

HAHV: Hierarchical Aggregation of Augmented Views for Image Captioning

Chia-Wen Kuo
Georgia Tech

albert.cwkuo@gatech.edu

Zsolt Kira
Georgia Tech

zkira@gatech.edu

Abstract

A great deal of progress has been made in image captioning, driven by research into how to encode the image using pre-trained models. This includes visual encodings (e.g. image grid features or detected objects) and more recently textual encodings (e.g. image tags or text descriptions of image regions). As more advanced encodings are available and incorporated, it is natural to ask: how to **efficiently** and **effectively** leverage the heterogeneous set of encodings? In this paper, we propose to regard the encodings as augmented views of the input image. The image captioning model encodes each view independently with a shared encoder efficiently, and a contrastive loss is incorporated across the encoded views in a novel way to improve their representation quality and the model’s data efficiency. Our proposed hierarchical decoder then adaptively weighs the encoded views according to their effectiveness for caption generation by first aggregating within each view at the token level, and then across views at the view level. We demonstrate significant performance improvements of +5.6% CIDEr on MS-COCO and +12.9% CIDEr on Flickr30k compared to state of the arts, and conduct rigorous analyses to demonstrate the importance of each part of our design.

1. Introduction

A large amount of progress has been made in vision-and-language (VL) tasks such as image captioning [1, 8], visual question answering [16, 23], and image-text retrieval. For these tasks, recent methods [31, 36, 45, 61] observe that encoding the input image by an object detector [42] pre-trained on Visual Genome [30] into a set of detected objects is not sufficient. To provide information complementary to detected objects, recent works proposed to encode an input image by different pre-trained models and into different modalities, and achieve substantial performance improvement by combining these heterogeneous encodings. For example, some works encode from the visual perspective (e.g. stronger object detector pre-trained on a larger vocabu-

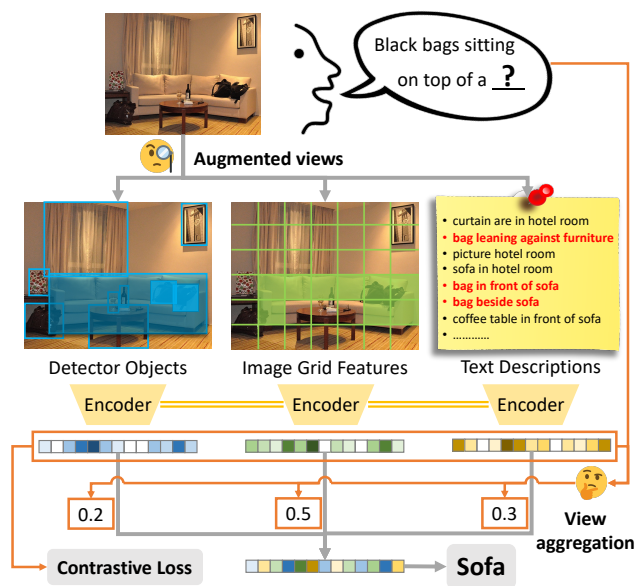


Figure 1. HAAV, Hierarchical Aggregation of Augmented Views, for image captioning at the step of predicting the word “sofa”. First, heterogeneous views such as detected objects [4], image grid features [45], and text descriptions [31] are generated from the input image by existing methods. We propose to regard these views as augmentations of the input image, and independently encode each view by a shared transformer encoder efficiently. A contrastive loss is incorporated to improve the representation quality of heterogeneous views. Finally, our proposed hierarchical decoder models the effectiveness of each view and adaptively weigh them according to their effectiveness for predicting the current word “sofa”.

lary and datasets [61] or global image features [24]), while other works encode from the textual perspective (e.g. image tags [36] and text descriptions of image regions [31]).

Given the great success of incorporating various heterogeneous encodings or “views”, one research question emerges naturally: *how to efficiently and effectively leverage these heterogeneous views for caption generation?* For **efficiency**, three factors are particularly important: computation, parameter count, and label efficiency. State-of-art VL and

image captioning models are typically a transformer encoder-decoder model [51], which has undesirable quadratic computational complexity with respect to the input sequence size. Therefore, as more views are incorporated, each represented by a sequence of tokens, we should carefully manage the computation and model size. Moreover, on the medium-scale MS-COCO image captioning benchmark [8] ($\sim 0.6M$ training samples), we should take label efficiency into consideration when training the data-hungry [11] transformer model to avoid negative effects such as overfitting. For **effectiveness**, different views contain some shared and some complementary information of the input image. Therefore, it is important to model the effectiveness of views and adaptively weigh them according to their effectiveness for predicting each word. Take image captioning in Figure 1 as an example, when predicting the current word “sofa”, for the incomplete caption “black bags sitting on top of a ?”, if say, the view of detected objects fails to detect sofa in the input image, the captioning model should down-weight the less effective view of detected objects and rely on other more effective views that properly encode the information about sofa. With these considerations in mind, we propose HAAV, **Hierarchical Aggregation of Augmented Views**. In HAAV, given a set of heterogeneous views of the input image from existing works such as detected objects [4], image grid features [45], and text descriptions [31], we propose to (1) regard heterogeneous views as augmentations of the input image, and (2) devise a hierarchical decoder layer to account for the effectiveness of heterogeneous views.

For (1), by regarding views as augmentations, we naturally choose to use a *shared* transformer encoder to encode each view *independently*. Compared to methods that concatenate all views into a long sequence as input [31, 36, 61], where the computational complexity scales up quadratically with respect to the number of views, our method scales up linearly. Compared to models that use entire models per view [10, 22, 26, 43] or methods that encode each view with unshared encoders [2, 3, 20, 34], our method is more parameter efficient. Furthermore, data augmentation increases data diversity and thus improves data efficiency, which is particularly important for training data-hungry transformer models. Last but not least, by regarding heterogeneous views generated by different pre-trained models as augmentations of the input image, we incorporate a contrastive loss in a novel way to help representation learning of heterogeneous views and increase data efficiency [6, 15, 17]. Different from how other VL methods [25, 34, 41, 57, 59] incorporate a contrastive loss, our formulation does **not** require annotated pairs (*e.g.* human annotated image-caption pairs in MS-COCO or image-text pairs scraped from the internet [44]) and can work with unlabeled image-only data to achieve better performance.

Also crucially, for (2) we devise a hierarchical decoder layer, which modifies the standard transformer decoder layer

by introducing two-tiered cross-attention modules. The hierarchical decoder first aggregates *within* each view at the token level and then aggregates *across* views at the view level. By introducing this hierarchical aggregating structure, we can better model the effectiveness of views and adaptively weigh them according to their effectiveness. For example, in Section 4 Experiment, we show that if we add noise to a certain view or mask out a prominent region of the input image, the proposed hierarchical decoder indeed down-weigh the noised and masked view when generating words and captions.

