

Separating Skills and Concepts for Novel Visual Question Answering

Spencer Whitehead^{1*}, Hui Wu², Heng Ji¹, Rogerio Feris², Kate Saenko^{2,3}
¹UIUC ²MIT-IBM Watson AI Lab, IBM Research ³Boston University
{srw5, hengji}@illinois.edu, {wuhu, rsferis}@us.ibm.com, saenko@bu.edu

Abstract

Generalization to out-of-distribution data has been a problem for Visual Question Answering (VQA) models. To measure generalization to novel questions, we propose to separate them into “skills” and “concepts”. “Skills” are visual tasks, such as counting or attribute recognition, and are applied to “concepts” mentioned in the question, such as objects and people. VQA methods should be able to compose skills and concepts in novel ways, regardless of whether the specific composition has been seen in training, yet we demonstrate that existing models have much to improve upon towards handling new compositions. We present a novel method for learning to compose skills and concepts that separates these two factors implicitly within a model by learning grounded concept representations and disentangling the encoding of skills from that of concepts. We enforce these properties with a novel contrastive learning procedure that does not rely on external annotations and can be learned from unlabeled image-question pairs. Experiments demonstrate the effectiveness of our approach for improving compositional and grounding performance.¹

1. Introduction

When humans answer questions, such as in Visual Question Answering (VQA), we first interpret the question, dissecting its content into parts (like concepts, relations, actions, question types), and then we select and execute the *skill* (or plan/program) necessary to produce an answer based on this information and the relevant knowledge base (e.g., the image) [26, 36, 59, 62]. The skills needed to produce an answer are general and can be applied to (composed with) many types of question-specific content. For example, if one can answer questions about “colors” for a variety of objects as well as recognize and answer questions about “cars”, then questions like “What color is the car?” should be straightforward to answer even if this specific composition has yet to be seen (Fig. 1). This ability of seam-

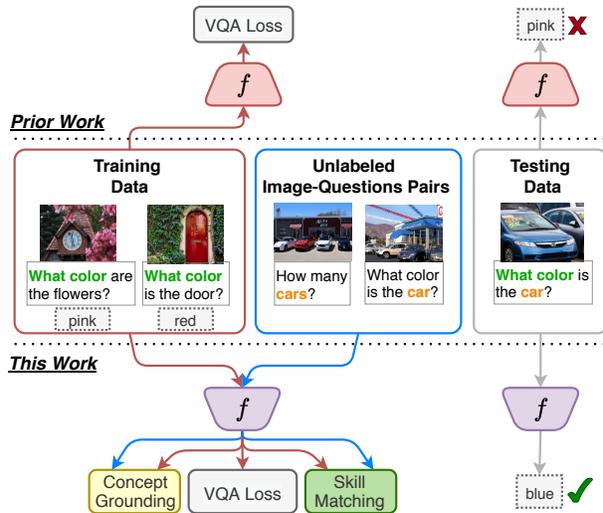


Figure 1. We propose a new view of compositionality in VQA that explores the ability to answer questions about unseen compositions of skills (e.g., color) and concepts (e.g., car). We present a method that learns to separate skills and concepts that can utilize both labeled and unlabeled image-question pairs in order to generalize to novel questions with new skill-concept compositions and new concepts.

lessly adapting and composing conceptual representations with skills is crucial to demonstrating true understanding of VQA and learning to generalize from less labeled data.

Compositionality is recognized as one of the essential properties of human cognition [35], but more research is still needed on incorporating compositionality into models and developing data-efficient, generalizable systems. While much progress has been made to achieve better performance on standard VQA test benchmarks [4, 17, 33, 61], most state-of-the-art models are still designed without any notion of built-in compositionality and tend to entangle skills and concepts in their learned representations. Some previous work has studied the lack of generalization ability of VQA models, and evaluated models using test splits with different answer distributions from the training data. However, this measurement only indirectly addresses the central issue

*Work was partly done as an intern at the MIT-IBM Watson AI Lab.

¹Code: <https://github.com/SpencerWhitehead/novelvqa>

(lack of compositionality), which manifests itself as poor generalization and over-reliance on language priors [2, 46].

To address these issues, our first contribution is a new view of VQA compositionality, called *skill-concept composition*, and a new evaluation setting that directly targets how VQA models can generalize to novel compositions of skills and concepts. This view is motivated by our observation that, to answer a natural question on real images requires the understanding of two distinct elements: 1) the visual concept referred to by the question; and 2) what information we need to extract from the referred concept. We elucidate this in Sec. 3 and evaluate a number of VQA architectures using this setting and demonstrate that the existing models have much to improve upon to answer novel questions.

We propose a novel approach to improve generalization that utilizes contrastive learning to separate skills and concepts within the internal representations of a model, while jointly learning to answer questions. We use grounding as a proxy to separate concepts so that the model learns to identify a concept in both the question and image, regardless of the specific context. Akin to weakly supervised grounding [3, 18], we train the model to recover a concept mentioned in a given image-question pair by contrasting the multi-modal representation of the masked concept word to the multi-modal representations of words in other questions. We utilize a new way to curate positive and negative examples for the contrastive loss so that the model learns to predict the concept based on relevant visual information rather than using superficial contextual cues. Additionally, our approach learns to separate skills from concepts by contrasting question representations that have the same or different skills. These properties are learned jointly alongside the VQA objective, on top of state-of-the-art models, and are generalizable to new architectures.

Some advantages of our approach are: 1) We learn grounding in a self-supervised manner using the VQA data *alone*, without external annotations. This is in contrast to previous approaches with similar goals that incur large expenses due to annotation requirements [48, 60]. 2) Our method does not rely on answer labels to learn skill-concept separation, so we are able to use *unlabeled* image-question pairs to learn these properties. Consequently, we are able to acquire new concepts and learn to answer questions about them without having labeled data with these concepts, which is pivotal for generalizing to a new domain or novel instances. Moreover, we focus on data-efficient methods and do not use prodigious amounts of data external to VQA, like pre-training approaches [9, 41, 54], which is expensive to obtain and can require prior knowledge of the domain and/or concepts in order to perform well [20, 51].

