIP from image- to pixel-level predictions. LSeg [29] learns pixel-wise visual embeddings that align with CLIP text embeddings while OpenSeg [21] learns class-agnostic segmentation proposals to pool visual features for region-text grounding. ZegFormer [16] and ZSseg [57] introduce a two-stage framework, which first learns class-agnostic segmentation mask predictions and classifies the corresponding region using a frozen CLIP. OVSeg [34] further finetunes CLIP on region-text pairs to compensate for the appearance shift of masked crops. To avoid the overhead of two stages, CAT-Seg [13] learns the aggregation of cost volumes between text embeddings and dense image embeddings from CLIP. Towards open vocabulary detection one family of methods, e.g. OWVLv2 [37], RegionCLIP [68], Detic [74], 3Ways [2], pseudolabel boxes for image caption data to use for localization pre-

training. An orthogonal family of methods including [32, 58, 59, 66] pretrain models to align class agnostic pseudoboxes to text as pretraining.

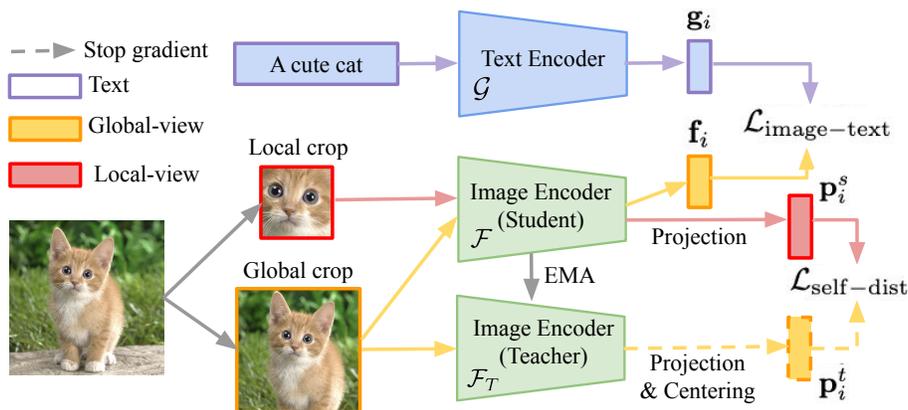


Fig. 2: **SILC** is a two-tower transformer based VLM. The first component of our training objective uses a global view of an image covering a large area and its paired caption to optimise a batch-wise contrastive loss for images and texts. The second component of our training objective enforces local-to-global consistency by self-distillation between the main model (the student) and an Exponential Moving Average (EMA)-based teacher. This local-to-global correspondence additionally allows the model to learn good visual features. Together the two objectives allow the model to excel at both traditional VLM tasks as well as tasks that require local understanding like segmentation and detection.

3 Method.

SILC builds on the contrastive pretraining framework of CLIP [47] and SigLIP [64]. SILC consists of a two-tower transformer model with a shared embedding space. We utilize a web-scale paired image-text dataset and rely on large-scale pretraining to learn the weights of the model. The first component of our pretraining objective focuses on aligning matching image-text pairs close together and away from other images and texts in the batch. This objective has been incredibly successful in recent literature [47, 64]. However, the contrastive objective in its current form does not focus on capturing rich local image semantics necessary for dense prediction tasks like segmentation and detection. Therefore, we propose to pair the contrastive pretraining objective with a local-to-global consistency objective that uses self-distillation as shown in Figure 2. **SILC** gets its name from the two training objectives consisting of **Self-Distillation from Images and Image-Language Contrastive Alignment from Image-Text** pairs.

3.1 Aligning Image and Text.

The contrastive pretraining objective relies on the Info-NCE framework [44]. It utilizes large amount of web-scale image-text dataset to learn an alignment between paired image and text. Given a minibatch $\mathcal{B} = \{(I_1, T_1), (I_2, T_2), \dots\}$, where (I_i, T_i) denotes a matching pair of image and text, the contrastive objective encourages matching image and text pairs to lie close together in a shared embedding space. The image I_i is processed by a learnable Vision Transformer \mathcal{F} to get its feature embedding. Similarly, the tokenized text T_i is processed by a learnable Text Transformer \mathcal{G} to get its feature embedding. These feature embeddings are normalized by their l_2 norm to get $\mathbf{f}_i = \frac{\mathcal{F}(I_i)}{\|\mathcal{F}(I_i)\|_2} \in \mathbb{R}^J$ for the image I_i and $\mathbf{g}_i = \frac{\mathcal{G}(T_i)}{\|\mathcal{G}(T_i)\|_2} \in \mathbb{R}^J$ for the paired text T_i where J is the feature dimension of the shared embedding space. The dot product of \mathbf{f}_i and \mathbf{g}_i computes their cosine similarity and is optimized with a pair of cross-entropy losses as proposed by CLIP [47] or a sigmoid loss as proposed by SigLIP [64]. The batch-wise contrastive losses of CLIP/ SigLIP, represented as $\mathcal{L}_{\text{image-text}}$, rely on a large batch size to align image-text pairs. This objective tuned over a large amount of data learns a shared embedding space between image and text and thus can be used for zero-shot transfer to multitude of computer vision tasks.

3.2 Distilling Local Image Features.

The image-text contrastive loss has shown to be very successful in learning zero-shot transfer models [26, 47]. Models learned with this objective have also been used to improve dense prediction tasks like open vocabulary segmentation and detection. However, the contrastive objective alone does not explicitly focus on learning good visual features for dense prediction tasks. These tasks require local image semantics to be sufficiently encoded in the output image and patch embeddings. Enforcing local-to-global consistency has emerged as a powerful technique to accomplish this on large unlabelled image data [9, 45, 73] in self-supervision literature. However, these methods can not be directly used for open vocabulary models as they are trained without any language information. In the second component of our training framework, we take inspiration from this subset of literature and additionally add local-to-global consistency as a training objective for images in our image-text dataset.

