

Training Vision-Language Models with Less Bimodal Supervision

Elad Segal

Ben Bogin

Jonathan Berant

*Blavatnik School of Computer Science
Tel Aviv University*

ELAD.SEGAL@GMAIL.COM

BEN.BOGIN@CS.TAU.AC.IL

JOBERANT@CS.TAU.AC.IL

Abstract

Standard practice in pertaining multimodal models, such as vision-language models, is to rely on pairs of aligned inputs from both modalities, for example, aligned image-text pairs. However, such pairs can be difficult to obtain in low-resource settings and for some modality pairs (e.g., structured tables and images). In this work, we investigate the extent to which we can reduce the reliance on such parallel data, which we term *bimodal supervision*, and use models that are pretrained on each modality independently. We experiment with a high-performing vision-language model, and analyze the effect of bimodal supervision on three vision-language tasks. We find that on simpler tasks, such as VQAv2 and GQA, one can eliminate bimodal supervision completely, suffering only a minor loss in performance. Conversely, for NLVR2, which requires more complex reasoning, training without bimodal supervision leads to random performance. Nevertheless, using only 5% of the bimodal data (142K images along with their captions), or leveraging weak supervision in the form of a list of machine-generated labels for each image, leads to only a moderate degradation compared to using 3M image-text pairs: 74% \rightarrow \sim 70%.

1. Introduction

Pretraining models on large amounts of raw data using self-supervision has revolutionized machine learning, and is now standard practice across a wide range of modalities [Liu et al., 2019, Raffel et al., 2020, Dosovitskiy et al., 2021b, Liu et al., 2021, Herzig et al., 2020, Schneider et al., 2019, Baevski et al., 2022]. While typically pretrained models are trained on data from a single modality (*unimodal data*), the success of pretraining has spread to the *bimodal* setup, where models are trained on pairs of inputs, each from a different modality (e.g. text and audio, Li et al., 2021a). Most notably, vision-language models, such as LXMERT [Tan and Bansal, 2019], ViLT [Kim et al., 2021a], METER [Dou et al., 2022], CLIP [Radford et al., 2021], and others [Li et al., 2019, Lu et al., 2019b, Li et al., 2021b], have been pretrained on manually or automatically collected parallel data that consists of aligned image-text pairs.

While effective, pretraining on bimodal data comes at a cost. First, gathering high-quality pairs can be challenging, especially in low-resource languages and domains, or for modality pairs where parallel data is scarce. Second, expanding this approach to more than two modalities (as in, e.g., MultimodalQA, Talmor et al., 2021) is challenging. Last, pretraining is computationally expensive [Strubell et al., 2019, Bommasani et al., 2021], and thus relying on pretraining for all modality pairs is inefficient.




Figure 1: The effect of unimodal and bimodal pretraining on downstream performance after finetuning. In VQAv2 and GQA, pretraining on unimodal data alone (without image-text pairs) is competitive with models pretrained on image-text pairs. On NLVR2, bimodal supervision is necessary, but one can reach reasonable performance using only 5% of the image-text pairs or training on machine-generated object labels. Random initialization leads to poor performance on all tasks.

Given these shortcomings, a natural question is how far can we get with models pretrained on unimodal data only (*unimodal models*), such as BERT [Devlin et al., 2019] and ViT [Dosovitskiy et al., 2021a], to reduce or obviate the need for *bimodal* pretraining. Can we align unimodal representations without resorting to pretraining over millions of input pairs? While past work [Dou et al., 2022, Li et al., 2021c, Zhai et al., 2022] used unimodal models as an initialization point before bimodal pretraining, it did not investigate its effect on the amount of necessary bimodal data.

In this work, we investigate to what extent we can reduce the burden of bimodal pretraining and finetune models on vision-language applications starting with models that were unimodally pretrained. We choose a high-performing architecture [Dou et al., 2022] – a transformer image encoder and a transformer text encoder, which pass their representations through additional transformer layers that capture the interaction between the image and the text, before performing a final classification task.

We test performance on visual question answering and visual reasoning tasks in the following setups: (a) randomly initialized image and text encoders, (b) unimodally-pretrained image and text encoders, and (c) unimodally-pretrained image and text encoders that are then pretrained with bimodal supervision. We test different sources for bimodal pretraining, which require different amounts of human effort: (a) automatically harvested image-caption pairs (Conceptual Captions, Sharma et al., 2018), (b) images paired with machine-generated object labels (CCIL, Ng et al., 2021), (c) manually annotated image-object pairs (ImageNet-1K, Russakovsky et al., 2015), and (d) image-question-answer triples from visual question answering tasks. We note that due to computational constraints the size of our pretraining corpus is smaller compared to those used by industry-based researchers [Li et al., 2022, Radford et al., 2021, Jia et al., 2021].

We find (Figure 1) that on some tasks, models that do not use any bimodal supervision are only slightly worse than models that are pretrained on large amounts of image-text pairs – 70.7→69.5 on VQAv2, and 56.1→53.6 on GQA. However, for a more complex reasoning task, such as NLVR2, bimodal supervision is *crucial*. Nevertheless, we show that one can

dramatically reduce the number of bimodal image-text pairs and still obtain reasonable performance – either by using only 5% of the pairs (74.3→70.2) or through machine-generated object labels (74.3→68.0). Our code is available at <https://github.com/eladsegal/less-bimodal-sup>.

2. Overview

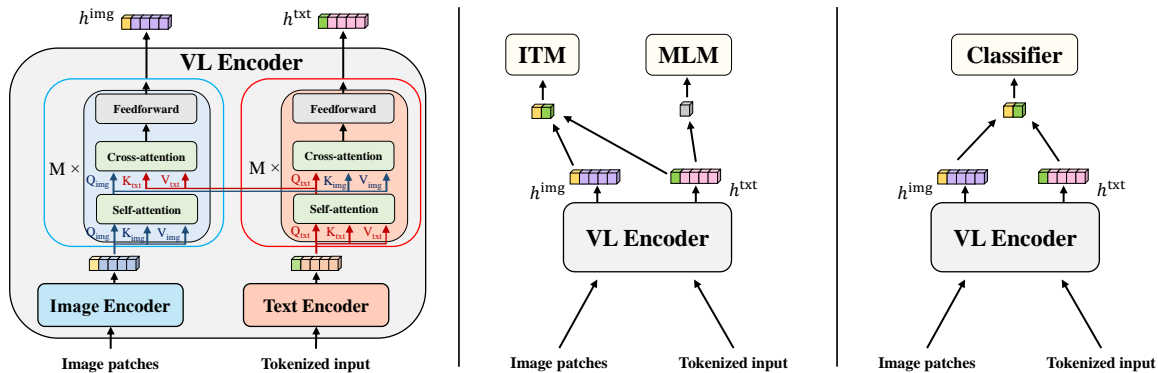


Figure 2: *Left*: Architecture overview: an image encoder and a text encoder followed by a few transformer fusion layers, capturing interaction between modalities through cross-attention. *Center*: We pretrain the VL encoder from bimodal supervision by taking contextualized representations of the image (h^{img}) and text (h^{txt}) and applying the image-text matching (ITM) and masked language modeling (MLM) loss functions. *Right*: We finetune the VL encoder on downstream classification tasks by concatenating the image and text representations and passing them through an MLP classifier.