Our main contributions in this paper are: 1) We present a novel view and evaluation setting for compositionality in VQA, called *skill-concept composition*, which enables

a more direct and interpretable evaluation of VQA models on real-image question answering. 2) We propose a novel contrastive learning approach, which combines the supervised VQA objective with self-supervised learning, to achieve skill-concept disentanglement at no additional annotation cost. 3) Our approach shows significant improvements over existing models on novel skill-concept compositions as well as generalization to unlabeled image-question pairs containing unseen concepts.

2. Related Work

VQA and Evaluations. Much progress has been made towards VQA [4, 33, 40, 54]. Large-scale pre-trained transformers [57], inspired by BERT [12], have recently become prevalent [9, 25, 38, 39, 40, 41, 53, 54, 61]. Progress is often measured by the widely adopted human-annotated VQA v2 [17] benchmark as well as other synthetic benchmarks [27, 30, 31]. Various datasets for reducing biases have been presented [2, 17], including VQA-CP [2] that creates distinct distributions of question prefixes and answers between the train and validation splits of VQA v2. Another dataset [31] evaluates on a conglomeration of data from VQA v1 [6] and synthetic data generated from question templates and image annotations, which they break down by vision task. The basic premise of considering questions by task is similar to our “skills”, but our setting examines compositional generalization of skills and concepts, where we evaluate on unseen compositions of skills and concepts. Existing work for evaluating generalization [2, 31] does not explore generalization to such novel compositions.

Compositionality and VQA. CLEVR [30] and GQA [27] are two VQA datasets that have received interest in recent years. Both offer *compositional questions*, which means that the questions involve various relational chains (e.g., “What color is the apple to the left of the bowl on the table?”). While GQA does not focus on novel compositions, CLEVR does investigate novel compositions of attributes and objects, where, for example, models see cubes of certain colors and cylinders of other colors during training, then the cubes and cylinders have their colors swapped in testing. This is analogous to our setting, although we propose to investigate skill-concept compositions and we experiment with natural questions about real images. Other efforts create compositional models to handle the relational reasoning chains [23, 24, 26, 28, 50]. Our approach implicitly learns compositional capabilities within state-of-the-art multi-modal transformer architectures, unlike these methods that do so explicitly. We also provide results on how well current compositional models (*i.e.*, neural module networks) can generalize to novel questions.

Grounding Visual Concepts. Visual grounding is typically studied on image-caption pairs. Previous work on visual grounding often learns grounding in a weakly supervised

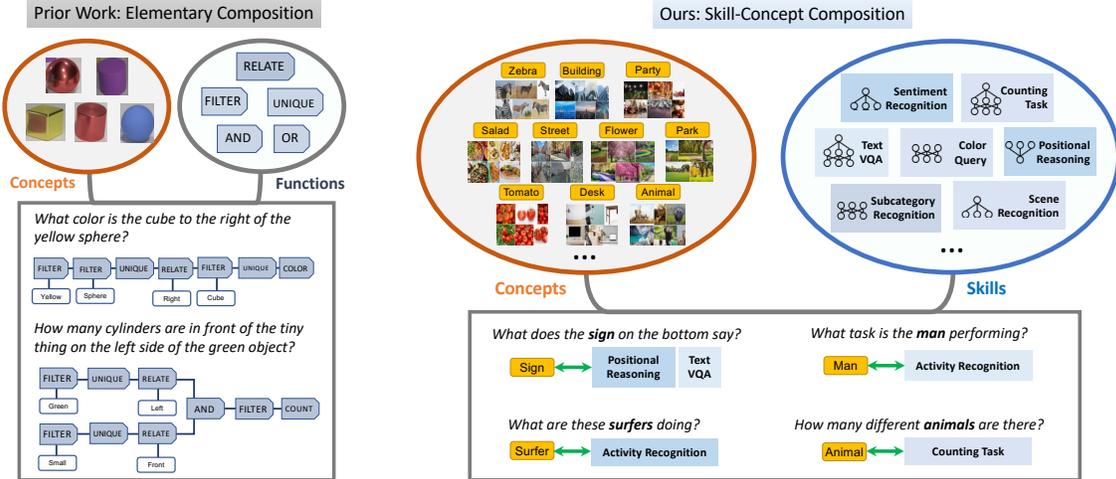


Figure 2. Illustration of skill-concept composition, a new view and evaluation setting for compositionality in VQA.

manner [3, 11, 15, 18], and sometimes leverages contrastive learning techniques [3, 18]. However, grounding concepts that appear in VQA questions has been less studied in previous work, which we explore in our paper. One direction that is somewhat related to visual grounding in VQA is learning VQA models which have interpretable visual attentions [48, 60]. However, these approaches usually rely on human annotations of the most influential regions for VQA answer prediction [10, 29], which are expensive to obtain and do not directly associate concept mentions with visual regions. Our work learns to ground concepts mentioned in questions, which facilitates compositional VQA modeling, and does so without using any additional annotations.

3. Skill-Concept Composition in VQA

We propose a novel, compositional view of VQA, called *skill-concept composition* (Fig. 2). *Concepts* are objects and other visually-grounded words or phrases. By *skills*, we refer to the collection of high-level vision understanding processes involved in answering common questions about real-world images. These skills operate on concepts and vary in terms of input/output representation complexity and the necessary reasoning processes. Our taxonomy of these skills is extracted from annotating a subset of the VQA v2 questions as well as taking inspiration from prior work on VQA.² Skills are generally standalone from each other and have been studied independently in the VQA literature (e.g., TextVQA [52], positional reasoning [27], or counting [1]).