The basic idea of this objective is as follows. A teacher network gets a global view of the image representing the scene as a whole and produces a feature embedding. A student model gets a partial view of the same image and produces a feature embedding. A self-distillation objective is introduced where the student needs to match the prediction of the teacher while only having partial information. This enforces the model to learn local semantics and their relation to global semantics of the scene. We add this criterion for the image encoder \mathcal{F} . We add a projection as a learnable MLP on top of the image encoder to map from the original shared embedding space of dimension J to K where $K > J$. The student \mathcal{F}_S is the main image encoder with a learnable projection head. Since

we rely on noisy web scale image-text data, we do not have an oracle teacher for the student to match. We therefore construct our teacher \mathcal{F}_T as an exponential moving average of the student \mathcal{F}_S from the previous training iterations to realize our self-distillation framework:

$$\mathcal{F}_T \leftarrow \lambda \mathcal{F}_T + (1 - \lambda) \mathcal{F}_S, \quad (1)$$

where λ controls the update step of the teacher. For a given image I_i , the teacher processes its global crop to produce $\mathbf{p}_i^t \in \mathbb{R}^K$ and the student processes its local crop to produce $\mathbf{p}_i^s \in \mathbb{R}^K$. To prevent the teacher from collapsing to a trivial solution, we apply sharpening on the outputs of teacher with τ_t and student with τ_s . To encourage each feature dimension to contribute to the output feature, we additionally introduce a centering operation on the prediction of the teacher. The centering term $\mathbf{c} \in \mathbb{R}^K$ is initialized with 0 and is updated by a momentum update with a factor of m with the first order batch statistics of the teacher’s prediction at each step as follows: $\mathbf{c} \leftarrow m\mathbf{c} + (1 - m) \frac{1}{|\mathcal{B}|} \sum_{i=1}^{|\mathcal{B}|} \mathbf{p}_i^t$.

To learn local-to-global correspondences, the student is faced with an information asymmetry. The student is given a local view of an image which is realized as a random crop over a small region of the image. The teacher, however, has access to a global view of the image containing more information about the scene. The student is tasked with matching the semantics of the teacher while only having partial information. Therefore, for a given image, the model needs to learn local semantics of the image and how it would fit in the global context of this image. This is realized as a knowledge-distillation loss where the student and the teacher’s feature vectors are first converted to a probability distribution by applying a softmax on the teacher prediction $\mathcal{P}_t(I_i^{gl}) = \text{softmax}((\mathbf{p}_i^t - \mathbf{c})/\tau_t)$ and student prediction $\mathcal{P}_s(I_i^{lc}) = \text{softmax}(\mathbf{p}_i^s/\tau_s)$. The student is optimized to match the teacher with a cross-entropy loss,

$$\mathcal{L}_{\text{self-dist}} = -\mathcal{P}_t(I_i^{gl}) \top \log(\mathcal{P}_s(I_i^{lc})). \quad (2)$$

This self-distillation objective incentivises the image encoder to learn local semantics of images over the large web scale dataset. Since the teacher is constructed with the student’s weights, and the image level features are pooled from patch embeddings in a Vision Transformer, this allows for richer local semantics to be captured in the image level as well as the patch level features.

While this objective has been explored in self-supervised learning [9, 45], to the best of our knowledge, we are the first work to show its complimentary nature to image-text contrastive learning on web-scale dataset. We show that when combined with text, this objective allows the model to develop a local understanding of the semantics of an image grounded in language. We find two important modifications compared to previous works that allows it to be complimentary to image-text contrastive learning. 1. Each global view used in $\mathcal{L}_{\text{self-dist}}$ needs to be aligned with text, otherwise the two objectives diverge. This is realized by computing the image-text contrastive loss for each global view while maintaining the same batch size. 2. The momentum scheduler of the EMA should not converge to 1.0. Otherwise the teacher stops learning from image-text loss as

Model	Zero-Shot Classification		Few-shot classification						Retrieval	
	ImageNet	CIFAR100	ImageNet			CIFAR100			COCO	
	T1	T1	1shot	5shot	10shot	1shot	5shot	10shot	I2T@1	T2I@1
CLIP (WebLI) [64]	74.1	68.4	42.8	63.2	67.3	39.4	59.6	64.6	61.7	43.9
SILC-C* (Ours)	<u>75.3</u>	<u>71.0</u>	<u>44.6</u>	<u>64.3</u>	<u>67.8</u>	<u>42.8</u>	<u>64.6</u>	<u>69.6</u>	<u>62.5</u>	<u>44.9</u>
SILC-C (Ours)	76.2	72.3	45.3	65.0	68.5	45.2	66.9	71.3	66.1	49.1
SigLIP [64]	75.1	<u>69.8</u>	44.0	64.2	68.4	39.0	61.7	66.3	62.6	<u>44.9</u>
SILC-S* (Ours)	<u>75.8</u>	69.2	<u>45.2</u>	<u>64.6</u>	<u>68.4</u>	<u>40.3</u>	<u>63.3</u>	<u>67.4</u>	<u>63.0</u>	44.6
SILC-S (Ours)	76.6	70.6	45.9	65.2	68.9	41.8	64.9	68.9	66.2	48.7

Table 1: Comparing SILC* with baselines, we observe that our pretraining framework results in a significant improvement over both CLIP and SigLIP objectives. We reproduce both CLIP and SigLIP on the same WebLI dataset [12] to quantify the improvements from our proposed training objective. We further finetune SILC* on a cleaner subset to get our final model SILC and see that it unlocks additional performance without significant extra retraining. The best performance for each variant is **bolded**, the second best is underlined.

the update step becomes too small in the later stage of the training. We therefore use a fixed momentum.

4 Experiments.

We compare our SILC pretraining framework with both CLIP [47] and SigLIP [64] on the same test bench and perform extensive experimentation. SILC models based on the CLIP objective are represented by **SILC-C** and the SigLIP versions are represented by **SILC-S**. We show that SILC sets a new state of the art on a variety of tasks: zero-shot classification, few-shot classification, retrieval, zero-shot segmentation and open vocabulary segmentation. We further show that SILC models also improve other local semantic understanding tasks including open vocabulary detection, captioning and VQA.