We provide an overview of the experimental settings explored in this work. As our architecture-of-choice, we leverage one that has been shown to perform well across multiple tasks [Dou et al., 2022], namely, a Vision-Language (VL) encoder, where a unimodal image encoder creates image representations, a unimodal text encoder creates text representations, and these two representations are passed through a few transformer [Vaswani et al., 2017] layers that capture cross-modal interactions (Figure 2, Left).

We experiment with three initializations of the image and text encoders. First, we use random initialization as a baseline. Second, we initialize from pretrained unimodal models (RoBERTa and ViT, Liu et al., 2019, Dosovitskiy et al., 2021a), which can potentially reduce the amount of bimodal pretraining. Last, we pretrain the entire VL encoder with bimodal supervision (Figure 2, Center), and compare different data sources for pretraining, each requiring different amounts of human effort.

In each experiment we finetune and evaluate the VL encoder on downstream VL applications (Figure 2, Right), focusing on classification tasks (visual question answering and visual reasoning).

3. Data

We now describe the datasets used during bimodal pretraining and finetuning. For downstream applications, we put an emphasis on tasks that require reasoning over image(s) and text. Table 1 provides an example from each dataset, and Appendix A provides key statistics and details on the composition of the training sets.







ImageNet	Conceptual Captions	Conceptual Captions Image Labels
 <p><i>Class label:</i> printer</p>	 <p><i>Caption:</i> snail on a branch isolated on white background</p>	 <p><i>Computer-generated labels:</i> room, interior design, furniture, blue, living room, green, property, turquoise, home, floor, yellow, table, building, wall, house</p>
VQAv2	GQA	NLVR2
 <p><i>Question:</i> How many chairs can you count? <i>Answer:</i> 2</p>	 <p><i>Question:</i> What vegetable is to the left of the bag? <i>Answer:</i> cauliflower</p>	 <p><i>Sentence:</i> The sink in one of the images is set into a brown wood hanging counter. <i>Label:</i> false</p>

Table 1: Examples from all datasets used in this work.

3.1 Pretraining Datasets

ImageNet-1K [Russakovsky et al., 2015] is a human-annotated dataset that consists of over 1.2M images, divided into 1,000 classes that are mapped to meaningful concepts according to the WordNet hierarchy [Fellbaum, 1998]. Each concept is described by one or more language phrases, and accompanied by ~ 1000 images to illustrate it. We consider ImageNet-1K as a source of lightweight bimodal supervision, relatively cheap to obtain, as images are paired with text describing a single concept rather than a full-sentence.

Conceptual Captions (CC) [Sharma et al., 2018] is a programmatically-generated dataset of image-text pairs that consists of 3.3M examples. Prior work has demonstrated that CC is an effective resource for vision-language pretraining [Kim et al., 2021a, Li et al., 2021c, Lu et al., 2019b, Hendricks et al., 2021]. We use CC as a primary source of bimodal supervision, since: (a) it does not involve manual annotations, (b) it is small enough to be used by resource-constrained researchers, and (c) its images are from a different origin than the downstream tasks. Therefore, it provides a suitable test bed for estimating models’ ability to generalize to new images.

Conceptual Captions Image Labels (CCIL) [Ng et al., 2021] is a subset of 2M images from CC that contains machine-generated labels using the Google Cloud image labelling API. While labels are cheap since they are automatically-generated, the API was presumably trained on large amounts of bimodal data. Nevertheless, we examine pretraining on images paired with sets of labels to investigate whether this provides a sufficiently rich

source of bimodal supervision despite lacking natural language sentences. Past work indeed showed that VL pretraining benefits from masking object labels [Bitton et al., 2021].

3.2 Downstream Tasks

VQAv2 [Goyal et al., 2017] VQAv2 is a human-authored visual question answering (VQA) dataset that consists of 1.1M natural language questions with 10 short answers per question over 204K images from COCO [Lin et al., 2014]. It is standard to treat VQAv2 as a classification task, by only keeping questions with the most common answers (3,129 classes) [Anderson et al., 2018, Tan and Bansal, 2019, Zhai et al., 2022].

GQA [Hudson and Manning, 2019] (**balanced**) is a VQA dataset whose public version contains 1.1M questions over 83K images from Visual Genome [Krishna et al., 2017]. Unlike VQAv2, questions are created programmatically from scene graphs created by human annotators. Using scene graphs allows GQA to generate questions that test various reasoning skills such as comparisons, logical inference, spatial reasoning, etc.

NLVR2 [Suhr et al., 2019] is a benchmark for testing models’ ability to reason over text and images. The dataset contains 107K examples, where each example contains an English sentence and two web images (see Table 1). The goal is to determine whether the sentence is true or false in the context of the pair of images, a binary classification task.

4. Method

Our goal is to develop a classifier $f : \mathcal{X} \times \mathcal{I} \rightarrow \mathcal{C}$ that given an utterance \mathbf{x} and an image \mathbf{i} predicts a class $c \in \mathcal{C}$.

4.1 Architecture

We use a VL architecture, adapted from Dou et al. [2022]. The tokens of the utterance $\mathbf{x} = (x_0, x_1, \dots, x_n)$ are fed into a transformer *text encoder*, where x_0 is the special symbol CLS_{txt} . Similarly, the image is broken into patches $\mathbf{i} = (i_0, i_1, \dots, i_m)$, where i_0 is a special symbol CLS_{img} , which are fed into a transformer *image encoder*.

The image and text encoders compute contextualized representations of the image and text $(\hat{h}_0^{\text{txt}}, \dots, \hat{h}_n^{\text{txt}})$ and $(\hat{h}_0^{\text{img}}, \dots, \hat{h}_m^{\text{img}})$, which are then linearly projected with projection matrices $W_{\text{proj}}^{\text{txt}} \in \mathbb{R}^{d_{\text{txt}} \times d}$, $W_{\text{proj}}^{\text{img}} \in \mathbb{R}^{d_{\text{img}} \times d}$, where $d_{\text{txt}}, d_{\text{img}}$ are the hidden state dimensions of the text and image encoders respectively. The projected representations of each modality are then passed through transformer *fusion* layers, which include both a self-attention sublayer, and a cross-attention sublayer. Namely, each modality performs cross-attention on the other modality to fuse information from its representations, capturing interaction between the modalities. Overall, the VL encoder outputs the image-and-text contextualized representations $\mathbf{h}^{\text{img}} = (h_0^{\text{img}}, \dots, h_n^{\text{img}})$ and $\mathbf{h}^{\text{txt}} = (h_0^{\text{txt}}, \dots, h_m^{\text{txt}})$. An overview of our architecture is given in Figure 2, Left.