We make an important yet intuitive observation about these VQA skills: to answer a question, it often requires the application of only a small number of skills (most often one) to one or more concepts in the image (Fig. 2). This ob-

ervation provides an interpretable view of a model’s generalization ability to out-of-distribution data: a model should learn that a skill is a separable process that can be applied to different concepts, and that the prediction process should not be tied to specific concepts co-occurring with this skill during training. This explicit notion of skill-concept separation underlies the contributions of this paper, including a new novel-VQA evaluation method which we will introduce next, as well as a new framework to weakly learn VQA models that can answer novel questions (Sec. 4).

Novel-VQA Evaluation. While conceptually intuitive, this skill-concept view offers natural ways to guide the evaluation of VQA models in terms of out-of-distribution data. In our experiments, we evaluate two novel-VQA settings: 1) answering questions on novel compositions of skills and concepts; 2) answering questions about concepts for which the model has not seen any answers before.

Comparison to Existing Evaluations. Our evaluation protocol is different from existing VQA benchmarks that also aim to measure VQA models’ generalization ability. VQA-CP [2] builds train-test splits from VQA v2/v1 [6, 17] with distinct answer distributions by greedily dividing the questions based on their annotated question types (i.e., first few words in the question: “how many”, “is the”, ...) and answers, but this does not capture skill-concept compositions because these question types do not necessarily correspond to skills (e.g., “Is the dog waiting?” requires action recognition and “Is the sky blue?” requires color recognition yet both have question type “is the”), and the same skill-concept composition can appear in training and testing, which violates our novel-VQA setting. TDIUC [31] evaluates VQA accuracy on different categories of task types, without regard for the concepts in the questions. [56] creates testing splits such that at least one word of a question is un-

²Please refer to Appendix A for a complete skill list.

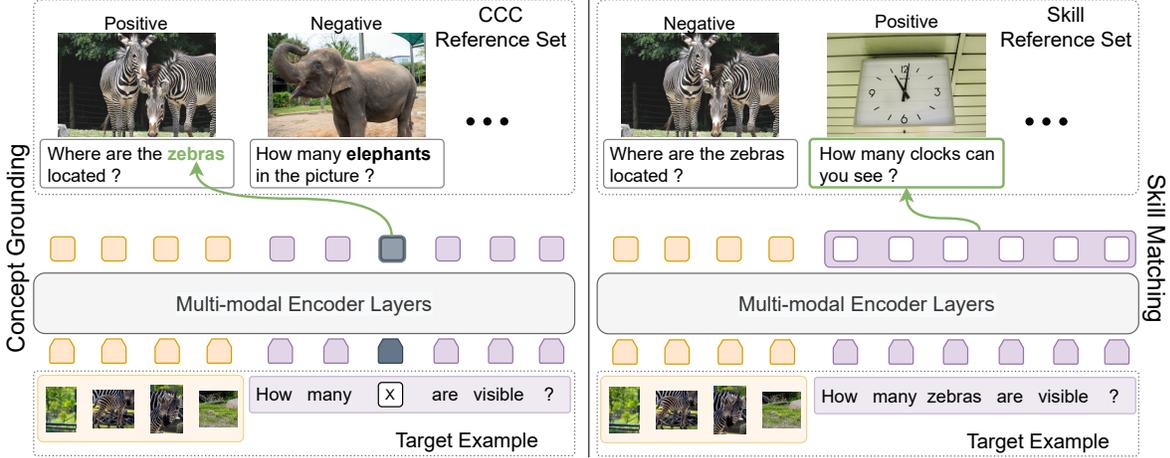


Figure 3. Overview of our approach. Left: We learn to ground concept representations by contrasting the multi-modal representations of a masked concept token in the target example and words in other questions. Right: We encode skills in the summary representations of the question by contrasting with summary representations of other questions with the same (positive) or different (negative) skills.

seen during training, which does not consider skills. None of these benchmarks directly address and evaluate skill-concept compositionality like our evaluation protocol.

Skill-Concept vs. Elementary Compositions. Existing compositional evaluations primarily define compositions as relational reasoning chains associated with questions [27, 30] (Fig. 2), which are suited for learning programs of elementary operations to answer questions. There are two main issues with applying this existing compositional view to real-image question answering. First, the concepts and their attributes are over-simplified and not representative of the diverse visual presentations in the real world. Second, the kinds of compositional questions in these synthetic datasets rarely appear in natural questions about real-world images. Our proposed *skill-concept composition* view is more applicable to real-world VQA and, as a result, can better represent the capabilities that people care about in real-image question answering.

4. Approach

Preliminaries. We assume that we are given a partially labeled dataset of tuples with image I , question Q , and answer labels A , where $(I^a, Q^a, A^a) \in \mathcal{D}^a$ has labels and $(I^u, Q^u) \in \mathcal{D}^u$ does not. Typically, models are trained with a VQA loss [55] using the labeled dataset, \mathcal{D}^a . Given an example (I, Q) , image region features, $g_v(I) = \{v_1, \dots, v_M\}$, and question token embeddings, $g_w(Q) = \{x_1, \dots, x_N\}$, are extracted and input to a multi-modal encoder to produce multi-modal representations of both modalities, $f(g_v(I), g_w(Q)) = (\{z_m\}_{m=1}^M, \{h_i\}_{i=1}^N)$, where z_m and h_i are the image and text multi-modal representations, respectively. An answer is predicted by pooling the encoded representations to a single representation (or

using a CLS token as input [12]), which is then input to a softmax output layer. We build upon this basic VQA setup to learn skill-concept separation, and uniquely take advantage of both labeled and unlabeled data.

