Vision-Language Intelligence: Tasks, Representation Learning, and Large Models

Feng Li[1,2*†] , Hao Zhang[1,2*†], Yi-Fan Zhang[3], Shilong Liu[2,4†] , Jian Guo[2], Lionel M. Ni[1,6], PengChuan Zhang[5],
Lei Zhang[2‡]
[1]The Hong Kong University of Science and Technology
[2]International Digital Economy Academy (IDEA)
[3]Chinese Academy of Science [4]Tsinghua University [5]Microsoft Research
[6]The Hong Kong University of Science and Technology (Guangzhou)

Abstract—**This paper presents a comprehensive survey of vision-language (VL) intelligence from the perspective of time. This survey is inspired by the remarkable progress in both computer vision and natural language processing, and recent trends shifting from single modality processing to multiple modality comprehension. We summarize the development in this field into three time periods, namely task-specific methods, vision-language pre-training (VLP) methods, and larger models empowered by large-scale weakly-labeled data. We first take some common VL tasks as examples to introduce the development of task-specific methods. Then we focus on VLP methods and comprehensively review key components of the model structures and training methods. After that, we show how recent work utilizes large-scale raw image-text data to learn language-aligned visual representations that generalize better on zero or few shot learning tasks. Finally, we discuss some potential future trends towards modality cooperation, unified representation, and knowledge incorporation. We believe that this review will be of help for researchers and practitioners of AI and ML, especially those interested in computer vision and natural language processing.**

I. INTRODUCTION

Computer vision (CV) and natural language processing (NLP) are sub-fields of artificial intelligence (AI) that focus on the simulation of human intelligence in vision and language. In the last decade, deep learning has greatly advanced single-modality learning in the two fields and led to state-of-the-art results on a series of tasks. At the core of the remarkable progress of deep learning lies the empowerment of rapidly evolving GPUs and the availability of large-scale datasets, which allows for accelerated training of deep models at scale.

Along with the advancement in deep learning, we have witnessed a series of powerful neural networks developed. Traditional neural networks are typically multi-layer perceptron (MLP) consisting of multiple stacked linear layers and non-linear activations (Rosenblatt, 1957, 1961). LeCun et al. (1998) proposed Convolutional Neural Network (CNN) to incorporate the shift-invariant property as a better inductive bias for 2D visual input, which inspired a large number of deep neural networks, including AlexNet (Krizhevsky et al., 2012), VGGNet (Simonyan and Zisserman, 2015a), GoogleNet (Szegedy et al., 2015), and ResNet (He et al., 2016a). Another prominent breakthrough is recurrent neural network (RNN) in the field of natural language processing (NLP), which proposed recurrent cells for sequential data modeling (Rumelhart et al., 1985; Hochreiter and Schmidhuber, 1997a). To mitigate the vanishing and exploding gradient issues in long sequence training, LSTM (Hochreiter and Schmidhuber, 1997a), a variant of RNN, and GRU (Chung et al., 2014), a more efficient version of LSTM, were proposed accordingly. Another great breakthrough in NLP is Transformer (Vaswani et al., 2017), which utilizes the attention mechanism to pursue better language representation. Using multiple stacked attention layers, Transformers can fuse information over language tokens globally with high parallelism, which facilitates both powerful representation and large-scale training.

Though inspiring progress has been achieved in single-modality domains, real-world problems often involve multiple modalities. For example, an autonomous vehicle should be able to process human orders (language), traffic signals (vision), and road conditions (vision and sounds). Even single modality learning benefits from multi-modality. For example, language learning needs perception which forms the basis of many semantic axioms (Bisk et al., 2020). Perception is the way humans understand the physical world and decides the assumption behind the human language. Since we all hear and see the same thing, we will leave some knowledge as common sense which is unwritten in our language (Bisk et al., 2020). Even restricted to language, speech contains more useful information than text-only, e.g., emotions can be implied through prosody. Noticing that multi-modal perception helps in both multi-modal and single-modal tasks, there comes a lot of research works. Within the field of multi-modality, the integration of vision and language gets much attention since vision is one of the most important perceptions for the human to understand the environment and language-aligned visual features can greatly improve the performances of both vision tasks and vision-language tasks. Moreover, the popularity of vision-language intelligence is also due to the availability of abundant datasets and benchmarks in this field.

The ambition to address many task-specific VL problems has fueled the initial development of VL learning. These VL

---

*Equal contribution.
†This work was done when Feng Li, Hao Zhang, and Shilong Liu were interns at IDEA.
‡Corresponding author.

problems include image captioning, visual question answering (VQA), image-text matching, etc. Xu et al. (2015); Karpathy et al. (2014); Vinyals et al. (2015) integrated a CNN image encoder and an RNN text decoder for image captioning. Antol et al. (2015); Yang et al. (2016); Anderson et al. (2018b) addressed the VQA task by mapping images and texts into the same latent space and predicting answers from the latent representations. Kiros et al. (2014); Karpathy et al. (2014); Huang et al. (2016); Lee et al. (2018) performed image-text matching by calculating the similarity between an image and a text either on sentence-level or token-level. These models are tailored for specific problems with various datasets and can only solve one task each.

Inspired by the prevalence of pre-training and fine-tuning in both language (Devlin et al., 2018) and vision, the inter-disciplinary field of vision and language embraces a new era: to learn a joint representation of vision and language by pre-training on image-text pairs. The surge of VLP models is mostly inspired by language models in both architecture design and training methods. For example, many recent studies (Li et al., 2019b; Lu et al., 2019; Zhang et al., 2021; Tan and Bansal, 2019; Li et al., 2020b; Yu et al., 2020; Chen et al., 2020) adopted BERT-like (Devlin et al., 2018) architectures and training methods. The development of VL learning meets a serious challenge due to the lack of sufficiently large scale manually labeled data. Recently, some studies (Radford et al., 2021; Jia et al., 2021; Wang et al., 2021; Li et al., 2021b) broke this limitation by adopting contrastive learning and making use of large-scale web-crawled data to learn visual-linguistic features which can be used for zero-shot learning.

The fast-evolving progress in the VL space urges a comprehensive survey of existing studies in this domain. This paper aims to provide a structured review of recent progress in the VL domain to help researchers gain a whole picture and better insight behind recent studies. We divide the development of VL learning into three eras. The first is from 2014 to 2018 where specialized models are designed for different tasks. The second era is from 2019 to 2021, during which joint representations of vision and language are learned by pre-training on well-labeled VL datasets. Finally, the third era began in 2021 with the appearance of CLIP (Shen et al., 2021), in which researchers seek to pre-train VL models on larger weakly-labeled datasets and to obtain a strong zero/few-shot vision model with VL pre-training.

When reviewing the whole development of VL intelligence, we find the general goal is to learn good visual representations. A good visual representation should have three attributes as summarized in (Li et al., 2021b), which are object-level, language-aligned, and semantic-rich. Object-level means the granularity of vision and language features should be as fine as in object and word-level, respectively. Language-aligned emphasizes that the vision feature aligned with language can help in vision tasks. Semantic-rich means the representation should be learned from large-scale data without domain restriction. Research works in the first era of VL aim to solve specific problems instead of learning the aforementioned good representations. In the second era, researchers train models on image-text pairs to obtain language-aligned visual features.

Some works in this era adopt detected regions for image representation to learn object-level features. Only in the third era can researchers deal with large-scale datasets and pre-train semantic-rich features.

There are also other survey paper discussing VL intelligence. Zhang et al. (2020) analyzed multi-modal deep learning from three perspectives: multimodal representation, fusion multimodal signals and multimodal applications. Mogadala et al. (2021) organized their survey by tasks. They reviewed some common tasks and corresponding methods. They also included datasets and metrics. Different from them, we view the VL intelligence from the perspective time. Further more, we include the general goal of this area and show how research works atain the goal step-by-step.

To the best of our knowledge, this is the first VL survey that summarizes studies from the viewpoint of the time period. The remainder of this paper is organized as follows. We start with some task-specific problems in VL such as image captioning, VQA, and image-text retrieval in Section II. Then we comprehensively explain the vision-language joint representation learning empowered by pre-training in Section III. Finally, we show some work learning language-aligned visual representations directly from raw image-text data with large-scale vision-language pre-training in Section VI.

II. TASK SPECIFIC PROBLEMS

Early VL methods are designed for specific tasks. The VL domain encompasses a broad range of tasks, including image captioning, VQA, image-text matching, visual grounding, and visual dialog, etc. Some common VL tasks are summarized in Table I, which shows the input, output, datasets, metrics, and mainstream methods of each task. In this section, we only introduce the three most common tasks in detail, including image captioning, VQA, and image-text matching. For these three tasks, we will introduce their task formulation and the development of mainstream methods. For the remaining tasks, a short description of each task is included. We summarize that the development of task-specific methods is from global representations to fine-grained object-centric representations. Most VL tasks experience three stages. The first stage is **gloabl vector representation and simple fusion**. The second stage is **grid feature representation and cross-modal attention**. The third stage is **object-centric feature representation and bottom-up top-down attention (Anderson et al., 2018b)**. The three stages and representative work are shown in Figure 1.
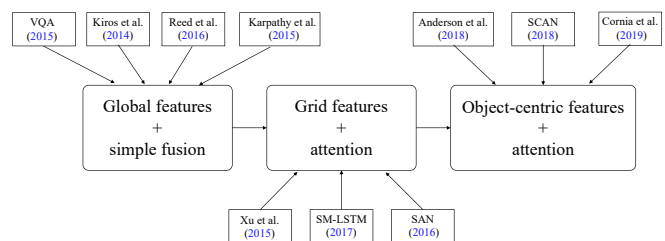

Fig. 1. The three stages of task specific methods. The main differences are the granularity of the visual representation and the way of fusing vision and language features.

	Tasks	Input	Output	Datasets	Metrics	Mainstream methods
Generation	Image Captioning (IC)	Image	Sentence	COCO (2014), Flickr30K (2014), Flickr8K (2013), CC3M (2018), CC12M (2021), SBU Captions (2011)	BLEU (2002), METEOR (2005), ROUGE (2004), CIDEr (2015), SPICE (2016)	Show and Tell (2015), Karpathy et al. (2015), Xu et al. (2015), m-RNN (2015), BUTD (2018b), Lu et al. (2017a), AoANet (2019), Lu et al. (2018), Cornia et al. (2019) AutoCaption (2020), ORT (2019), CPTR (2021)
	Text-to-Image Generation	Text	Image	COCO (2014), CUB (2011)	Inception score (2016) FID (2017) R-precision	StackGAN (2017), AttnGAN (2018), ObjGAN (2019c)
Understanding	Visual Question and Answer (VQA)	Image+ Text	Phrase	VQA (2015), VQA v2 (2017), DAQUAR (2014), Visual Genome (2017b), COCO QA (2015)	VQA Accuracy (2015)	Antol et al. (2015), Kim et al. (2016), Ren et al. (2015), Malinowski et al. (2015), Gao et al. (2015), SAN (2016), BUTD (2018a), MCB (2016), MUTAN (2017)
	Visual Dialog (VisDial)	Image+ Dialog+ Sentence	Sentence	VisDial (2017), GuessWhat?! (2017), Image Chat (2018), CLEVER (2019)	Mean Rank, MRR, Recall	VisDial (2017)
	Visual Reasoning (VR)	Image+ Text+ Graph	Text	GQA (2019), CLEVER (2019), NLVR (2017), $NLVR^2$ (2019), VCR (2019)	Accuracy	NMN (2017), N2NMN (2017), PG+EE (2017), TbD-net (2018), StackNMN (2019), NS-VQA (2019), XNM-Det (2019b)
	Visual Entailment (VE)	Image+ Text	label	SNLI-VE (2019)	Accuracy	EVE-ROI (2019)
Retrieval	Image-text Retrieval (IR)/ Text-image Retrieval (TR)	Text/Image	Image/Text	COCO (2014), Flickr30K (2014), Flickr8K (2013)	Recall@K, Median r	Frome et al. (2013), Socher et al. (2014), Deep Fragment (2014), MNLM (2014), Xu et al. (2015), m-CNN (2015), Klein et al. (2015), Wang et al. (2016), m-RNN (2015), Show and Tell (2015), Huang et al. (2016) Nam et al. (2017), SCAN (2018), Faghri et al. (2018)
Grounding	Phrase Grounding (PG)	Image+ Phrases	Bounding Boxes	Visual Genome (2017a), Flicker30K (2014), Flicker30K Entities (2016)	Phrase localization Accuracy, Recall@K	MDETR (2021), Gupta et al. (2020) Align2Ground (2019)
	Reference Expression Comprehension (RE)	Image+ Text	Bounding Boxes	RefCOCO(+/g) (2016; 2014; 2016), Talk2Car (2019) Visual7W (2016)	Accuracy	MDETR (2021), MAttNet (2018), MAC (2018)

TABLE I

A COMPARISON OVER TASK-SPECIFIC PROBLEMS. TASKS ARE CLASSIFIED INTO FOUR CATEGORIES. FOR EACH TASK, WE HAVE SUMMARIZED THE INPUT, OUTPUT, DATASETS, METRICS, AND MAINSTREAM METHODS.