All model parameters are jointly trained by defining loss functions over classification heads, which we describe next. Since some model parameters are initialized from a pre-trained model, while other are randomly initialized, we use a higher learning rate for randomly initialized weights compared to pretrained weights, similar to Dou et al. [2022].

4.2 Pretraining Objectives

For pretraining, we use two objectives: masked language modeling (MLM) [Devlin et al., 2019, Tan and Bansal, 2019], and image-text matching (ITM) [Tan and Bansal, 2019], which are the most common objectives for VL pretraining and lead to state-of-the-art performance [Dou et al., 2022]. During training, we sum the ITM loss and the MLM loss for each training instance.

In MLM, given a masked token x_i the goal is to maximize the probability of the gold token given the representation h_i^{txt} , using cross-entropy loss. In ITM, given a image-text pair (\mathbf{x}, \mathbf{i}) , we concatenate the special CLS tokens h_0^{img} and h_0^{txt} , and use a sigmoid layer to predict if the image matches the text or not. We train with binary cross-entropy loss.

When pretraining on Conceptual Captions, we use the same masking scheme employed by Dou et al. [2022], that is, randomly masking 15% of the tokens. For ImageNet, we are given an image and a text label and mask all of its tokens. For CCIL, we are given an image and a list of text labels, concatenated with commas as separators, ordered by their machine-generated confidence scores. We then mask all tokens of a randomly-sampled label.

In ITM, in 50% of the examples, given a positive pair (\mathbf{x}, \mathbf{i}) , we substitute the true image with a random one and label it as a negative example.

4.3 Finetuning

Since the downstream applications in §3.2 can all be framed as classification tasks, we add a classification head to finetune the VL encoder. The classification head is a two-layer MLP, as in Kim et al. [2021b]. Specifically, we take as input the concatenation of all the image and text CLS representations, i.e., $[h_0^{\text{img}}, h_0^{\text{txt}}]$, and use the MLP to map them to $|\mathcal{C}|$ logits based on the number of task classes. The objective during training is to maximize the probability of the correct class(es), and we use standard cross-entropy loss. At inference time, we return the top-scoring class for all downstream tasks.

In NLVR2, where each example has two images, we consider each example as two image-text pairs, duplicating the text, and pass them separately through the VL encoder (dubbed ‘the pair setup’ in Chen et al. [2020]). We then pass four CLS representations (two for the images, two for the text) to the MLP to obtain the prediction.

5. Experiments

Experimental Setup We use ViT-Base [Dosovitskiy et al., 2021a] as the image encoder, pretrained and finetuned on ImageNet-21K at a resolution of 224x224 with a patch size of 16x16. We use RoBERTa-Base [Liu et al., 2019] as the text encoder. For the cross-modal transformer, we use only two layers to save computational resources, as previous work [Lu et al., 2019a, Hendricks et al., 2021], as well as our own preliminary findings, have shown that the effect of depth is small after finetuning.

We run pretraining (§4.2) for a maximum of 7,400 steps, and finetune each downstream task for 10 epochs. We specify batch sizes and learning rates for each case in Appendix B.1.

The evaluation score for VQAv2 is VQA score, and accuracy for GQA and NLVR2. Each result for VQAv2 and GQA is a 3-run average on the test-dev split, and for NLVR2 a 10-run average on the public test split.

Limitations Our work is performed within a limited compute budget. Therefore, we choose our largest pretraining dataset to be CC although there are datasets orders of magnitude larger. Compared to other work, we use images in a lower resolution, which has been shown to decrease performance [Dou et al., 2022]. Also, Dou et al. [2022] showed that better image encoders can significantly improve performance even before bimodal pretraining, but we did not experiment with different text and image encoders nor with larger models. Additionally, even though further pretraining in some setups results in small performance improvements, we decided the computational cost was unjustified. Bugliarello et al. [2021] showed pretraining variance exists when training on CC, but we were only able to pretrain once in each setup, due to the high computational costs. All of the above means that our work is self-contained, but cannot be directly compared in numbers to other works.

5.1 Main Results

	VQAv2	GQA	NLVR2
Random init.	52.3±0.1	43.1±0.1	random
Vision Random init.	54.2±0.0	44.3±0.2	random
Language Random init.	66.3±0.1	51.2±0.1	random
Unimodally-pretrained	69.5±0.1	53.6±0.1	random
Bimodally-pretrained with CCIL	69.6±0.3	52.9±0.5	68.0±0.7
Bimodally-pretrained with CC	70.7±0.0	56.1±0.3	74.3±0.3

Table 2: Main results for all downstream tasks.

Table 2 shows the results of finetuning on all downstream tasks for different initializations of the image and text encoders.

In addition to finetuning a model that is initialized with ViT and RoBERTa (‘Unimodally-pretrained’), and in order to verify the importance of unimodal pretraining, we finetune our model when the image encoder, text encoder, or both encoders are randomly initialized. We find that pretraining the vision model is essential for good performance, and observe a smaller drop in performance when the text encoder is randomly initialized, similar to Zhai et al. [2022].

Comparing the unimodally-pretrained model to one that was further pretrained on CC (‘Bimodally-pretrained with CC’), we see that for VQAv2 the gap is only 1 point, and for GQA it is just 2.5 points. However, on the more challenging NLVR2, which requires complex reasoning operations, bimodal pretraining is essential, and the model achieves random performance without it. Nevertheless, training with a weaker form of supervision, namely, a list of machine-generated object labels from CCIL is sufficient for non-random and reasonably high performance on NLVR2 (but has no effect on VQAv2 and GQA).

5.2 Effect of CC Size on Pretraining

Table 2 showed that bimodal pretraining is essential for obtaining non-random results on NLVR2. A natural question is whether this can be obtained with fewer pretraining examples. To this end, we pretrain on different fractions of CC and present the results after finetuning in Table 8 (in the Appendix) and Figure 3.