4.1 Implementation Details.

We implement our model in jax in the `big_vision` codebase [4, 5], following the contrastive pretraining setups from [64], and use the WebLI dataset [12] for our experiments. We utilize two global views cropped between (0.4 – 1.0) of the original image area and eight local views cropped between (0.05 – 0.4) of the original image area for the self-distillation loss. The global views are resized to (256 × 256) and the local views are resized to (96 × 96). The teacher momentum λ is kept fixed at 0.966 and the center update momentum m is kept fixed at 0.9 through the training. The teacher temperature τ_t is fixed at 0.04 and the student temperature τ_s is fixed at 0.1. K is 65536. We resize the original image to (256 × 256) for the contrastive loss between image-text pairs. We trained with a batch size of 16k on Google TPUs. We use *example-seen* to represent how many image and text pairs are drawn from the dataset throughout the training. We train all baselines in our main comparisons in Table 1 for 20 Billion example-seen on the WebLI dataset [12] following [64]. Our models trained on WebLI

are marked as **SILC***. We use a rsqrt learning scheduler [63] with base learning rate of 0.001 with 50000 warm up and 50000 cooldown steps. Additional training details including compute cost are discussed in the supplementary.

We additionally finetune our model using a smaller but cleaner WebLI subset [12] for 1 Billion additional example-seen and represent this model as **SILC**. The smaller WebLI subset contains 100 million image-text pairs with finer-grained text filters etc.

4.2 Classification and Retrieval.

We compare our pretraining framework with CLIP and SigLIP under the same training and evaluation protocol in Table 1. We compare at ViT/B16 and see that the introduction of self-distillation to both consistently improve their performance on zero-shot classification, few shot classification and retrieval. On zero-shot classification on ImageNet, SILC-C* improves on CLIP (WebLI) by 1.2 points, similarly we notice an improvement of 2.6 points on CIFAR-100 showing the benefit of local feature self-distillation. Similar improvements are noted for few-shot classification where SILC-C* improves over CLIP (WebLI) by 1.8, 1.1 and 0.5 points on ImageNet 1 shot, 5 shot and 10 shot classification respectively. We make similar observation on retrieval where SILC-C* shows improvements on image to text as well as text to image retrieval. Moving to SigLIP versions of the model, we see a similar trend where the introduction of self-distillation objective allows SILC-S* to consistently improve almost all metrics over the evaluated tasks. We therefore conclude that capturing better local semantics results in learning stronger visual features which also helps tasks that require global understanding of the image.

Comparing SILC* models with SILC, we notice that the finetuning on the cleaner subset unlocks additional performance for the model without significant extra training. For the CLIP based SILC-C, We notice another 0.9 point improvement over SILC-C* on zero-shot ImageNet classification. We observe improvements of the same magnitude on few-shot classification. Comparing retrieval performance, we see a significant increase in retrieval performance on COCO where SILC-C achieves a 3.6 and 4.2 points improvement on Image to Text and Text to Image Recall@1. The SigLIP based SILC-S follows a similar trend and consistently improves on SILC-S* on all metrics. SILC models set a new state-of-the-art for these tasks at ViT/B16 model size. We additionally provide system level comparison with open-source CLIP variants in Supplementary Table 1, and performance of SILC trained with open source datasets (LAION and Conceptual Captions) in Supplementary Table 2. We show that SILC provides consistent improvements over all CLIP variants.

4.3 Zero-Shot Semantic Segmentation.

Zero-shot semantic segmentation aims to measure the grounding performance of a VLM usually from its patch embeddings. MaskCLIP [71] and CLIPpy [49] found that this grounding naturally emerges as a consequence of image-text

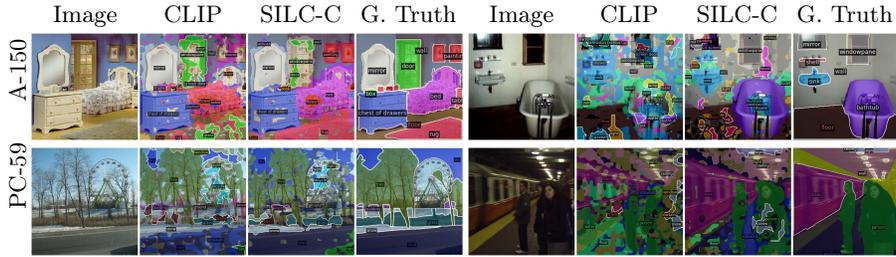


Fig. 3: Qualitative results on zero-shot segmentation show that SILC-C achieves significant improvements over CLIP (WebLI). SILC-C produces less noisy segmentation and better distinguishes semantic classes. This semantic segmentation emerges without any segmentation supervision.

Model	A-150	PC-59	Cityscapes	VOC-20	COCO-Stuff
GroupViT [56]	9.2	23.4	11.1	79.7	11.1
MaskCLIP [71]	9.8	26.4	12.6	74.9	16.4
ReCo [50]	11.2	22.3	21.1	57.7	14.8
TCL [10]	14.9	30.3	23.1	77.5	19.6
CLIP (WebLI) [64]	15.0	24.0	22.6	69.5	15.0
SILC-C* (Ours)	17.2	29.3	25.1	73.5	18.2
SILC-C (Ours)	19.3	31.6	26.9	77.5	20.8
SigLIP [64]	13.6	22.9	20.8	64.7	13.4
SILC-S* (Ours)	16.7	28.6	23.4	72.1	17.3
SILC-S (Ours)	18.6	30.9	25.2	76.3	19.7

Table 2: Comparing Zero-Shot Segmentation performance we see that SILC* models trained on noisy web image-text data already outperform several ZS segmentation baselines that use cleaner image-text data. When we tune our model on a cleaner subset of image-text data to get SILC-C, we see that it sets the absolute state-of-the-art on 4 out of 5 datasets.

contrastive training. We use a Vision Transformer with a MAP pooling head [63]. We observe that grounding for our model emerges in the values of the MAP head instead of the last encoder block. For a given set of possible classes in a segmentation dataset, we obtain the corresponding text embeddings by querying our text encoder with a standard prompt. We compute the cosine similarity between the image patch embeddings and the text features of each class name to generate a segmentation map in zero-shot. We report the mean-IOU (mIOU) performance of our model in Table 2 and compare with baselines at ViT/B16 similar to previous works. We follow the evaluation protocol of TCL [10] without the background class. However, we do not use any post-refinement e.g. PAMR as we argue that the raw segmentation of a VLM is the true depiction of its zero-shot performance.