To sum up, in this paper, given a set of heterogeneous views of the input image from existing works, we focus on how to **efficiently** and **effectively** leverage these views and make the following contributions: (1) regard heterogeneous views as augmentations of the input image and propose a novel use of contrastive loss to improve computation, parameter, and data efficiency; (2) devise a hierarchical decoder layer to model the effectiveness of each view and weigh each view accordingly for caption generation; (3) achieve significant improvement of +5.6% CIDEr on MS-COCO over state of the art, and achieve comparable or often better performance compared with methods using large-scale transformer pre-training even though we do not do so; and (4) provide thorough ablations and rigorous analyses to validate our proposed method for efficiency and effectiveness.

2. Related Works

Image captioning. The goal of image captioning is to generate text descriptions for an image. It can be roughly divided into two settings: (1) trained-from-scratch and (2) with pre-training, depending on whether the model is pre-trained on a large image-text corpus or not. For the trained-from-scratch setting [4, 10, 12, 22, 26, 28, 33, 40, 43, 57], researchers train the model only on the image captioning dataset such as MS-COCO [8, 37]. They mainly focus on model architecture improvements and/or better training losses, and tackle the image captioning task specifically. On the other hand, for the pre-training setting [9, 21, 34–36, 38, 48, 49, 54, 60, 61], researchers first pre-train the model on a large image-text corpus and then fine-tune to various downstream VL tasks. They mainly focus on how to effectively pre-train the model so that it transfers well to a broad set of downstream VL tasks. In this paper, we work on the trained-from-scratch setting for image captioning specifically.

Image encodings. One important aspect of VL approaches is to properly encode relevant information from the input image, on which the captioning model is conditioned for captions generation. [4] proposed to encode finer-grained information of the input image into a sequence of objects detected by an object detector pre-trained on Visual Genome [30]. This method achieved great success and soon became the dominant approach in many VL

tasks [9, 35, 36, 38, 48, 49]. However, recent works found that the object detector pre-trained on visual genome fails to encode important information for image captioning. To tackle this problem VinVL [61] proposed to train the object detector on much larger training corpora that combine multiple publicly annotated object detection datasets. Thanks to advances in foundation models [25, 29, 41] trained on large-scale datasets, some works [12, 34, 45, 55] propose to encode the input image by a foundation model into image grid features and achieved good performance.

Nevertheless, these works still encode an input image from a single point of view by a single pre-trained model, and may still fail to encode some important information for image captioning. For example, a pre-trained object detector may fail to encode important information such as scene information or object relationships [31]. To encode complementary information, recent works proposed to include other heterogeneous “*views*” of the input image along with the detected objects. For example, OSCAR [36] use image-level object tags along with detected object for VL pre-training. Some works [24, 48, 58, 62] include an extra token of global image features encoded by an ImageNet pre-trained image encoder. Xmodal-Ctx [31] include text descriptions of the input image and image sub-regions by CLIP cross-modal retrieval. In this paper, we use heterogeneous views from existing works including detected objects, image grid features, and text descriptions, and focus on how to leverage these views efficiently and effectively.

Heterogeneous view image captioning. To leverage heterogeneous views encoded by different pre-trained models and potentially encoded into different modalities, common approaches include (1)

zka separate model for each view [10, 22, 26, 43], which trains $|V|$ (number of views) models and average the predicted word probability from each model, and (2) concatenation of views [31, 36, 61], which concatenates the views, each represented by a sequence of tokens, into a long single view, and feeds the concatenate view into an image captioning model for caption generation. However, having a separate model per view is parameter inefficient as the number of models and parameters grows linearly with $|V|$. Concatenated views are computationally inefficient as the computational complexity from the transformer model is quadratic with respect to $|V|$. In this paper, in pursuit of efficiency and effectiveness, we encode the views independently with a shared encoder, and devise a novel hierarchical decoder to first aggregate within each view at the token level and then aggregate across views at the view level.

Representation learning. To properly encode heterogeneous views, existing methods [2, 3, 20, 34] typically encode each view with a dedicated encoder. This is parameter inefficient as the number of encoders grows linearly with respect to $|V|$. This may also not be as data efficient and

be prone to overfitting on relatively smaller-scale datasets such as MS-COCO. To learn a better representation of heterogeneous views, we propose to incorporate a contrastive loss, which facilitates superior self-/un-supervised representation learning [5, 7, 17, 50], and is beneficial in low-label settings [6, 15, 17]. Driven by the contrastive loss, the representation of views (augmentations) from the same image are pulled together in an embedding space while that from different images are pushed apart. Different from existing multi-modal works [25, 34, 41, 57, 59] that incorporate contrastive loss only in the pre-training stage, or require annotated pairs (e.g. human annotated image-caption pairs in MS-COCO), our method incorporates contrastive loss together with the target image captioning task, and requires **no** annotated pairs.

3. Method

3.1. Overview

In image captioning, given heterogeneous views of the input image from existing works (see Figure 1) such as detected objects [4], image grid features [45], and text descriptions [31], we aim to **efficiently** and **effectively** leverage these views for the target image captioning (IC) task. Following the successful paradigm of image captioning [4, 10, 31], an input image x is first encoded by frozen pre-trained models into heterogeneous views $\{v_1, v_2, \dots, v_{|V|}\}$, each v represented by a sequence of d -dimensional tokens. These views v are then fed into an IC transformer encoder-decoder model, where the IC-encoder encodes important information from v into u for the later decoding step, and the IC-decoder generates captions y conditioned on the encoded views u . In the rest of the paper, since the input image is encoded into v by some frozen pre-trained models from existing works and is not the focus of this paper, for the sake of brevity, we use “*input views*” or simply “*views*” for v , “*encoded views*” for u , and “*encode*” for the encoding process by the IC-encoder.

In Section 3.2 we revisit other successful paradigms in heterogeneous view image captioning: methods that use a separate entire model per view [10, 22, 26, 43] and concatenation of views [31, 36, 61]. By comparing the pros and cons of these two methods, we aim to propose our method that can efficiently account for the effectiveness of input heterogeneous views for caption generation. In Section 3.3, we first propose to regard input views v as augmentations of the input image, and thus naturally choose to encode each v *independently* with a *shared* IC-encoder into u . This formulation is more efficient in terms of computation, parameter, and label. It also allows us to add a contrastive loss in a novel way to help representation learning of the encoded views u , and increase label efficiency. In Section 3.4, we then propose a *hierarchical decoder layer* to first aggregate *within* each encoded view u at the token level and then aggregate *across* all u at the view level. By introducing this

