

Single-Stream Multi-Level Alignment for Vision-Language Pretraining

Zaid Khan¹, Vijay Kumar B G², Xiang Yu², Samuel Schuster²,
Manmohan Chandraker^{2,3}, and Yun Fu¹

¹ Northeastern University

² NEC Labs America

³ UC San Diego

khan.za@northeastern.edu, vijay.kumar@nec-labs.com, xiangyu@nec-labs.com,
samuel@nec-labs.com, mkchandraker@eng.ucsd.edu, yunfu@ece.neu.edu

Abstract. Self-supervised vision-language pretraining from pure images and text with a contrastive loss is effective, but ignores fine-grained alignment due to a dual-stream architecture that aligns image and text representations only on a global level. Earlier, supervised, non-contrastive methods were capable of finer-grained alignment, but required dense annotations that were not scalable. We propose a single stream architecture that aligns images and language at multiple levels: global, fine-grained patch-token, and conceptual/semantic, using two novel tasks: symmetric cross-modality reconstruction (XMM) and a pseudo-labeled key word prediction (PSL). In XMM, we mask input tokens from one modality and use cross-modal information to reconstruct the masked token, thus improving fine-grained alignment between the two modalities. In PSL, we use attention to select keywords in a caption, use a momentum encoder to recommend other important keywords that are missing from the caption but represented in the image, and then train the visual encoder to predict the presence of those keywords, helping it learn semantic concepts that are essential for grounding a textual token to an image region. We demonstrate competitive performance and improved data efficiency on image-text retrieval, grounding, visual question answering/reasoning against larger models and models trained on more data. Code and models available at zaidkhan.me/SIMLA.

Keywords: Vision-Language Modeling, Cross-Modality Learning

1 Introduction

To learn a joint representation of images and language, early work [6,33,48] follows a supervised approach, using a pre-trained object detector to extract image regions, which are then aligned with corresponding image captions or dense annotations. Such approaches are limited by the amount of available densely annotated data and the semantic concepts the pretrained object detector can represent. A recent alternative approach is to directly align image representations with the corresponding text representations using a contrastive loss [42,29,17,55,14,35],

sidestepping the need for a pretrained object detector or dense annotations. Such approaches can learn from image-text pairs alone, which can be scraped from the web at large scales. However, the image-text contrastive learning paradigm is data hungry, using 1b+ [17,57] or 100m+ [42,62] pairs to overcome the noisiness of web-scraped image-text pairs. Second, the standard image-text contrastive learning architecture and objective uses a dual-stream architecture that aligns the global image and text representations, making it difficult to learn fine-grained details [56]. Third, contrastive learning does not explicitly align visual and language *concepts*, only features. Because the data complexity of images is greater than that of short captions, it can be challenging for the vision model to learn a representation that captures modality-invariant instance information corresponding to coherent natural language concepts rather than vision-specific semantics irrelevant to the modality-invariant image content.

We propose an approach that aligns image and language representations on multiple levels using a single stream transformer-only architecture that enables early, local interactions between image regions and language tokens, without the need for a pretrained object detector or dense annotations. We design a symmetric cross-modality reconstruction task to teach fine-grained alignment between image patches and language tokens, and construct a concept prediction task that extracts pseudo labels for each image without supervision and trains the visual encoder to detect concepts that are missing from the caption but present in the image. This allows us to align vision and language on multiple levels: fine-grained (cross-modality reconstruction), coarse (contrastive learning) and discrete (concept-level supervision). We empirically evaluate our proposed model, SIMLA (**S**ingle-Stream **M**ulti-**L**evel **A**lignment) on several downstream tasks, following prior work [24]. The entirely self-supervised SIMLA achieves state-of-the-art results on image-text retrieval and grounding, while outperforming larger models trained with supervision on downstream vision-language tasks, and demonstrates greater data efficiency compared to prior work in an ablation study. Our contributions, summarized:

1. A symmetric cross-modality reconstruction task to learn fine-grained alignment between image patches and language tokens.
2. A natural language, pseudo-labeling approach to align concept-level semantics without dense annotations.
3. A single-stream architecture to enable the proposed multi-level alignment.
4. Extensive experiments on image-text retrieval, vision-language reasoning, and visual grounding to demonstrate effectiveness of the proposed modules.

2 Method

Images are dense, unstructured, and require significant processing to extract useful semantic information. In contrast, language is highly structured, and contains directly available semantic information. Because of this asymmetry, attempting to align image features with language features too early will be futile, because the image features are too low-level to be matched with the more abstract language features. Contemporary architectures thus employ a symmetric encoder

design, in which both image and text are processed by equally deep encoders before late fusion through alignment of global image and text representations. This approach wastes model capacity, as high-level image semantics often correspond directly to low-level language semantics, so processing language to same depth as images is wasteful. In addition, both language and images contain a semantic pyramid of concepts, with some concepts being highly localized (e.g. a small image patch / single word) while other concepts are highly abstract (e.g. multiple interrelated image patches / multi-token sequences). Cross-modal concepts can exist at different levels of the semantic pyramid for each modality (e.g. the singular token 'throwing' describes a complex spatial scene / the phrase 'bird walking on rocky ground' may describe a small local image region). Thus, the problems in vision-language learning are twofold:

1. Asymmetry in inherent semantic abstraction between image and text data.
2. Semantic concepts appear at disparate levels in the abstraction hierarchy across modalities.

We propose an asymmetric architecture with a multi-task loss to address the above issues. Concretely, our architecture consists of a deep stack of transformer encoder layers that can be interpreted as a transformer language model [50] stacked atop a visual transformer [11]. During the forward pass, an image is fed through the bottom of the stack, while language tokens are injected at the middle of the stack, into the bottom of the language model. This design allows processing of the image to an appropriate level of semantic abstraction before fusion with language. Our multi-task loss consists of four tasks, engineered to align vision and language representations at multiple levels. We begin with an image-text matching task for very coarse instance-level alignment, and add a contrastive loss for global feature-level alignment. Next, we add a patch-level reconstruction task for fine-grained region-level alignment. Finally, we add a pseudo-label supervision task to the visual encoder to explicitly ensure the level of abstraction between the visual and language tokens is synchronized prior to fine-grained fusion.