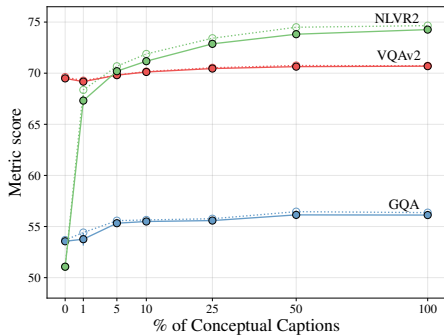


Figure 3: Effect of the fraction of examples from CC on downstream task performance. Solid/dashed line – average/maximum score over seeds.

Max # of labels	Unique Labels	NLVR2
1	5.3K	52.6±1.4 (55.3)
3	8.0K	67.8±0.5 (68.6)
15	14.3K	68.0±0.7 (68.9)

Table 3: Performance on NLVR2 when restricting the number of labels per image during pretraining on CCIL (max. value is in the parentheses).

Surprisingly, even when using only $\sim 1\%$ of CC (30K examples), performance on NLVR2 is far from random – 67.3. When using 5% of the data, performance is only moderately lower than when using CC in its entirety – 70.2 vs. 74.3. When using 25% of the data for pretraining, performance on all three datasets is less than two points lower than when using 100%, showing that indeed the amount of bimodal supervision can be considerably reduced with only a small hit in performance.

The aforementioned results were obtained by finetuning on all of the downstream data per task. However, an interesting variant to consider is a low-resource setting where we have only *some* of the downstream data – what is the importance of bimodal pretraining then? Table 9 (in the Appendix) and Figure 4 show for VQAv2 and GQA that when less data is used for finetuning, the benefit of pretraining with 5% or more of CC is greater than the benefit observed when 100% of the downstream data is used for finetuning. For NLVR2, we see that pretraining is still very helpful even with 100% of the downstream data. The reason for the difference might be that VQAv2 and GQA are much larger than NLVR2.

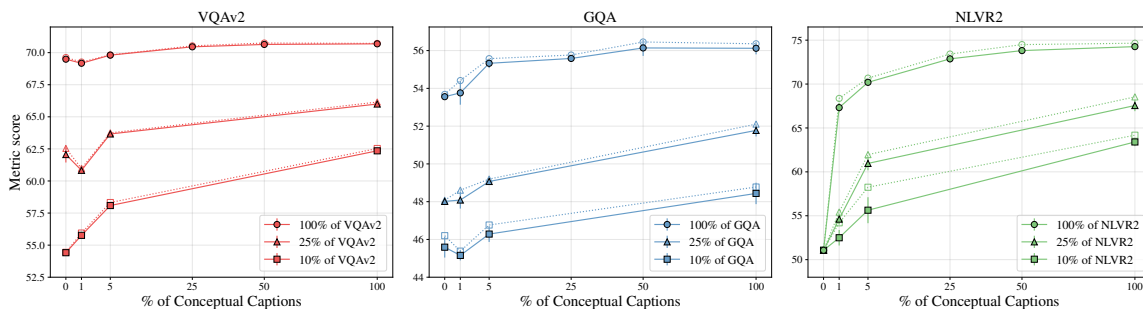


Figure 4: Effect of the fraction of examples from CC on downstream task performance when finetuning on less downstream data. Solid/dashed line – average/max. over seeds.

5.3 Pretraining with ImageNet Labels

We have seen in §5.1 that image-caption pairs are useful for pretraining VL models. Here, we investigate if a weaker source of language supervision, namely image labels only, suffices for aligning text and vision representations. Specifically, we pair each ImageNet image with its label, treating it as a caption, and pretrain with MLM and ITM as described in §4.2.

We observe *no difference* in results compared to unimodally-pretrained models (Table 10 in the Appendix) – performance remains random for NLVR2, and similar for VQAv2 and GQA. This suggests that ImageNet labels do not provide adequate signal for VL pretraining.

5.4 Pretraining with CCIL

One hypothesis for the lack of improvement when pretraining with ImageNet is that a single label per image is too limiting, since images typically contain many objects. To test this, we pretrain with CCIL, where each image is paired with machine-generated labels, providing a richer image representation. We pretrain with MLM and ITM as described in §4.2.

While pretraining on CCIL does not improve performance on VQAv2 and GQA, it leads to dramatic improvement on NLVR2, reaching an average accuracy of 68.0 ± 0.7 and a maximum accuracy of 68.9. This shows that providing a set of object labels lets the model better align image and text representations. Table 3 further validates this by showing results when restricting the maximal number of labels per image. We observe that having multiple labels per image is crucial, as performance is roughly random when using a single label. Using 3 labels is already sufficient for bootstrapping the model, and performance is barely lower compared to using all 15 labels.

5.5 Transfer Learning

Finally, we test whether a model finetuned on a source downstream task (VQAv2 and GQA) can improve performance on a target task, i.e., in a transfer learning setup, where we vary the amount of annotated data in the source task.

Table 11 (in the Appendix) and Figure 5 (left) show results when VQAv2 is the source task and GQA and NLVR2 are the target tasks. VQAv2 appears to be an effective source of bimodal supervision for both tasks – when using all of VQAv2, performance on GQA is even slightly higher compared to pretraining on CC data, and 3 points lower on NLVR2 (74.3→71.1). Nevertheless, the amount of data in the source task is important, and performance on NLVR2 is much lower when using 5%-25% of the data.

Table 12 (in the Appendix) and Figure 5 (right) show results when GQA is the source task and VQAv2 and NLVR2 are the target tasks. We observe that VQAv2 is a better source task compared to GQA – GQA does not improve performance on VQAv2, and its effect on NLVR2 is much more moderate. A possible explanation is that VQAv2 has natural language questions, while questions in GQA are automatically generated. Another potential factor is the fact that VQAv2 typically require less reasoning steps compared to GQA.

Overall, in both cases we find transfer learning on downstream tasks is useful, and can even perform closely to bimodally-pretrained models.

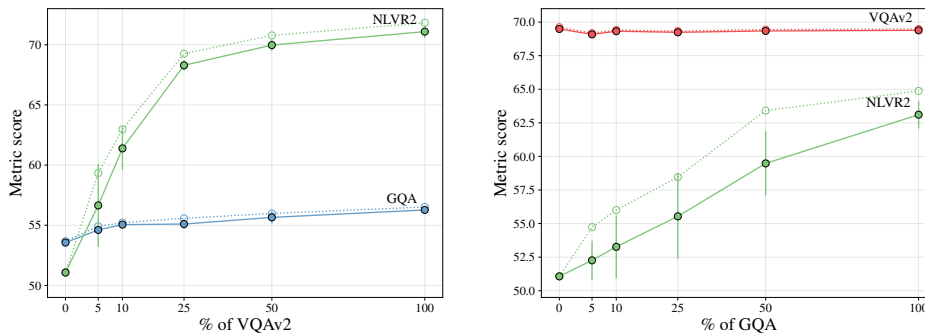


Figure 5: Effect of the fraction of examples from VQAv2 (left) and GQA (right) on downstream task performance. Solid/dashed line – average/maximum score over seeds.