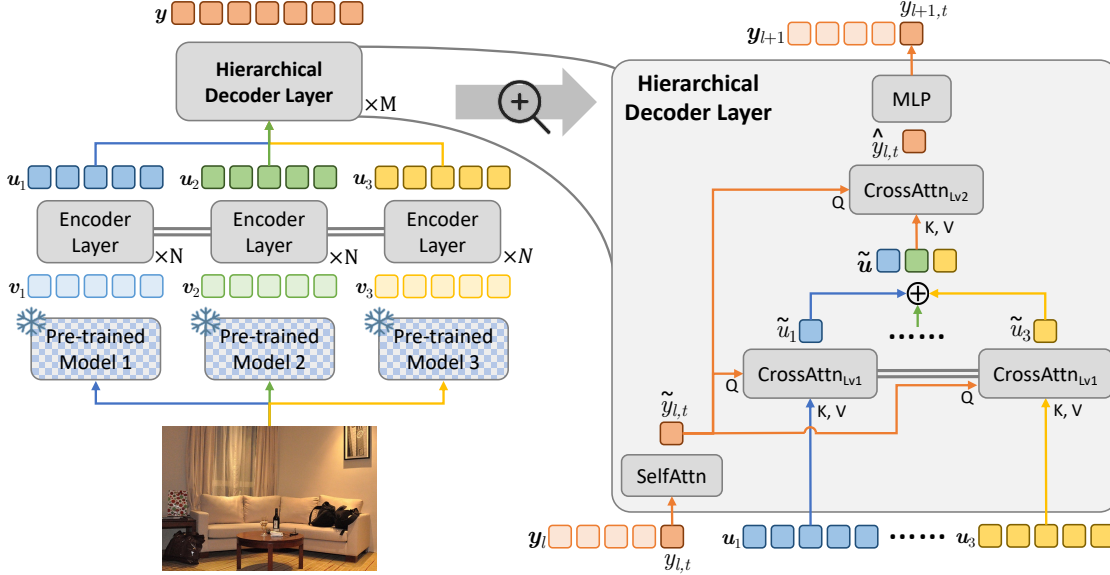


Figure 2. Model architecture. **(Left)** An input image is first encoded by some frozen pre-trained models into a set of heterogeneous views $\{v_1, v_2, v_3\}$, each represented by a sequence of d -dimensional tokens \square . These input views v are regarded as augmentations of the input image, and thus are encoded independently with a shared transformer encoder into $\{u_1, u_2, u_3\}$. **(Right)** To model the effectiveness of the encoded views u for current word prediction, we propose a hierarchical decoder layer that first aggregates within each encoded view u at the token level with a shared CrossAttn_{Lv1} module, and then aggregates across the encoded views at the view level with a CrossAttn_{Lv2} module. For clarity of illustration, we only show operations for the current $y_{l,t}$ token and not the $y_{l,1:t-1}$ past tokens.

hierarchical aggregating structure, the decoder better models the effectiveness of encoded views u , and adaptively weighs them according to their effectiveness. The overall model architecture is illustrated in Figure 2.

3.2. Heterogeneous Views Image Captioning

We start by reviewing two closely related formulations in existing works: (1) combining separate models for each view [10, 22, 26, 43] and (2) concatenation of input views [31, 36, 61]. For (1), $|V|$ captioning models are trained independently for each input view. The word probability predicted by these models is averaged as the final prediction. For (2), all input views, each represented by a sequence of tokens, are concatenated along the sequence dimension to form a single long view. The concatenated view is then fed into a single image captioning model for caption generation.

These two formulations are two extreme cases. In (1), views are encoded and decoded independently, and the effectiveness of each view is not accounted for to generate the final prediction. The computational complexity scales linearly, but the number of parameters also scales linearly with respect to $|V|$. In contrast, in (2), views are encoded and decoded jointly, and the effectiveness of views for caption generation is modeled internally by the attention module within the encoder and decoder layers. The number of parameters remains constant, but this formulation has undesirable quadratic computational complexity with respect to $|V|$. We ask: *can we*

find a middle ground between these two formulations that enjoys the advantages of both? Specifically, we would like to model the effectiveness of views for caption generation in (2), but at the same time maintain linear computational complexity in (1) and a constant number of parameters in (2).

3.3. Heterogeneous View Encoding

We first observe that each input view v is generated from the input image x and thus contains partial information of x . This means that input views v can be regarded as augmentations of the input image x [50] that describe x from different perspectives. Just like how we use data augmentation in computer vision, we would not concatenate all of the augmented images into one single image before sending it into the model or use different encoders to encode different augmented images. Therefore, we propose to decouple the roles of the IC-encoder and the IC-decoder. Each v is encoded *independently* using a *shared* IC-encoder into u efficiently for the later decoding step. The IC-decoder then models the effectiveness of encoded views u and adaptively weighs them according to their effectiveness. We will focus on the IC-encoder in this section and the IC-decoder in the next section.

By encoding each view independently, the computational complexity scales up linearly ($\mathcal{O}(|V|)$), which is computationally efficient. By encoding each view with a shared IC-encoder, the model size stays constant ($\mathcal{O}(1)$), which is parameter efficient. Data augmentation also increases data

diversity and thus improves label efficiency. Furthermore, since we regard heterogeneous input views as augmentations of the input image, we propose to add a contrastive loss [5, 7, 17, 18] to improve the representation quality of encoded heterogeneous views u . In contrastive learning, augmented views from the same input image are positive pairs and those from different images are negative pairs. The representation of positive pairs is pulled together in an embedding space while the representation of negative pairs is pushed apart. To further increase diversity across heterogeneous views for the contrastive loss [5, 32, 46, 50], we add channel-wise and sequence-wise dropout during training. For channel-wise dropout, each channel across views and sequences is randomly zero-ed out with probability p_c . For sequence-wise dropout, each token, which is a d -dimensional feature vector, across views is randomly zero-ed out with probability p_s . Note that our goal is not to propose a new contrastive loss here but to incorporate an existing one in a novel way in the VL setting. Compared to how most other VL pre-training methods incorporate a contrastive loss [25, 34, 41, 57, 59], where paired annotations such as image and text pairs annotated by human or scraped from the internet are required, our formulation only requires input image *without* paired annotations to construct positive and negative pairs.

3.4. Heterogeneous View Decoding

Given the encoded views u from the IC-encoder, in this section we shift our attention to the IC-decoder to account for the effectiveness of u for predicting each word. Specifically, the IC-decoder should adaptively weigh each encoded view u according to their effectiveness at each word prediction step conditioned on previously generated words and all other encoded views u . For example, in Figure 1, when predicting the current word “sofa” for the incomplete sentence “Black bags sitting on top of a ?”, if say, the object detector fails to detect sofa, the decoder should down-weight the view of detected objects and should leverage other views that properly encode sofa information to correctly predict the next word. On the other hand, we spent great efforts in Section 3.3 to encode the views efficiently. Therefore, the proposed decoding strategy should be at least as efficient as the proposed encoding strategy.