2.1 Preliminary Architectures

Our model is a 24-deep stack of transformer [50] layers that can be decomposed into a vision encoder E_v , a language encoder E_l , and a multimodal encoder E_{mm} . Specifically, we stack the language encoder E_l atop the vision encoder E_v . We then add cross-attention layers after each self-attention layer in the language encoder E_l , allowing us to use it as a multimodal encoder E_{mm} when an image-text pair is passed in, and as a unimodal language encoder when language tokens are passed in. To obtain a multimodal embedding, we first use the bottom half of the transformer encoder stack (E_v) to encode an input image I into a sequence of embeddings $E_v(I) = \{\vec{v}_{\text{cls}}, \vec{v}_1, \dots, \vec{v}_N\}$ where v_{cls} is the embedding of the [CLS] token. We then pass the sequence of image embeddings $\{\vec{v}_{\text{cls}}, \vec{v}_1, \dots, \vec{v}_N\}$ into the top half of the transformer encoder stack, corresponding to the language

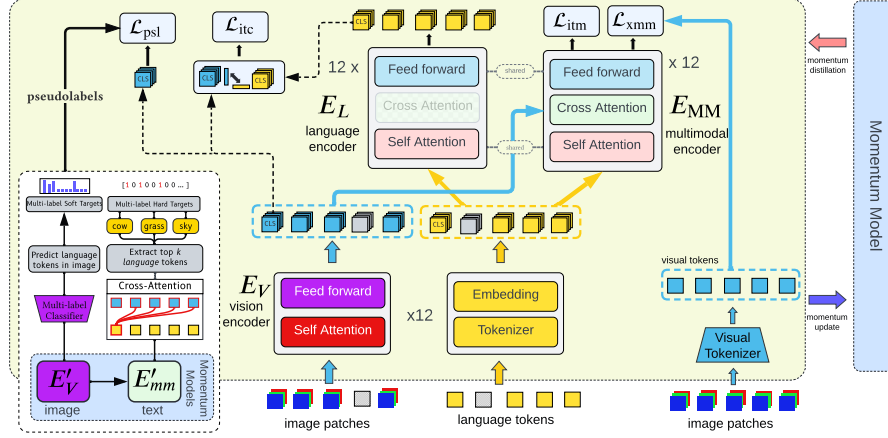


Fig. 1. SIMLA architecture. A language encoder E_l is stacked atop a vision encoder E_v . We add cross attention to E_l , allowing us to reuse it as a multimodal encoder E_{mm} by consuming image embeddings from E_v . Four tasks align images and language at multiple levels, exploiting a momentum model for additional supervision. A D-VAE tokenizes image patches for the cross-modality reconstruction task.

model with cross-attention, while concurrently injecting the associated caption, so the image embeddings $\{\vec{v}_{cls}, \vec{v}_1, \dots, \vec{v}_N\}$ from bottom half of the stack and the input *tokens* $\{[cls], t_1, \dots, t_N\}$ are consumed simultaneously and fused through cross-attention after each self attention layer to yield a sequence of multimodal embeddings $\{\vec{m}_{cls}, \vec{m}_1, \dots, \vec{m}_N\} = E_{mm}(\{\vec{v}_{cls}, \vec{v}_1, \dots, \vec{v}_N\}, \{[cls], t_1, \dots, t_N\})$.

2.2 Coarse Cross-Modality Alignment

Image-Text Contrastive Learning The simplest level of alignment is coarse, global alignment between image and text representations. Global alignment is useful training signal for two reasons: (i) it is robust to mismatches in fine-grained details between an image and caption (ii) it is an easier task than fine-grained alignment and enables faster learning during the earlier stages of training, when fine-grained alignment is infeasible due to the large domain gap between images and text. Coarse, global alignment requires learning image and text representations which capture modality-invariant information. A simple, effective and scalable [42, 17] approach to learning modality invariant representations is multi-view contrastive learning [49]. The multi-view contrastive objective pushes embeddings of matched image-text pairs together while pulling those of unmatched image-text pairs apart. Our contrastive loss follows the InfoNCE [37] formulation. Contrastive losses benefit from larger batch sizes, but batch sizes are bounded by GPU memory. To increase effective batch size, we follow MoCo [15] by using memory queues of size M for the unimodal image (Q^{img}) and text (Q^{txt}) features, as well as maintaining momentum (time-averaged) versions of

the text and image encoders. The normalized image-to-text and text-to-image similarity are calculated as

$$\begin{aligned} p_m^{i2t}(I, Q^{\text{txt}}) &= \frac{\exp(S(I, Q_m^{\text{txt}})/\tau)}{\sum_{m=1}^M \exp(S(I, Q_m^{\text{txt}})/\tau)} \\ p_m^{t2i}(T, Q^{\text{img}}) &= \frac{\exp(S(T, Q_m^{\text{img}})/\tau)}{\sum_{m=1}^M \exp(S(T, Q_m^{\text{img}})/\tau)} \end{aligned} \quad (1)$$

where τ is a learnable temperature parameter, $S(I, T) = g_v(\vec{v}_{\text{cls}})g'_l(\vec{l}'_{\text{cls}})$ and $S(T, I) = g_l(\vec{l}_{\text{cls}})^T g'_v(\vec{v}'_{\text{cls}})$ are raw similarity scores between image and text [CLS] tokens, obtained by $E_v(I)$ and $E_l(T)$ respectively. The functions g_v and g_l are linear transformations that project the unimodal [CLS] embeddings of the image and text, respectively, to lower-dimensional representations, followed by normalization to unit length. We use $g'_v(\vec{v}'_{\text{cls}})$ and $g'_l(\vec{l}'_{\text{cls}})$ to denote the momentum features, retrieved from the memory queues. The boolean one-hot vectors $\vec{y}^{i2t}(I)$ and $\vec{y}^{t2i}(T)$ represent the ground-truth similarity, with the positive pair indicated by a 1 and a 0 for all negatives. Then, the image-text contrastive loss is defined as the cross-entropy H between \vec{p} and \vec{y} :

$$\mathcal{L}_{\text{itc}} = \frac{1}{2} \mathbb{E}_{(I, T) \sim D} [\text{H}(\vec{y}^{i2t}(I), \vec{p}^{i2t}(I)) + \text{H}(\vec{y}^{t2i}(T), \vec{p}^{t2i}(T))] \quad (2)$$

The one-hot labels $\vec{y}^{i2t}(I)$ and $\vec{y}^{t2i}(T)$ penalize all predictions which do not match each image to the text it came paired with, and vice versa. However, one caption can potentially describe many different images, and similarly, many captions may match an image. To avoid this noisy penalization, we soften the hard targets with soft targets generated by the momentum model, corresponding to knowledge distillation [16] with the momentum model as a teacher. The complete loss can then be written as

$$\mathcal{L}_{\text{itc}}^{\text{mod}} = (1 - \alpha) \mathcal{L}_{\text{itc}} + \alpha \mathcal{L}'_{\text{itc}} \quad (3)$$

$$\mathcal{L}'_{\text{itc}} = \frac{1}{2} \mathbb{E}_{(I, T) \sim D} [\text{H}(p_m^{i2t}(I), p^{i2t}(I)) + \text{H}(p_m^{t2i}(T), p^{t2i}(T))] \quad (4)$$

where $p_m^{i2t}(I)$ and $p_m^{t2i}(T)$ is Equation 1 using only the momentum encoders.

Image-Text Matching is a binary classification task to predict if an image-text pair is matched. We define the ITM loss to be

$$\mathcal{L}_{\text{itm}} = \mathbb{E}_{(I, T) \sim D} \text{H}(\mathbf{y}^{\text{itm}}, \mathbf{p}^{\text{itm}}(I, T)) \quad (5)$$

where \mathbf{y}^{itm} is a one-hot vector indicating whether the pair is matched or not, and \mathbf{p}^{itm} is a two-class probability vector predicted by a single fully connected layer on top of the multimodal [CLS] token. We mine in-batch hard negatives for each image and text in a pair following ALBEF [24].