6. Analysis

To better understand what data properties are important for pretraining, we train on small subsets of CC (1% of the data) and VQAv2 (5% of the data), with particular characteristics:

- *Min/max length*: We create subsets that minimize/maximize the average input length.
- *Min/max vocabulary size* - We create subsets that minimize/maximize the size of the vocabulary. To do so we use a greedy procedure, where (a) we initialize an empty set of examples, and at each step (b) randomly sample a candidate set of 10K examples, and (c) choose the example that minimizes/maximizes the current vocabulary size.

	1% CC			5% VQAv2		
Subset	Length	Vocab.	NLVR2	Length	Vocab.	NLVR2
Min length	5.0	8.0K	67.1±0.3 (67.4)	4.45	3.6K	57.4±3.1 (60.8)
Max length	30.3	19.1K	64.1±1.4 (65.7)	12.7	7.6K	53.9±2.0 (57.0)
Min vocab.	6.5	0.3K	64.8±1.2 (65.9)	5.8	0.2K	57.4±3.8 (62.2)
Max vocab.	14.0	44.4K	67.3±0.3 (67.7)	7.7	16.4K	55.1±3.1 (58.2)
Random	10.3	13.3K	67.3±0.5 (68.4)	7.3	5.8K	56.6±3.5 (59.3)

Table 4: Analyzing the effect of pretraining on CC/VQAv2 subsets with particular properties. After training on each subset, we finetune on NLVR2.

Results are in Table 4. No subset is noticeably better than a random subset. For CC, results are similar. For VQAv2, while performance when minimizing length and vocabulary is better on average, the differences seem negligible, given the high standard deviation.

Effect of length on pretraining Table 4 shows that pretraining on long inputs substantially hurts performance – results are reduced by at least 3 points for both CC and VQAv2. This is surprising as one might hypothesize that longer inputs should be better since they contain more information. A possible explanation is that simple examples are necessary to bootstrap the pretraining procedure and align the text and image representations.

Effect of vocabulary size on pretraining Pretraining on a subset with higher lexical diversity should expose the model to more concepts, both in images and texts, and therefore improve its performance. While for CC this is indeed the case, for VQAv2 results for the

max vocabulary size setup with 16.4K words are lower than the min vocabulary size setup with only 0.2K words. A possible explanation is the amount of yes/no questions in the min/max vocabulary size subsets which is 80.7% and 44.5%, respectively – Since NLVR2 is a yes/no task, training on more yes/no questions might be closer to its distribution.

7. Related Work

Dou et al. [2022] investigated unimodally-pretrained models, finetuning different image and text encoders on multiple VL tasks, recognizing it as efficient and performant. However, they did not consider the effects of the types and amount of bimodal supervision. Past work investigated bimodal supervision on VL models, but for models that use *frozen* features from an object detection model [Singh et al., 2020, Hendricks et al., 2021], which (a) cannot be adapted to unseen concepts [Zhang et al., 2021], (b) require heavily annotated object-level data for the training of the object detection model [Krishna et al., 2017, Anderson et al., 2018], and (c) result in an architectural inductive bias towards objects (which is very beneficial for VQA tasks). Singh et al. [2020] compared performance between multiple pretraining datasets, varying their sizes. Unlike us, for all tasks, the effect of different usage of bimodal supervision was small, compared to our NLVR2 experiments. Hendricks et al. [2021] assessed the contribution of pretraining datasets from a set of standard VL pretraining datasets, but focused on zero-shot image retrieval tasks.

Li et al. [2021d] and Zhou et al. [2022] also share the motivation of reducing bimodal pretraining for VL models. With some similarity to our CCIL experiments in §5.4, they avoid pretraining on collected parallel image-text data altogether by utilizing predictions of regions and tags from an object detection model to create VL-specialized training objectives.

Opposite to our setup, a current trend is to pretrain models on vast amounts of bimodal data [Radford et al., 2021, Zhai et al., 2022, Alayrac et al., 2022], and perform zero/few-shot evaluation. While remarkable results were achieved, performance is lower than finetuned models pretrained on less bimodal data, which is relatively cheap to obtain.

8. Conclusion

A current obstacle on the road to multimodal models is reliance on bimodal supervision. In this work, we go in an opposite direction from current trends, and instead of using increasing amounts of bimodal data, we examine whether one can use *less* of it. We find that indeed this is the case, where for simple tasks just finetuning unimodally-pretrained models leads to performance that is similar to bimodally-pretrained models, at a much lower cost. For complex tasks, while bimodal pretraining is still necessary, its amount (100%→5%) and source quality (CC→CCIL) can be significantly reduced with only a moderate degradation in performance. We also find that models finetuned on one downstream task are useful in a transfer learning setup, achieving results close to bimodally-pretrained models.

Acknowledgements

This research was partially supported by The Yandex Initiative for Machine Learning, and the European Research Council (ERC) under the European Union Horizons 2020 research and innovation programme (grant ERC DELPHI 802800).