We propose a **hierarchical decoder layer** shown in Figure 2, which modifies the standard transformer decoder layer by introducing the two-tiered cross-attention structure. First, current word $y_{l,t}$ is contextualized with all previously generated words $y_{l,1:t-1}$ into $\tilde{y}_{l,t}$ by the SelfAttn module. Next, to model the effectiveness of each encoded view for current word prediction, the hierarchical decoder aggregates *within* each encoded view u independently at the token level by a shared CrossAttn_{L_{V1}} module. The outputs of CrossAttn_{L_{V1}} from all views are collected into \tilde{u} . Finally, to adaptively weigh all the views according to their effectiveness for cur-

rent word prediction, the hierarchical decoder then aggregates *across* views \tilde{u} at the view level by the CrossAttn_{L_{V2}} module. We also add a view-wise dropout to randomly drop an encoded view u with probability p_v during training to encourage the decoder to better leverage all encoded views instead of focusing on a few certain ones. By introducing this hierarchical aggregation structure, we can better model the effectiveness of the encoded views u by CrossAttn_{L_{V1}} and adaptively weigh them by CrossAttn_{L_{V2}}.

In terms of computation and parameter efficiency, the hierarchical decoder layer uses a shared CrossAttn_{L_{V1}} module to first aggregate within each view. The computational complexity scales up linearly, and the model size is constant with respect to $|V|$. It then aggregates across views using a CrossAttn_{L_{V2}} module with linear computational complexity, and constant model size with respect to $|V|$. Overall, the computational complexity and model size does not exceed our proposed view encoding strategy in Section 3.3.

4. Experiment

4.1. Implementation Details

Our proposed HAAV can be easily incorporated into existing transformer encoder-decoder IC models. In this paper, we choose the state-of-art Xmodal-Ctx [31] as our base model. Following the conventional training procedure [10, 36, 52], the model is first trained with cross-entropy loss and then fine-tuned with SCST [43] loss, which optimizes CIDEr score by reinforcement learning (note this is done by all compared methods). The dropout rates p_c , p_s , and p_t are all set to 0.1 following [14] without tuning. We reuse the heterogeneous views from previous methods without modifications, including detected objects from [4], CLIP ViT-B/32 [41] image grid features from [45], and text descriptions of image regions from [31]. For the contrastive loss, we prepend a learnable [CLS] token for each input view v before sending it into the IC-encoder and use the encoded [CLS] token as a view-level representation for computing the contrastive loss. We use MoCo-v2 [24] with an exponential moving average (EMA) transformer encoder and a memory buffer to compute the contrastive loss. More implementation details and hyper-parameters can be found in the supplementary.

4.2. Main Results

In Table 1, we show the results of our HAAV on the test set of MS-COCO Karpathy split [27]. It is worth noting that HAAV is **not** pre-trained on external image-and-text corpus and instead is trained from scratch only on the MS-COCO image captioning dataset [8]. Therefore, for a fair comparison, we separate these two different settings in Table 1, where methods on the top block use pre-training while those on the bottom are trained from scratch. In the same trained-from-scratch setting, our HAAV outperforms pre-

Table 1. Image captioning results on the test set of MS-COCO Karpathy split [27]. Since our HAAV is trained from scratch, we separately compare methods with large-scale transformer pre-training on the top block, and those trained from scratch on the bottom block. Our HAAV outperforms previous trained-from-scratch methods by a large margin and achieves comparable or often better performance compared to methods with large-scale transformer pre-training.

Method	Pretrain Data	B-4	M	C	S
VLP [63]	3M	39.5	29.3	129.3	23.2
X-VLM [60]	16M	40.4	-	139.3	-
Oscar [36]	6.5M	40.5	29.7	137.6	22.8
VinVL [61]	8.8M	40.9	30.9	140.4	25.1
GIT [53]	4M	41.3	30.4	139.1	24.3
ViTCap [13]	4M	41.2	30.1	138.1	24.1
SimVLM _{base}	39.0	32.9	134.8	24.0	
SCST [43]	None	34.2	26.7	114.0	-
Up-Down [4]	None	36.3	27.7	120.1	21.4
AoANet [22]	None	38.9	29.2	129.8	22.4
\mathcal{M}^2 [10]	None	39.1	29.1	131.2	22.6
CLIP-ViL [45]	None	40.2	29.7	134.2	23.8
X-LAN [40]	None	39.5	29.5	132.0	23.4
DLCT [39]	None	39.8	29.5	133.8	23.0
Xmodal-Ctx [31]	None	39.7	30.0	135.9	23.7
HAAV (ours)	None	41.0	30.2	141.5	23.9

Table 2. Image captioning results on a smaller-scale Flickr30K Karpathy split [27]. We also demonstrate semi-supervised learning (SSL) for image captioning with labeled data from Flickr30K and unlabeled data from MS-COCO. Compared to previous state of the arts, our HAAV achieves significant performance improvement, which may come from the label efficiency of our method.

Method	B-4	M	C	S
Show & Tell [52]	21.5	18.3	41.7	12.2
Show, Attend & Tell [56]	23.6	19.2	49.1	13.3
Up-Down [4]	28.3	21.6	63.3	15.9
\mathcal{M}^2 [10]	29.8	22.4	68.4	16.2
ORT [19]	30.1	22.8	68.8	16.9
HAAV (ours)	34.3	24.6	81.7	18.0
HAAV + SSL (ours)	34.3	25.1	85.6	19.0

vious state-of-art Xmodal-Ctx [31] by 5.6% in CIDEr and 1.3% in BLEU-4. On the other hand, when compared with methods with transformer pre-training, our HAAV, despite being only trained on MS-COCO, achieves comparable or often better performance.

To further demonstrate the label efficiency of HAAV, we train it on the Karpathy split [27] of Flickr30K, which only has ~ 0.15 M training data (4x smaller than MS-COCO), and

show the results in Table 2. The results of other methods are taken from [47]. We can see that our HAAV outperforms previous state-of-art ORT substantially by 12.9% in CIDEr and 4.3% in BLEU-4. The significant improvement may come from the label efficiency of our method, which is particularly important when trained on the smaller-scale dataset of Flickr30K.

4.3. Ablations and Analyses

The goal of this paper is to propose an *efficient* and *effective* way to leverage heterogeneous views. Therefore, in this section we closely examine whether our proposed HAAV achieves these goals in Section 4.3.1 for efficiency and Section 4.3.2 for effectiveness. Following the convention in [19, 22, 31], we only train the model with cross-entropy loss for all ablations and analyses.

4.3.1 Computation, parameter, and label efficient

In Section 3.3, we propose to regard input views v as augmentations of the input image and encode the views independently with a shared IC-encoder. To demonstrate the efficiency of our method, in Table 3, we show the theoretical computation and parameter complexity as well as the actual training speed and trainable parameters. Since we do not want to trade performance in pursuit of efficiency, we also show in Table 3 that our method achieves better performance despite being more efficient.