A. Image Captioning

Task definition: The goal of image captioning is to generate a caption for a given image. A caption is a sentence to summarize the content of the image. The caption usually contains objects of interest, what they are doing, and the positional relations among them.

Methods: Early image captioning methods before deep learning (Kulkarni et al., 2013; Farhadi et al., 2010) are mainly rule-based. They first recognize objects and their relations and then generate captions based on predefined rules. Early methods have limited impact since their visual recognizers have limited vocabulary and the rule-based method can not deal with complex scenarios in human language.

The breakthrough in deep learning has greatly empowered image captioning. Seq2Seq (Sutskever et al., 2014) achieved great success in machine translation by utilizing a text encoder to encode text from the source language and a text decoder to generate text from the target language. Following the encoder-decoder structure in Seq2Seq, Xu et al. (2015) proposed to replace the text encoder with an image encoder using GoogleNet (Szegedy et al., 2014) and achieved state-of-the-art performance at that time. The encoder-decoder structure became popular and was widely adopted by later works. This structure is called img2seq as shown in Fig.2. Early studies (Xu et al., 2015; Karpathy and Fei-Fei, 2015; Vinyals et al., 2015) adopted a CNN model as the image encoder to extract
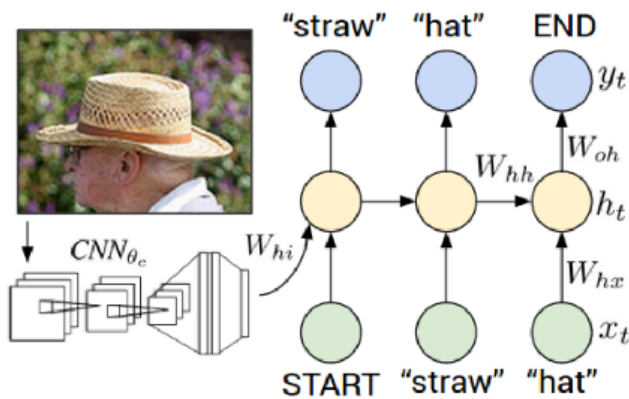
Fig. 2. The img2seq structure contains an image encoder such as a CNN and a language decoder such as an LSTM.

a global CNN feature, which is fed into the text decoder as the initial hidden state. m-RNN (Mao et al., 2015) and LRCN (Donahue et al., 2015) proposed to feed global CNN feature to each step of the LSTM decoder.

Global CNN feature has an obvious weakness because the decoder cannot focus on important regions of an image as humans do. To address this problem, attention mechanisms was introduced. Xu et al. (2015) proposed the first method leveraging attention over gird features. Assume the output feature map of a CNN feature extractor has a shape (H, W, C) where H, W are the height and width of the feature map and C is the feature dimension. The feature map can be flattened along spatial dimensions to $H \times W$ grid features with dimension C. For each cell of the LSTM decoder, the hidden state attends to grid features to decide which grids to focus on.

Compared to convolution, the attention mechanism has the following benefits. It allows the model to focus on certain parts of an image by giving high attention weights to important grid features. Moreover, the model learns alignments that correspond strongly with human intuition. The explainability of the model can also be improved through visualization of the attention scores, which potentially helps to troubleshoot.

However, splitting an image into equally-sized grids is a naive way to perform attention because grids align poorly with objects. To address this problem, some researchers seek to align attention with more meaningful regions. Anderson et al. (2018a) proposed a bottom-up and top-down attention (BUTD) approach to aligning attention with salient regions that are acquired with a detection model. BUTD extracts region features with a Faster-RCNN (Ren et al., 2016) model pre-trained on Visual Genome (Krishna et al., 2017a). Because detected object regions normally contain meaningful visual concepts and align better with human language, BUTD significantly improves the performance of both image captioning and VQA. Therefore, the pre-trained detector is widely adopted by subsequent VL studies.

There are also some variants of how to attend. For example, Lu et al. (2017a) claim that the decoder does not need to always attend to visual features, since some words are irrelevant to visual features. Therefore, they proposed to use a gate to decide whether to attend or not. AoA (Huang et al., 2019) designed

a special "Attention on attention" for image captioning task. After standard attention, they concatenate the attended vector and the query. Then they generate an information vector and an attention gate from the concatenated vector and multiply the gate and information vector to obtain the output.

Except for works with attention mechanism, there are also works that do not use attention. For example, Neural Baby Talk (Lu et al., 2018) first generates a sentence template and then fill it with concepts detected in the image. Cornia et al. (2019) generate a sentence by predicting a sequence of noun chunks. They first detect regions and then sort the regions with a sorting network. Finally, each region will be converted to a noun chunk to form the sentence.

In summary, there are two main aspects in the development of early image captioning methods, i.e. visual representation and language decoding. Visual representation develops from image-level global features to fine-grained and object-level region features and language decoding develops from LSTM to attention-based models.

B. VQA

Task definition: Given an image-question pair, VQA requires answering a question based on the image. Most studies treat VQA as a classification problem on a predefined answer set. For example, VQA v2 (Goyal et al., 2017) has around 2K predefined answers.

Methods: The vanilla VQA (Antol et al., 2015) is a combination of an LSTM (Hochreiter and Schmidhuber, 1997b) question encoder and a VGG (Simonyan and Zisserman, 2015b) image encoder. The output image embedding and question embedding are simply fused by point-wise multiplication. Then the fused vector goes through a linear layer followed by a softmax layer to output the probability of choosing each candidate answer. The architecture of the model is shown in Figure 3. Follow-up studies in VQA usually adopt the same
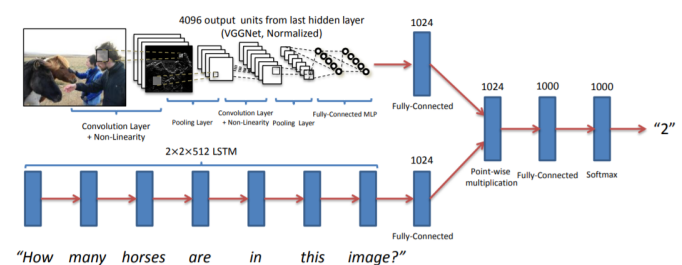

Fig. 3. The architecture of vanilla VQA (Antol et al., 2015) contains a CNN model to encode the input images and an LSTM model to encode the input question. The encoded image and question features are merged with dot product and then go through a fully connected layer to predict the probability over candidate answers.

prototype.

Early studies (Antol et al., 2015; Malinowski et al., 2015; Gao et al., 2015) normally adopted global image representation and simple fusion. Malinowski et al. (2015) proposed to feed the CNN image feature into each LSTM cell of the question encoder. Gao et al. (2015) used a shared LSTM to encode the question and decode the answer. They fuse the CNN image feature with the output of each decoder cell and generate the answer word by word.

Question answering is usually only related to some regions of the image. Therefore, global representation only leads to a sub-optimal solution due to the noise introduced by irrelevant regions. (Yang et al., 2016) proposed Stacked Attention Network (SAN) to stack multiple question-guided attention layers. In each layer, the semantic representation of the question is used as a query to attend to image grids. SAN is the first work that verifies the effectiveness of attention in VQA. Fukui et al. (2016) also adopted grid features, while they fuse image and language features through bilinear pooling (Lin et al., 2017).

Grid features have limited power as we have illustrated in the image captioning task. Shih et al. (2016) proposed to use region features as visual representations. They adopted Edge Boxes (Zitnick and Dollár, 2014) to locate regions. BUTD (Anderson et al., 2018b) pre-trained a powerful detector and uses the question features as queries to attend to region features. Lu et al. (2017b) argued that attention on text is of equal importance as on the image. Therefore, they developed a co-attention method that jointly performs text-guided image attention and image-guided text attention.

Except for attention, there are other strategies for modality fusion. Ren et al. (2015) treat an image feature as a language token. They concatenates the image embeddings with language tokens as the input to LSTM. Kim et al. (2016) proposed an iterative way of element-wise multiplication for modality fusion named Multimodal Residual Networks (MRN). MUTAN (Ben-Younes et al., 2017) presented parameterized bilinear interactions between modalities. Although there are many ways to fuse image and language features, attention is the most widely used one.

The core of VQA is to obtain a joint representation of image and language (the question). Researchers in this field pursued various ways to better encode and fuse image and language, which have laid the foundation for the following VLP methods. Most works in this field encode image and language independently and then fuse them, which is similar to dual-stream methods for VLP. Ren et al. (2015) treat image embedding as a language token, which is similar to single-stream methods.

C. Image Text Matching

Task definition: Image-text matching (ITM), or say image-text retrieval, is one of the fundamental topics in vision. Given a query in a certain modality (vision or language), it aims to find the semantically closest target from another modality. Depending on the query and the target modality, it contains two sub-tasks: image-to-text retrieval and text-to-image retrieval.

Methods: The core of image-text matching is to calculate the similarity or distance between an image and a piece of text. A widely adopted prototype is to map image and text into a shared embedding space and then calculate their similarity. The matched image and sentence are expected to have the highest similarity.

Early methods (Frome et al., 2013; Socher et al., 2014; Kiros et al., 2014) mainly adopted global feature to encode image and text. Kiros et al. (2014) proposed to learn cross-view representation with a hinge-based triplet ranking loss.

Faghri et al. (2018) considered hard negatives to improve the performance.

Karpathy et al. (2014) proposed Deep Fragment which is the first attempt to use fine-grained representation on both the image side and text side. The architecture of Deep Fragment is shown in Fig. 4. Instead of directly representing the whole image and sentence, they map each image fragment and sentence fragment into the cross-modality embedding space. Then they align fragments between different modalities. Since one image region may be related to several words, they find the most similar region embedding for each word embedding. The similarity between the image and the sentence is the sum of the similarities between aligned word and region pairs.
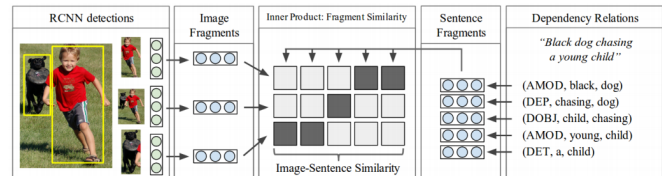

Fig. 4. The overview of the Deep fragment(Karpathy et al., 2014) architecture. **Left:** Detected objects are mapped to fragment embedding space. **Right:** Dependence tree relations are encoded to fragment embedding space.

As the attention mechanism has shown great success in other VL tasks, Huang et al. (2016) proposed to introduce attention into image-text matching (ITM). They developed a context-modulated attention scheme to attend to instance pairs appearing in both image and text. Nam et al. (2017) proposed a dual attention framework that attends to specific regions in images and words in the text through multiple steps and gathers essential information from both modalities. These methods proved the effectiveness of attention in the ITM task. However, as a limitation, they are multi-step methods and can only focus on one semantic part at a time.

Lee et al. (2018) proposed a cross-attention algorithm called SCAN to calculate the similarity between image and sentence. To enable cross attention, they represent an image as a set of regions and a sentence as a set of words. The core idea of cross attention is to not only use the sentence as a query to attend to image regions but also use the image as a query to attend to words.

Essentially, image-text matching is a problem of calculating the similarity between image and text. Early works encode image and text into global features and calculate their cosine similarity through dot product. Subsequent works adopt fine-grained features—object-level features for images and word-level features for language. They also develop more complex algorithms for calculating similarities such as cross attention.