Comparing against CLIP and SigLIP, we see that both SILC-C* and SILC-S* show significantly superior zero-shot semantic segmentation performance. In fact, both variants achieve multiple mIOU points improvements over all 5 datasets. This validates our hypothesis that the combination of image-text con-

VLM	Method	A-847	PC-459	A-150	PC-59	VOC-20	VOC-21
CLIP-B/16	ZegFormer [16]	5.6	10.4	18.0	45.5	89.5	65.5
CLIP-B/16	ZSseg [57]	7.0	-	20.5	47.7	88.4	-
CLIP-B/16	OVSeg [34]	7.1	11.0	24.8	53.3	92.6	-
CLIP-B/16	CAT-Seg [13]	8.4	16.6	27.2	57.5	93.7	78.3
SILC-C-B/16	CAT-Seg [13]	13.4 (+5.0)	22.0 (+5.4)	36.6 (+9.4)	61.2 (+3.7)	95.9 (+2.2)	80.4 (+2.1)
SILC-S-B/16	CAT-Seg [13]	13.5 (+5.1)	21.9 (+5.3)	37.0 (+9.8)	61.2 (+3.7)	96.1 (+2.4)	80.9 (+2.6)
CLIP-L/14	ZSseg [57]	7.1	10.2	21.7	52.2	92.3	-
CLIP-L/14	OVSeg [34]	9.0	12.4	29.6	55.7	94.5	-
CLIP-L/14	CAT-Seg [13]	10.8	20.4	31.5	62.0	96.6	81.8
SILC-C-L/16	CAT-Seg [13]	15.0 (+4.2)	25.8 (+5.4)	37.7 (+6.2)	63.5 (+1.5)	97.6 (+1.0)	82.5 (+0.7)
CLIP-G/14	CAT-Seg [13]	13.3	21.4	36.2	61.5	97.1	81.4

Table 3: Comparing Open Vocabulary Semantic Segmentation performance, we observe that SILC models improve over CLIP by significant margins on all unseen test sets. SILC particularly improves the performance for challenging test sets with large vocabularies. SILC-L/16 even outperforms the much larger CLIP-G/14. All models are trained on COCO-Stuff.

trastive learning and local-to-global correspondence learning allows the model to develop better understanding of local semantics of the image grounded in language. From Table 2, we observe that the CLIP objective in general results in superior zero-shot segmentation than the SigLIP objective. This is also apparent as we compare SILC-C* with SILC-S*. Moreover, we observe that finetuning on a cleaner subset further improves the zero-shot segmentation performance of both SILC-C and SILC-S. We observe that the CLIP variant SILC-C also outperforms SILC-S here. We show the improvements of SILC on CLIP (WebLI) qualitatively in Figure 3. We can observe that SILC is better at segmenting and labeling semantic classes in images. We would like to emphasize that SILC achieves this without any segmentation ground truth.

Comparing with current SOTA, we observe that SILC-C consistently beats all specialized zero-shot segmentation baselines on 4/5 datasets to set a new state of the art (SOTA). Compared to the previous state of the art TCL, SILC achieves a remarkable 4.3 mIOU points improvement on A-150 [69, 70], 2.9 points improvement on PC-59 [39], and 4.9 points improvement on CityScapes [15]. Similar improvements are noted on VOC-20 [19] and COCO-Stuff [6], however GroupViT maintains the best result on VOC-20. These methods use relatively cleaner image captioning datasets for their segmentation specific training objectives. We noticed that the improvements in zero-shot segmentation are achievable by just finetuning on a cleaner subset of data. We did not observe superior performance by learning an expensive patch-wise attention as proposed by PACL [41]. We show in the supplementary that SILC models also outperform PACL trained on WebLI. Methods like TCL and ReCo, which are designed to improve the zero-shot segmentation performance of a frozen VLM, can in theory further improve the performance of our model. However, since we aim to improve vision-language pretraining over all tasks, this is out of the scope of this work.

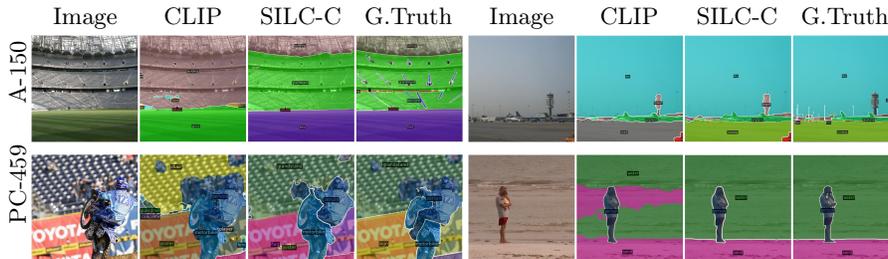


Fig. 4: Comparing qualitative examples for open vocabulary segmentation, we observe that SILCw/ CAT-Seg better distinguishes semantically similar classes such as field/grass, runway/road, grandstand/chair and sand/water than CLIP.

Initialization	Training data	COCO		LVIS	
		AP	AP _{all}	AP _{rare}	
CLIP (WebLI)	WebLI N-grams	40.4	31.9	29.2	
SILC-C*(Ours)	WebLI N-grams	41.8	33.3	30.4	
SigLIP	WebLI N-grams	40.9	32.8	30.4	
SILC-S*(Ours)	WebLI N-grams	42.7	34.2	32.4	

Table 4: Training OWLv2 for Object Detection with SILC models offers consistent improvement over CLIP and SigLIP for open vocabulary object detection. These models are trained with pseudo labels from WebLI N-grams [37] and evaluated zero-shot on COCO and LVIS.