Comparison with common aggregation approaches.

As described in Section 3.2, other common approaches include (1) using a separate model per view [4, 10, 22, 43], which trains $|V|$ IC models, one for each input view, and averages the predictions of word probability from each IC model as the final word prediction, and (2) concatenation of views [31, 36], which concatenates all input views along the sequence dimension into a long single view, and feeds the concatenated input views into the IC model for caption generation. In Table 3, using a separate model per view has to train $|V|$ models and thus is parameters inefficient ($\mathcal{O}(|V|)$). Concatenation of views concatenates all views into one long single view and is computationally expensive ($\mathcal{O}(|V|^2)$). Compared to these two approaches, our HAAV has linear computation ($\mathcal{O}(|V|)$) and constant parameter ($\mathcal{O}(1)$) complexity, and performs consistently better across all metrics by large margins.

Shared v.s. unshared encoder. Another design choice is whether or not to use a shared encoder for each view. In Table 3, in the case of training without a contrastive loss (*w/o* \mathcal{L}_{con}), using unshared encoders [2, 3, 20, 34] for each view does not bring any performance benefits compared to using a shared encoder, but with the cost of higher parameter complexity ($\mathcal{O}(|V|)$). Furthermore, we ablate the proposed contrastive loss \mathcal{L}_{con} , and found that the unshared encoder

Table 3. Comparison with other encoding strategies in terms of efficiency and performance. We train the models on MS-COCO Karpathy split [27] and report the validation results. All models are trained with cross-entropy loss only (*w/o* SCST fine-tuning) following the convention in [19, 22, 31]. Our strategy of shared encoder trained with an additional contrastive loss \mathcal{L}_{con} achieves the best performance and is computation and parameter efficient. Please see Section 4.3.1 for more details of different conditions.

Conditions	\mathcal{L}_{con}	Computation		Parameter		Performance			
		complexity	iter/sec \uparrow	complexity	#params \downarrow	B-4	M	C	S
Model per View		$\mathcal{O}(V)$	2.23	$\mathcal{O}(V)$	52.5M	38.5	28.6	121.1	21.3
Concatenated views		$\mathcal{O}(V ^2)$	4.20	$\mathcal{O}(1)$	13.1M	38.5	28.7	122.8	21.7
Unshared encoders		$\mathcal{O}(V)$	5.33	$\mathcal{O}(V)$	20.8M	39.7	29.1	125.4	22.1
Unshared encoders	\checkmark	$\mathcal{O}(V)$	3.96	$\mathcal{O}(V)$	22.7M	39.7	29.3	125.8	22.2
Shared encoder		$\mathcal{O}(V)$	5.97	$\mathcal{O}(1)$	13.5M	39.7	29.1	125.6	22.1
Ours (Shared encoder)	\checkmark	$\mathcal{O}(V)$	4.83	$\mathcal{O}(1)$	15.4M	40.5	29.4	127.6	22.3




Figure 3. Data efficiency of HAAV. HAAV (blue) only needs about 40-50% labeled data to achieve similar performance as other common methods for heterogeneous view image captioning trained on 100% of data such as using concatenated views (red) and a model per view (gold). All models are trained with cross-entropy loss only.

does not benefit from \mathcal{L}_{con} as much as the shared encoder. Although \mathcal{L}_{con} introduces additional overhead such as extra parameters from the projection heads and extra computation from pairwise similarity during the training time (not for inference), the increase in complexity is moderate but the improvement in performance is significant when incorporated into our shared-encoder strategy.

Label efficiency with fewer training data. To test the label efficiency of HAAV, in Figure 3 we train the model with $\{10, 20, \dots, 90\}\%$ of labeled data, and compare the CIDEr score with other common approaches for heterogeneous view image captioning, using concatenation of views (red) and a separate model per view (gold), trained on 100% of data. We can see that HAAV only needs about 40-50% labeled data to achieve comparable performance. With the same input views and similar model architectures, this indicates that our method is more label efficient, likely due to our novel use of heterogeneous views as augmentations and the contrastive

loss \mathcal{L}_{con} . In the supplementary, we also show that HAAV suffers less from overfitting on MS-COCO in the supplementary, indicating that our model is more label efficient.

Semi-supervised training. To demonstrate label efficiency brought by the contrastive loss \mathcal{L}_{con} even further, we train a semi-supervised image captioning model. Specifically, the model is trained on the labeled data (image-caption pairs) from Flickr30K, and unlabeled data (image only) from MS-COCO. Since our contrastive loss is applied differently than other VL methods and does not require annotated pairs (*e.g.* human-annotated image-caption pairs in MS-COCO), it can be applied on the unlabeled image-only data to aid representation learning of encoded views u . In Table 2, with semi-supervised training, the already strong HAAV achieves an even better performance of +3.9% CIDEr and +1% SPICE.

4.3.2 Effectiveness of hierarchical decoder layer

In Section 3.4, we propose a hierarchical decoder layer to model the effectiveness of encoded views and adaptively weigh them according to their effectiveness for caption generation. To verify if the proposed hierarchical decoder achieves this goal, we monitor the view-level attention weights of $\text{CrossAttn}_{L_{v2}}$ in two control studies. To show the architectural design of the proposed hierarchical decoder layer is effective, we compare with other common designs in Table 4.

Adaptive view-level attention weights. We design two control studies in Figure 4 to show how the view-level attention weights of $\text{CrossAttn}_{L_{v2}}$ vary adaptively according to the effectiveness of an input view. In the first experiment, we add noise to a view by randomly zeroing out tokens in a view to make a view *less effective*, and expect a drop of weights toward that noised view. To measure the weights, we take the multi-head attention weights of $\text{CrossAttn}_{L_{v2}}$ at the last decoder layer and average the attention weights across heads. In Figure 4a, the weights for the noised view drops consistently at each word prediction step compared to the same view *without* added noise. This means that our hierarchical