Overview. We aim to learn separable skills and concepts, such that we can compose them to answer novel questions. To do so, the model should recognize that concepts mentioned in the question are manifested by their appearances in the image (*i.e.*, grounding) and that skills should be identifiable regardless of the concepts in the question or image. Gathering supervision for identifying concepts in the question, grounding them in the image, and labeling questions with skills would be very costly. Therefore, we propose to learn skill-concept separation in a self-supervised manner using contrastive learning [8, 44]. Illustrated in Fig. 3, we train the model with two additional contrastive objectives jointly with the VQA objective: *concept grounding* (Sec. 4.1), which learns grounded concept representations, and *skill matching* (Sec. 4.2), which encodes concept-agnostic representations of skills. For each of our objectives, the model is presented with a *target example* and a *reference set* of positive and negative examples sampled from carefully curated *candidate references*. Each objective trains the model to make the representation from the target example similar to those of the positive ones. We expound our training procedure for learning these objectives jointly with VQA in Sec. 4.3. In the following, for brevity, we put specific details of functions/settings in Appendix B.

4.1. Concept Grounding

To learn grounded representations of concepts, we mask the concept mention from the target question and then train the model to recover this concept mention, using the multi-

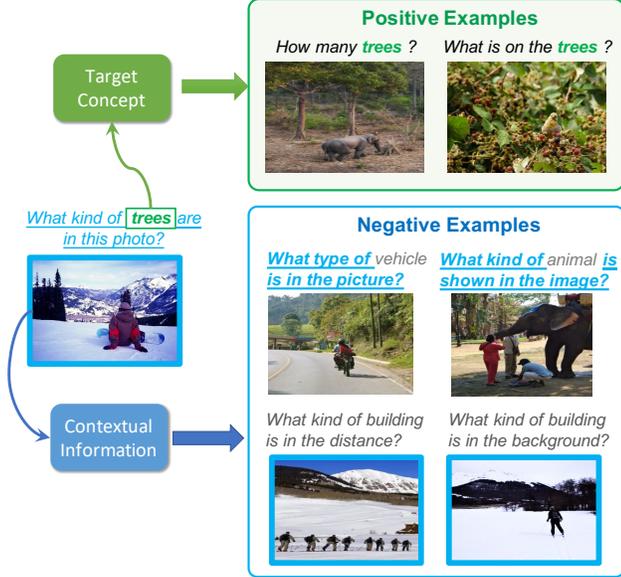


Figure 4. Target example (left) and a reference set of positive/negative examples (right) from our CCC references.

modal contextual information, by pointing to the same concept mention in examples in the reference set (Fig. 3).

Concept Discovery. We first identify the concept words that can be grounded in images. While this can be done with different methods [13, 21, 32], we simply use heuristics. We use POS tagging and lemmatization [22] to identify the 400 most frequent nouns in VQA v2 and then we filter out concepts that cannot be grounded (e.g., “time”). For a given question, Q , we want to find examples that have co-occurring mentions of a concept and the appearance of that concept in the image. It is likely that if a question about an image mentions a concept, then that concept may appear in the image [16]. Therefore, we identify the set of questions that mention the same concept, c , call it $\tilde{\mathcal{R}}_g^+(I, Q, c)$, which we consider as candidate positive references for Q . The set of questions not mentioning any of the same concepts are considered as candidate negative references, call it $\tilde{\mathcal{R}}_g^-(I, Q, c)$. To increase the likelihood of the concept appearing in the image, we employ a set of NLP-based heuristics to remove questions whose images may not contain the concept, such as counting questions with an answer of “0”.

Concept-Context Contrastive (CCC) References. Given a target question, Q , and a target concept mention, c , in Q , we could simply create reference sets by randomly sampling positive and negative examples [3, 44] from $\tilde{\mathcal{R}}_g^+(I, Q, c)$ and $\tilde{\mathcal{R}}_g^-(I, Q, c)$, respectively, based solely on whether the question contains c or not. However, we propose a novel reference example filtering strategy to encourage concept grounding. Our motivation is that, during VQA training, a concept often co-occurs with certain types of vi-

sual scenes or language priors. So the positive and negative examples should force the model to not rely on superficial cues when contrasted against the target example and, instead, look at the correct visual regions. Our solution is to build sets of refined reference candidates, $\mathcal{R}_g^+(I, Q, c)$ and $\mathcal{R}_g^-(I, Q, c)$, for each (I, Q, c) tuple to ensure that the co-occurrence factor present in the dataset can be reduced. As shown in Fig. 4, we want to find positive examples that also contain the concept “tree”, but with distinct visual scenes and questions from the target. For negative examples, we seek distractors that are similar to the target in terms of the question or visual scene (e.g., mountains with skiers in Fig. 4), but do not reference “tree”. To achieve this, we first represent the context of c by masking out c in the question and inputting the masked question and the image into off-the-shelf feature extractors to obtain question context representation q and image representation v .³ We measure the contextual similarity by:

$$\xi = \beta \cos(q, q') + (1 - \beta) \cos(v, v'), \quad (1)$$

where β is a scalar and (v, q) and (v', q') are the representations from target and candidate examples, respectively.

To select positive examples from $\tilde{\mathcal{R}}_g^+(I, Q, c)$, we use $\beta = 0.6$ and sample a set, $\mathcal{R}_g^+(I, Q, c)$, of N^+ examples that minimize ξ as our candidate positive examples for (I, Q, c) . For negatives, we apply two settings of β that maximize ξ : $\beta = 0.7$, which favors examples with more textual similarity, and $\beta = 0.3$, which prioritizes images with similar visual context. We select N^- examples from each setting as our candidate negative examples, $\mathcal{R}_g^-(I, Q, c)$. Illustrated in Fig. 4, when sampling reference sets from these two sets of candidates, the examples encourage the model to learn the specific correspondence between the concept mention in the question and its appearance in the image. Intuitively, the model must learn to ground the concept mention in the presence of the distractors.