D. Other tasks

There are a broad variety of tasks in the interdisciplinary field of vision and language that we can not elaborate on in detail. Therefore, we list some of the important tasks in Table I, including:

Text-to-Image Generation: Given a piece of text, generate an image containing the content of the text. We leave the details of text-to-image generation to Section IV-B.
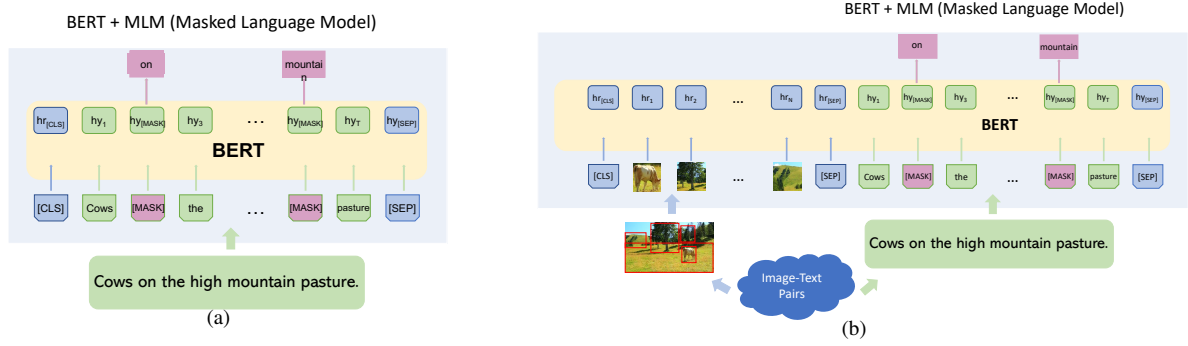
Fig. 5. (a) Original BERT with single-modality, where some language tokens are masked for prediction to train language representation. (b) Modified BERT with multi-modality, where both image and language tokens are fed into a BERT-like Transformer model.

Visual Dialog: Given an image, a dialog history, and a question about the image, answer the question.

Visual Reasoning: Similar to VQA, which requires answering a question about an input image, visual reasoning requires a further ability to understand the image. Visual reasoning task usually contains adequate annotations about the objects in an image, the structure of the questions, etc.

Visual Entailment: Given an image and a text, decide whether the image semantically entails the input text.

Phrase Grounding and Reference Expression Comprehension: These two tasks require a model to output bounding boxes corresponding to the text. For phrase grounding, the text is a set of phrases and for reference expression comprehension, the text is an expression.

In the era of task-specific methods, researchers design specific models for different tasks. Although the models for different tasks vary significantly, they follow a similar trajectory. They all have three stages as shown in Fig. 1. The technological development of this era laid the foundation for the VLP era.

III. VISION LANGUAGE JOINT REPRESENTATION

The pre-training and fine-tuning paradigm has been adopted across a wide range of domains and advanced various downstream tasks. Among the most prominent factors that leverage the prevailing large-scale pre-training is the availability of abundant datasets along with the rapid evolution of GPUs. Motivated by the success in single-modal language/vision pre-training, researchers started to explore the joint representation of language and vision (Sun et al., 2019a; Lu et al., 2019), giving birth to the cross-modal VLP models.

The recent surge of VLP models is mostly inspired by language models in both architecture design and training methods. One of the most important breakthroughs is Transformer which was developed by Vaswani et al. (2017) for improving language representation. Using multiple stacked attention layers, Transformers can fuse information over language tokens globally with high parallelism, which facilitates both powerful representation and large-scale training. A successful application of Transformer is BERT (Devlin et al., 2018) which leverages Transformer encoders and introduces a bidirectional masking technique that allows each language token to attend to other tokens bidirectionally. As shown in Figure 5(a), training is conducted by replacing some text tokens with a special [MASK] token and predicting each [MASK] using its context information. With this technique, language representation training can be considered as a denoising process, in which an input sentence learns to reconstruct itself with some contaminated tokens. This denoising training forces the masked tokens to utilize all the unmasked information, hence resulting in contextualized representations. The architecture design and mask training technique developed for Transformer-based language models is the main principle behind a broad variety of cross-modal developments that contributes to the recent surge of VLP models. Figure 5(b) shows a simple cross-modal BERT (Anderson et al., 2018b). Similar to language training, image is tokenized and embedded along with language tokens with certain techniques, which will be elaborated on later. Usually, the tokenized visual features and textual features together are fed into a Transformer encoder with masked language training to learn a joint representation.

In this section, we will go through the main components of VLP models. As shown in Figure 6, there are primarily three components in VLP models, namely the *visual embedding (VE), textual embedding (TE), and modality fusion (MF)* modules. VE and TE are normally pre-trained with images and texts respectively, whereas MF takes the features extracted from VE and TE, and fuses them with image-text pre-training. The goal of VLP is to learn *object-level, language-aligned, and semantic-rich* visual representations. Object-level means the learned representation is fine-grained and aligned with objects rather than for a whole image. Works that use features of detected objects to represent images (Tan and Bansal, 2019; Lu et al., 2019; Yu et al., 2020; Li et al., 2019b,a, 2020b; Hu et al., 2020; Li et al., 2021b) are object-level. Language-aligned aims to learn visual features that are well-aligned with language words, which is the goal for most VLP methods. Semantic-rich strives for a representation that can be generalized to a broad range of semantic concepts and needs to be learned from a large-scale dataset.

Pre-training on a massive dataset is crucial for improving the performance on downstream tasks with smaller datasets, as the learned representation can be transferred in downstream tasks. VLP models have been proven very effective to empower downstream tasks.
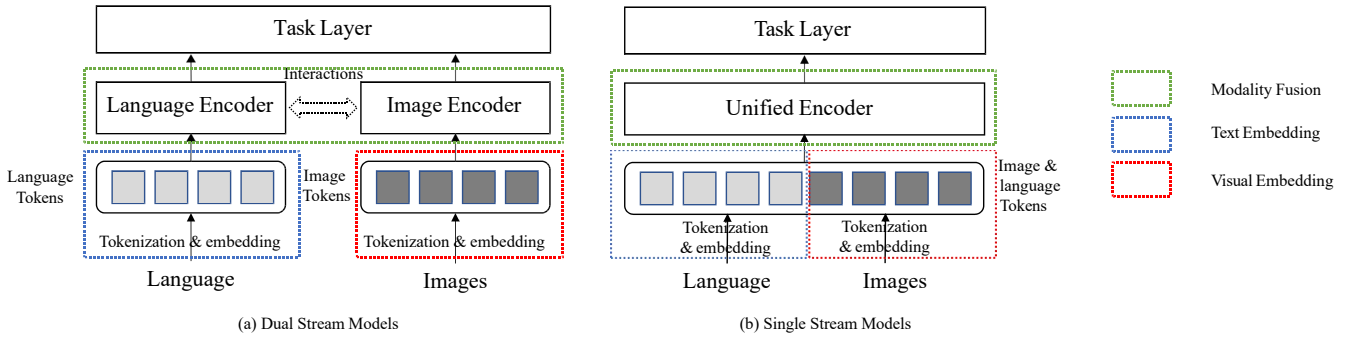
Fig. 6. The architecture of VLP models usually contains Visual Embedding (VE), Textual Embedding (TE), and Modality Fusion (MF). (a) is the illustration of a dual-stream model, and (b) shows a single stream model. In a dual-stream model, modality fusion is optional and is done by the interactions (usually cross attention) between language and image encoder. In a single stream model, modality fusion is done in a unified encoder (usually multi-layer Transformers).

A. Why Pre-training Is Needed?

Deep learning is essentially a statistical data-driven approach, which aims to learn a mapping function from seen data so that to make predictions on unseen data using the learned mapping function. Note that the ultimate goal is to achieve a good performance on unseen data. In terms of statistical learning, such a goal is represented as minimizing an *expected* loss over the whole data space which follows a fixed but unknown distribution. However, such an expected loss minimization is not tractable as the distribution is unknown. In practice, one has to sample data from this distribution and define an *empirical* loss as a proxy to the expected loss. This may sound strange, but is actually a commonly used practice in machine learning. For example, for an image classification problem of telling whether an input image has a cat or not, the most practical approach is to collect training images of cat and no-cat and then train a classifier by minimizing an empirical loss defined on this training set. However, the distribution of cat and no-cat images is indeed unknown.

Statistical learning theory shows that for a sufficiently large number of independent and identically distributed (i.i.d.) data sampled from such an unknown distribution, the empirical loss minimization result converges to the expected loss minimization result. That is, asymptotically, one can use i.i.d. samples to approximate a loss function defined by an unknown distribution. However, in practice, data is never sufficient to represent the unknown distribution and thus leads to many deficiencies, such as low performance on uncovered data, being vulnerable to adversarial attacks, etc.

Pre-training allows one to leverage an unlimited amount of data without labels (or with weak labels) to learn features that conform with downstream tasks. Such a large-scale data set helps define a better approximation to the expected loss for learning more robust and non-spurious patterns from data. Thanks to the shared model between the pre-training and fine-tuning stages, with very limited (e.g., few-shot) supervision, the learned feature can lead to high accuracy on downstream tasks after fine-tuning. This makes the paradigm of pre-training and fine-tuning an effective solution to solving (or mitigating) the data shortage problem. For more discussions, see (Zhang and Shum, 2022).

B. Modality Embedding

Text and image are different levels of information in nature concerning dimensionality and structure. To resolve this modality discrepancy, modality embedding is normally utilized to extract features from each modality independently and then map the features into a shared feature space. As shown in Figure 6, modality embedding involves visual embedding and textual embedding, both encompassing a tokenization process and an embedding process. Visual embedding aims to follow textual embedding to convert an image into a number of tokens with the level of representation as text tokens. Ablation studies conducted by Bugliarello et al. (2021) demonstrate that training datasets and hyperparameters are responsible for most performance improvements across many different VLP models and also emphasize the importance of modality embeddings.

1) Text Tokenization and Embedding:
Before text embedding, the text should be tokenized. Considering the discretization nature of language, early work simply regards each word as a token. A pioneering study is Word2Vec (Mikolov et al., 2013), which proposed a continuous CBOW and a skip-gram model to train word vector representation. Word2Vec is computationally efficient to scale to large corpus and produces high-quality embeddings. However, although its vocabulary size is as large as about one million, this method suffers from out-of-vocabulary issues due to rare or unseen words, making it difficult to learn word sub-units such as 'est'. To resolve this problem, *Sennrich et. al* (Sennrich et al., 2015) proposed a subword tokenization approach, which segments words into smaller units with byte pair encoding (BPE) (Gage, 1994). Subword tokenization is widely used in many language models including BERT.

Most VLP models adopt text embeddings from pre-trained BERT (Devlin et al., 2018). As BERT is trained with masked token learning using Transformer encoders, it has a strong bidirectional representation ability.

2) Visual Tokenization and Embedding:
Different from language tokens that are discrete and arranged in a single dimension, images are from a high dimensional space and have co-related pixel values. Therefore, image tokenization is usually more complicated than text tokenization. Basically, image tokenization can be categorized into *region based*, *grid based*, and *patch based*, which are introduced as follows.

1) Grid features are directly extracted from equally sized image grids with a convolution feature extractor as aforementioned. For example, Huang et al. (2020, 2021) adopted grid features as the image embedding of their VLP models. The advantages of grid features are mainly two folds. The first is convenient as it does not require a pre-trained object detector. The second is that besides salient objects, grid features also contain background which may be useful for downstream tasks.

2) Region features are extracted by a pre-trained object detector. Most recent VLP models adopt region features to learn object-level joint representations. Especially, most VLP models adopt Faster R-CNN trained on the Visual Genome (VG) dataset as region feature embedding following the work of BUTD (Anderson et al., 2018b). There are three essential components of region features, which are bounding boxes, object tags, and RoI features (feature vectors after RoI pooling). Bounding boxes are commonly used in VLP as position indicators, which are encoded through a transformation into the same dimensional space as RoI features and added to RoI features. Object tags are widely utilized in training methods such as Masked Region Classification, which will be elaborated later in III-D3. The advantage of region features is that they help a VLP model focus on meaningful regions of the image. These regions are usually closely related to downstream tasks.

3) Patch features are usually extracted by a linear projection on evenly divided image patches. The main difference between patch and grid features is that grid features are extracted from the feature map of a convolutional model while patch features directly utilize a linear projection. Patch features were first introduced by Vision Transformer (ViT) (Dosovitskiy et al., 2021a) and then adopted by VLP models (Kim et al., 2021; Xue et al., 2021). The advantage of using patch features is efficiency. For example, ViLT accelerates the pre-training by 10 times with competitive results (Kim et al., 2021).