4.4 Open-Vocabulary Semantic Segmentation.

Open Vocabulary Semantic Segmentation aims to develop segmentation models that can segment novel classes beyond the training vocabulary. Most recent methods in this area rely on a pretrained CLIP due to its open-vocabulary capabilities and adapt it for segmentation task. To evaluate the open vocabulary segmentation potential of SILC, we take the current state-of-the-art model CAT-Seg [13] and replace the CLIP model used by the authors with SILC. The models are trained on COCO-Stuff-164k with 172 classes and tested on unseen datasets with different vocabularies: ADE-20k with 847 or 150 classes (A-847/A-150), Pascal Context (PC-459/PC-59), and Pascal VOC (VOC-20/VOC-21).

From Table 3, we observe that SILC significantly improves over CLIP [47]. In fact, SILC-C-B/16 performs on par with the much bigger CLIP-G/14 on the three most challenging test datasets A-847, PC-459 and A-150. Moreover, we observe that while SILC-S performed slightly worse than SILC-C in zero-shot segmentation, it achieves slightly better performance when trained for open vocabulary segmentation. SILC-S-B/16 further improve on the performance of SILC-C-B/16. The observed improvements of SILC-C also transfer to the larger ViT-L variant, where CAT-Seg with SILC-C-L/16 outperforms CAT-Seg with CLIP-L/14 on all datasets by a significant margin. In particular, it achieves more than +4 mIOU improvement on the challenging A-847, PC-459, and A-150. SILC-L/16 even significantly outperforms the much bigger CLIP-G/14 on all tested datasets. The improvements of SILC-C over CLIP are also reflected

Model	Classification		Captioning	Question Ans.	
	ImageNet	SUN397	COCO	GQA	VQAv2
CLIP (WebLI)	82.3	82.4	118.1	52.5	63.5
SILC-C*(Ours)	83.8	83.4	120.8	53.1	64.6
SigLIP	82.5	82.2	117.5	51.9	63.0
SILC-S*(Ours)	83.7	82.9	121.2	53.2	64.5

Table 5: Evaluating SILC visual representation with LiT-Decoder in a multi-task setup, we observe consistent improvements on all tasks compared to CLIP and SigLIP. These improvements are especially apparent for tasks that require local understanding of the image i.e. Captioning and Question Answering.

Model	ImageNet 0 shot	ImageNet Few shot			COCO Retrieval		ZS Segmentation			Open Vocab Seg		
	T1	1shot	5shot	10shot	I2T@1	T2I@1	A-150	Stuff	PC-59	PC-459	A-150	PC-59
CLIP (WebLI)	71.7	36.4	57.7	62.5	59.1	42.9	11.8	12.9	20.1	18.6	30.5	57.7
+ additional views	73.6	38.7	60.8	65.7	60.6	43.2	11.7	13.0	20.0	19.2	32.1	57.8
+ EMA	73.7	38.4	60.7	65.5	61.3	43.1	11.9	13.3	20.5	19.0	32.2	57.5
+ Self Dist (SILC-C*)	74.3	39.9	61.2	65.7	62.7	43.9	12.2	15.3	21.1	21.0	33.3	60.7

Table 6: We ablate over each component of our model to verify our design choices. The addition of image augmentation and EMA to CLIP (WebLI) improves classification and retrieval metrics while only slightly impact the segmentation. Adding local-to-global consistency by self-distillation, we observe an improvement across the board especially on segmentation metrics.

in the qualitative examples in Fig. 4. We observe that SILC-C better distinguishes semantically similar classes such as grandstand/building, field/grass, runway/road and grandstand/chair. Further, it improves segmentation in difficult cases and better handles transparent segments as shown in supplementary. Results on additional WebLI models are also provided in supplementary with similar conclusions.

4.5 SILC for Open-Vocabulary Detection.

We utilize OWLv2 [37] as a framework to test the open vocabulary detection potential of SILC models. OWLv2 initializes a detection model with a contrastive image-text model’s weights and utilizes pseudo labelled boxes from WebLI (WebLI N-grams [37]) to learn an open vocabulary detection model. We utilize the test bench of the authors and retrain OWLv2 initialized from our VLM baselines to report results in Table 4. We evaluate these models zero-shot on COCO and LVIS without doing any finetuning on respective dataset to test their open vocabulary performance. From Table 4 we observe that SILC models also benefit open vocabulary detection thanks to learning better local semantics. SILC-S* achieves an improvement of +1.8AP on COCO. The improvements are also consistent on the challenging LVIS benchmark where SILC-S* achieves an improvement of +1.4AP on all classes and a remarkable +2.0AP on rare classes. We make similar observations as we compare SILC-C* with CLIP (WebLI) where SILC-C* offers consistent improvements. This further validates that SILC models offer better performance for dense tasks.

4.6 Evaluating SILC features with LiT-Decoder.

LiT-Decoder [3] proposes to utilise a frozen image encoder and train a single autoregressive decoder to learn a multi-task model for Classification, Captioning and Visual Question Answering. We use LiT-Decoder as a framework to evaluate the quality of visual representations learned by SILC models against baseline CLIP (WebLI) and SigLIP. We use the authors’ implementation and only replace the ViT with the respective baselines. We report results in Table 5. We observe that SILC models offer consistent improvements in this multi-task setup as well. Compared to SigLIP, SILC-S* improves classification by +1.2 points on ImageNet and +1.0 point on SUN397. The improvements are even more profound on captioning (+3.6 CIDEr score on COCO) and Visual Question Answering (+1.3 on GQA and +1.5 on VQAv2). These tasks greatly benefit from SILC features’ ability to better encode local semantics. Similar improvements are noted for CLIP based baselines. This further validates that SILC features can simultaneously benefit multiple computer vision problems.