References

- Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel Lenc, Arthur Mensch, Katie Millican, Malcolm Reynolds, Roman Ring, Eliza Rutherford, Serkan Cabi, Tengda Han, Zhitao Gong, Sina Samangooei, Marianne Monteiro, Jacob Menick, Sebastian Borgeaud, Andy Brock, Aida Nematzadeh, Sahand Sharifzadeh, Mikolaj Binkowski, Ricardo Barreira, Oriol Vinyals, Andrew Zisserman, and Karen Simonyan. Flamingo: a visual language model for few-shot learning. *arXiv*, abs/2204.14198, 2022.
- 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 *2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018*, pages 6077–6086. IEEE Computer Society, 2018. doi: 10.1109/CVPR.2018.00636. URL http://openaccess.thecvf.com/content_cvpr_2018/html/Anderson_Bottom-Up_and_Top-Down_CVPR_2018_paper.html.
- Alexei Baevski, Wei-Ning Hsu, Qiantong Xu, Arun Babu, Jiatao Gu, and Michael Auli. Data2vec: A general framework for self-supervised learning in speech, vision and language. *arXiv*, abs/2202.03555, 2022.
- Yonatan Bitton, Michael Elhadad, Gabriel Stanovsky, and Roy Schwartz. Data efficient masked language modeling for vision and language. In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 3013–3028. Association for Computational Linguistics, 2021. doi: 10.18653/v1/2021.findings-emnlp.259. URL <https://aclanthology.org/2021.findings-emnlp.259>.
- Rishi Bommasani, Drew A. Hudson, Ehsan Adeli, Russ Altman, Simran Arora, Sydney von Arx, Michael S. Bernstein, Jeannette Bohg, Antoine Bosselut, Emma Brunskill, Erik Brynjolfsson, S. Buch, Dallas Card, Rodrigo Castellon, Niladri S. Chatterji, Annie S. Chen, Kathleen Creel, Jared Davis, Dora Demszky, Chris Donahue, Moussa Doumbouya, Esin Durmus, Stefano Ermon, John Etchemendy, Kawin Ethayarajh, Li Fei-Fei, Chelsea Finn, Trevor Gale, Lauren E. Gillespie, Karan Goel, Noah D. Goodman, Shelby Grossman, Neel Guha, Tatsunori Hashimoto, Peter Henderson, John Hewitt, Daniel E. Ho, Jenny Hong, Kyle Hsu, Jing Huang, Thomas F. Icard, Saahil Jain, Dan Jurafsky, Pratyusha Kalluri, Siddharth Karamcheti, Geoff Keeling, Fereshte Khani, O. Khattab, Pang Wei Koh, Mark S. Krass, Ranjay Krishna, Rohith Kuditipudi, Ananya Kumar, Faisal Ladhak, Mina Lee, Tony Lee, Jure Leskovec, Isabelle Levent, Xiang Lisa Li, Xuechen Li, Tengyu Ma, Ali Malik, Christopher D. Manning, Suvir P. Mirchandani, Eric Mitchell, Zanele Munyikwa, Suraj Nair, Avanika Narayan, Deepak Narayanan, Benjamin Newman, Allen Nie, Juan Carlos Niebles, Hamed Nilforoshan, J. F. Nyarko, Giray Ogut, Laurel Orr, Isabel Papadimitriou, Joon Sung Park, Chris Piech, Eva Portelance, Christopher Potts, Aditi Raghunathan, Robert Reich, Hongyu Ren, Frieda Rong, Yusuf H. Roohani, Camilo Ruiz, Jack Ryan, Christopher R’e, Dorsa Sadigh, Shiori Sagawa, Keshav Santhanam, Andy Shih, Krishna Parasuram Srinivasan, Alex Tamkin, Rohan Taori, Armin W. Thomas, Florian Tramèr, Rose E. Wang, William Wang, Bohan Wu, Jiajun

- Wu, Yuhuai Wu, Sang Michael Xie, Michihiro Yasunaga, Jiaxuan You, Matei A. Zaharia, Michael Zhang, Tianyi Zhang, Xikun Zhang, Yuhui Zhang, Lucia Zheng, Kaitlyn Zhou, and Percy Liang. On the opportunities and risks of foundation models. *arXiv*, abs/2108.07258, 2021.
- Emanuele Bugliarelli, 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. doi: 10.1162/tacl_a_00408. URL <https://aclanthology.org/2021.tacl-1.58>.
- 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*, 2020.
- Ekin Dogus Cubuk, Barret Zoph, Jon Shlens, and Quoc Le. Randaugment: Practical automated data augmentation with a reduced search space. In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020. URL <https://proceedings.neurips.cc/paper/2020/hash/d85b63ef0ccb114d0a3bb7b7d808028f-Abstract.html>.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186. Association for Computational Linguistics, 2019. doi: 10.18653/v1/N19-1423. URL <https://aclanthology.org/N19-1423>.
- 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. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net, 2021a. URL <https://openreview.net/forum?id=YicbFdNTTy>.
- 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. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net, 2021b. URL <https://openreview.net/forum?id=YicbFdNTTy>.
- Zi-Yi Dou, Yichong Xu, Zhe Gan, Jianfeng Wang, Shuohang Wang, Lijuan Wang, Chenguang Zhu, Nanyun Peng, Zicheng Liu, and Michael Zeng. An empirical study of training end-to-end vision-and-language transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022.
- Christiane Fellbaum. *WordNet: An Electronic Lexical Database*. Bradford Books, 1998.

- 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 *2017 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017, Honolulu, HI, USA, July 21-26, 2017*, pages 6325–6334. IEEE Computer Society, 2017. doi: 10.1109/CVPR.2017.670. URL <https://doi.org/10.1109/CVPR.2017.670>.
- Lisa Anne Hendricks, John Mellor, Rosalia Schneider, Jean-Baptiste Alayrac, and Aida Nematzadeh. Decoupling the role of data, attention, and losses in multimodal transformers. *Transactions of the Association for Computational Linguistics*, 9:570–585, 2021. doi: 10.1162/tacl.a.00385. URL <https://aclanthology.org/2021.tacl-1.35>.
- Jonathan Herzig, Pawel Krzysztof Nowak, Thomas Müller, Francesco Piccinno, and Julian Eisenschlos. TaPas: Weakly supervised table parsing via pre-training. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4320–4333. Association for Computational Linguistics, 2020. doi: 10.18653/v1/2020.acl-main.398. URL <https://aclanthology.org/2020.acl-main.398>.
- Drew A. Hudson and Christopher D. Manning. GQA: A new dataset for real-world visual reasoning and compositional question answering. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019, Long Beach, CA, USA, June 16-20, 2019*, pages 6700–6709. Computer Vision Foundation / IEEE, 2019. doi: 10.1109/CVPR.2019.00686. URL http://openaccess.thecvf.com/content_CVPR_2019/html/Hudson_GQA_A_New_Dataset_for_Real-World_Visual_Reasoning_and_Compositional_CVPR_2019_paper.html.
- Chao Jia, Yinfei Yang, Ye Xia, Yi-Ting Chen, Zarana Parekh, Hieu Pham, Quoc V. Le, Yun-Hsuan Sung, Zhen Li, and Tom Duerig. Scaling up visual and vision-language representation learning with noisy text supervision. In *ICML*, 2021.
- Wonjae Kim, Bokyung Son, and Ildoo Kim. Vilt: Vision-and-language transformer without convolution or region supervision. In *ICML*, 2021a.
- Wonjae Kim, Bokyung Son, and Ildoo Kim. Vilt: Vision-and-language transformer without convolution or region supervision. In *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pages 5583–5594. PMLR, 2021b. URL <http://proceedings.mlr.press/v139/kim21k.html>.
- Ranjay Krishna, Yuke Zhu, Oliver Groth, Justin Johnson, Kenji Hata, Joshua Kravitz, Stephanie Chen, Yannis Kalantidis, Li-Jia Li, David A. Shamma, Michael S. Bernstein, and Li Fei-Fei. Visual genome: Connecting language and vision using crowdsourced dense image annotations. *International Journal of Computer Vision*, 123(1):32–73, 2017. ISSN 0920-5691. doi: 10.1007/s11263-016-0981-7. URL <https://doi.org/10.1007/s11263-016-0981-7>.
- Hang Li, Wenbiao Ding, Yu Kang, Tianqiao Liu, Zhongqin Wu, and Zitao Liu. CTAL: Pre-training cross-modal transformer for audio-and-language representations. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 3966–3977. Association for Computational Linguistics, 2021a. doi: 10.18653/v1/2021.emnlp-main.323. URL <https://aclanthology.org/2021.emnlp-main.323>.