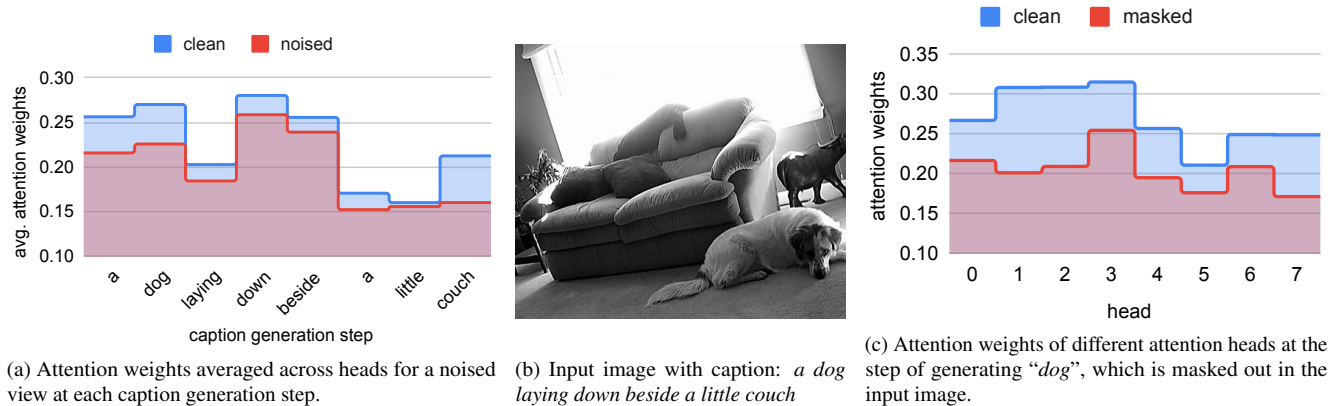


Figure 4. The view-level attention weights of $\text{CrossAttn}_{L_{v2}}$ vary adaptively according to the effectiveness of a view at the view level and at the word level when generating a caption for the center image. (a) We add random noise to a view and show that the attention weights averaged across different attention heads drop consistently at each step of caption generation. (c) We mask out dog in the input image and show that the attention weights of different attention heads drop consistently at the step of generating the word “dog”.

Table 4. Ablations of design choices for the proposed hierarchical decoder layer. Our proposed two-tiered cross-attention architecture works the best. The models are trained with cross-entropy only.

Method	B-4	M	C	S
Concatenate	39.2	29.2	125.1	22.1
Mean pool	39.4	29.2	125.2	22.2
Max pool	38.7	28.8	122.9	21.5
Sigmoid	39.4	29.2	125.5	22.2
Tanh	39.4	29.3	124.8	22.0
Ours	40.5	29.4	127.6	22.3

decoder indeed learns to adaptively weigh the views according to their effectiveness at the view level. In the second experiment, we randomly mask out a prominent region of the input image for a view. For example, we mask out dog in the input image (Figure 4b) with caption “*a dog laying down beside a little couch*” to make a view *less effective* at the step of generating the word “dog”. We expect a drop of the weights toward the masked view at the step of generating the word “dog”. To measure the weights, we take the multi-head attention weights of $\text{CrossAttn}_{L_{v2}}$ at the last decoder layer and measure the attention weights of each head at the step of generating the word “dog”. In Figure 4c, the weights for the masked view drops consistently across all attention heads compared to the same view without masking. This means that our hierarchical decoder indeed learns to adaptively weigh the input views according to their usefulness at the word level. The example in Figure 4 is randomly chosen and more examples can be found in the supplementary.

Design choices of the hierarchical decoder. In Table 4, we ablate different design choices of the hierarchical

decoder. One simple alternative is to concatenate the encoded views u along the sequence dimension into a long single view, and use a single cross-attention module to jointly aggregate the views at the token level. One can also mean-/max-pool to aggregate across views in place of the $\text{CrossAttn}_{L_{v2}}$ module. Other works also propose to use a sigmoid/tanh gating mechanism [22] to aggregate across layers [10], which can be generalized to aggregating across views. In Table 4, we can see that our proposed hierarchical decoder layer with a two-tiered cross-attention structure achieves the best performance compared to all other designs.

5. Conclusion

In this paper, we focus on the problem of how to *efficiently* and *effectively* leverage heterogeneous views v . To tackle this problem, we propose HAAV to (1) regard heterogeneous views as augmentations of the input image, and naturally encode each view independently with a shared encoder in an efficient manner, (2) incorporate a contrastive loss across encoded views u to improve representation quality and label efficiency (especially semi-supervised training to leverage image-only unlabeled data), and (3) carefully design a hierarchical decoder layer that first aggregates within each view at the token level and then across views at the view level to account for the effectiveness of each encoded view u . Through rigorous analysis, we show HAAV is computation, parameter, and data efficient, and outperforms other less efficient designs and existing approaches for heterogeneous views image captioning. We also demonstrate that our hierarchical decoder successfully models the effectiveness of views and weigh them adaptively according to their effectiveness. We demonstrate significant performance improvements of +5.6% CIDEr compared to state-of-art *w/o* transformer pre-training on MS-COCO and +16.8% CIDEr with SSL on Flickr30K.