4.7 Ablation on Model Components.

We ablate on the various design choices of our model and their impact on various tasks. We train all models for 5 Billion example-seen and report the performance in Table 6. Since our method processes additional image augmentations in the contrastive loss, we first test if our improvements are a consequence of processing more augmentations. We observe that the introduction of additional image augmentations (second row) improve the classification and retrieval metrics but their impact on zero-shot segmentation and open vocabulary segmentation is not as significant. When we add an EMA over this model’s weights similar to our model (third row), we notice a slight improvement as seen in previous SSL literature. Finally when we add the self-distillation from local crops, we see an improvement across the board on all tasks. In particular, we observe the strongest improvement on segmentation tasks highlighting our proposal’s impact on them.

5 Conclusion.

We propose to integrate local-to-global correspondence learning by self-distillation as a complementary objective to the popular VLM contrastive objective of CLIP [47] and SigLIP [64]. We show that the introduction of this results in remarkable performance improvements on several computer vision tasks. We see a consistent performance improvement on zero-shot classification, few-shot classification, and retrieval. We further test our VLM on zero-shot segmentation and show that our training framework results in significant improvements without using any dense ground truth. Finally we show that SILC models as pretrained backbones significantly improve a model’s performance on open vocabulary segmentation, open vocabulary detection, captioning and VQA. SILC models set a new state of the art in Vision-Language Foundational Models.

Bibliography

- [1] Alayrac, J.B., Donahue, J., Luc, P., Miech, A., Barr, I., Hasson, Y., Lenc, K., Mensch, A., Millican, K., Reynolds, M., et al.: Flamingo: a visual language model for few-shot learning. In: NeurIPS (2022) [3](#)
- [2] Arandjelović, R., Andonian, A., Mensch, A., Hénaff, O.J., Alayrac, J.B., Zisserman, A.: Three ways to improve feature alignment for open vocabulary detection. arXiv preprint arXiv:2303.13518 (2023) [4](#)
- [3] Beyer, L., Wan, B., Madan, G., Pavetic, F., Steiner, A., Kolesnikov, A., Pinto, A.S., Bugliarello, E., Wang, X., Yu, Q., et al.: A study of autoregressive decoders for multi-tasking in computer vision. arXiv preprint arXiv:2303.17376 (2023) [14](#)
- [4] Beyer, L., Zhai, X., Kolesnikov, A.: Better plain vit baselines for imagenet-1k (2022) [8](#)
- [5] Beyer, L., Zhai, X., Kolesnikov, A.: Big vision. https://github.com/google-research/big_vision (2022) [8](#)
- [6] Caesar, H., Uijlings, J., Ferrari, V.: Coco-stuff: Thing and stuff classes in context. In: CVPR (2018) [11](#)
- [7] Caron, M., Bojanowski, P., Joulin, A., Douze, M.: Deep clustering for unsupervised learning of visual features. In: ECCV (2018) [4](#)
- [8] Caron, M., Misra, I., Mairal, J., Goyal, P., Bojanowski, P., Joulin, A.: Unsupervised learning of visual features by contrasting cluster assignments. In: NeurIPS (2020) [4](#)
- [9] Caron, M., Touvron, H., Misra, I., Jégou, H., Mairal, J., Bojanowski, P., Joulin, A.: Emerging properties in self-supervised vision transformers. In: ICCV (2021) [1](#), [2](#), [4](#), [6](#), [7](#)
- [10] Cha, J., Mun, J., Roh, B.: Learning to generate text-grounded mask for open-world semantic segmentation from only image-text pairs. In: CVPR (2023) [4](#), [10](#)
- [11] Chen, T., Kornblith, S., Norouzi, M., Hinton, G.: A simple framework for contrastive learning of visual representations. In: ICML (2020) [1](#), [4](#)
- [12] Chen, X., Wang, X., Changpinyo, S., Piergiovanni, A., Padlewski, P., Salz, D., Goodman, S., Grycner, A., Mustafa, B., Beyer, L., Kolesnikov, A., Puigcerver, J., Ding, N., Rong, K., Akbari, H., Mishra, G., Xue, L., Thapliyal, A.V., Bradbury, J., Kuo, W., Seyedhosseini, M., Jia, C., Ayan, B.K., Ruiz, C.R., Steiner, A.P., Angelova, A., Zhai, X., Houlsby, N., Soricut, R.: PaLI: A jointly-scaled multilingual language-image model. In: ICLR (2023) [3](#), [8](#), [9](#)
- [13] Cho, S., Shin, H., Hong, S., An, S., Lee, S., Arnab, A., Seo, P.H., Kim, S.: Cat-seg: Cost aggregation for open-vocabulary semantic segmentation. arXiv preprint arXiv:2303.11797 (2023) [2](#), [4](#), [11](#), [12](#)
- [14] Chowdhery, A., Narang, S., Devlin, J., Bosma, M., Mishra, G., Roberts, A., Barham, P., Chung, H.W., Sutton, C., Gehrmann, S., et al.: Palm: Scaling