- Junnan Li, Ramprasaath Selvaraju, Akhilesh Gotmare, Shafiq Joty, Caiming Xiong, and Steven Chu Hong Hoi. Align before fuse: Vision and language representation learning with momentum distillation. *Advances in Neural Information Processing Systems*, 34, 2021b.
- Junnan Li, Ramprasaath R. Selvaraju, Akhilesh Deepak Gotmare, Shafiq Joty, Caiming Xiong, and Steven Hoi. Align before fuse: Vision and language representation learning with momentum distillation. In *NeurIPS*, 2021c.
- Junnan Li, Dongxu Li, Caiming Xiong, and Steven Hoi. Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation. In *ICML*, 2022.
- 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*, abs/1908.03557, 2019.
- Liunian Harold Li, Haoxuan You, Zhecan Wang, Alireza Zareian, Shih-Fu Chang, and Kai-Wei Chang. Unsupervised vision-and-language pre-training without parallel images and captions. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 5339–5350. Association for Computational Linguistics, 2021d. doi: 10.18653/v1/2021-naacl-main.420. URL <https://aclanthology.org/2021.naacl-main.420>.
- Tsung-Yi Lin, Michael Maire, Serge J. Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C. Lawrence Zitnick. Microsoft coco: Common objects in context. In *ECCV*, 2014.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. Roberta: A robustly optimized bert pretraining approach. *arXiv*, abs/1907.11692, 2019.
- Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 10012–10022, 2021.
- Jiasen Lu, Dhruv Batra, Devi Parikh, and Stefan Lee. Vilbert: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks. In *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pages 13–23, 2019a. URL <https://proceedings.neurips.cc/paper/2019/hash/c74d97b01eae257e44aa9d5bade97baf-Abstract.html>.
- Jiasen Lu, Dhruv Batra, Devi Parikh, and Stefan Lee. Vilbert: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks. In *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pages 13–23, 2019b. URL <https://proceedings.neurips.cc/paper/2019/hash/c74d97b01eae257e44aa9d5bade97baf-Abstract.html>.

- Edwin G. Ng, Bo Pang, Piyush Sharma, and Radu Soricut. Understanding guided image captioning performance across domains. In *Proceedings of the 25th Conference on Computational Natural Language Learning*, pages 183–193. Association for Computational Linguistics, 2021. doi: 10.18653/v1/2021.conll-1.14. URL <https://aclanthology.org/2021.conll-1.14>.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International Conference on Machine Learning*, pages 8748–8763. PMLR, 2021.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of Machine Learning Research*, 21(140): 1–67, 2020. URL <http://jmlr.org/papers/v21/20-074.html>.
- Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C. Berg, and Li Fei-Fei. ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision (IJCV)*, 115(3):211–252, 2015. doi: 10.1007/s11263-015-0816-y.
- Steffen Schneider, Alexei Baevski, Ronan Collobert, and Michael Auli. wav2vec: Unsupervised pre-training for speech recognition. *arXiv*, abs/1904.05862, 2019.
- Piyush Sharma, Nan Ding, Sebastian Goodman, and Radu Soricut. Conceptual captions: A cleaned, hypernymed, image alt-text dataset for automatic image captioning. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2556–2565. Association for Computational Linguistics, 2018. doi: 10.18653/v1/P18-1238. URL <https://aclanthology.org/P18-1238>.
- Amanpreet Singh, Vedanuj Goswami, and Devi Parikh. Are we pretraining it right? digging deeper into visio-linguistic pretraining. *arXiv*, abs/2004.08744, 2020.
- Emma Strubell, Ananya Ganesh, and Andrew McCallum. Energy and policy considerations for deep learning in NLP. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3645–3650. Association for Computational Linguistics, 2019. doi: 10.18653/v1/P19-1355. URL <https://aclanthology.org/P19-1355>.
- Alane Suhr, Stephanie Zhou, Ally Zhang, Iris Zhang, Huajun Bai, and Yoav Artzi. A corpus for reasoning about natural language grounded in photographs. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 6418–6428. Association for Computational Linguistics, 2019. doi: 10.18653/v1/P19-1644. URL <https://aclanthology.org/P19-1644>.
- Alon Talmor, Ori Yoran, Amnon Catav, Dan Lahav, Yizhong Wang, Akari Asai, Gabriel Ilharco, Hannaneh Hajishirzi, and Jonathan Berant. Multimodalqa: Complex question answering over text, tables and images. In *International Conference on Learning Representations (ICLR)*, 2021.

- Hao Tan and Mohit Bansal. LXMERT: Learning cross-modality encoder representations from transformers. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 5100–5111. Association for Computational Linguistics, 2019. doi: 10.18653/v1/D19-1514. URL <https://aclanthology.org/D19-1514>.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pages 5998–6008, 2017. URL <https://proceedings.neurips.cc/paper/2017/hash/3f5ee243547dee91fbd053c1c4a845aa-Abstract.html>.
- Xiaohua Zhai, Xiao Wang, Basil Mustafa, Andreas Steiner, Daniel Keysers, Alexander Kolesnikov, and Lucas Beyer. Lit: Zero-shot transfer with locked-image text tuning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 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 (CVPR)*, pages 5579–5588, 2021.
- Mingyang Zhou, Licheng Yu, Amanpreet Singh, Mengjiao Wang, Zhou Yu, and Ning Zhang. Unsupervised vision-and-language pre-training via retrieval-based multi-granular alignment. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 16485–16494, 2022.

Appendix for “Training Vision-Language Models with Less Bimodal Supervision”

Appendix A. Training Data

Since for some of datasets the official training splits aren’t used as-is, we provide the exact details of the training data composition for each dataset and also key statistics for all of the datasets in Table 5.

Dataset	Training instances	Unique texts	Training images
ImageNet	656K	738	656K
Conceptual Captions (CC)	2.84M	2M	2.84M
Conceptual Captions Image Labels (CCIL)	1.84M	1.79M	1.84M
VQAv2	620K	210K	118K
GQA	943K	538K	72K
NLVR2	86K	23K	103K

Table 5: Key statistics for the training datasets.

ImageNet-1K [Russakovsky et al., 2015] Since ImageNet classes are often too fine-grained, we manually collapse fine-grained classes into an ancestor WordNet class,¹ e.g., dog breeds are collapsed to “dog”. Then, we create a balanced training set according to the updated classes of the images.