Concept Grounding Loss. Let (I, Q) and c be the target example and target concept mention, respectively, and let $\mathcal{X} = \{(I_k, Q_k)\}_{k=1}^K$ be a corresponding reference set. Let k^* be the index of the positive example in \mathcal{X} sampled from $\mathcal{R}_g^+(I, Q, c)$, while the other $K - 1$ examples are negative examples from $\mathcal{R}_g^-(I, Q, c)$. Let w_i be the token in Q that refers to the concept c . We mask out w_i and input this masked version of the question along with the corresponding image into the model, f , which outputs multi-modal representations from which we extract the representation of the masked concept token, h_i . Next, we individually feed the examples from \mathcal{X} into the model to obtain each token representation $\hat{h}_{k,j}$, where j is the index of a token in Q_k . Let \hat{h}_{k^*,j^*} be the representation of the concept mention in the positive example’s question. Our grounding loss is an

³We use BERT [12] for questions and ResNet101 [19] for images.

NCE objective [8, 44] that requires the model to match the multi-modal representation of the masked concept mention to the representation of the same concept mention in the reference set:

$$\mathcal{L}_g = -\log \frac{\exp(\text{sim}(\phi_g(h_i), \phi_g(\hat{h}_{k^*,j^*})))}{\sum_{k,j} \exp(\text{sim}(\phi_g(h_i), \phi_g(\hat{h}_{k,j})))}, \quad (2)$$

where ϕ_g is a learned projection function and $\text{sim}(\cdot, \cdot)$ is a similarity function (e.g., dot product or cosine similarity). To correctly match h_i with \hat{h}_{k^*,j^*} , the model must encode the visual features that match between the images of these examples in both token representations. Our CCC references encourage these representations of the concept mention in the positive example and the masked concept mention to be grounded to the right visual regions as the model cannot rely on superficial textual or visual co-occurrences.

4.2. Skill Matching

Contrary to concepts, the essential skill needed to answer a certain question is largely independent of image appearances and mentions of concepts in the question. For example, *counting* questions should share a similar process to produce an answer: image areas associated with the subjects of counting are summarized to make the count prediction. This process should be independent of the type of objects being asked about. In other words, we seek to learn summary representations of questions that share the essential steps to infer the answer and are invariant to concepts.

Skill References. A straightforward approach to learn skills is to annotate questions which explicitly require the same reasoning steps. This annotation can be readily available on synthetic datasets [27, 30], but not available on datasets involving real-world images and questions. Instead, we propose to mine sets of contrasting examples to learn which questions require the same/different skills, matching questions with the same skills. Since the skills required for the question are typically indicated by the words of the question, we identify questions that are semantically similar. Essentially, questions that require the same skill (e.g., “*What color ...*”) should be related to one another, regardless of the specific concept mentions in the question. So, for each question, we mask out the concept words and we compute their BERT [12] representations. For a given (I, Q) , the set of positive reference examples, $\mathcal{R}_s^+(I, Q)$, are sampled from the top-200 most similar questions using BERT representation, and the set of negative examples, $\mathcal{R}_s^-(I, Q)$, are randomly chosen from the rest of the dataset.

Skill Matching Loss. For a given target example, (I, Q) , let h be a summary representation of the target question. This can be computed using a special input token like BERT [12] or via a pooling operation on all question token representations output from the encoder. We sample a reference set of image-question pairs, $\{(I_l, Q_l)\}_{l=1}^L$, where

the positive example, Q_{l^*} from $\mathcal{R}_s^+(I, Q)$, shares the same skill as the target question, and the rest of the reference set are negative examples from $\mathcal{R}_s^-(I, Q)$. Let \hat{h}_l be a summary representation for a question in the reference set. Shown in Fig. 3, our skill matching loss is defined as

$$\mathcal{L}_s = -\log \frac{\exp(\text{sim}(\phi_s(h), \phi_s(\hat{h}_{l^*})))}{\sum_l \exp(\text{sim}(\phi_s(h), \phi_s(\hat{h}_l)))}, \quad (3)$$

where \hat{h}_{l^*} is the positive example representation and ϕ_s is another learned projection function. This loss makes the representations of questions with the same skill more similar, regardless of the concepts mentioned.

4.3. Training Procedure

With our losses, we use a multi-tasking learning procedure [14, 42], where at each step we employ our objectives with probability p_{sep} or not with probability $1 - p_{\text{sep}}$. During training, we always first sample an instance from the labeled data, \mathcal{D}^a , and update the model by minimizing the VQA objective. If at the current iteration we do not use our skill and concept objectives, then we only use the VQA objective. Otherwise, we first use the VQA objective and then apply our other objectives. Both objectives are computed in the same fashion: for \mathcal{L}_g (or \mathcal{L}_s), we sample a target example from $\mathcal{D}^a \cup \mathcal{D}^u$ along with N_r^+ positive examples from \mathcal{R}_g^+ (or \mathcal{R}_s^+) as well as N_r^- negative examples from \mathcal{R}_g^- (or \mathcal{R}_s^-), combine the sampled references to form the current reference set, and compute the loss term. We then sum \mathcal{L}_g and \mathcal{L}_s , and update the model to minimize the negative sum.

5. Experiments

Data and Settings. We run our experiments on VQA v2 [17], which contains real images, human-written questions, and a variety of skills required to answer the questions. Since the goal of this work is to examine a model’s performance on different types of novel questions, it requires the availability of answer annotations for the test data. Since the annotations of test-dev and test-std sets of VQA v2 are not publicly available, we use questions from the validation set for testing as is common [2, 49]. We do not train or tune hyperparameters with the validation set; it is strictly used for evaluation. We compare performance using the VQA accuracy [6] on different splits of novel questions. Details are provided in Appendix C.

Model Comparisons. We select a set of recent VQA models to benchmark their novel-VQA performance. The first category is compositional models [5, 23, 24, 50]. We use StackNMN [23] and XNM [50], which are designed to handle compositional questions, like those in CLEVR [30], and have state-of-the-art performance on these datasets while also being applicable to real images without supervision from functional programs or image scene graphs.