- language modeling with pathways. arXiv preprint arXiv:2204.02311 (2022) [1](#)
- [15] Cordts, M., Omran, M., Ramos, S., Rehfeld, T., Enzweiler, M., Benenson, R., Franke, U., Roth, S., Schiele, B.: The cityscapes dataset for semantic urban scene understanding. In: CVPR (2016) [11](#)
- [16] Ding, J., Xue, N., Xia, G.S., Dai, D.: Decoupling zero-shot semantic segmentation. In: CVPR (2022) [2](#), [4](#), [11](#)
- [17] Dong, X., Bao, J., Zheng, Y., Zhang, T., Chen, D., Yang, H., Zeng, M., Zhang, W., Yuan, L., Chen, D., et al.: Maskclip: Masked self-distillation advances contrastive language-image pretraining. In: CVPR (2023) [3](#), [4](#)
- [18] Dou, Z.Y., Kamath, A., Gan, Z., Zhang, P., Wang, J., Li, L., Liu, Z., Liu, C., LeCun, Y., Peng, N., et al.: Coarse-to-fine vision-language pre-training with fusion in the backbone. In: NeurIPS (2022) [3](#)
- [19] Everingham, M., Van Gool, L., Williams, C.K., Winn, J., Zisserman, A.: The pascal visual object classes (voc) challenge. IJCV **88** (2010) [11](#)
- [20] Fang, A., Jose, A.M., Jain, A., Schmidt, L., Toshev, A., Shankar, V.: Data filtering networks. arXiv preprint arXiv:2309.17425 (2023) [3](#)
- [21] Ghiasi, G., Gu, X., Cui, Y., Lin, T.Y.: Scaling open-vocabulary image segmentation with image-level labels. In: ECCV. Springer (2022) [2](#), [4](#)
- [22] Gidaris, S., Singh, P., Komodakis, N.: Unsupervised representation learning by predicting image rotations. In: ICLR (2018) [4](#)
- [23] Grill, J.B., Strub, F., Althé, F., Tallec, C., Richemond, P., Buchatskaya, E., Doersch, C., Avila Pires, B., Guo, Z., Gheshlaghi Azar, M., et al.: Bootstrap your own latent-a new approach to self-supervised learning. NeurIPS (2020) [1](#)
- [24] He, K., Chen, X., Xie, S., Li, Y., Dollár, P., Girshick, R.: Masked autoencoders are scalable vision learners. In: CVPR (2022) [4](#)
- [25] He, K., Fan, H., Wu, Y., Xie, S., Girshick, R.: Momentum contrast for unsupervised visual representation learning. In: CVPR (2020) [4](#)
- [26] Jia, C., Yang, Y., Xia, Y., Chen, Y.T., Parekh, Z., Pham, H., Le, Q., Sung, Y.H., Li, Z., Duerig, T.: Scaling up visual and vision-language representation learning with noisy text supervision. In: ICML (2021) [1](#), [3](#), [6](#)
- [27] Kuo, W., Cui, Y., Gu, X., Piergiovanni, A., Angelova, A.: F-vlm: Open-vocabulary object detection upon frozen vision and language models. ICLR (2023) [2](#)
- [28] Li, B., Weinberger, K.Q., Belongie, S., Koltun, V., Ranftl, R.: Language-driven semantic segmentation. In: International Conference on Learning Representations (2022), <https://openreview.net/forum?id=RriDjddCLN> [2](#)
- [29] Li, B., Weinberger, K.Q., Belongie, S., Koltun, V., Ranftl, R.: Language-driven semantic segmentation. In: International Conference on Learning Representations (2022), <https://openreview.net/forum?id=RriDjddCLN> [4](#)
- [30] Li, J., Li, D., Xiong, C., Hoi, S.: Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation. In: ICML (2022) [3](#)

- [31] Li, J., Selvaraju, R., Gotmare, A., Joty, S., Xiong, C., Hoi, S.C.H.: Align before fuse: Vision and language representation learning with momentum distillation. In: NeurIPS (2021) [3](#)
- [32] Li, L.H., Zhang, P., Zhang, H., Yang, J., Li, C., Zhong, Y., Wang, L., Yuan, L., Zhang, L., Hwang, J.N., et al.: Grounded language-image pre-training. In: CVPR (2022) [3](#), [5](#)
- [33] Li, Y., Liang, F., Zhao, L., Cui, Y., Ouyang, W., Shao, J., Yu, F., Yan, J.: Supervision exists everywhere: A data efficient contrastive language-image pre-training paradigm. In: ICML (2021) [3](#)
- [34] Liang, F., Wu, B., Dai, X., Li, K., Zhao, Y., Zhang, H., Zhang, P., Vajda, P., Marculescu, D.: Open-vocabulary semantic segmentation with mask-adapted clip. In: CVPR (2023) [2](#), [4](#), [11](#)
- [35] Lu, J., Batra, D., Parikh, D., Lee, S.: Vilbert: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks. In: NeurIPS (2019) [3](#)
- [36] Luo, H., Bao, J., Wu, Y., He, X., Li, T.: Segclip: Patch aggregation with learnable centers for open-vocabulary semantic segmentation. In: ICLR (2023) [3](#)
- [37] Minderer, M., Gritsenko, A., Hounsby, N.: Scaling open-vocabulary object detection. arXiv preprint arXiv:2306.09683 (2023) [4](#), [12](#), [13](#)
- [38] Misra, I., Maaten, L.v.d.: Self-supervised learning of pretext-invariant representations. In: CVPR (2020) [4](#)
- [39] Mottaghi, R., Chen, X., Liu, X., Cho, N.G., Lee, S.W., Fidler, S., Urtasun, R., Yuille, A.: The role of context for object detection and semantic segmentation in the wild. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2014) [11](#)
- [40] Mu, N., Kirillov, A., Wagner, D., Xie, S.: Slip: Self-supervision meets language-image pre-training. In: ECCV (2022) [4](#)
- [41] Mukhoti, J., Lin, T.Y., Poursaeed, O., Wang, R., Shah, A., Torr, P.H., Lim, S.N.: Open vocabulary semantic segmentation with patch aligned contrastive learning. In: CVPR (2023) [11](#)
- [42] Naeem, M.F., Khan, M.G.Z.A., Xian, Y., Afzal, M.Z., Stricker, D., Van Gool, L., Tombari, F.: I2mvformer: Large language model generated multi-view document supervision for zero-shot image classification. In: CVPR (2023) [3](#)
- [43] Naeem, M.F., Xian, Y., Gool, L.V., Tombari, F.: I2dformer: Learning image to document attention for zero-shot image classification. NeurIPS (2022) [3](#)
- [44] Oord, A.v.d., Li, Y., Vinyals, O.: Representation learning with contrastive predictive coding. arXiv preprint arXiv:1807.03748 (2018) [6](#)
- [45] Oquab, M., Darcet, T., Moutakanni, T., Vo, H., Szafraniec, M., Khaidov, 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) [1](#), [2](#), [6](#), [7](#)
- [46] Pathak, D., Krahenbuhl, P., Donahue, J., Darrell, T., Efros, A.A.: Context encoders: Feature learning by inpainting. In: CVPR (2016) [4](#)