References

- [1] Harsh Agrawal, Karan Desai, Yufei Wang, Xinlei Chen, Rishabh Jain, Mark Johnson, Dhruv Batra, Devi Parikh, Stefan Lee, and Peter Anderson. Nocaps: Novel object captioning at scale. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8948–8957, 2019. [1](#)
- [2] Hassan Akbari, Liangzhe Yuan, Rui Qian, Wei-Hong Chuang, Shih-Fu Chang, Yin Cui, and Boqing Gong. Vatt: Transformers for multimodal self-supervised learning from raw video, audio and text. *Advances in Neural Information Processing Systems*, 34:24206–24221, 2021. [2](#), [3](#), [6](#)
- [3] Jean-Baptiste Alayrac, Adria Recasens, Rosalia Schneider, Relja Arandjelović, Jason Ramapuram, Jeffrey De Fauw, Lucas Smaira, Sander Dieleman, and Andrew Zisserman. Self-supervised multimodal versatile networks. *Advances in Neural Information Processing Systems*, 33:25–37, 2020. [2](#), [3](#), [6](#)
- [4] Peter Anderson, Xiaodong He, Chris Buehler, Damien Teney, Mark Johnson, Stephen Gould, and Lei Zhang. Bottom-up and top-down attention for image captioning and visual question answering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 6077–6086, 2018. [1](#), [2](#), [3](#), [5](#), [6](#)
- [5] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, pages 1597–1607. PMLR, 2020. [3](#), [5](#)
- [6] Ting Chen, Simon Kornblith, Kevin Swersky, Mohammad Norouzi, and Geoffrey E Hinton. Big self-supervised models are strong semi-supervised learners. *Advances in neural information processing systems*, 33:22243–22255, 2020. [2](#), [3](#)
- [7] Xinlei Chen, Haoqi Fan, Ross Girshick, and Kaiming He. Improved baselines with momentum contrastive learning. *arXiv preprint arXiv:2003.04297*, 2020. [3](#), [5](#)
- [8] Xinlei Chen, Hao Fang, Tsung-Yi Lin, Ramakrishna Vedantam, Saurabh Gupta, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco captions: Data collection and evaluation server. *arXiv preprint arXiv:1504.00325*, 2015. [1](#), [2](#), [5](#)
- [9] Yen-Chun Chen, Linjie Li, Licheng Yu, Ahmed El Kholy, Faisal Ahmed, Zhe Gan, Yu Cheng, and Jingjing Liu. Uniter: Universal image-text representation learning. In *European conference on computer vision*, pages 104–120. Springer, 2020. [2](#), [3](#)
- [10] Marcella Cornia, Matteo Stefanini, Lorenzo Baraldi, and Rita Cucchiara. Meshed-memory transformer for image captioning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10578–10587, 2020. [2](#), [3](#), [4](#), [5](#), [6](#), [8](#)
- [11] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. *ICLR*, 2021. [2](#)
- [12] Zi-Yi Dou, Yichong Xu, Zhe Gan, Jianfeng Wang, Shuohang Wang, Lijuan Wang, Chenguang Zhu, Nanyun (Violet) Peng, Zicheng Liu, and Michael Zeng. An empirical study of training end-to-end vision-and-language transformers. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022., June 2022. [2](#), [3](#)
- [13] Zhiyuan Fang, Jianfeng Wang, Xiaowei Hu, Lin Liang, Zhe Gan, Lijuan Wang, Yezhou Yang, and Zicheng Liu. Injecting semantic concepts into end-to-end image captioning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18009–18019, 2022. [6](#)
- [14] Tianyu Gao, Xingcheng Yao, and Danqi Chen. SimCSE: Simple contrastive learning of sentence embeddings. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 6894–6910, Online and Punta Cana, Dominican Republic, Nov. 2021. Association for Computational Linguistics. [5](#)
- [15] Priya Goyal, Mathilde Caron, Benjamin Lefaudeux, Min Xu, Pengchao Wang, Vivek Pai, Mannat Singh, Vitaliy Liptchinsky, Ishan Misra, Armand Joulin, et al. Self-supervised pretraining of visual features in the wild. *arXiv preprint arXiv:2103.01988*, 2021. [2](#), [3](#)
- [16] Yash Goyal, Tejas Khot, Douglas Summers-Stay, Dhruv Batra, and Devi Parikh. Making the v in vqa matter: Elevating the role of image understanding in visual question answering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 6904–6913, 2017. [1](#)
- [17] Jean-Bastien Grill, Florian Strub, Florent Altché, Corentin Tallec, Pierre Richemond, Elena Buchatskaya, Carl Doersch, Bernardo Avila Pires, Zhaohan Guo, Mohammad Gheshlaghi Azar, et al. Bootstrap your own latent—a new approach to self-supervised learning. *Advances in Neural Information Processing Systems*, 33:21271–21284, 2020. [2](#), [3](#), [5](#)
- [18] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. *arXiv preprint arXiv:1911.05722*, 2019. [5](#)
- [19] Simao Herdade, Armin Kappeler, Kofi Boakye, and Joao Soares. Image captioning: Transforming objects into words. *Advances in Neural Information Processing Systems*, 32, 2019. [6](#), [7](#)
- [20] Ronghang Hu and Amanpreet Singh. Unit: Multimodal multitask learning with a unified transformer. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1439–1449, 2021. [2](#), [3](#), [6](#)
- [21] Xiaowei Hu, Zhe Gan, Jianfeng Wang, Zhengyuan Yang, Zicheng Liu, Yumao Lu, and Lijuan Wang. Scaling up vision-language pre-training for image captioning. *arXiv preprint arXiv:2111.12233*, 2021. [2](#)
- [22] Lun Huang, Wenmin Wang, Jie Chen, and Xiao-Yong Wei. Attention on attention for image captioning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4634–4643, 2019. [2](#), [3](#), [4](#), [6](#), [7](#), [8](#)
- [23] Drew A Hudson and Christopher D Manning. Gqa: A new dataset for real-world visual reasoning and compositional question answering. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 6700–6709, 2019. [1](#)
- [24] Jiayi Ji, Yunpeng Luo, Xiaoshuai Sun, Fuhai Chen, Gen Luo, Yongjian Wu, Yue Gao, and Rongrong Ji. Improving

- image captioning by leveraging intra-and inter-layer global representation in transformer network. In *Proceedings of the AAAI conference on artificial intelligence*, volume 35, pages 1655–1663, 2021. 1, 3, 5
- [25] Chao Jia, Yinfei Yang, Ye Xia, Yi-Ting Chen, Zarana Parekh, Hieu Pham, Quoc Le, Yun-Hsuan Sung, Zhen Li, and Tom Duerig. Scaling up visual and vision-language representation learning with noisy text supervision. In *International Conference on Machine Learning*, pages 4904–4916. PMLR, 2021. 2, 3, 5
- [26] Wenhao Jiang, Lin Ma, Yu-Gang Jiang, Wei Liu, and Tong Zhang. Recurrent fusion network for image captioning. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 499–515, 2018. 2, 3, 4
- [27] Andrej Karpathy and Li Fei-Fei. Deep visual-semantic alignments for generating image descriptions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3128–3137, 2015. 5, 6, 7
- [28] Wonjae Kim, Bokyung Son, and Ildoo Kim. Vilt: Vision-and-language transformer without convolution or region supervision. In *International Conference on Machine Learning*, pages 5583–5594. PMLR, 2021. 2
- [29] Alexander Kolesnikov, Lucas Beyer, Xiaohua Zhai, Joan Puigcerver, Jessica Yung, Sylvain Gelly, and Neil Houlsby. Big transfer (bit): General visual representation learning. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part V 16*, pages 491–507. Springer, 2020. 3
- [30] Ranjay Krishna, Yuke Zhu, Oliver Groth, Justin Johnson, Kenji Hata, Joshua Kravitz, Stephanie Chen, Yannis Kalantidis, Li-Jia Li, David A Shamma, et al. Visual genome: Connecting language and vision using crowdsourced dense image annotations. *International journal of computer vision*, 123(1):32–73, 2017. 1, 2
- [31] Chia-Wen Kuo and Zsolt Kira. Beyond a pre-trained object detector: Cross-modal textual and visual context for image captioning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022. 1, 2, 3, 4, 5, 6, 7
- [32] Chia-Wen Kuo, Chih-Yao Ma, Jia-Bin Huang, and Zsolt Kira. Featmatch: Feature-based augmentation for semi-supervised learning. In *European Conference on Computer Vision*, pages 479–495. Springer, 2020. 5
- [33] Junnan Li, Dongxu Li, Caiming Xiong, and Steven Hoi. Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation. *arXiv preprint arXiv:2201.12086*, 2022. 2
- [34] Junnan Li, Ramprasaath Selvaraju, Akhilesh Gotmare, Shafiq Joty, Caiming Xiong, and Steven Chu Hong Hoi. Align before fuse: Vision and language representation learning with momentum distillation. *Advances in Neural Information Processing Systems*, 34, 2021. 2, 3, 5, 6
- [35] Liunian Harold Li, Mark Yatskar, Da Yin, Cho-Jui Hsieh, and Kai-Wei Chang. Visualbert: A simple and performant baseline for vision and language. *arXiv preprint arXiv:1908.03557*, 2019. 2, 3
- [36] Xiujun Li, Xi Yin, Chunyuan Li, Pengchuan Zhang, Xiaowei Hu, Lei Zhang, Lijuan Wang, Houdong Hu, Li Dong, Furu Wei, et al. Oscar: Object-semantics aligned pre-training for vision-language tasks. In *European Conference on Computer Vision*, pages 121–137. Springer, 2020. 1, 2, 3, 4, 5, 6
- [37] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *European conference on computer vision*, pages 740–755. Springer, 2014. 2
- [38] Jiasen Lu, Dhruv Batra, Devi Parikh, and Stefan Lee. Vilbert: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks. *arXiv preprint arXiv:1908.02265*, 2019. 2, 3
- [39] Yunpeng Luo, Jiayi Ji, Xiaoshuai Sun, Liujuan Cao, Yongjian Wu, Feiyue Huang, Chia-Wen Lin, and Rongrong Ji. Dual-level collaborative transformer for image captioning. In *Proceedings of the AAAI conference on artificial intelligence*, volume 35, pages 2286–2293, 2021. 6
- [40] Yingwei Pan, Ting Yao, Yehao Li, and Tao Mei. X-linear attention networks for image captioning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10971–10980, 2020. 2, 6
- [41] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning transferable visual models from natural language supervision. In *ICML*, 2021. 2, 3, 5
- [42] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. *Advances in neural information processing systems*, 28:91–99, 2015. 1
- [43] Steven J Rennie, Etienne Marcheret, Youssef Mroueh, Jerret Ross, and Vaibhava Goel. Self-critical sequence training for image captioning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7008–7024, 2017. 2, 3, 4, 5, 6
- [44] Piyush Sharma, Nan Ding, Sebastian Goodman, and Radu Soricut. Conceptual captions: A cleaned, hypernymed, image alt-text dataset for automatic image captioning. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2556–2565, 2018. 2
- [45] Sheng Shen, Liunian Harold Li, Hao Tan, Mohit Bansal, Anna Rohrbach, Kai-Wei Chang, Zhewei Yao, and Kurt Keutzer. How much can clip benefit vision-and-language tasks? *arXiv preprint arXiv:2107.06383*, 2021. 1, 2, 3, 5, 6
- [46] Kihyuk Sohn, David Berthelot, Nicholas Carlini, Zizhao Zhang, Han Zhang, Colin A Raffel, Ekin Dogus Cubuk, Alexey Kurakin, and Chun-Liang Li. Fixmatch: Simplifying semi-supervised learning with consistency and confidence. *Advances in Neural Information Processing Systems*, 33:596–608, 2020. 5
- [47] Matteo Stefanini, Marcella Cornia, Lorenzo Baraldi, Silvia Cascianelli, Giuseppe Fiamenì, and Rita Cucchiara. From show to tell: a survey on deep learning-based image captioning. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2022. 6