Model	Counting				Color		Subcat.	Overall
	animal	{animals}	{vehicles}	{electronics}	animal	{dishware}	vegetable	
XNM [50]	56.02	48.32	44.35	<u>51.94</u>	77.22	<u>65.73</u>	57.33	<u>57.27</u>
StackNMN [23]	<u>54.22</u>	<u>47.56</u>	46.10	52.83	76.57	69.22	<u>57.17</u>	57.67
X-Att [54]	58.94	56.28	46.30	57.05	<u>73.15</u>	67.29	57.25	59.47
X-BERT [9]	63.58	54.58	<u>42.34</u>	56.84	75.88	70.31	58.96	60.36
Base	62.57	59.19	48.33	61.84	76.57	72.91	58.33	62.82
Ours	65.16	59.87	50.75	62.21	77.45	73.76	61.04	64.32

Table 1. VQA accuracy on novel skill-concept compositions. The highest and the lowest numbers of each experiment are emphasized.

Model	lamp	fruit	fridge	surfer	flag	skateb.	oven	sheep	banana	zebra	Overall
XNM [50]	53.69	50.23	57.98	72.68	36.58	70.16	53.49	54.96	52.35	61.50	56.36
StackNMN [23]	54.27	46.10	58.97	74.10	41.31	74.11	56.30	57.12	50.98	61.25	57.45
X-Att [54]	46.26	33.10	51.52	67.68	31.53	69.73	51.69	<u>49.83</u>	<u>41.71</u>	64.93	50.80
X-BERT [9]	<u>44.43</u>	<u>30.72</u>	<u>50.83</u>	<u>61.60</u>	32.46	<u>66.05</u>	48.10	<u>50.32</u>	43.10	57.06	<u>48.47</u>
Base	55.14	52.99	59.06	74.12	39.05	71.67	56.60	63.31	49.83	<u>56.05</u>	57.78
Ours	57.40	54.40	60.92	74.36	40.15	75.27	59.91	64.04	50.78	60.77	59.80

Table 2. VQA Accuracy on individual novel-concept split. *skateb.* refers to *skateboarder*.

The second type of model we experiment with is transformer-based [9, 12, 37, 40, 41, 54, 57, 61]. We use two top-performing transformer architectures from this model family: 1) a two-stream, cross-attention model [54] (**X-Att**), which has modality specific branches and cross-attentions in early layers followed by multi-modal layers later in the network; and 2) a vision-and-language transformer model [9] (**X-BERT**) that acts as multi-modal encoder throughout the entire network. For fair comparison, we do not use pre-training, same as our model, since we are specifically interested in the generalization ability of data-efficient models without requiring large-scale (e.g., 9M+ image-text pairs), in-domain data external to VQA [20, 51].

Lastly, for our base model, we employ a variant of the standard multi-modal transformer where image features can be attended by both the CLS token and language features. When the base model is trained without the proposed skill-concept contrastive losses, it serves as a baseline model (denoted by **Base**). Details are in Appendix C and Appendix D.

5.1. Novel Skill-Concept Composition VQA

We select three prevalent and common skills present in VQA v2: counting, color querying and subcategory recognition. For each skill, we remove the data labels for its co-occurring questions with one concept or a set of multiple concepts which can form a distinct category from training, and then test on these compositions. The concepts (or concept groups) are randomly sampled with two criteria: each skill-concept composition contains reliable amount of test data to measure accuracy and the compositions have diverse coverage across the dataset (more details in Appendix A).

Tab. 1 shows the VQA accuracy on each of the novel compositional subset. Interestingly, although neural module networks are designed to explicitly break down the question answering process into subtasks, which in principle

should help with adapting these subtasks to new questions and thus generalize better, they yield lower performance than transformer models. This may be due to the effective feature learning capacity of self-attention mechanisms. Among all transformer models, our base encoder achieves competitive performance to existing networks, demonstrating that it is a strong baseline among multi-modal transformers. Finally, our contrastive learning framework outperforms the baseline and all other approaches across each novel composition set. This supports the effectiveness of our framework for generalization to new compositions.

5.2. Novel-Concept VQA

For this experiment, we are interested in the setting where models are never trained to answer questions about a concept but can make use of the unlabeled image-question pairs, and then are tested on questions that have mentions of this given concept. Similar to the previous experiment, these concepts were sampled to maximize coverage as well as maintaining a reasonable test size. This setting is more challenging than the previous experiment since the model misses the VQA training supervision on any questions that have the given concept, as opposed to any questions that have both the given concept and a certain skill.

We provide qualitative examples of novel-concept VQA in Appendix E and report quantitative results in Tab. 2. For this more challenging setting of novel question answering, on average, two of the existing transformer architectures underperform other models by a noticeable margin. This may suggest that the transformer architectures, which perform well on large-scale vision and language pre-training, may have difficulty specializing to the VQA task. The Base model slightly outperforms neural module networks. Lastly, our framework again outperforms all models on average, demonstrating its value in improving VQA general-

ization ability on novel concepts.

5.3. Analysis

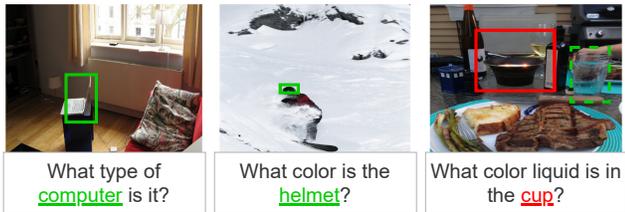


Figure 5. Correct and incorrect grounding examples. We visualize the most similar visual region to the concept in the question.

Concept Grounding. Since our approach learns to ground concepts without strong supervision, we would like to test its grounding abilities directly. To obtain an evaluation set, we manually annotate 320 image-question-concept tuples with the visual regions in the image that corresponds to the concept in each tuple. Candidate visual regions are found using Faster-RCNN [47]. We use recall@5 as our grounding metric, considering a grounding correct if the correct visual region falls within the top 5 most similar visual regions to the target concept token. The model trained with our framework achieves a grounding recall of **59.12**, compared to **43.71** of Base. Note that our framework obtained this improvement with no additional training data for grounding. As shown in Fig. 5, our model can often correctly ground a variety of objects, but can be fooled by ambiguous looking concepts like the candle in the incorrect example. Further, it is challenging to learn to differentiate concepts that almost always co-occur (e.g., “*shirt*” and “*person*”).