Image embedding methods normally vary for different tokenization schemes. Grid features and region features are usually from pre-trained convolutional models, whereas patch features can be simply embedded by a linear layer.

C. Modality Fusion

At the core of VLP models lies the modality fusion, which models intra-modality and inter-modality fusion to produce contextualized joint representations of image and text. MF schemas can be categorized into dual-stream modeling and single-stream modeling. The general structure of VLP is shown in Figure 6.

1) Dual stream modeling:
Dual-stream modeling aims to map vision and language into the same semantic space. It is the seminal method for modality fusion (Tan and Bansal, 2019; Yu et al., 2020; Li et al., 2021a). As shown in Fig. 6(a), it adopts two separate encoders to learn high-level representations for vision and language, respectively. The dual-stream design allows variable network depth and architecture to be adaptive to each modality. Apart from intra-modal fusion within each modality, some studies (Lu et al., 2019; Tan and Bansal, 2019) also explicitly design inter-modal interactions between two encoders to enable modality fusion in different encoding stages.

2) Single stream modeling:
Single stream modeling aims to learn one joint representation. Image and text tokens are concatenated and inputted into Transformers as shown in Fig. 6(b). Most VLP models adopt this modality fusion scheme (Alberti et al., 2019; Chen et al., 2020; Li et al., 2019a; Huang et al., 2021; Jia et al., 2021). Single-stream modeling performs implicit intra-modal and inter-modal fusion, free from the architecture design of the fusion stage in dual-stream modeling.

D. Training

To learn a joint representation of vision and language, vision language pre-training methods usually use several self-supervised learning losses to pre-train the model on a large dataset. As studied in (Chen et al., 2020; Li et al., 2019a; Lu et al., 2019; Su et al., 2019; Tan and Bansal, 2019; Zhou et al., 2020), there are mainly three pre-training methods, which are Image Text Matching (ITM), Masked Language Modeling (MLM), and Masked Visual Modeling (MVM).

1) Image Text Matching:
The goal of ITM is to predict whether a pair of images and text is matched. ITM can be formulated as a binary classification task. Previous work (Chen et al., 2020) applies a sigmoid function on the output of the special token $[CLS]$ to predict whether the input image and text are matched. The loss function is

$$\mathcal{L}_{\text{ITM}} = -\mathbb{E}_{(\mathcal{W},V)\sim D} \log p(y \mid \mathcal{W}, V) \tag{1}$$

where $\mathcal{W} = \{w_1, w_2, ..., w_n\}$ denotes a sequence of language tokens and V denotes the visual content. $y = 0$ or 1 indicates whether the image is matched ($y = 1$) or not ($y = 0$).

2) Masked Language Modeling:
According to (Chen et al., 2020), MLM is utilized to encourage the model to learn the implicit relation between language tokens and visual content. The goal is to reconstruct the masked language tokens from the known language tokens and visual contents. This goal can be formulated as

$$\mathcal{L}_{\text{MLM}} = -\mathbb{E}_{(\mathcal{W},V)\sim D} \log p\left(w_i \mid \mathcal{W}_{\setminus i}, V\right), \tag{2}$$

where $\mathcal{W}_{\setminus i}$ denotes the sentence without the i-th word. Note that, although BPE is normally adopted for language tokenization, the minimal masked unit is a whole word instead of a subword. The reason is that a subword can be easily predicted from its surrounding subwords due to information leakage. There are also improved versions of MLM. For example, Sun et al. (2019b) proposed Knowledge Masked Language Modeling, which performs phrase-level masking and entity-level masking to integrate phrase and entity-level knowledge into the language representation. For entity level masking, they treat a named entity as a whole. For example, J. K. Rowling, which contains three tokens, is a person name and should be masked together in entity-level masking. The phrase level masking treats a group of words as a conceptual unit. They mask all the tokens belonging to a phrase and predict them simultaneously.
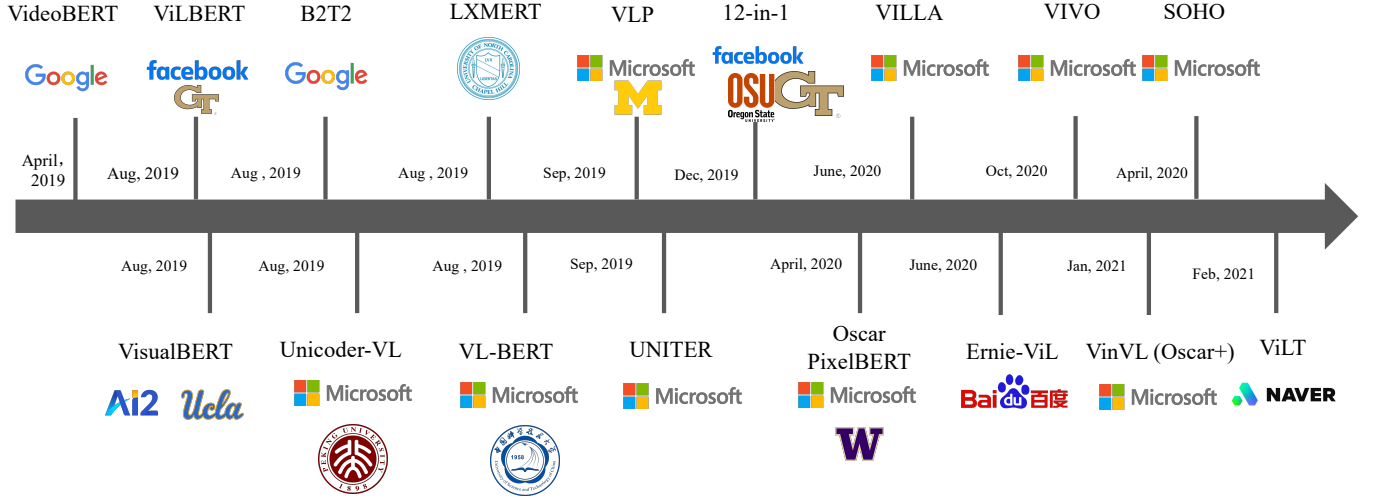
Fig. 7. A Landscape of VLP methods. Works are sorted according to the time they were published. We also show the logos of main institutions where each work is from.

3) Masked Vision Modeling:

Inspired by MLM, MVM is designed to learn contextualized visual representation by reconstructing masked visual contents. MVM is more challenging than MLM since the information density of image is lower than that of language. When reconstructing a missing word, sophisticated language understanding is required. On the contrary, a missing image patch can be recovered from neighboring patches without cross-modality understanding (He et al., 2021). To overcome this gap, most works mask detected object regions that have relatively high information density. Other works such as SOHO (Huang et al., 2021) use a visual dictionary (VD) to represent more comprehensive and compact semantics in the visual domain so that they can apply MVM in the same way as MLM. In summary, there are mainly four MVM schemes.

1) Masked Region Prediction (MRP): MRP Minimizes the distance between the predicted feature and the feature output by a pre-trained object detector. The distance metric is usually L2 (Tan and Bansal, 2019). We denote masked image regions as $\mathbf{v_m} = \{v_m^1, ..., v_m^M\}$, h_θ^i as the model prediction corresponding to v_m^i, and $r(v_m^i)$ as the ROI-pooled feature of v_m^i. The loss function is

$$\mathcal{L}_{MRP} = \sum_{i=1}^{M} \left\| h_\theta^i - r\left(v_m^i\right) \right\|_2^2 \qquad (3)$$

2) Masked Region Classification (MRC): MRC requires a model to predict the object semantic class for each masked region. Since there is no ground-truth label, the label $c(v_m^i)$ predicted by a pre-trained object detector of v_m^i is used as the target label of v_m^i. Denoting the predicted label of the VLP model as $g_\theta(v_m^i)$, the loss fuction is

$$\mathcal{L}_{MRC} = \sum_{i=1}^{M} \text{CE}\left(c\left(v_m^i\right), g_\theta\left(v_m^i\right)\right) \qquad (4)$$

3) Masked Region Classification with KL-Divergence (MRC-KL): As the target label for MRC is inaccurate, MRC-KL adopts the soft label $\tilde{c}(v_m^i)$ of v_m^i as the supervision signal, which is the raw output of the object detector after softmax. The loss function is

$$\mathcal{L}_{MRC-kl} = \sum_{i=1}^{M} D_{KL}\left(\tilde{c}\left(v_m^i\right) \| g_\theta^i\right) \qquad (5)$$

where g_θ^i denotes the soft label of v_m^i predicted by the VLP model.

4) Masked Visual Modeling with Visual Dictionary (MVMVD): Similar to language models which have a vocabulary dictionary, MVMVD requires a visual vocabulary dictionary (VD). The goal of MVMVD is to reconstruct the masked VD tokens (Huang et al., 2021). The loss function is

$$\mathcal{L}_{\text{MVM}} = -\mathbb{E}_{(\mathcal{W}, f(\mathcal{V})) \sim D} \log p\left(f\left(v_j\right) \mid \mathcal{W}, f(\mathcal{V})_{\backslash j}\right) \qquad (6)$$

where $f(\cdot)$ denotes the mapping from an image grid to a visual token in the VD and j denotes the index of the masked token in the VD.

There are two points that are worth noting. Firstly, to encourage inter-modal fusion, some works such as UNITER-VL (Chen et al., 2019) only mask tokens in one modality each time during training to encourage the masked tokens to attend to another modality for missing information (Chen et al., 2019). Secondly, for MVMVD, neighboring image grids tend to map to the same VD token as they are highly co-related. When performing reconstruction, the model may directly copy the surrounding tokens. Therefore, all visual embedding vectors mapped to the same VD token are masked together in SOHO. Despite all the above mentioned MVM methods, effective vision modeling remains a challenging problem. The results of some ablation studies in VLP models such as SOHO (Huang et al., 2021) indicate that adding MVM task only yields small additional improvement to the performance. Cao et al. (2020) reveal that VLP models exhibit the propensity to focus on textual information than visual information in downstream tasks.

Architecture	Method	Visual Embedding	Pre-training Tasks	Pre-training Datasets	Downstream Tasks
Dual-Stream	ViLBERT (2019)	BUTD	ITM, MLM MRC-kl	CC3M	VQA, VR, RE IR, Zero-shot IR
	LXMERT (2019)	BUTD	ITM, MLM, MRP, MRC	COCO, VG (2017a), VQA v2 GQA, Visual7W (2016)	VQA, VR
	12-in-1 (2020)	BUTD	ITM,MLM,MRC-kl	VQA v2, Flickr30k SNLI-VE, COCO GuessWhat, VG RefCOCO, RefCOCO+,RefCOCOG Visual 7W, GQA, NLVR^2	VQA ,IR, RE, VE, VR
	Ernie-ViL (2020)	BUTD	Object Prediction Attribute Prediction Relationship Prediction ITM, MLM,MRC-kl	CC3M, SBU (out-of-domain) COCO, VG (in-domain)	VQA, VR, IT, TR, RE
Single-Stream	VideoBERT (2019a)	S3D (2017b)	ITM, MLM, MVM	YouTube cooking videos (2017b)	Zero-shot Action prediction (2017b) Video Captioning (2017b)
	VisualBERT (2019b)	Pre-trained Fast R-CNN	ITM, MLM	COCO	VQA, VR, PG
	B2T2 (2019)	ResNet-152 (2016b)	ITM, MLM	CC3M	VR
	Unicoder-VL (2019a)	Pre-trained Faster-RCNN (2018)	ITM, MLM, MRC	CC3M, SBU	IR, TR, VR Zero-shot IR/TR
	VL-BERT (2019)	BUTD	MLM, MRP	CC3M, English Wikipedia BooksCorpus (2015)	VQA, VR, RE
	Unified VLP (2020)	BUTD variant (with ResNext-101 backbone) (2017a)	MLM (Sequentially and bidirectionally)	CC3M	IC, VQA
	UNITER (2019)	BUTD	ITM, MLM, MRP, MRC MRC-kl	COCO, VG, CC3M, SBU	VQA, VR, VE, IR, TR, RE
	Oscar (2020b)	BUTD+tags	MLM (include tags) ITM (pollute tags)	COCO, VG, CC3M, SBU, fliker30k, GQA	VQA, VR, VE, IR, TR RE, IC
	PixelBERT (2020)	Pixel feature embedding (2015)	MLM, ITM	COCO, VG	VQA, IR, TR, VR
	VILLA (2020)	BUTD	ITM, MLM, MRC-kl	COCO, VG, CC3M, SBU	VQA, VR, VE, IR, TR, RE
	VIVO (2020)	BUTD	Mask Tag Prediction (Hungarian match loss)	Open Images V5 (2018; 2019)	IC
	SOHO (2021)	VD	ITM, MLM MVMVD	COCO, VG	IR, TR, VQA VR, VE (based on VD)
	ViLT (2021)	Patch Projection	ITM, MLM	COCO, VG, SBU, CC	VQA, VR, IR, TR, Zero-shot IR&TR
Hybrid	SemVLP (2021a)	Pre-tained Faster-RCNN	MLM, MRP, VQA	COCO, VG, VQA v2, GQA, Visual7W	VQA, VR, IR, TR

TABLE II

PRE-TRAINING WORK COMPARISON. THE PRE-TRAINING TASKS CORRESPOND TO THE TASKS DESCRIBED IN SECTION III-D. WE ALSO LIST THE PRE-TRAINING DATASETS AND DOWNSTREAM TASKS OF THESE WORKS. THE DATASETS AND DOWNSTREAM TASKS ARE DESCRIBED IN TABLE I

E. Landscape of General Pre-training Studies

After introducing the general pipeline of VLP models, in this section, we summarize some pioneering works in the cross-domain of VLP.