2.3 Finer-Grained Cross-Modality Alignment

A contrastive loss such as \mathcal{L}_{itc} aligns the global image and text representations. However, solely aligning the global representations while simultaneously fusing the image and text at the last possible opportunity makes it difficult to learn fine-grained correspondences, such as those between subregions of an image and subsequences of a caption. We design a reconstruction task to teach a model fine-grained alignment between images and patches. We mask the image, and force the model to reconstruct the masked image region from the remaining portion of the image using the caption as context. We then reverse the reconstruction task, forcing the model to reconstruct masked language tokens from the remaining portion of the caption using the image as context.

Concretely, (I, T) be an image text pair. Let \mathcal{M}_I be a mask for the image, generated following the masking strategy of BEiT [4], and let \mathcal{M}_T be the mask for the language tokens, generated following the masking strategy of BERT [9]. We then generate⁴ a masked image as $\hat{I} = I \odot \mathcal{M}_I$ and masked text as $\hat{T} = T \odot \mathcal{M}_T$. Then, the loss to be minimized is

$$\mathcal{L}_{\text{xmm}} = \mathbb{E}_{(I, \hat{T}) \sim D} H(\mathbf{y}^{\text{MLM}}, \mathbf{p}^{\text{MLM}}(I, \hat{T})) + \mathbb{E}_{(\hat{I}, T) \sim D} H(\mathbf{y}^{\text{MIM}}, \mathbf{p}^{\text{MIM}}(\hat{I}, T)) \quad (6)$$

The cross-modality masked language modeling loss \mathcal{L}_{xmm} is a sum of two cross-entropy losses, where \mathbf{y}^{MLM} and \mathbf{y}^{MIM} indicate the ground-truth value of the masked language token and masked image token respectively, and $\mathbf{p}^{\text{MLM}}(I, \hat{T})$, $\mathbf{p}^{\text{MIM}}(\hat{I}, T)$ represents the model’s probability estimates of the masked language and image tokens respectively. Because images are continuous, use the strategy of [4] to discretize the images into a sequence of tokens and mask them. We divide each image into patches and tokenize each patch with a discrete VAE [38] that maps each patch to one of 8192 visual tokens from a learned codebook.

In many cases, the ground-truth visual or language token can be plausibly replaced with an alternative. However, the ground truth target vectors are one-hot encoded and penalize any predictions that do not exactly match for the ground truth, even if they are plausible. Furthermore, the image masking and language masking are random, so it is possible for non-content tokens (e.g. *the*, *it*) or tokens that cannot be predicted well based on context to be masked. To allow the model to learn even when the ground-truth target for the masked token cannot be reasonably predicted from context, we again use the momentum distillation strategy. Specifically, we decompose \mathcal{L}_{xmm} into

$$\mathcal{L}_{\text{xmm}}^{\text{mod}} = (1 - \alpha)\mathcal{L}_{\text{MIM}} + \alpha\mathcal{L}'_{\text{MIM}} + (1 - \alpha)\mathcal{L}_{\text{MLM}} + \alpha\mathcal{L}'_{\text{MLM}} \quad (7)$$

where $\mathcal{L}'_{\text{MIM}} = H(\mathbf{p}_m^{\text{MIM}}, \mathbf{p}^{\text{MIM}}(I, \hat{T}))$, $\mathcal{L}'_{\text{MLM}} = H(\mathbf{p}_m^{\text{MLM}}, \mathbf{p}^{\text{MLM}}(\hat{I}, T))$ and $\mathbf{p}_m^{\text{MLM}}$, $\mathbf{p}_m^{\text{MIM}}$ are the softmax-normalized outputs of the MIM and MLM momentum prediction heads over the visual and language token distributions, respectively.

⁴ We depict the masking as a boolean operation for notational simplicity. The implementation follows the strategy of BEiT[4] and BERT[9] for I , T respectively.

2.4 Concept-Level Alignment

Semantic concepts may appear at disparate levels in the abstraction hierarchy across modalities. A concept may be highly complex in the visual modality, while being expressible with a single token in the language modality, and vice versa. This results in a concept-level mismatch between images and text. Although an asymmetric architecture that subjects image inputs to greater processing than text inputs prior to fusion addresses the intrinsic disparity in the semantic abstraction between image and text data, it does not guarantee that the visual embeddings $E_v(I) = \{\vec{v}_{\text{cls}}, \vec{v}_1, \dots, \vec{v}_N\}$ express concepts that are commonly described with language, or even possible to describe with language. Furthermore, it is possible that during the alignment process, the *unimodal* representations may degrade, because the emphasis is only on alignment.

To address this, we design a high-level alignment task in which the visual representation is aligned to represent concepts expressible by the language encoder by teaching it to label images with language concepts associated to the image, which also maintains the quality of the unimodal visual representation. We use the self-attention map of the multimodal [CLS] token to determine which language tokens within the text are most salient to the image-text pair. We choose k of the most salient tokens as pseudo-labels for the image, and generate a "hard" 2-D binary target vector $\mathbf{y}^{\text{PSL}} \in \mathbb{R}^V$, where V is the number of tokens known to the language model, and a 1 in the $[0][i]$ -th position indicates the i -th token is a target pseudo-label and a 1 in the $[1][j]$ -th position indicates the j -th token is not a target. We seek to minimize

$$\mathcal{L}_{\text{PSL}} = -\frac{1}{V} \sum_{i=1}^V \mathbf{y}_i^{\text{PSL}} \cdot \log(\sigma(\mathbf{p}_i^{\text{PSL}})) + (1 - \mathbf{y}_i^{\text{PSL}}) \cdot \log(1 - \sigma(\mathbf{p}_i^{\text{PSL}})) \quad (8)$$

where \mathbf{p}^{PSL} is the output of a single fully-connected layer placed atop the unimodal image [CLS] token, $\sigma(\cdot)$ is a sigmoid function used to clamp the output of the fully-connected layer between 0 and 1, and V is the number of tokens in the vocabulary of the tokenizer. This corresponds to multi-label loss where the model is trained to predict which language concepts (corresponding to tokens) are present in the image, using only the image context. However, the binary pseudolabels \mathbf{y}^{PSL} may fail to capture relevant concepts in the image, because the caption typically only describes a small number of aspects of an image. To provide a stronger self-supervisory signal, we use the momentum model as a teacher and minimize the K-L divergence between the predicted pseudolabels and the momentum pseudolabels. This can be expressed as a distillation loss where \mathbf{p}^{PSL} is the vector of momentum pseudolabel predictions.

$$\mathcal{L}_{\text{PSL}}^{\text{mod}} = (1-\alpha)\mathcal{L}_{\text{PSL}} - \frac{\alpha}{V} \sum_{i=1}^V \mathbf{p}_i^{\text{PSL}} \cdot \log(\sigma(\mathbf{p}_i^{\text{PSL}})) + (1 - \mathbf{p}_i^{\text{PSL}}) \cdot \log(1 - \sigma(\mathbf{p}_i^{\text{PSL}})) \quad (9)$$

The full pre-training objective can be expressed as

$$\mathcal{L} = \mathcal{L}_{\text{itc}}^{\text{mod}} + \mathcal{L}_{\text{xmm}}^{\text{mod}} + \mathcal{L}_{\text{itm}} + \mathcal{L}_{\text{PSL}}^{\text{mod}} \quad (10)$$

2.5 Implementation Details

We initialize the bottom 12 layers of the transformer encoder stack dedicated to vision (corresponding to the visual encoder) with the weights and architecture ViT/B-16 vision transformer [11], which is equipped with self-attention only. We initialize the top 12 multimodal layers of the transformer encoder stack (corresponding to the shared text / multimodal encoder) with the weights and architecture of BERT [9], with cross-attention. We pre-train the model for 30 epochs on 8 NVIDIA A100 GPUs with a batch size of 512. During pre-training, random 256×256 crops of images are used as input, and RandAugment [8] is applied to the images with color changes removed, following [24]. We set the momentum coefficient to 0.995, and linearly scale the distillation coefficient α from $0 \rightarrow 0.4$ in the first epoch. We use an $M = 65,536$ length memory queue. The AdamW [32] optimizer is used to train the model, with a weight decay of 0.02 and a cosine learning rate scheduler with a linear warmup to $1e^{-4}$ followed by a decay to $1e^{-5}$ in the subsequent epochs.