Loss Ablation. We ablate our losses by sampling three novel compositions and three novel concepts and report their average performance in Tab. 3. Adding our losses leads to consistent gains, with top performance achieved with our full framework. When used alone, our grounding loss seems to contribute a larger benefit compared to the skill loss. Nonetheless, the best performance is achieved by combining the two components, further supporting the value of skill and concept separation. We also experiment with a masked language modeling (MLM) objective [12] that replaces our losses. Our objectives perform better than the MLM objective, implying that the improvements our objectives offer are not simply due to additional data.

CCC Reference Sets. To study the effects of our CCC reference set selection strategy, we compare it with the commonly used random sampling method [3, 44] and report novel-concept VQA results in Tab. 4. We train both models with our full framework, the only difference being the reference set construction method for the concept loss. Both models improve upon the Base model, with our reference set construction method offering more consistent gains.

Model	Avg. Novel Count	Avg. Novel Concepts
Base	58.03	60.42
Base+MLM	58.41	60.25
Base+ \mathcal{L}_s	58.83	61.85
Base+ \mathcal{L}_g	59.80	62.06
Ours	60.71	63.19

Table 3. Effect of using different losses on novel skill-concept composition and novel-concept VQA.

Model	lamp	fruit	fridge	surfer	flag
Base	55.14	52.99	59.06	74.12	39.05
Random	+2.80	+0.49	-0.16	-0.14	+0.50
CCC (Ours)	+2.26	+1.41	+0.98	+0.24	+1.10

Table 4. Comparing different reference set construction schemes for concept learning across five different concepts.

Model	VQA-CP	Test-dev	Test-std
Base	40.98	69.60	69.99
Ours	41.71	69.78	70.09

Table 5. Single model VQA performance on VQA-CP [2] and VQA v2 test-dev and test-std splits. Both models see the exact same training data (no compositions/concepts are removed).

Existing Benchmarks. We also evaluate on VQA-CP [2] and the test-dev/test-std splits of VQA v2 (Tab. 5). While we see gains in general, notably, our approach is able to improve on VQA-CP without extra annotations, ensembling/tuning, or a performance drop on VQA v2.

6. Conclusions

We propose a new setting for generalization in VQA: measuring the ability to compose the skills needed to answer a question and the visual concepts that should be grounded to the image. We show that existing approaches have difficulty generalizing to unseen compositions of these two factors. We present a novel approach that implicitly disentangles skills and concepts, while grounding concepts visually, using a contrastive learning procedure. Our approach is able to learn from unlabeled VQA data in order to answer questions about previously unseen concepts. Results on the VQA v2 show that the proposed framework can achieve state-of-the-art performance on novel skill-concept compositions as well as generalize from unlabeled data.

Acknowledgements: We thank David Cox for the helpful discussions. From the UIUC side: This work was in part supported by the U.S. DARPA AIDA Program No. FA8750-18-2-0014. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation here on.