Inspired by the success of pre-training in NLP and CV, a boosting number of research works in the domain of VLP has recently surged to pursue a unified cross-modality representation. A landscape of VLP works is shown in Figure 7. A more detailed comparison of related works is shown in Table II. We elaborate on some representative studies in this section.

Single Stream Models: VideoBERT (Sun et al., 2019a) is a pioneering work to learn joint representation for video and language. The primary idea is to feed visual and textual tokens into a single-stream model built upon BERT (Devlin et al., 2018). Textual tokens are extracted by converting video speech into text with an automatic speech recognition approach, and visual tokens are acquired by extracting features from video clips using a convolutional backbone. VideoBERT is capable of performing a wide range of downstream classification and generation tasks, including video captioning and zero-shot mask verbs/nouns prediction. Note that VideoBERT is pre-trained on cooking videos, where the contents are instructional and of high-quality. It assumes that spoken words are well aligned with visual content, which limits its application to only certain

videos (e.g. instructional videos). Another issue confining its scalability is its delicately designed captioning text template, for example, *now let's [MASK] the [MASK] to the [MASK], and then [MASK] the [MASK]*, which only works for cooking videos.

Li et al. (2019b) proposed a simple single-stream VLP model named VisualBERT. The extracted visual and textual tokens are directly combined and fed into Transformers, where cross-modality fusion can be performed implicitly. Similar to VisualBERT, several concurrent studies such as Unicoder-VL (Li et al., 2019a), VL-BERT (Su et al., 2019), and UNITER (Chen et al., 2020) also adopt the single-stream architecture. These VLP studies are similar in the following aspects: 1) They all utilize an object detection backbone to compute image embedding. 2) Masked language modeling task is adopted by all of them. 3) They all adopt the single-stream BERT architecture. They differ from each other in their pre-training methods and datasets, as shown in Table II.

Dual Stream Models: ViLBERT (Lu et al., 2019) and LXMBERT (Tan and Bansal, 2019) are pioneering works to extend BERT to dual-stream VLP models. They are pre-trained on the Conceptual Captions dataset (Sharma et al., 2018) and leverage a pre-trained Faster R-CNN model (Ren et al., 2017) to detect regions as visual tokens. ViLBERT processes visual and textual tokens separately with two parallel streams which can

fuse cross-modality information through cross-attention layers when needed. In other words, ViLBERT assumes the different processing architectures for vision and language. Its cross-modal fusion is designed to be sparse and explicit between the two processing pipelines. LXMBERT differs from ViLBERT by decoupling intra-modal and inter-modal processing. More specifically, visual and textual tokens are encoded separately in the first phase and then fed into a cross-modality encoder to produce the joint representation.

Other Fusion Methods: Fundamentally, single-stream modeling and dual-stream modeling differ in the fusion time, where single-stream fuses different modalities in an earlier stage while dual-stream prefers to extract high-level features of each modality before fusion. SemVLP (Li et al., 2021a) proposed to combine the two prevalent modeling architectures by training them iteratively. Such an approach takes advantage of both architectures and performs cross-modality semantic alignment on both low-level and high-level. Especially, the Transformer encoder is shared between both modeling methods with an additional cross-modal attention module in the dual-stream encoder, which is found to contribute to the semantic alignment and reduce parameters.

Most VLP models attempt to encode vision and language into separate tokens that interact with each other explicitly or implicitly through modality fusion. Another line of VLP models alternatively attaches visual tokens to textual tokens based on object detection models. B2T2 (Alberti et al., 2019) proposed to fuse the features of detected objects in textual tokens, based on which MLM and ITM are performed in pre-training. In B2T2, a token T can be expressed as

$$T = t + \sum_{i=1}^{n} (h_t heta(b_i) + b_i) \tag{7}$$

where t is the original textual embedding, n is the number of detected objects whose label is token t, b_i is the embedding of the i-th object's bounding box, and $h_t heta(b_i)$ denotes the extracted visual feature from the bounding box. B2T2 also analyzes the stages of fusing objects and textual tokens. The result indicates the effectiveness of early fusion.

Early Attempts to Bridge Modality Gap: To enable both generation and understanding tasks, Zhou et al. (2020) proposed a unified vision-language pre-training approach. It introduces two mask schemes namely bidirectional attention mask and sequence-to-sequence mask to empower understanding and generation tasks, respectively. It is worth noting that this unified VLP approach only adopts MLM during pre-training and achieves a competitive performance on image captioning and VQA. 12-in-1 (Lu et al., 2020) extended multi-task training to four broad tasks pre-trained on 12 datasets. The experimental results indicate multi-task training can consistently improve the performance of downstream tasks and yield a more lightweight model with fewer parameters.

VILLA (Gan et al., 2020) introduced adversarial training at the embedding level of visual and textual tokens based on the design of UNITER (Chen et al., 2019). It performs adversarial training by adding perturbations in the embedding space as regularization and yields decent performance improvement.

Motivated by the knowledge masking scheme of ERNIE (Sun et al., 2019b), structured knowledge is first incorporated in the VLP models in ERNIE-ViL (Yu et al., 2020). To develop better cross-modality semantic alignments by constructing scene graphs, ERNIE-ViL proposes scene graph prediction tasks to model objects, attributes, and relationships in the graph to learn object-level and attribute-aware representation. Incorporating knowledge in cross-modality training is challenging and remains an open problem.

Grid & Patch features: While the prevalence of region feature embedding facilitates the training of VLP models, it also restricts the scalability and generalization capability of VLP models. We can analyze the weakness of region features from Faster R-CNN as follows.

- Limited categories: Visual feature is limited by object detection models which are trained on relatively small datasets with predefined object categories. For example, the widely adopted Faster R-CNN model in BUTD (Anderson et al., 2018a) is trained on VG with a fixed number of 1594 object classes and 524 attributes.
- Low quality: Region features often suffer from low quality as the Faster R-CNN models are trained on small well-labeled datasets (Anderson et al., 2018b).
- Lack context: Region features extract RoI features that belong to certain categories without any background information, neglecting semantic relationships between these region features. In reality, these semantic relationships are important.

PixelBERT (Huang et al., 2020) attempted to break this limitation and fully utilizes visual information by directly learning from pixel features. Instead of utilizing all the pixels as visual features, to reduce computation cost and improve model robustness, a fixed number of 100 pixels are randomly sampled during pre-training. However, the experimental results indicate that random sampling only slightly improve the performance, less than 0.5 VQA score in downstream tasks.

SOHO (Huang et al., 2021) is another pioneering work that leverages grid features for cross-modality understanding. To learn a semantically comprehensive representation for visual context, SOHO proposes to learn a VD for visual tokenization. VD is learned by first obtaining high-level features from a convolutional network, which are then grouped according to feature similarity and fed into a moving-averaged encoder to dynamically update VD. As visual embeddings are trainable, SOHO is an end-to-end pre-training framework that directly learns from pixels, eliminating the need for bounding boxes. With this dynamic VD updating during training, the serial number of each token in the VD can be considered as a label just like language tokens, making it natural to perform masked vision modeling. For pre-training tasks, SOHO proposes a novel MVMVD method (described in III-D3) to mask all the visual tokens of the same label simultaneously in an image to avoid any information leakage.

Image embeddings based on regions or grids as aforementioned are computationally heavy and the extracted high-level features prevent early fusion of cross-modality information. Inspired by ViT (Dosovitskiy et al., 2021b), ViLT (Kim et al., 2021) adopts simple linear projection of image patches as

visual embedding, which greatly accelerates pre-training by 10 times with competitive results. It implies that instead of designing novel visual embedding, designing better modality-fusion could be the key to improving the representation of VLP models.

Improve Aligned Representation: Vision-language-aligned representation is a fundamental goal in VLP. To achieve this goal, some works propose to adopt additional object-level data in VLP. For example, many VLP methods adopt RoI region features with detection models. However, the detected object tags as an important component are not explicitly modeled in VLP models. To leverage this additional information, Oscar (Li et al., 2020b) introduced object tags as anchor points to help learn cross-modality-aligned representation. This learning process is empirically natural as the detected object tags often appear in the image-paired text, which helps align vision and language. In addition, training with object tags contributes to learning co-occurrence of objects. Therefore, Oscar yields significant improvement in the downstream understanding and generation tasks. However, the drawback of Oscar is also obvious as it relies on well-labeled image-caption datasets, making it hard to scale.

As VLP models are limited by inadequate well-aligned (image, caption) pairs, VIVO (Hu et al., 2020) proposed to scale up pre-training using a large amount of (image, tag) pairs. VIVO adopts a Hungarian matching loss to perform masked tag prediction, which enables visual vocabulary learning and improves the model generalization ability to describe novel objects in downstream tasks. As a result, it surpassed human performance for the first time on the Novel Object Captioning at Scale (NoCaps) (Agrawal et al., 2019) benchmark. Scaling VLP models based on RoI features also calls for more powerful visual representations. VinVL (Zhang et al., 2021) followed BUTD (Anderson et al., 2018b) and develops an improved object detection model for VLP with a larger training dataset. More specifically, it adopts ResNeXt152-C4 and merges four public datasets including VG, COCO, Objects365, and OpenImagesV5 for large-scale training. VinVL yields significant improvement on VLP models like VIVO and Oscar and achieves top results on the NoCaps, image captioning, and VQA leaderboards.

IV. SCALE UP MODELS AND DATA

Though inspiring progress has been made in the vision-language joint representation, most aforementioned studies primarily focus on object-level representation to pursue good cross-modal alignment. However, they take a strong assumption: image and text pairs are well labeled, which restricts training datasets to relatively small "gold-labeled" datasets. For example, the largest public dataset widely used for VL pre-training is Conceptual Captions (Sharma et al., 2018) with three million image-text pairs. To obtain richer semantics and stronger generalization capability, larger weakly-labeled datasets such as web-crawled datasets are greatly desired. CLIP (Radford et al., 2021) and DALL-E (Gu et al., 2021) are the first successful practice to utilize large-scale (400M image-text pairs for CLIP and 250M for DALL-E) web-crawled data for pre-training.

Motivated by the success of CLIP and DALL-E, several recent works further built more powerful models with even larger datasets.

This section aims to introduce models trained with large-scale weakly-labeled datasets. The section is divided into two parts. The first part includes works utilizing large-scale datasets for visual understanding such as CLIP, ALIGN (Jia et al., 2021), SimVLM (Wang et al., 2021) and Florence (Yuan et al., 2021). The second part contains visual generation models empowered by large-scale datasets such as DALL-E, GODIVA (Wu et al., 2021a) and NÜWA (Wu et al., 2021b).