3 Experiments

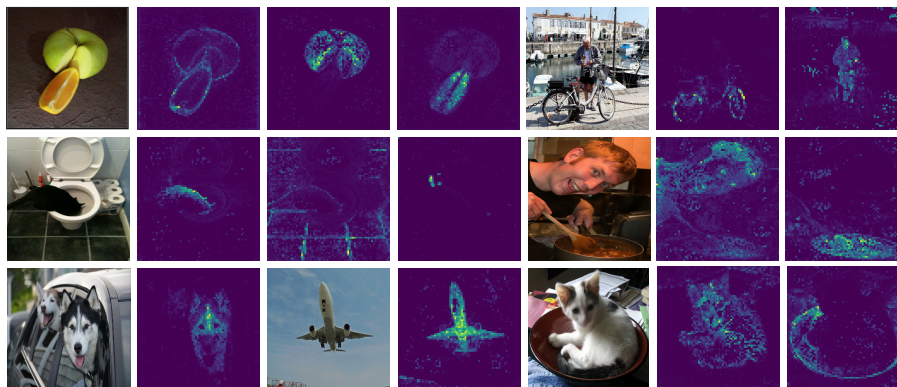


Fig. 2. Self-attention maps of the visual [CLS] token from different heads.

3.1 Experimental Setup

Pretraining Data is constructed by concatenating four image-text datasets: Conceptual Captions [45], SBU Captions [39], COCO[30] and Visual Genome [22], for a total of 4M image-text pairs, identical to [6,24].

Image-Text Retrieval The goal of text retrieval (TR) is to retrieve texts matching a query image. Image retrieval (IR) reverses the roles of the modalities. We evaluate retrieval on MSCOCO [30] and Flickr30k [41]. We use the

Karpathy[18] train/val/test splits for finetuning: 113k/5k/5k for MSCOCO and 29k/1k/1k for Flickr. For 0-shot retrieval on Flickr30k, we use the model finetuned on COCO, following [24]. We use ITC (Eq. 2) and ITM losses (Eq. 5) during fine-tuning. We finetune using a learning rate of $1e-5$ for 10 epochs. For fashion image retrieval, we use FashionGen[43], following the protocol of [13]. Fashion retrieval results and evaluation details are in the appendix.

Visual Question Answering (VQA) requires predicting an answer from an (image, question) pair. Following [24], we treat the task as a text generation problem using a auto-regressive decoder atop the multimodal encoder. For answer generation, we use [CLS] as the start of sequence token and [SEP] as the end of sequence token. The decoder is initialized from the multimodal encoder’s weights and finetuned with a language modeling loss. We restrict the decoder to generate answers from a predefined set $3k$ of candidate answers [20].

Visual Entailment is a visual reasoning task where a model must decide whether an image (the premise) entails a sentence (hypothesis), contradicts it, or is neutral. We stack a multi-layer perceptron atop the [CLS] token of the multi-modal encoder and treat the task as a 3-way classification problem.

Visual Grounding requires localizing the image region corresponding to a text description (referring expression). We use the RefCOCO+ dataset [58] with 141k referring expressions for 20k images from the COCO dataset. Following [24], we simulate a weakly supervised setting where the bounding box annotations are not used during finetuning. The model is finetuned for 5 epochs in manner similar to image-text retrieval.

Table 1. An ablation study on the components of the proposed approach. ITM: image-text matching. ITC: image-text contrastive learning. MLM: masked language modeling. MIM: masked image modeling. PLS: pseudo-label supervision.

	Components					Flickr30k True 0-shot (1k test set)					
	ITM	ITC	MLM	MIM	PLS	TR@1	TR@5	TR@10	IR@1	IR@5	IR@10
(a)	✓					6.1	9.3	11.6	7.3	10.4	11.7
(b)	✓	✓				73.1	85.9	88.5	56.6	79.0	83.6
(c)	✓	✓	✓			84.0	96.4	97.8	69.5	89.2	93.9
(d)	✓	✓	✓	✓		85.1	97.1	99.2	70.1	89.3	94.6
(d)	✓	✓	✓	✓	✓	86.2	97.2	98.7	69.5	90.2	94.7

3.2 Results and Discussion

We run each fine-tuning experiment five times with different random seeds, and report the mean and standard deviation in the following tables.

Zero-shot Retrieval Table 2 reports results on zero-shot image-text retrieval. Our SIMLA model outperforms both CLIP [42] and ALIGN [17], which were

trained on 100x and 300x more pairs respectively. We achieve better Rank-1 performance on both text and image retrieval compared to ALBEF [24].

Table 2. Zero-shot image-text retrieval results on Flickr30K.

Method	# Pre-train Images	Flickr30K (1K test set)					
		TR			IR		
		R@1	R@5	R@10	R@1	R@5	R@10
UNITER	4M	83.6	95.7	97.7	68.7	89.2	93.9
CLIP	400M	88.0	98.7	99.4	68.7	90.6	95.2
ALIGN	1.2 B	88.6	98.7	99.7	75.7	93.8	96.8
ALBEF	4M	90.5	98.8	99.7	76.8	93.7	96.7
SIMLA	4M	91.9	98.6	99.1	78.1	93.9	96.7
Std. Dev		± 0.4	± 0.4	± 0.2	± 0.4	± 0.3	± 0.2

3.3 Image-Text Retrieval

Table 3 reports results on fine-tuned image-text retrieval. Our SIMLA model outperforms all other approaches on Rank-1 retrieval across both modalities and dataset, with a substantial (3%) increase over ALBEF [24] and a (6%) increase over OSCAR [28] on Rank-1 MSCOCO text retrieval.

3.4 VQA/NLVR/SNLI-VE

Table 4 compares the performance of SIMLA to existing methods on vision-language understanding tasks. SIMLA achieves state of the art performance, outperforming methods that use object annotations [28] adversarial training [12], and dual stream architectures [24].

3.5 Weakly-Supervised Visual Grounding

We show results on RefCOCO+ in Table 5. We outperform ALBEF [24], which itself outperforms existing methods by $\approx 10\% - 30\%$, by 1.5% and 1.2% on TestA and TestB respectively. We ground a referring expression in an image using Grad-CAM [44] on the cross-attention maps in the 8th layer of the multimodal encoder, using the gradients of the image-text matching score $\mathbf{p}^{\text{itm}}(I, T)$ for a text-image pair (I, T) . In Figure 3, we visualize the Grad-CAM to show the grounding and fine-grained alignment ability of the model.