- [47] Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al.: Learning transferable visual models from natural language supervision. In: ICLR (2021) [1](#), [2](#), [3](#), [5](#), [6](#), [8](#), [12](#), [14](#)
- [48] Radford, A., Narasimhan, K., Salimans, T., Sutskever, I., et al.: Improving language understanding by generative pre-training. In: OpenAI (2018) [1](#)
- [49] Ranasinghe, K., McKinzie, B., Ravi, S., Yang, Y., Toshev, A., Shlens, J.: Perceptual grouping in contrastive vision-language models. In: ICCV (2023) [4](#), [9](#)
- [50] Shin, G., Xie, W., Albanie, S.: Reco: Retrieve and co-segment for zero-shot transfer. In: NeurIPS (2022) [4](#), [10](#)
- [51] Singh, A., Hu, R., Goswami, V., Couairon, G., Galuba, W., Rohrbach, M., Kiela, D.: Flava: A foundational language and vision alignment model. In: CVPR (2022) [3](#)
- [52] Sun, Q., Fang, Y., Wu, L., Wang, X., Cao, Y.: Eva-clip: Improved training techniques for clip at scale. arXiv preprint arXiv:2303.15389 (2023) [3](#)
- [53] Tschannen, M., Kumar, M., Steiner, A., Zhai, X., Houlsby, N., Beyer, L.: Image captioners are scalable vision learners too. arXiv preprint arXiv:2306.07915 (2023) [3](#)
- [54] Wang, J., Yang, Z., Hu, X., Li, L., Lin, K., Gan, Z., Liu, Z., Liu, C., Wang, L.: Git: A generative image-to-text transformer for vision and language. TMLR (2022)
- [55] Wang, Z., Yu, J., Yu, A.W., Dai, Z., Tsvetkov, Y., Cao, Y.: Simvlm: Simple visual language model pretraining with weak supervision. In: ICLR (2022) [3](#)
- [56] Xu, J., De Mello, S., Liu, S., Byeon, W., Breuel, T., Kautz, J., Wang, X.: Groupvit: Semantic segmentation emerges from text supervision. In: CVPR (2022) [4](#), [10](#)
- [57] Xu, M., Zhang, Z., Wei, F., Lin, Y., Cao, Y., Hu, H., Bai, X.: A simple baseline for open-vocabulary semantic segmentation with pre-trained vision-language model. In: ECCV. Springer (2022) [2](#), [4](#), [11](#)
- [58] Yao, L., Han, J., Liang, X., Xu, D., Zhang, W., Li, Z., Xu, H.: Detclipv2: Scalable open-vocabulary object detection pre-training via word-region alignment. In: CVPR (2023) [5](#)
- [59] Yao, L., Han, J., Wen, Y., Liang, X., Xu, D., Zhang, W., Li, Z., Xu, C., Xu, H.: Detclip: Dictionary-enriched visual-concept paralleled pre-training for open-world detection. NeurIPS (2022) [5](#)
- [60] Yao, L., Huang, R., Hou, L., Lu, G., Niu, M., Xu, H., Liang, X., Li, Z., Jiang, X., Xu, C.: Filip: Fine-grained interactive language-image pre-training. In: ICLR (2021) [3](#)
- [61] Yu, J., Wang, Z., Vasudevan, V., Yeung, L., Seyedhosseini, M., Wu, Y.: Coca: Contrastive captioners are image-text foundation models. In: TMLR (2022) [3](#)
- [62] Yu, Q., He, J., Deng, X., Shen, X., Chen, L.C.: Convolutions die hard: Open-vocabulary segmentation with single frozen convolutional clip. NeurIPS (2023) [2](#)

- [63] Zhai, X., Kolesnikov, A., Houlsby, N., Beyer, L.: Scaling vision transformers. In: CVPR (2022) [9](#), [10](#)
- [64] Zhai, X., Mustafa, B., Kolesnikov, A., Beyer, L.: Sigmoid loss for language image pre-training. In: ICCV (2023) [1](#), [3](#), [5](#), [6](#), [8](#), [10](#), [14](#)
- [65] Zhai, X., Wang, X., Mustafa, B., Steiner, A., Keysers, D., Kolesnikov, A., Beyer, L.: Lit: Zero-shot transfer with locked-image text tuning. In: CVPR (2022) [3](#)
- [66] Zhang, H., Zhang, P., Hu, X., Chen, Y.C., Li, L., Dai, X., Wang, L., Yuan, L., Hwang, J.N., Gao, J.: Glipv2: Unifying localization and vision-language understanding. NeurIPS (2022) [5](#)
- [67] Zhang, R., Isola, P., Efros, A.A.: Colorful image colorization. In: ECCV (2016) [3](#)
- [68] Zhong, Y., Yang, J., Zhang, P., Li, C., Codella, N., Li, L.H., Zhou, L., Dai, X., Yuan, L., Li, Y., et al.: Regionclip: Region-based language-image pretraining. In: CVPR (2022) [3](#), [4](#)
- [69] Zhou, B., Zhao, H., Puig, X., Fidler, S., Barriuso, A., Torralba, A.: Scene parsing through ade20k dataset. In: CVPR (2017) [11](#)
- [70] Zhou, B., Zhao, H., Puig, X., Xiao, T., Fidler, S., Barriuso, A., Torralba, A.: Semantic understanding of scenes through the ade20k dataset. IJCV (2019) [11](#)
- [71] Zhou, C., Loy, C.C., Dai, B.: Extract free dense labels from clip. In: ECCV (2022) [4](#), [9](#), [10](#)
- [72] Zhou, J., Dong, L., Gan, Z., Wang, L., Wei, F.: Non-contrastive learning meets language-image pre-training. In: CVPR (2023) [4](#)
- [73] Zhou, J., Wei, C., Wang, H., Shen, W., Xie, C., Yuille, A., Kong, T.: Image BERT pre-training with online tokenizer. In: ICLR (2022) [2](#), [6](#)
- [74] Zhou, X., Girdhar, R., Joulin, A., Krähenbühl, P., Misra, I.: Detecting twenty-thousand classes using image-level supervision. In: ECCV. Springer (2022) [4](#)