A. Visual Understanding

The core idea of CLIP is the training method. Instead of training to predict masked visual or textual tokens as in other VLP methods, CLIP learns to recognize paired image and text. Given a batch of N (image-text) pairs, the goal is to predict which of the $N \times N$ possible pairs are matched pairs (positive samples) and which are unmatched pairs (negative samples). After pre-training, CLIP can perform zero-shot image classification by using phrases such as "a photo of" plus a category name as prompts to tell the model which categories an input image is the most similar to. Compared with fully supervised baselines, zero-shot CLIP outperforms the baseline on 16 of 27 datasets.

Similar to CLIP, ALIGN (Jia et al., 2021) also adopts a dual encoder model with a contrastive loss for zero-shot tasks. It utilizes a larger raw dataset with $1.8B$ image-text pairs. ALIGN outperforms CLIP on many zero-shot visual tasks, which proves that a larger dataset leads to better performance. Except for vision tasks, ALIGN also outperforms previous work on image-text retrieval tasks. SimVLM (Wang et al., 2021) developed a new approach to VL pre-training. It follows a simple prefix language modeling objective to predict the next token in an autoregressive way. It achieves competitive results on multiple VL tasks and has the ability for text-guided zero-shot learning. Unlike previous works that adopt coarse (image-level) representations and static (image) data, Florence (Yuan et al., 2021) adopts fine-grained (object-level) representations and extends to dynamic (video) data. For object-level representations, Florence adds an adaptor Dynamic Head (Dai et al., 2021) to the image encoder and trains with an extra object detection dataset. Through pre-training on $900M$ pairs of image-text pairs, Florence achieves new state-of-the-art results in a majority of 44 representative benchmarks.

Apart from zero-shot classification, CLIP can also help detection. For example, ViLD (Gu et al., 2021) proposes a zero-shot detector via CLIP distillation. Other studies show that CLIP can learn multi-modal features which are more like neurons in the human brain (Goh et al., 2021) and can help VL tasks (Shen et al., 2021).

B. Visual Generation

Besides visual understanding, large-scale weakly-labeled image-text-paired data can also assist text-to-image generation. Ramesh et al. (2021) developed an image generation system called DALL-E. DALL-E converts images into discrete visual

tokens using a discrete variational auto encoder (dVAE) so that a (text, image) pair can be viewed as a single stream of data. During training, the text-image stream is fed into a decoder only Transformer. As for the attention mask, each image token can see all the text tokens. The attention mask among text tokens is standard causal mask. And image-to-image attention use either row, column or convolutional attention mask. In inference time, given text tokens, the generation process is to predict image tokens in an auto-regressive way as in GPT. DALL-E shows impressive results in four aspects: creating anthropomorphized versions of animals and objects, combining unrelated concepts, rendering text, and applying a transformation to existing images.

Inspired by the training method in DALL-E, Wu et al. (2021a) proposed a method named GODIVA to generate videos from the text. Similar to DALL-E, GODIVA tokenizes each frame of the video and concatenates the text and visual tokens sequentially as a stream to train the model. DALL-E and GODIVA are designed for text-to-image generation and text-to-video generation, respectively, while Wu et al. (2021b) proposed a unified visual generation model which achieves state-of-the-art results on 8 downstream tasks including text-to-image, text-to-video, video prediction, etc. They proposed a 3D Transformer that is able to encode all three data formats including text (1D), images (2D), and videos (3D). To better attend on videos, they designed a 3D Nearby Attention to apply attention along both spatial and temporal axes.

V. FUTURE TRENDS

In the last few years, we have witnessed how VLP models scale to use large quantities of weakly-labeled and more diverse data. In the future, the models and data will continue to scale up to achieve stronger modality cooperation and even unified representation. In addition, incorporating knowledge can further empower VLP models to gain better generalization abilities. In this section, we will discuss these future trends.

A. Toward Modality Cooperation

Besides improving cross-modal tasks with VL datasets, modal cooperation is emerging in pre-training to boost the performance of both single-modal tasks and multi-modal tasks. Modal cooperation is to help different modalities to help each other and learn better representation. For example, improve language tasks with vision data and improve cross-modal tasks with single-modality data.

1) Improve Language Tasks with Vision Data :
Improving language learning with visual information has been explored on a broad range of language tasks, including machine translation (Ive et al., 2019; Zhang et al., 2019), semantic parsing (Shi et al., 2019a; Kojima et al., 2020), and language grounding (Bordes et al., 2019; Kiela et al., 2018). These studies are tailored for specific language tasks and may suffer from modality discrepancies. Tan and Bansal (2020) proposed a general pre-training model for language representation with a visual aid, in which a "vokenization" model was introduced to extrapolate vision-language-alignment from image captioning datasets to pure-language corpora. More

specifically, the "vokenization" model is trained with image-text matching to construct a visual image vocabulary, which is leveraged to map text tokens in language-only datasets to the retrieved images with the highest score. The experimental results show that it can yield additional improvement over self-supervised language models.

2) Improve Cross-Modal Tasks with Single-Modality Data.: To address the data shortage issue, some VLP models utilize extra single-modal data to improve the representation capability. For example, in an image-text dataset, text is usually short with several tokens, which restricts the textual representation. Therefore, VL-BERT (Su et al., 2019) adds additional linguistic corpora to improve the language part in cross-modal tasks.

B. Toward General Unified-Modality

Thanks to the Transformer architecture, researchers have achieved remarkable progress in both single-modal and multi-modal representation learning. In previous sections, we have discussed multi-modal representation and modal cooperation, which connect vision and language in different ways. A more ambitious goal is to build a general representation model which can unify multiple modalities. As a pioneering work, UNIMO (Li et al., 2020a) proposed a unified pre-training model, which can handle both single-modal and multi-modal downstream tasks including understanding and generation. Trained with a huge amount of single-modal as well as cross-modal data including BookWiki (Zhu et al., 2015) and OpenWebText (language data), OpenImages (Krasin et al., 2017) and COCO (Lin et al., 2014) (image data), and COCO (Lin et al., 2014), Visual Genome (Krishna et al., 2016) , Conceptual Captions (Sharma et al., 2018) and SBU (Ordonez et al., 2011) (image-text data). As a result, UNIMO improves many single-modal and multi-modal downstream tasks by a large margin. Another interesting work is a general-purpose vision system developed by Gupta et al. (2021) for a series of vision and cross-modal tasks.

C. VL+Knowledge

Many VL tasks require common sense and factual information beyond training datasets. However, most VLP models do not have a mechanism to consume extra knowledge. ERNIE (Sun et al., 2019b) proposed a multi-stage knowledge-based masking strategy. Instead of directly adding knowledge embedding, it masks language in three levels, which are basic-level, phrase-level, and entity-level masking. For entity-level masking, the model masks a whole entity rather than a sub-word. Such entities include persons, locations, organizations, products, etc. There are also ways of integrating knowledge into VLP models. Shevchenko et al. (2021) proposed to directly inject knowledge embeddings into a vision-language Transformer. They first build a knowledge base (KB) with knowledge embeddings and then match the sentences in the training data with knowledge embeddings. During training, they use an auxiliary loss to encourage the learned representation to be aligned with knowledge embeddings. Although there have been some works seeking to integrate knowledge into VLP models, there are still many challenges to be addressed such as how to

efficiently utilize large wikidata with high noise and how to learn from knowledge in an explainable way.

REFERENCES

Harsh Agrawal, Karan Desai, Xinlei Chen, Rishabh Jain, Dhruv Batra, Devi Parikh, Stefan Lee, and Peter Anderson. nocaps: novel object captioning at scale. In *ICCV*, 2019.

Chris Alberti, Jeffrey Ling, Michael Collins, and David Reitter. Fusion of detected objects in text for visual question answering. In *EMNLP*, pages 2131–2140, 2019.

Peter Anderson, Basura Fernando, Mark Johnson, and Stephen Gould. SPICE: Semantic Propositional Image Caption Evaluation. In *ECCV*, 2016.

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 *CVPR*, pages 6077–6086, 2018a.

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, 2018b.

Jacob Andreas, Marcus Rohrbach, Trevor Darrell, and Dan Klein. Neural module networks, 2017.

Stanislaw Antol, Aishwarya Agrawal, Jiasen Lu, Margaret Mitchell, Dhruv Batra, C Lawrence Zitnick, and Devi Parikh. Vqa: Visual question answering. In *Proceedings of the IEEE international conference on computer vision*, pages 2425–2433, 2015.

Satanjeev Banerjee and Alon Lavie. METEOR: An automatic metric for MT evaluation with improved correlation with human judgments. 2005.

Hedi Ben-Younes, Rémi Cadene, Matthieu Cord, and Nicolas Thome. Mutan: Multimodal tucker fusion for visual question answering. In *ICCV*, pages 2612–2620, 2017.

Rodrigo Benenson, Stefan Popov, and Vittorio Ferrari. Large-scale interactive object segmentation with human annotators. In *CVPR*, 2019.

Yonatan Bisk, Ari Holtzman, Jesse Thomason, Jacob Andreas, Yoshua Bengio, Joyce Chai, Mirella Lapata, Angeliki Lazaridou, Jonathan May, Aleksandr Nisnevich, Nicolas Pinto, and Joseph Turian. Experience grounds language, 2020.

Patrick Bordes, Éloi Zablocki, Laure Soulier, Benjamin Piwowarski, and Patrick Gallinari. Incorporating visual semantics into sentence representations within a grounded space. In *2019 Conference on Empirical Methods in Natural Language Processing and 9th International Joint Conference on Natural Language Processing*, pages 696–707. Association for Computational Linguistics, 2019.

Emanuele Bugliarello, Ryan Cotterell, Naoaki Okazaki, and Desmond Elliott. Multimodal pretraining unmasked: A meta-analysis and a unified framework of vision-and-language berts. *Transactions of the Association for Computational Linguistics*, 9:978–994, 2021.

Jize Cao, Zhe Gan, Yu Cheng, Licheng Yu, Yen-Chun Chen, and Jingjing Liu. Behind the scene: Revealing the secrets of pre-trained vision-and-language models. In *European Conference on Computer Vision*, pages 565–580. Springer, 2020.

Soravit Changpinyo, Piyush Sharma, Nan Ding, and Radu Soricut. Conceptual 12M: Pushing Web-Scale Image-Text Pre-Training To Recognize Long-Tail Visual Concepts. In *CVPR*, 2021.

Yen-Chun Chen, Linjie Li, Licheng Yu, Ahmed El Kholy, Faisal Ahmed, Zhe Gan, Yu Cheng, and Jingjing Liu. Uniter: Learning universal image-text representations. 2019.

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 *ECCV*, pages 104–120, 2020.

Junyoung Chung, Caglar Gulcehre, KyungHyun Cho, and Yoshua Bengio. Empirical evaluation of gated recurrent neural networks on sequence modeling, 2014.

Marcella Cornia, Lorenzo Baraldi, and Rita Cucchiara. Show, Control and Tell: A Framework for Generating Controllable and Grounded Captions. In *CVPR*, 2019.

Xiyang Dai, Yinpeng Chen, Jianwei Yang, Pengchuan Zhang, Lu Yuan, and Lei Zhang. Dynamic detr: End-to-end object detection with dynamic attention, October 2021.

Abhishek Das, Satwik Kottur, Khushi Gupta, Avi Singh, Deshraj Yadav, José MF Moura, Devi Parikh, and Dhruv Batra. Visual dialog. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 326–335, 2017.

Samyak Datta, Karan Sikka, Anirban Roy, Karuna Ahuja, Devi Parikh, and Ajay Divakaran. Align2ground: Weakly supervised phrase grounding guided by image-caption alignment, 2019.

Harm De Vries, Florian Strub, Sarath Chandar, Olivier Pietquin, Hugo Larochelle, and Aaron Courville. Guesswhat?! visual object discovery through multi-modal dialogue. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5503–5512, 2017.

Thierry Deruyttere, Simon Vandenhende, Dusan Grujicic, Luc Van Gool, and Marie-Francine Moens. Talk2car: Taking control of your self-driving car. *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, 2019. doi: 10.18653/v1/d19-1215. URL http://dx.doi.org/10.18653/v1/D19-1215.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: Pre-training of deep bidirectional transformers for language understanding. 2018.

Jeffrey Donahue, Lisa Anne Hendricks, Sergio Guadarrama, Marcus Rohrbach, Subhashini Venugopalan, Kate Saenko, and Trevor Darrell. Long-term recurrent convolutional networks for visual recognition and description. In *CVPR*, 2015.

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,

2021a.

Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale. *ICLR*, 2021b.