Data Efficiency The additional pretraining tasks of SIMLA result in a stronger training signal that allow the model to learn faster with fewer training steps. In Figure 4, we show zero-shot image-text retrieval accuracy as both ALBEF and SIMLA train. SIMLA’s zero-shot accuracy smoothly and quickly rises in the beginning stages of training, compared to the more gradual and rocky climb of ALBEF.

Table 3. Fine-tuned image-text retrieval results on Flickr30K and MSCOCO.

Method	Pairs	Flickr30K (1k test set)						MSCOCO (5k test set)					
		TR			IR			TR			IR		
		R@1	R@5	R@10	R@1	R@5	R@10	R@1	R@5	R@10	R@1	R@5	R@10
UNITER	4M	87.3	98.0	99.2	75.6	94.1	96.8	65.7	88.6	93.8	52.9	79.9	88.0
VILLA	4M	87.9	97.5	98.8	76.3	94.2	96.8	—	—	—	—	—	—
OSCAR	4M	—	—	—	—	—	—	70.0	91.1	95.5	54.0	80.8	88.5
ALBEF	4M	94.3	99.4	99.8	82.8	96.7	98.4	73.1	91.4	96.0	56.8	81.5	89.2
SIMLA	4M	94.7	99.5	99.7	83.3	96.5	98.2	75.8	92.9	96.2	57.7	81.9	92.0
Std. Dev		± 0.2	± 0.1	± 0.1	± 0.2	± 0.2	± 0.1	± 0.3	± 0.3	± 0.1	± 0.2	± 0.4	± 0.55

Table 4. Comparison with state-of-the-art methods on downstream vision-language tasks. ALBEF results are from our reproduction, due to expired URLs in NLVR. SNLI-VE results may be noisy due to label errors [10].

Model	VQA		NLVR ²		SNLI-VE	
	test-dev	test-std	dev	test-P	val	test
VisualBERT [25]	70.8	71.0	67.4	67.0	-	-
VL-BERT [47]	71.16	-	-	-	-	-
LXMERT [48]	72.4	72.5	74.9	74.5	-	-
12-in-1 [34]	73.2	-	-	78.9	-	77.0
UNITER [6]	72.7	72.9	77.2	77.9	78.6	78.3
VL-BART/T5 [7]	-	71.3	-	73.6	-	-
ViLT [21]	70.9	-	75.2	76.2	-	-
OSCAR [28]	73.2	73.4	78.1	78.4	-	-
VILLA [12]	73.6	73.7	78.4	79.3	79.5	79.0
ALBEF [24]	74.5	74.7	79.2	80.0	79.1	80.1
SIMLA	74.5	74.8	79.8	79.5	79.6	80.2
Std. Dev	± 0.1	± 0.1	± 0.4	± 0.5	± 0.2	± 0.3

Table 5. Weakly-supervised visual grounding on RefCOCO+ [59] dataset.

Method	Val	TestA	TestB
ARN [31]	32.8	34.4	32.1
CCL [63]	34.3	36.9	33.6
ALBEF [24]	58.5	65.9	46.3
SIMLA	58.1	67.4	47.5
Std. Dev	± 0.5	± 0.29	± 0.33

Ablation Study In Table 1, we study the effect of the various losses on image-text retrieval performance. Training with only the image-text matching (ITM) loss provides only a weak supervisory signal. Explicit alignment is crucial, and each level of alignment provides an increase in performance. Global alignment (\mathcal{L}_{itc}) provides the largest boost in performance, but fine-grained alignment ($\mathcal{L}_{\text{mim}} + \mathcal{L}_{\text{mlm}}$) is crucial for increasing performance, and pseudo-label supervi-



Fig. 3. Examples of fine-grained alignment learned by SIMLA. The model can ground abstract concepts (e.g. metal flying object) in addition to simple concepts (e.g shoe).




Fig. 4. True zero-shot fast retrieval accuracy on the Flickr30K test set as a function of training time. SIMLA achieves higher accuracy in less training time than ALBEF.

sion (\mathcal{L}_{psl}) successfully exploits additional supervisory signals to learn a better-aligned representation.

Qualitative Results In Figure 5, we show examples of pseudolabels generated by the momentum models. When the captions are nondescriptive, the pseudolabels provide a strong surrogate supervisory signal that grounds the content of the image in natural language concepts. Even when the captions *are* descriptive (bottom middle of Figure 5), the pseudolabels provide additional supervision by requiring the visual representation to reflect concepts present in the image but not in the caption. We show self-attention maps obtained from the [CLS] token of the visual encoder in Figure 2. Different heads of the visual encoder work together to decompose a scene, with heads focusing on various parts of the scene. The subjects of attention correspond to the visual concepts humans are most likely to notice, even in cluttered or dense scenes. The attention map segments objects well, despite having no access to object-level annotations and receiving supervision from oftentimes noisy captions.

3.6 Parameter Count and Inference Speed

In Tables 2, 3 and 4, we report results against the **large** sized versions of UNITER [6], OSCAR [28] and VILLA [12], which have $\approx 335\text{M}$ parameters and a depth of 24 transformer encoder layers. SIMLA has equivalent depth (24 layers) with fewer parameters (223M) due to parameter sharing. SIMLA is also substantially faster at inference time (7.2 pairs / second vs ≈ 1.1 pairs / second for UNITER/VILLA/OSCAR) due to the dual-encoder design shared with ALBEF [24], in which the text/image encoders can be used separately to quickly

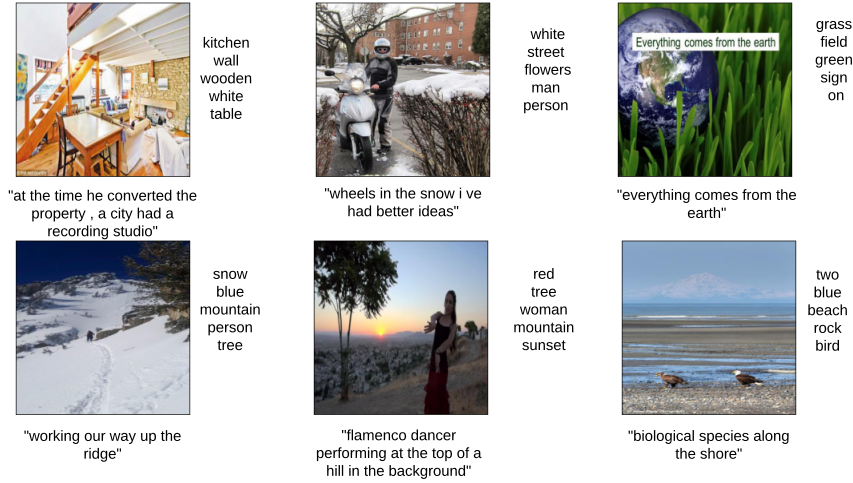


Fig. 5. Examples of pseudolabels used for pseudolabel supervision, obtained by decoding high-probability concept head logits from the momentum encoders.

retrieve the top- k candidates matching a query, and re-ranking them using the slower multimodal encoder.