- [48] Weijie Su, Xizhou Zhu, Yue Cao, Bin Li, Lewei Lu, Furu Wei, and Jifeng Dai. Vl-bert: Pre-training of generic visual-linguistic representations. In *International Conference on Learning Representations*, 2020. 2, 3
- [49] Hao Tan and Mohit Bansal. Lxmert: Learning cross-modality encoder representations from transformers. *arXiv preprint arXiv:1908.07490*, 2019. 2, 3
- [50] Yonglong Tian, Chen Sun, Ben Poole, Dilip Krishnan, Cordelia Schmid, and Phillip Isola. What makes for good views for contrastive learning? *Advances in Neural Information Processing Systems*, 33:6827–6839, 2020. 3, 4, 5
- [51] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017. 2
- [52] Oriol Vinyals, Alexander Toshev, Samy Bengio, and Dumitru Erhan. Show and tell: A neural image caption generator. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3156–3164, 2015. 5, 6
- [53] Jianfeng Wang, Zhengyuan Yang, Xiaowei Hu, Linjie Li, Kevin Lin, Zhe Gan, Zicheng Liu, Ce Liu, and Lijuan Wang. Git: A generative image-to-text transformer for vision and language. *arXiv preprint arXiv:2205.14100*, 2022. 6
- [54] Peng Wang, An Yang, Rui Men, Junyang Lin, Shuai Bai, Zhikang Li, Jianxin Ma, Chang Zhou, Jingren Zhou, and Hongxia Yang. Unifying architectures, tasks, and modalities through a simple sequence-to-sequence learning framework. *arXiv preprint arXiv:2202.03052*, 2022. 2
- [55] Zirui Wang, Jiahui Yu, Adams Wei Yu, Zihang Dai, Yulia Tsvetkov, and Yuan Cao. SimVLM: Simple visual language model pretraining with weak supervision. In *International Conference on Learning Representations*, 2022. 3
- [56] Kelvin Xu, Jimmy Ba, Ryan Kiros, Kyunghyun Cho, Aaron Courville, Ruslan Salakhudinov, Rich Zemel, and Yoshua Bengio. Show, attend and tell: Neural image caption generation with visual attention. In *International conference on machine learning*, pages 2048–2057. PMLR, 2015. 6
- [57] Jinyu Yang, Jiali Duan, Son Tran, Yi Xu, Sampath Chanda, Liqun Chen, Belinda Zeng, Trishul Chilimbi, and Junzhou Huang. Vision-language pre-training with triple contrastive learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15671–15680, 2022. 2, 3, 5
- [58] Fei Yu, Jiji Tang, Weichong Yin, Yu Sun, Hao Tian, Hua Wu, and Haifeng Wang. Ernie-vil: Knowledge enhanced vision-language representations through scene graph. In *AAAI*, 2021. 3
- [59] Jiahui Yu, Zirui Wang, Vijay Vasudevan, Legg Yeung, Mojtaba Seyedhosseini, and Yonghui Wu. Coca: Contrastive captioners are image-text foundation models. *arXiv preprint arXiv:2205.01917*, 2022. 2, 3, 5
- [60] Yan Zeng, Xinsong Zhang, and Hang Li. Multi-grained vision language pre-training: Aligning texts with visual concepts. *arXiv preprint arXiv:2111.08276*, 2021. 2, 6
- [61] Pengchuan Zhang, Xiujun Li, Xiaowei Hu, Jianwei Yang, Lei Zhang, Lijuan Wang, Yejin Choi, and Jianfeng Gao. Vinvl: Revisiting visual representations in vision-language models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5579–5588, 2021. 1, 2, 3, 4, 6
- [62] Xuying Zhang, Xiaoshuai Sun, Yunpeng Luo, Jiayi Ji, Yiyi Zhou, Yongjian Wu, Feiyue Huang, and Rongrong Ji. Rstnet: Captioning with adaptive attention on visual and non-visual words. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15465–15474, 2021. 3
- [63] Luowei Zhou, Hamid Palangi, Lei Zhang, Houdong Hu, Jason Corso, and Jianfeng Gao. Unified vision-language pre-training for image captioning and vqa. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 13041–13049, 2020. 6