Fartash Faghri, David J. Fleet, Jamie Ryan Kiros, and Sanja Fidler. Vse++: Improving visual-semantic embeddings with hard negatives, 2018.

Ali Farhadi, Mohsen Hejrati, Mohammad Amin Sadeghi, Peter Young, Cyrus Rashtchian, Julia Hockenmaier, and David Forsyth. Every picture tells a story: Generating sentences from images. In *ECCV*, 2010.

Andrea Frome, Greg Corrado, Jonathon Shlens, Samy Bengio, Jeffrey Dean, Marc'Aurelio Ranzato, and Tomas Mikolov. Devise: A deep visual-semantic embedding model. In *Neural Information Processing Systems (NIPS)*, 2013.

Akira Fukui, Dong Huk Park, Daylen Yang, Anna Rohrbach, Trevor Darrell, and Marcus Rohrbach. Multimodal compact bilinear pooling for visual question answering and visual grounding, 2016.

Philip Gage. A new algorithm for data compression. *C Users Journal*, 12(2):23–38, 1994.

Zhe Gan, Yen-Chun Chen, Linjie Li, Chen Zhu, Yu Cheng, and Jingjing Liu. Large-scale adversarial training for vision-and-language representation learning. *arXiv preprint arXiv:2006.06195*, 2020.

Haoyuan Gao, Junhua Mao, Jie Zhou, Zhiheng Huang, Lei Wang, and Wei Xu. Are you talking to a machine? dataset and methods for multilingual image question answering, 2015.

Gabriel Goh, Nick Cammarata, Chelsea Voss, Shan Carter, Michael Petrov, Ludwig Schubert, Alec Radford, and Chris Olah. Multimodal neurons in artificial neural networks. *Distill*, 6(3):e30, 2021.

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 *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.

Xiuye Gu, Tsung-Yi Lin, Weicheng Kuo, and Yin Cui. Zero-shot detection via vision and language knowledge distillation, 2021.

Tanmay Gupta, Arash Vahdat, Gal Chechik, Xiaodong Yang, Jan Kautz, and Derek Hoiem. Contrastive learning for weakly supervised phrase grounding, 2020.

Tanmay Gupta, Amita Kamath, Aniruddha Kembhavi, and Derek Hoiem. Towards general purpose vision systems, 2021.

Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition, 2015.

Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016a.

Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Identity mappings in deep residual networks. In *European conference on computer vision*, pages 630–645. Springer,

2016b.

Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners, 2021.

Simao Herdade, Armin Kappeler, Kofi Boakye, and Joao Soares. Image Captioning: Transforming Objects into Words. In *NeurIPS*, 2019.

Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. *Advances in neural information processing systems*, 30, 2017.

Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997a.

Sepp Hochreiter and Jürgen Schmidhuber. Long Short-Term Memory. *Neural Computation*, 9(8):1735–1780, 11 1997b. ISSN 0899-7667. doi: 10.1162/neco.1997.9.8.1735. URL https://doi.org/10.1162/neco.1997.9.8.1735.

Micah Hodosh, Peter Young, and Julia Hockenmaier. Framing image description as a ranking task: Data, models and evaluation metrics. 2013.

Ronghang Hu, Jacob Andreas, Marcus Rohrbach, Trevor Darrell, and Kate Saenko. Learning to reason: End-to-end module networks for visual question answering, 2017.

Ronghang Hu, Jacob Andreas, Trevor Darrell, and Kate Saenko. Explainable neural computation via stack neural module networks, 2019.

Xiaowei Hu, Xi Yin, Kevin Lin, Lijuan Wang, Lei Zhang, Jianfeng Gao, and Zicheng Liu. Vivo: Surpassing human performance in novel object captioning with visual vocabulary pre-training. *arXiv e-prints*, pages arXiv–2009, 2020.

Lun Huang, Wenmin Wang, Jie Chen, and Xiao-Yong Wei. Attention on attention for image captioning, 2019.

Yan Huang, Wei Wang, and Liang Wang. Instance-aware image and sentence matching with selective multimodal lstm, 2016.

Zhicheng Huang, Zhaoyang Zeng, Bei Liu, Dongmei Fu, and Jianlong Fu. Pixel-bert: Aligning image pixels with text by deep multi-modal transformers. *arXiv:2004.00849*, 2020.

Zhicheng Huang, Zhaoyang Zeng, Yupan Huang, Bei Liu, Dongmei Fu, and Jianlong Fu. Seeing out of the box: End-to-end pre-training for vision-language representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12976–12985, 2021.

Drew A. Hudson and Christopher D. Manning. Compositional attention networks for machine reasoning, 2018.

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.

Julia Ive, Pranava Madhyastha, and Lucia Specia. Distilling translations with visual awareness. *arXiv preprint arXiv:1906.07701*, 2019.

Chao Jia, Yinfei Yang, Ye Xia, Yi-Ting Chen, Zarana Parekh, Hieu Pham, Quoc V Le, Yunhsuan Sung, Zhen Li, and Tom Duerig. Scaling up visual and vision-language representation learning with noisy text supervision. *arXiv preprint arXiv:2102.05918*, 2021.

Justin Johnson, Bharath Hariharan, Laurens van der Maaten,

Judy Hoffman, Li Fei-Fei, C. Lawrence Zitnick, and Ross Girshick. Inferring and executing programs for visual reasoning, 2017.

Aishwarya Kamath, Mannat Singh, Yann LeCun, Gabriel Synnaeve, Ishan Misra, and Nicolas Carion. Mdetr – modulated detection for end-to-end multi-modal understanding, 2021.

Andrej Karpathy and Li Fei-Fei. Deep visual-semantic alignments for generating image descriptions. In *CVPR*, 2015.

Andrej Karpathy, Armand Joulin, and Li Fei-Fei. Deep fragment embeddings for bidirectional image sentence mapping, 2014.

Sahar Kazemzadeh, Vicente Ordonez, Mark Matten, and Tamara Berg. Referitgame: Referring to objects in photographs of natural scenes. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, pages 787–798, 2014.

Douwe Kiela, Alexis Conneau, Allan Jabri, and Maximilian Nickel. Learning visually grounded sentence representations. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 408–418, 2018.

Jin-Hwa Kim, Sang-Woo Lee, Dong-Hyun Kwak, Min-Oh Heo, Jeonghee Kim, Jung-Woo Ha, and Byoung-Tak Zhang. Multimodal residual learning for visual qa, 2016.

Wonjae Kim, Bokyung Son, and Ildoo Kim. Vilt: Vision-and-language transformer without convolution or region supervision. *arXiv preprint arXiv:2102.03334*, 2021.

Ryan Kiros, Ruslan Salakhutdinov, and Richard S. Zemel. Unifying visual-semantic embeddings with multimodal neural language models, 2014.

Benjamin Klein, Guy Lev, Gil Sadeh, and Lior Wolf. Associating neural word embeddings with deep image representations using fisher vectors. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4437–4446, 2015.

Noriyuki Kojima, Hadar Averbuch-Elor, Alexander Rush, and Yoav Artzi. What is learned in visually grounded neural syntax acquisition. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, 2020.

Satwik Kottur, José MF Moura, Devi Parikh, Dhruv Batra, and Marcus Rohrbach. Clevr-dialog: A diagnostic dataset for multi-round reasoning in visual dialog. *arXiv preprint arXiv:1903.03166*, 2019.

Ivan Krasin, Tom Duerig, Neil Alldrin, Vittorio Ferrari, Sami Abu-El-Haija, Alina Kuznetsova, Hassan Rom, Jasper Uijlings, Stefan Popov, Andreas Veit, et al. Openimages: A public dataset for large-scale multi-label and multi-class image classification. *Dataset available from https://github.com/openimages*, 2(3):18, 2017.

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. *arXiv preprint arXiv:1602.07332*, 2016.

Ranjay Krishna, Yuke Zhu, Oliver Groth, Justin Johnson, Kenji Hata, Joshua Kravitz, Stephanie Chen, Yannis Kalantidis, Li-Jia Li, David A Shamma, Michael Bernstein, and Li Fei-Fei. Visual Genome: Connecting Language and Vision Using Crowdsourced Dense Image Annotations. *IJCV*, 123(1): 32–73, 2017a.

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, 2017b.

Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. *Advances in neural information processing systems*, 25:1097–1105, 2012.

Girish Kulkarni, Visruth Premraj, Vicente Ordonez, Sagnik Dhar, Siming Li, Yejin Choi, Alexander C Berg, and Tamara L Berg. Babytalk: Understanding and generating simple image descriptions. 2013.

Alina Kuznetsova, Hassan Rom, Neil Alldrin, Jasper Uijlings, Ivan Krasin, Jordi Pont-Tuset, Shahab Kamali, Stefan Popov, Matteo Malloci, Tom Duerig, and Vittorio Ferrari. The open images dataset v4: Unified image classification, object detection, and visual relationship detection at scale. *arXiv:1811.00982*, 2018.

Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.

Kuang-Huei Lee, Xi Chen, Gang Hua, Houdong Hu, and Xiaodong He. Stacked Cross Attention for Image-Text Matching. In *ECCV*, 2018.

Chenliang Li, Ming Yan, Haiyang Xu, Fuli Luo, Wei Wang, Bin Bi, and Songfang Huang. Semvlp: Vision-language pre-training by aligning semantics at multiple levels. *arXiv preprint arXiv:2103.07829*, 2021a.

Gen Li, Nan Duan, Yuejian Fang, Daxin Jiang, and Ming Zhou. Unicoder-VL: A universal encoder for vision and language by cross-modal pre-training. In *AAAI*, pages 11336–11344, 2019a.

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:1908.03557*, 2019b.

Liunian Harold Li, Pengchuan Zhang, Haotian Zhang, Jianwei Yang, Chunyuan Li, Yiwu Zhong, Lijuan Wang, Lu Yuan, Lei Zhang, Jenq-Neng Hwang, Kai-Wei Chang, and Jianfeng Gao. Grounded language-image pre-training, 2021b.

Wei Li, Can Gao, Guocheng Niu, Xinyan Xiao, Hao Liu, Jiachen Liu, Hua Wu, and Haifeng Wang. Unimo: Towards unified-modal understanding and generation via cross-modal contrastive learning. *arXiv preprint arXiv:2012.15409*, 2020a.

Wenbo Li, Pengchuan Zhang, Lei Zhang, Qiuyuan Huang, Xiaodong He, Siwei Lyu, and Jianfeng Gao. Object-driven text-to-image synthesis via adversarial training. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12174–12182, 2019c.

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 *ECCV*, 2020b.

Chin-Yew Lin. Rouge: A package for automatic evaluation of summaries. 2004.

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.

Tsung-Yu Lin, Aruni RoyChowdhury, and Subhransu Maji. Bilinear cnns for fine-grained visual recognition, 2017.

Wei Liu, Sihan Chen, Longteng Guo, Xinxin Zhu, and Jing Liu. CPTR: Full Transformer Network for Image Captioning. *arXiv preprint arXiv:2101.10804*, 2021.

Jiasen Lu, Caiming Xiong, Devi Parikh, and Richard Socher. Knowing when to look: Adaptive attention via a visual sentinel for image captioning, 2017a.

Jiasen Lu, Jianwei Yang, Dhruv Batra, and Devi Parikh. Hierarchical question-image co-attention for visual question answering, 2017b.

Jiasen Lu, Jianwei Yang, Dhruv Batra, and Devi Parikh. Neural baby talk. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7219–7228, 2018.

Jiasen Lu, Dhruv Batra, Devi Parikh, and Stefan Lee. Vilbert: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks. In *NeurIPS*, 2019.

Jiasen Lu, Vedanuj Goswami, Marcus Rohrbach, Devi Parikh, and Stefan Lee. 12-in-1: Multi-task vision and language representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10437–10446, 2020.

Lin Ma, Zhengdong Lu, Lifeng Shang, and Hang Li. Multimodal convolutional neural networks for matching image and sentence, 2015.

Mateusz Malinowski and Mario Fritz. A multi-world approach to question answering about real-world scenes based on uncertain input. *Advances in neural information processing systems*, 27:1682–1690, 2014.

Mateusz Malinowski, Marcus Rohrbach, and Mario Fritz. Ask your neurons: A neural-based approach to answering questions about images, 2015.