4 Related Work

4.1 Vision Language Pretraining

Early transformer-based [50,9] vision language pretraining techniques [25,33,48,6] required an pretrained object detector and were limited to visual categories the object detector could identify. Recent contrastive-image text learners such as CLIP[42] do not rely on object detectors and understand a far wider range of visual concepts. However, CLIP relies on a massive dataset (400m pairs) to overcome label noise. Several methods [24,29,62,55,14] have been proposed for data-efficient pretraining. DeCLIP [29] exploits inter/intra-modality supervision to train a CLIP-like model with less data, similar to [55]. ALBEF[24] proposes a contrastive alignment followed by deeper fusion with a multimodal encoder. Methods such as BLIP[23], CoCa[57], SimVLM [52], UNIMO [26,27] incorporate a decoder and add image-to-text generation as an auxillary task. Other lines of work on vision-language foundation models [5] are multi-task models [1,60,51,46] or foundation model ensembles [61,54,5]. We propose a data-efficient, detector-free pretraining approach, architecturally similar to [24], but with additional supervision for the visual encoder to learn key words (high level concepts) that are present in the image, as well as a symmetric cross-modality reconstruction task inspired by the masked image modeling techniques of [4,64].

4.2 Fusion Methods

Existing fusion techniques can be broadly classified into three categories: early [40],[53], middle [19,36,36], and late fusion [3,42,33,17]. Late fusion (e.g. CLIP) is the dominant approach due to its scalability and encodes input modalities separately using unimodal encoders and fuses the resulting representations at the end. However, each modality can have different levels of information density. For example, [2] shows audio and video to have fine-grained information while text has coarse-grained information, and [62] draws the conclusion that in dual-stream contrastive architectures, the strength of the visual encoder matters more than that of the language encoder. It is thus essential to consider the information density of the input modalities for fusion. Compared to language, images require significant processing to extract useful semantic information, but current dual-stream approaches [42,62,29] apply the same amount of processing to both. In contrast, we use a single stream architecture where the input modalities undergo asymmetric processing before fusion. OSCAR[28], UNITER[6] and ViLBERT[33] also include a similar patch-level concept prediction task, but they use region labels produced by a pretrained object detector as prediction targets. SIMLA is fully self-supervised, and needs no labels, bounding boxes, or object detectors.

5 Conclusion

We propose SIMLA, a framework for vision-language pretraining. In contrast to contemporary dual-stream approaches that employ symmetric encoders and introduce multimodal interactions after unimodal representation learning, SIMLA uses a single-stream architecture with asymmetric depth of processing for each modality and enables earlier multimodal interactions. SIMLA aligns images and text on multiple levels, and explicitly enriches the visual modality with pseudo-label supervision to ensure similar levels of conceptual abstraction in the representations before fusion. We empirically verify the strength of the approach and achieve state-of-the-art results on image-text retrieval, natural language visual grounding, and vision language reasoning tasks. Finally, we show that the additional training tasks provide additional supervision that increases the data efficiency of SIMLA relative to other state-of-the-art approaches.

References

1. Alayrac, J.B., Donahue, J., Luc, P., Miech, A., Barr, I., Hasson, Y., Lenc, K., Mensch, A., Millican, K., Reynolds, M., Ring, R., Rutherford, E., Cabi, S., Han, T., Gong, Z., Samangooei, S., Monteiro, M., Menick, J., Borgeaud, S., Brock, A., Nematzadeh, A., Sharifzadeh, S., Binkowski, M., Barreira, R., Vinyals, O., Zisserman, A., Simonyan, K.: Flamingo: a visual language model for few-shot learning. ArXiv **abs/2204.14198** (2022) [13](#)
2. Alayrac, J., Recasens, A., Schneider, R., Arandjelovic, R., Ramapuram, J., Fauw, J.D., Smaira, L., Dieleman, S., Zisserman, A.: Self-supervised multimodal versatile networks. CoRR **abs/2006.16228** (2020) [14](#)
3. Alwassel, H., Mahajan, D., Torresani, L., Ghanem, B., Tran, D.: Self-supervised learning by cross-modal audio-video clustering. CoRR **abs/1911.12667** (2019) [14](#)
4. Bao, H., Dong, L., Wei, F.: Beit: BERT pre-training of image transformers. CoRR **abs/2106.08254** (2021), <https://arxiv.org/abs/2106.08254> [6](#), [13](#)
5. Bommasani, R., Hudson, D.A., Adeli, E., Altman, R., Arora, S., von Arx, S., Bernstein, M.S., Bohg, J., Bosselut, A., Brunskill, E., Brynjolfsson, E., Buch, S., Card, D., Castellon, R., Chatterji, N.S., Chen, A.S., Creel, K.A., Davis, J., Demszky, D., Donahue, C., Doumbouya, M., Durmus, E., Ermon, S., Etchemendy, J., Ethayarajh, K., Fei-Fei, L., Finn, C., Gale, T., Gillespie, L.E., Goel, K., Goodman, N.D., Grossman, S., Guha, N., Hashimoto, T., Henderson, P., Hewitt, J., Ho, D.E., Hong, J., Hsu, K., Huang, J., Icard, T.F., Jain, S., Jurafsky, D., Kalluri, P., Karamcheti, S., Keeling, G., Khani, F., Khattab, O., Koh, P.W., Krass, M.S., Krishna, R., Kudithipudi, R., Kumar, A., Ladhak, F., Lee, M., Lee, T., Leskovec, J., Levent, I., Li, X.L., Li, X., Ma, T., Malik, A., Manning, C.D., Mirchandani, S.P., Mitchell, E., Munyikwa, Z., Nair, S., Narayan, A., Narayanan, D., Newman, B., Nie, A., Niebles, J.C., Nilforoshan, H., Nyarko, J.F., Ogut, G., Orr, L., Papadimitriou, I., Park, J.S., Piech, C., Portelance, E., Potts, C., Raghunathan, A., Reich, R., Ren, H., Rong, F., Roohani, Y.H., Ruiz, C., Ryan, J., R'e, C., Sadigh, D., Sagawa, S., Santhanam, K., Shih, A., Srinivasan, K.P., Tamkin, A., Taori, R., Thomas, A.W., Tramèr, F., Wang, R.E., Wang, W., Wu, B., Wu, J., Wu, Y., Xie, S.M., Yasunaga, M., You, J., Zaharia, M.A., Zhang, M., Zhang, T., Zhang, X., Zhang, Y., Zheng, L., Zhou, K., Liang, P.: On the opportunities and risks of foundation models. ArXiv **abs/2108.07258** (2021) [13](#)
6. Chen, Y., Li, L., Yu, L., Kholy, A.E., Ahmed, F., Gan, Z., Cheng, Y., Liu, J.: UNITER: universal image-text representation learning. In: ECCV. vol. 12375, pp. 104–120 (2020) [1](#), [8](#), [11](#), [12](#), [13](#), [14](#)
7. Cho, J., Lei, J., Tan, H., Bansal, M.: Unifying vision-and-language tasks via text generation. arXiv preprint arXiv:2102.02779 (2021) [11](#)
8. Cubuk, E.D., Zoph, B., Shlens, J., Le, Q.V.: Randaugment: Practical automated data augmentation with a reduced search space. In: CVPR Workshops. pp. 702–703 (2020) [8](#)
9. Devlin, J., Chang, M., Lee, K., Toutanova, K.: BERT: pre-training of deep bidirectional transformers for language understanding. In: Burstein, J., Doran, C., Solorio, T. (eds.) NAACL. pp. 4171–4186 (2019) [6](#), [8](#), [13](#)
10. Do, V., Camburu, O.M., Akata, Z., Lukasiewicz, T.: e-snli-ve-2.0: Corrected visual-textual entailment with natural language explanations. ArXiv **abs/2004.03744** (2020) [11](#)
11. Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., Uszkoreit, J., Houlsby, N.:

- An image is worth 16x16 words: Transformers for image recognition at scale. In: ICLR (2021) [3](#), [8](#)
12. Gan, Z., Chen, Y., Li, L., Zhu, C., Cheng, Y., Liu, J.: Large-scale adversarial training for vision-and-language representation learning. In: Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M., Lin, H. (eds.) NeurIPS (2020) [10](#), [11](#), [12](#)
 13. Gao, D., Jin, L., Chen, B., Qiu, M., Wei, Y., Hu, Y., Wang, H.J.: Fashionbert: Text and image matching with adaptive loss for cross-modal retrieval. Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval (2020) [9](#)
 14. Goel, S., Bansal, H., Bhatia, S.K., Rossi, R.A., Vinay, V., Grover, A.: Cyclip: Cyclic contrastive language-image pretraining. ArXiv [abs/2205.14459](#) (2022) [1](#), [13](#)
 15. He, K., Fan, H., Wu, Y., Xie, S., Girshick, R.: Momentum contrast for unsupervised visual representation learning. In: CVPR (2020) [4](#)
 16. Hinton, G., Vinyals, O., Dean, J.: Distilling the knowledge in a neural network. arXiv preprint arXiv:1503.02531 (2015) [5](#)
 17. Jia, C., Yang, Y., Xia, Y., Chen, Y.T., Parekh, Z., Pham, H., Le, Q.V., Sung, Y., Li, Z., Duerig, T.: Scaling up visual and vision-language representation learning with noisy text supervision. arXiv preprint arXiv:2102.05918 (2021) [1](#), [2](#), [4](#), [9](#), [14](#)
 18. Karpathy, A., Li, F.: Deep visual-semantic alignments for generating image descriptions. In: CVPR. pp. 3128–3137 (2015) [9](#)
 19. Kazakos, E., Nagrani, A., Zisserman, A., Damen, D.: Epic-fusion: Audio-visual temporal binding for egocentric action recognition. CoRR [abs/1908.08498](#) (2019) [14](#)
 20. Kim, J., Jun, J., Zhang, B.: Bilinear attention networks. In: Bengio, S., Wallach, H.M., Larochelle, H., Grauman, K., Cesa-Bianchi, N., Garnett, R. (eds.) NIPS. pp. 1571–1581 (2018) [9](#)
 21. Kim, W., Son, B., Kim, I.: Vilt: Vision-and-language transformer without convolution or region supervision. arXiv preprint arXiv:2102.03334 (2021) [11](#)
 22. Krishna, R., Zhu, Y., Groth, O., Johnson, J., Hata, K., Kravitz, J., Chen, S., Kalanidis, Y., Li, L., Shamma, D.A., Bernstein, M.S., Fei-Fei, L.: Visual genome: Connecting language and vision using crowdsourced dense image annotations. IJCV **123**(1), 32–73 (2017) [8](#)
 23. Li, J., Li, D., Xiong, C., Hoi, S.C.H.: Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation. In: ICML (2022) [13](#)
 24. Li, J., Selvaraju, R.R., Gotmare, A.D., Joty, S., Xiong, C., Hoi, S.: Align before fuse: Vision and language representation learning with momentum distillation. In: NeurIPS (2021) [2](#), [5](#), [8](#), [9](#), [10](#), [11](#), [12](#), [13](#)
 25. Li, L.H., Yatskar, M., Yin, D., Hsieh, C., Chang, K.: Visualbert: A simple and performant baseline for vision and language. arXiv preprint arXiv:1908.03557 [abs/1908.03557](#) (2019) [11](#), [13](#)
 26. Li, W., Gao, C., Niu, G., Xiao, X., Liu, H., Liu, J., Wu, H., Wang, H.: Unimo: Towards unified-modal understanding and generation via cross-modal contrastive learning. ArXiv [abs/2012.15409](#) (2021) [13](#)
 27. Li, W., Gao, C., Niu, G., Xiao, X., Liu, H., Liu, J., Wu, H., Wang, H.: UNIMO-2: End-to-end unified vision-language grounded learning. In: Findings of the Association for Computational Linguistics: ACL 2022. pp. 3187–3201. Association for Computational Linguistics, Dublin, Ireland (May 2022). <https://doi.org/10.18653/v1/2022.findings-acl.251>, <https://aclanthology.org/2022.findings-acl.251> [13](#)