Following is a list of the classes we collapse sub-classes into:

dog, fox, wild dog, wolf, coyote, domestic cat, bear, monkey, snake, lizard, turtle, frog, salamander, lobster, crab, beetle, butterfly, spider, rabbit, bird, fungus

Conceptual Captions (CC) [Sharma et al., 2018] Out of the 3.3M examples in the official training set, we were able to download 2.84M examples from the provided image URLs.

Conceptual Captions Image Labels (CCIL) [Ng et al., 2021] Out of the 2M examples in the official training set, we were able to download 1.84M examples from the provided image URLs.

VQAv2 [Goyal et al., 2017] We create our training set similar to previous works on VQAv2 [Tan and Bansal, 2019, Dou et al., 2022], and use the same validation set as Tan and Bansal [2019], which was constructed from the official validation set based on 5,000 randomly chosen images.

To create the training set, we first create an answer set that contains only majority answers that occurred at least 9 times on the official training and validation sets together. Then, out of the official training and validation sets, we filter out all of the examples that doesn’t have any answer in the created answer set. Finally, out of the remaining examples, we discard every example that appears in our validation set.

GQA [Hudson and Manning, 2019] We use the official training set.

1. <https://observablehq.com/@mbostock/imagenet-hierarchy>.

NLVR2 [Suhr et al., 2019] We use the official training set.

Appendix B. Experimental Setup

B.1 Additional Implementation Details

Image Preprocessing Both in pretraining and finetuning, we apply center crop on the image and resize it to 224x224. When training, we additionally use RandAugment [Cubuk et al., 2020] as in Kim et al. [2021a] with the exclusion of color-changing strategies (Invert, Posterize, Solarize, SolarizeAdd) and the coutout strategy.

Model Architecture We use the model from Dou et al. [2022], but we simplify it with the removal of two of its components since we didn’t observe a performance difference: the single-layer feedforward network before the feeding of the [CLS] representations to a task-specific head (e.g. ITM, MLM, classifier), and the image token type embeddings.

Pretraining We run pretraining for 7,400 steps, except when training on 1%, 5% and 10% of CC, as more training results in an increase of the validation loss. We train for 1850 steps on 1% and %5 of CC, and 3700 steps for 10% of CC.

The batch size is 3,840 and learning rates of $1e^{-4}$ and $5e^{-4}$ are used for the pretrained and randomly initialized weights respectively. The learning rate is warmed up from zero during the first 10% steps, and then linearly decays back to zero throughout the remaining steps.

We use 8 NVIDIA V100 GPUs, and training takes about 16 hours for 100% of CC.

Finetuning For finetuning, we use a batch size of 96 for VQAv2 and GQA, and 48 for NLVR2. We specify the learning rates for finetuning before and after bimodal pretraining in Tables 6 and 7 respectively. The learning rate is warmed up from zero during the first 10% steps, and then linearly decays back to zero throughout the remaining steps.

We use a single NVIDIA RTX 3090 GPU, and training takes 10, 15 and 4 hours for VQAv2, GQA and NLVR2 respectively.

Weights	VQAv2	GQA	NLVR2
Image encoder, Text encoder	$2e - 5$	$1e - 5$	$1e - 5$
Cross-modal transformer	$2e - 4$	$1e - 4$	$1e - 4$
Classifier head	$2e - 4$	$1e - 4$	$1e - 4$

Table 6: Learning rates per weights for finetuning *before* bimodal pretraining for each downstream task.

Weights	VQAv2	GQA	NLVR2
Image encoder, Text encoder	$2e-5$	$1e-5$	$1e-5$
Cross-modal transformer	$1e-4$	$1e-4$	$5e-5$
Classifier head	$1e-3$	$1e-4$	$5e-4$

Table 7: Learning rates per weights for finetuning *after* bimodal pretraining for each downstream task.

Appendix C. Results

CC Data	VQAv2	GQA	NLVR2
0%	69.5±0.1	53.6±0.1	random
1%	69.2±0.1	53.8±0.6	67.3±0.5
5%	69.8±0.0	55.3±0.3	70.2±0.3
10%	70.1±0.1	55.5±0.2	71.2±0.4
25%	70.5±0.1	55.6±0.2	72.9±0.4
50%	70.6±0.1	56.1±0.4	73.8±0.4
100%	70.7±0.0	56.1±0.3	74.3±0.3

Table 8: Effect of the fraction of examples from CC on downstream task performance. Visualized with Fig. 3.

CC Data	VQAv2			GQA			NLVR2		
	10%	25%	100%	10%	25%	100%	10%	25%	100%
0%	54.4±0.0	62.1±0.6	69.5±0.1	45.6±0.5	48.0±0.1	53.6±0.1	random	random	random
1%	55.8±0.2	60.8±0.1	69.2±0.1	45.2±0.3	48.1±0.5	53.8±0.6	52.5±0.9	54.6±0.6	67.3±0.5
5%	58.1±0.2	63.7±0.1	69.8±0.0	46.3±0.4	49.1±0.1	55.3±0.3	55.6±1.5	61.0±0.8	70.2±0.3
100%	62.4±0.2	66.0±0.1	70.7±0.0	48.4±0.6	51.8±0.3	56.1±0.3	63.4±0.5	67.6±0.6	74.3±0.3

Table 9: Effect of the fraction of examples from CC on downstream task performance when finetuning on less downstream data. Visualized with Fig. 4.

	VQAv2	GQA	NLVR2
Unimodally-pretrained	69.5±0.1	53.6±0.1	random
Bimodally-pretrained with ImageNet	69.3±0.0	53.5±0.2	random

Table 10: Performance on all downstream tasks, with and without ImageNet pretraining. No difference is observed.

VQAv2 Data	GQA	NLVR2
0%	53.6±0.1	random
5%	54.6±0.3	56.6±3.5
10%	55.1±0.2	61.4±1.8
25%	55.1±0.4	68.3±0.4
50%	55.7±0.5	70.0±0.5
100%	56.3±0.2	71.1±0.5
Bimodally-pretrained with CC	56.1±0.3	74.3±0.3

Table 11: Effect of the fraction of examples from VQAv2 on downstream task performance. Visualized with Fig. 5 (left).

GQA Data	VQAv2	NLVR2
0%	69.5±0.1	random
5%	69.1±0.1	52.3±1.5
10%	69.3±0.1	53.3±2.3
25%	69.2±0.1	55.5±3.2
50%	69.3±0.1	59.5±2.4
100%	69.4±0.1	63.1±1.0
Bimodally-pretrained with CC	70.7±0.0	74.3±0.3

Table 12: Effect of the fraction of examples from GQA on downstream task performance. Visualized with Fig. 5 (right).