Junhua Mao, Wei Xu, Yi Yang, Jiang Wang, Zhiheng Huang, and Alan Yuille. Deep captioning with multimodal recurrent neural networks (m-rnn), 2015.

Junhua Mao, Jonathan Huang, Alexander Toshev, Oana Camburu, Alan L Yuille, and Kevin Murphy. Generation and comprehension of unambiguous object descriptions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 11–20, 2016.

David Mascharka, Philip Tran, Ryan Soklaski, and Arjun Majumdar. Transparency by design: Closing the gap between performance and interpretability in visual reasoning. *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, Jun 2018. doi: 10.1109/cvpr.2018.00519. URL http://dx.doi.org/10.1109/CVPR.2018.00519.

Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781*, 2013.

Aditya Mogadala, Marimuthu Kalimuthu, and Dietrich Klakow. Trends in integration of vision and language research: A survey of tasks, datasets, and methods. *Journal of Artificial Intelligence Research*, 71:1183–1317, 2021.

Hyeonseob Nam, Jung-Woo Ha, and Jeonghee Kim. Dual attention networks for multimodal reasoning and matching. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 299–307, 2017.

Vicente Ordonez, Girish Kulkarni, and Tamara Berg. Im2text: Describing images using 1 million captioned photographs. *NeurIPS*, 2011.

Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. BLEU: a method for automatic evaluation of machine translation. 2002.

Bryan A. Plummer, Liwei Wang, Chris M. Cervantes, Juan C. Caicedo, Julia Hockenmaier, and Svetlana Lazebnik. Flickr30k entities: Collecting region-to-phrase correspondences for richer image-to-sentence models, 2016.

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. *arXiv preprint arXiv:2103.00020*, 2021.

Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. Zero-shot text-to-image generation, 2021.

Mengye Ren, Ryan Kiros, and Richard Zemel. Exploring models and data for image question answering, 2015.

Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks, 2016.

Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster R-CNN: towards real-time object detection with region proposal networks. 39(6):1137–1149, 2017.

Frank Rosenblatt. *The perceptron, a perceiving and recognizing automaton Project Para*. Cornell Aeronautical Laboratory, 1957.

Frank Rosenblatt. Principles of neurodynamics. perceptrons and the theory of brain mechanisms. Technical report, Cornell Aeronautical Lab Inc Buffalo NY, 1961.

David E Rumelhart, Geoffrey E Hinton, and Ronald J Williams. Learning internal representations by error propagation. Technical report, California Univ San Diego La Jolla Inst for Cognitive Science, 1985.

Tim Salimans, Ian Goodfellow, Wojciech Zaremba, Vicki Cheung, Alec Radford, and Xi Chen. Improved techniques for training gans. *Advances in neural information processing systems*, 29:2234–2242, 2016.

Rico Sennrich, Barry Haddow, and Alexandra Birch. Neural machine translation of rare words with subword units. *arXiv preprint arXiv:1508.07909*, 2015.

Piyush Sharma, Nan Ding, Sebastian Goodman, and Radu Soricut. Conceptual captions: A cleaned, hypernymed, image alt-text dataset for automatic image captioning. 2018.

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?, 2021.

Violetta Shevchenko, Damien Teney, Anthony Dick, and Anton van den Hengel. Reasoning over vision and language: Exploring the benefits of supplemental knowledge, 2021.

Haoyue Shi, Jiayuan Mao, Kevin Gimpel, and Karen Livescu. Visually grounded neural syntax acquisition. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1842–1861, 2019a.

Jiaxin Shi, Hanwang Zhang, and Juanzi Li. Explainable and explicit visual reasoning over scene graphs, 2019b.

Kevin J. Shih, Saurabh Singh, and Derek Hoiem. Where to look: Focus regions for visual question answering, 2016.

Kurt Shuster, Samuel Humeau, Antoine Bordes, and Jason Weston. Image chat: Engaging grounded conversations. *arXiv preprint arXiv:1811.00945*, 2018.

Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. In *ICLR*, 2015a.

Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition, 2015b.

Amanpreet Singh, Vedanuj Goswami, Vivek Natarajan, Yu Jiang, Xinlei Chen, Meet Shah, Marcus Rohrbach, Dhruv Batra, and Devi Parikh. Pythia-a platform for vision & language research. In *NeurIPS, SysML Workshop*, volume 2018, 2018.

Richard Socher, Andrej Karpathy, Quoc V. Le, Christopher D. Manning, and Andrew Y. Ng. Grounded compositional semantics for finding and describing images with sentences. *Transactions of the Association for Computational Linguistics*, 2:207–218, 2014. doi: 10.1162/tacl_a_00177.

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 *ICLR*, 2019.

Alane Suhr, Mike Lewis, James Yeh, and Yoav Artzi. A corpus of natural language for visual reasoning. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 217–223, 2017.

Alane Suhr, Stephanie Zhou, Ally Zhang, Iris Zhang, Huajun Bai, and Yoav Artzi. A corpus for reasoning about natural language grounded in photographs. In *ACL*, pages 6418–6428, 2019.

Chen Sun, Austin Myers, Carl Vondrick, Kevin Murphy, and Cordelia Schmid. Videobert: A joint model for video and language representation learning. In *ICCV*, 2019a.

Yu Sun, Shuohuan Wang, Yukun Li, Shikun Feng, Xuyi Chen, Han Zhang, Xin Tian, Danxiang Zhu, Hao Tian, and Hua Wu. Ernie: Enhanced representation through knowledge integration. *arXiv preprint arXiv:1904.09223*, 2019b.

Ilya Sutskever, Oriol Vinyals, and Quoc V. Le. Sequence to sequence learning with neural networks, 2014.

Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. Going deeper with convolutions, 2014.

Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. Going deeper with convolutions. In *CVPR*, 2015.

Hao Tan and Mohit Bansal. Lxmert: Learning cross-modality encoder representations from transformers. 2019.

Hao Tan and Mohit Bansal. Vokenization: improving language understanding with contextualized, visual-grounded supervision. *arXiv preprint arXiv:2010.06775*, 2020.

Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *NeurIPS*, 2017.

Ramakrishna Vedantam, C Lawrence Zitnick, and Devi Parikh. CIDEr: Consensus-based Image Description Evaluation. In *CVPR*, 2015.

Oriol Vinyals, Alexander Toshev, Samy Bengio, and Dumitru Erhan. Show and tell: A neural image caption generator. In *CVPR*, 2015.

Catherine Wah, Steve Branson, Peter Welinder, Pietro Perona, and Serge Belongie. The caltech-ucsd birds-200-2011 dataset. 2011.

Liwei Wang, Yin Li, and Svetlana Lazebnik. Learning deep structure-preserving image-text embeddings, 2016.

Zirui Wang, Jiahui Yu, Adams Wei Yu, Zihang Dai, Yulia Tsvetkov, and Yuan Cao. Simvlm: Simple visual language model pretraining with weak supervision. *arXiv preprint arXiv:2108.10904*, 2021.

Chenfei Wu, Lun Huang, Qianxi Zhang, Binyang Li, Lei Ji, Fan Yang, Guillermo Sapiro, and Nan Duan. Godiva: Generating open-domain videos from natural descriptions, 2021a.

Chenfei Wu, Jian Liang, Lei Ji, Fan Yang, Yuejian Fang, Daxin Jiang, and Nan Duan. Nüwa: Visual synthesis pre-training for neural visual world creation, 2021b.

Ning Xie, Farley Lai, Derek Doran, and Asim Kadav. Visual entailment task for visually-grounded language learning, 2019.

Saining Xie, Ross Girshick, Piotr Dollár, Zhuowen Tu, and Kaiming He. Aggregated residual transformations for deep neural networks. In *CVPR*, pages 1492–1500, 2017a.

Saining Xie, Chen Sun, Jonathan Huang, Zhuowen Tu, and Kevin Murphy. Rethinking spatiotemporal feature learning for video understanding. *arXiv preprint arXiv:1712.04851*, 1(2):5, 2017b.

Kelvin Xu, Jimmy Ba, Ryan Kiros, Kyunghyun Cho, Aaron Courville, Ruslan Salakhutdinov, Richard S Zemel, and Yoshua Bengio. Show, attend and tell: Neural image caption generation with visual attention. 2015.

Tao Xu, Pengchuan Zhang, Qiuyuan Huang, Han Zhang, Zhe Gan, Xiaolei Huang, and Xiaodong He. Attngan: Fine-grained text to image generation with attentional generative adversarial networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1316–1324, 2018.

Hongwei Xue, Yupan Huang, Bei Liu, Houwen Peng, Jianlong Fu, Houqiang Li, and Jiebo Luo. Probing inter-modality: Visual parsing with self-attention for vision-language pre-training. *arXiv preprint arXiv:2106.13488*, 2021.

Zichao Yang, Xiaodong He, Jianfeng Gao, Li Deng, and Alex Smola. Stacked attention networks for image question answering, 2016.

Kexin Yi, Jiajun Wu, Chuang Gan, Antonio Torralba, Pushmeet

Kohli, and Joshua B. Tenenbaum. Neural-symbolic vqa: Disentangling reasoning from vision and language understanding, 2019.

Peter Young, Alice Lai, Micah Hodosh, and Julia Hockenmaier. From image descriptions to visual denotations: New similarity metrics for semantic inference over event descriptions. 2014.

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. *arXiv preprint arXiv:2006.16934*, 1:12, 2020.

Licheng Yu, Patrick Poirson, Shan Yang, Alexander C. Berg, and Tamara L. Berg. Modeling context in referring expressions, 2016.

Licheng Yu, Zhe Lin, Xiaohui Shen, Jimei Yang, Xin Lu, Mohit Bansal, and Tamara L. Berg. Mattnet: Modular attention network for referring expression comprehension, 2018.

Lu Yuan, Dongdong Chen, Yi-Ling Chen, Noel Codella, Xiyang Dai, Jianfeng Gao, Houdong Hu, Xuedong Huang, Boxin Li, Chunyuan Li, Ce Liu, Mengchen Liu, Zicheng Liu, Yumao Lu, Yu Shi, Lijuan Wang, Jianfeng Wang, Bin Xiao, Zhen Xiao, Jianwei Yang, Michael Zeng, Luowei Zhou, and Pengchuan Zhang. Florence: A new foundation model for computer vision, 2021.

Rowan Zellers, Yonatan Bisk, Ali Farhadi, and Yejin Choi. From recognition to cognition: Visual commonsense reasoning. In *CVPR*, pages 6720–6731, 2019.

Chao Zhang, Zichao Yang, Xiaodong He, and Li Deng. Multimodal intelligence: Representation learning, information fusion, and applications. *IEEE Journal of Selected Topics in Signal Processing*, 14(3):478–493, 2020.

Han Zhang, Tao Xu, Hongsheng Li, Shaoting Zhang, Xiaogang Wang, Xiaolei Huang, and Dimitris Metaxas. Stackgan: Text to photo-realistic image synthesis with stacked generative adversarial networks, 2017.

Lei Zhang and Heung-Yeung Shum. Statistical foundation behind machine learning and its impact on computer vision. *Manuscript under review*, 2022.

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.

Zhuosheng Zhang, Kehai Chen, Rui Wang, Masao Utiyama, Eiichiro Sumita, Zuchao Li, and Hai Zhao. Neural machine translation with universal visual representation. In *International Conference on Learning Representations*, 2019.

Luowei Zhou, Hamid Palangi, Lei Zhang, Houdong Hu, Jason J Corso, and Jianfeng Gao. Unified vision-language pretraining for image captioning and VQA. In *AAAI*, pages 13041–13049, 2020.

Xinxin Zhu, Weining Wang, Longteng Guo, and Jing Liu. AutoCaption: Image Captioning with Neural Architecture Search. *arXiv preprint arXiv:2012.09742*, 2020.

Yuke Zhu, Oliver Groth, Michael Bernstein, and Li Fei-Fei. Visual7w: Grounded question answering in images. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4995–5004, 2016.

Yukun Zhu, Ryan Kiros, Rich Zemel, Ruslan Salakhutdinov, Raquel Urtasun, Antonio Torralba, and Sanja Fidler. Aligning books and movies: Towards story-like visual explanations by watching movies and reading books. In *Proceedings of the IEEE international conference on computer vision*, pages 19–27, 2015.

C Lawrence Zitnick and Piotr Dollár. Edge boxes: Locating object proposals from edges. In *European conference on computer vision*, pages 391–405. Springer, 2014.