28. Li, X., Yin, X., Li, C., Zhang, P., Hu, X., Zhang, L., Wang, L., Hu, H., Dong, L., Wei, F., Choi, Y., Gao, J.: Oscar: Object-semantics aligned pre-training for vision-language tasks. In: ECCV. pp. 121–137 (2020) [10](#), [11](#), [12](#), [14](#)
29. 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: International Conference on Learning Representations (2022), <https://openreview.net/forum?id=zq1iJkNk3uN> [1](#), [13](#), [14](#)
30. Lin, T., Maire, M., Belongie, S.J., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft COCO: common objects in context. In: Fleet, D.J., Pajdla, T., Schiele, B., Tuytelaars, T. (eds.) ECCV. vol. 8693, pp. 740–755 (2014) [8](#)
31. Liu, X., Li, L., Wang, S., Zha, Z., Meng, D., Huang, Q.: Adaptive reconstruction network for weakly supervised referring expression grounding. In: ICCV. pp. 2611–2620 (2019) [11](#)
32. Loshchilov, I., Hutter, F.: Decoupled weight decay regularization. arXiv preprint arXiv:1711.05101 (2017) [8](#)
33. Lu, J., Batra, D., Parikh, D., Lee, S.: Vilbert: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks. In: Wallach, H.M., Larochelle, H., Beygelzimer, A., d’Alché-Buc, F., Fox, E.B., Garnett, R. (eds.) NeurIPS. pp. 13–23 (2019) [1](#), [13](#), [14](#)
34. Lu, J., Goswami, V., Rohrbach, M., Parikh, D., Lee, S.: 12-in-1: Multi-task vision and language representation learning. In: CVPR. pp. 10434–10443 (2020) [11](#)
35. Mu, N., Kirillov, A., Wagner, D.A., Xie, S.: Slip: Self-supervision meets language-image pre-training. ArXiv [abs/2112.12750](#) (2021) [1](#)
36. Nagrani, A., Yang, S., Arnab, A., Jansen, A., Schmid, C., Sun, C.: Attention bottlenecks for multimodal fusion. CoRR [abs/2107.00135](#) (2021) [14](#)
37. van den Oord, A., Li, Y., Vinyals, O.: Representation learning with contrastive predictive coding. CoRR [abs/1807.03748](#) (2018), <http://arxiv.org/abs/1807.03748> [4](#)
38. van den Oord, A., Vinyals, O., Kavukcuoglu, K.: Neural discrete representation learning. CoRR [abs/1711.00937](#) (2017), <http://arxiv.org/abs/1711.00937> [6](#)
39. Ordonez, V., Kulkarni, G., Berg, T.L.: Im2text: Describing images using 1 million captioned photographs. In: Shawe-Taylor, J., Zemel, R.S., Bartlett, P.L., Pereira, F.C.N., Weinberger, K.Q. (eds.) NIPS. pp. 1143–1151 (2011) [8](#)
40. Owens, A., Efros, A.A.: Audio-visual scene analysis with self-supervised multisensory features. CoRR [abs/1804.03641](#) (2018) [14](#)
41. Plummer, B.A., Wang, L., Cervantes, C.M., Caicedo, J.C., Hockenmaier, J., Lazebnik, S.: Flickr30k entities: Collecting region-to-phrase correspondences for richer image-to-sentence models. In: ICCV. pp. 2641–2649 (2015) [8](#)
42. 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. arXiv preprint arXiv:2103.00020 (2021) [1](#), [2](#), [4](#), [9](#), [13](#), [14](#)
43. Rostamzadeh, N., Hosseini, S., Boquet, T., Stokowiec, W., Zhang, Y., Jauvin, C., Pal, C.J.: Fashion-gen: The generative fashion dataset and challenge. ArXiv [abs/1806.08317](#) (2018) [9](#)
44. Selvaraju, R.R., Das, A., Vedantam, R., Cogswell, M., Parikh, D., Batra, D.: Grad-cam: Visual explanations from deep networks via gradient-based localization. International Journal of Computer Vision **128**, 336–359 (2017) [10](#)
45. Sharma, P., Ding, N., Goodman, S., Soricut, R.: Conceptual captions: A cleaned, hypernymed, image alt-text dataset for automatic image captioning. In: Gurevych, I., Miyao, Y. (eds.) ACL. pp. 2556–2565 (2018) [8](#)

46. Singh, A., Hu, R., Goswami, V., Couairon, G., Galuba, W., Rohrbach, M., Kiela, D.: Flava: A foundational language and vision alignment model. ArXiv **abs/2112.04482** (2021) [13](#)
47. Su, W., Zhu, X., Cao, Y., Li, B., Lu, L., Wei, F., Dai, J.: Vi-bert: Pre-training of generic visual-linguistic representations. In: ICLR (2020) [11](#)
48. Tan, H., Bansal, M.: LXMERT: learning cross-modality encoder representations from transformers. In: Inui, K., Jiang, J., Ng, V., Wan, X. (eds.) EMNLP. pp. 5099–5110 (2019) [1](#), [11](#), [13](#)
49. Tian, Y., Krishnan, D., Isola, P.: Contrastive multiview coding. CoRR **abs/1906.05849** (2019), <http://arxiv.org/abs/1906.05849> [4](#)
50. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L., Polosukhin, I.: Attention is all you need. In: Guyon, I., von Luxburg, U., Bengio, S., Wallach, H.M., Fergus, R., Vishwanathan, S.V.N., Garnett, R. (eds.) NIPS. pp. 5998–6008 (2017) [3](#), [13](#)
51. Wang, P., Yang, A., Men, R., Lin, J., Bai, S., Li, Z., Ma, J., Zhou, C., Zhou, J., Yang, H.: Ofa: Unifying architectures, tasks, and modalities through a simple sequence-to-sequence learning framework. CoRR **abs/2202.03052** (2022) [13](#)
52. Wang, Z., Yu, J., Yu, A.W., Dai, Z., Tsvetkov, Y., Cao, Y.: Simvlm: Simple visual language model pretraining with weak supervision. ArXiv **abs/2108.10904** (2021) [13](#)
53. Xiao, F., Lee, Y.J., Grauman, K., Malik, J., Feichtenhofer, C.: Audiovisual slowfast networks for video recognition. CoRR **abs/2001.08740** (2020) [14](#)
54. Xie, Y., Zhou, L., Dai, X., Yuan, L., Bach, N., Liu, C., Zeng, M.: Visual clues: Bridging vision and language foundations for image paragraph captioning. ArXiv **abs/2206.01843** (2022) [13](#)
55. YANG, J., Duan, J., Tran, S., Xu, Y., Chanda, S., Chen, L., Zeng, B., Chilimbi, T., HUANG, J.: Vision-language pre-training with triple contrastive learning. In: CVPR 2022 (2022), <https://www.amazon.science/publications/vision-language-pre-training-with-triple-contrastive-learning> [1](#), [13](#)
56. 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. ArXiv **abs/2111.07783** (2021) [2](#)
57. Yu, J., Wang, Z., Vasudevan, V., Yeung, L., Seyedhosseini, M., Wu, Y.: Coca: Contrastive captioners are image-text foundation models. ArXiv **abs/2205.01917** (2022) [2](#), [13](#)
58. Yu, L., Poirson, P., Yang, S., Berg, A.C., Berg, T.L.: Modeling context in referring expressions. ArXiv **abs/1608.00272** (2016) [9](#)
59. Yu, L., Poirson, P., Yang, S., Berg, A.C., Berg, T.L.: Modeling context in referring expressions. In: Leibe, B., Matas, J., Sebe, N., Welling, M. (eds.) ECCV. pp. 69–85 (2016) [11](#)
60. Yuan, L., Chen, D., Chen, Y.L., Codella, N.C.F., Dai, X., Gao, J., Hu, H., Huang, X., Li, B., Li, C., Liu, C., Liu, M., Liu, Z., Lu, Y., Shi, Y., Wang, L., Wang, J., Xiao, B., Xiao, Z., Yang, J., Zeng, M., Zhou, L., Zhang, P.: Florence: A new foundation model for computer vision. ArXiv **abs/2111.11432** (2021) [13](#)
61. Zeng, A., Wong, A.S., Welker, S., Choromanski, K., Tombari, F., Purohit, A., Ryoo, M.S., Sindhvani, V., Lee, J., Vanhoucke, V., Florence, P.R.: Socratic models: Composing zero-shot multimodal reasoning with language. ArXiv **abs/2204.00598** (2022) [13](#)
62. Zhai, X., Wang, X., Mustafa, B., Steiner, A., Keysers, D., Kolesnikov, A., Beyer, L.: Lit: Zero-shot transfer with locked-image text tuning. In: Proceedings of the

- IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 18123–18133 (June 2022) [2](#), [13](#), [14](#)
63. Zhang, Z., Zhao, Z., Lin, Z., Zhu, J., He, X.: Counterfactual contrastive learning fo weakly-supervised vision-language grounding. In: Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M., Lin, H. (eds.) NeurIPS (2020) [11](#)
64. Zhou, J., Wei, C., Wang, H., Shen, W., Xie, C., Yuille, A., Kong, T.: ibot: Image bert pre-training with online tokenizer. International Conference on Learning Representations (ICLR) (2022) [13](#)