References

- [1] Manoj Acharya, Kushal Kafle, and Christopher Kanan. Tal-lyqa: Answering complex counting questions. In *AAAI*, 2019. 3
- [2] Aishwarya Agrawal, Dhruv Batra, Devi Parikh, and Anirudha Kembhavi. Don't just assume; look and answer: Overcoming priors for visual question answering. In *CVPR*, 2018. 2, 3, 6, 8, 12
- [3] Hassan Akbari, Svebor Karaman, Surabhi Bhargava, Brian Chen, Carl Vondrick, and Shih-Fu Chang. Multi-level multi-modal common semantic space for image-phrase grounding. In *CVPR*, 2019. 2, 3, 5, 8
- [4] Peter Anderson, Xiaodong He, Chris Buehler, Damien Teney, Mark Johnson, Stephen Gould, and Lei Zhang. Bottom-up and top-down attention for image captioning and visual question answering. In *CVPR*, 2018. 1, 2, 12
- [5] Jacob Andreas, Marcus Rohrbach, Trevor Darrell, and Dan Klein. Neural module networks. In *CVPR*, 2016. 6
- [6] Stanislaw Antol, Aishwarya Agrawal, Jiasen Lu, Margaret Mitchell, Dhruv Batra, C Lawrence Zitnick, and Devi Parikh. Vqa: Visual question answering. In *ICCV*, 2015. 2, 3, 6
- [7] Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. Neural machine translation by jointly learning to align and translate. In *ICLR*, 2015. 11
- [8] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *ICML*, 2020. 4, 6, 12
- [9] Yen-Chun Chen, Linjie Li, Licheng Yu, Ahmed El Kholy, Faisal Ahmed, Zhe Gan, Yu Cheng, and Jingjing Liu. Uniter: Universal image-text representation learning. In *ECCV*, 2020. 2, 7, 12, 13
- [10] Abhishek Das, Harsh Agrawal, C. Lawrence Zitnick, Devi Parikh, and Dhruv Batra. Human Attention in Visual Question Answering: Do Humans and Deep Networks Look at the Same Regions? In *EMNLP*, 2016. 3
- [11] Samyak Datta, Karan Sikka, Anirban Roy, Karuna Ahuja, Devi Parikh, and Ajay Divakaran. Align2ground: Weakly supervised phrase grounding guided by image-caption alignment. In *ICCV*, 2019. 3
- [12] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: Pre-training of deep bidirectional transformers for language understanding. In *NAACL*, 2019. 2, 4, 5, 6, 7, 8, 13
- [13] Jesse Dodge, Amit Goyal, Xufeng Han, Alyssa Mensch, Margaret Mitchell, Karl Stratos, Kota Yamaguchi, Yejin Choi, Hal Daumé III, Alexander Berg, et al. Detecting visual text. In *NAACL*, 2012. 5
- [14] Daxiang Dong, Hua Wu, Wei He, Dianhai Yu, and Haifeng Wang. Multi-task learning for multiple language translation. In *ACL*, 2015. 6
- [15] Hao Fang, Saurabh Gupta, Forrest Iandola, Rupesh K Srivastava, Li Deng, Piotr Dollár, Jianfeng Gao, Xiaodong He, Margaret Mitchell, John C Platt, et al. From captions to visual concepts and back. In *CVPR*, 2015. 3
- [16] Siddha Ganju, Olga Russakovsky, and Abhinav Gupta. What's in a question: Using visual questions as a form of supervision. In *CVPR*, 2017. 5
- [17] 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 *CVPR*, 2017. 1, 2, 3, 6, 12
- [18] Tanmay Gupta, Arash Vahdat, Gal Chechik, Xiaodong Yang, Jan Kautz, and Derek Hoiem. Contrastive learning for weakly supervised phrase grounding. In *ECCV*, 2020. 2, 3
- [19] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016. 5
- [20] Lisa Anne Hendricks, John Mellor, Rosalia Schneider, Jean-Baptiste Alayrac, and Aida Nematzadeh. Decoupling the role of data, attention, and losses in multimodal transformers. *arXiv preprint arXiv:2102.00529*, 2021. 2, 7
- [21] Jack Hessel, David Mimno, and Lillian Lee. Quantifying the visual concreteness of words and topics in multimodal datasets. In *NAACL*, 2018. 5
- [22] Matthew Honnibal and Ines Montani. spaCy 2: Natural language understanding with Bloom embeddings, convolutional neural networks and incremental parsing. To appear, 2017. 5
- [23] Ronghang Hu, Jacob Andreas, Trevor Darrell, and Kate Saenko. Explainable neural computation via stack neural module networks. In *ECCV*, 2018. 2, 6, 7, 12
- [24] Ronghang Hu, Jacob Andreas, Marcus Rohrbach, Trevor Darrell, and Kate Saenko. Learning to reason: End-to-end module networks for visual question answering. In *ICCV*, 2017. 2, 6
- [25] Zhicheng Huang, Zhaoyang Zeng, Bei Liu, Dongmei Fu, and Jianlong Fu. Pixel-bert: Aligning image pixels with text by deep multi-modal transformers. *arXiv preprint arXiv:2004.00849*, 2020. 2
- [26] Drew A Hudson and Christopher D Manning. Compositional attention networks for machine reasoning. In *ICLR*, 2018. 1, 2
- [27] Drew A Hudson and Christopher D Manning. Gqa: A new dataset for real-world visual reasoning and compositional question answering. In *CVPR*, 2019. 2, 3, 4, 6, 11
- [28] Drew A Hudson and Christopher D Manning. Learning by abstraction: The neural state machine. In *NeurIPS*, 2019. 2
- [29] Dong Huk Park, Lisa Anne Hendricks, Zeynep Akata, Anna Rohrbach, Bernt Schiele, Trevor Darrell, and Marcus Rohrbach. Multimodal explanations: Justifying decisions and pointing to the evidence. In *CVPR*, 2018. 3
- [30] Justin Johnson, Bharath Hariharan, Laurens van der Maaten, Li Fei-Fei, C Lawrence Zitnick, and Ross Girshick. Clevr: A diagnostic dataset for compositional language and elementary visual reasoning. In *CVPR*, 2017. 2, 4, 6
- [31] Kushal Kafle and Christopher Kanan. An analysis of visual question answering algorithms. In *ICCV*, 2017. 2, 3, 11
- [32] Gitit Kehat and James Pustejovsky. Integrating vision and language datasets to measure word concreteness. In *IJCNLP*, 2017. 5
- [33] Jin-Hwa Kim, Jaehyun Jun, and Byoung-Tak Zhang. Bilinear attention networks. In *NeurIPS*, 2018. 1, 2
- [34] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *ICLR*, 2015. 12

- [35] Brenden M Lake, Tomer D Ullman, Joshua B Tenenbaum, and Samuel J Gershman. Building machines that learn and think like people. *Behavioral and brain sciences*, 2017. 1
- [36] Wendy Lehnert. Human and computational question answering. *Cognitive Science*, 1977. 1
- [37] Gen Li, Nan Duan, Yuejian Fang, Ming Gong, Daxin Jiang, and Ming Zhou. Unicoder-vl: A universal encoder for vision and language by cross-modal pre-training. In *AAAI*, 2020. 7
- [38] 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*, 2020. 2
- [39] Liunian Harold Li, Mark Yatskar, Da Yin, Cho-Jui Hsieh, and Kai-Wei Chang. Visualbert: A simple and performant baseline for vision and language. *arXiv preprint arXiv:1908.03557*, 2019. 2
- [40] Xiujun Li, Xi Yin, Chunyuan Li, Pengchuan Zhang, Xiaowei Hu, Lei Zhang, Lijuan Wang, Houdong Hu, Li Dong, Furu Wei, et al. Oscar: Object-semantic aligned pre-training for vision-language tasks. In *ECCV*, 2020. 2, 7
- [41] Jiasen Lu, Dhruv Batra, Devi Parikh, and Stefan Lee. Vilbert: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks. In *NeurIPS*, 2019. 2, 7
- [42] Minh-Thang Luong, Quoc V Le, Ilya Sutskever, Oriol Vinyals, and Lukasz Kaiser. Multi-task sequence to sequence learning. In *ICLR*, 2016. 6
- [43] Thang Luong, Hieu Pham, and Christopher D. Manning. Effective approaches to attention-based neural machine translation. In *EMNLP*, 2015. 11
- [44] Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*, 2018. 4, 5, 6, 8, 12
- [45] Jeffrey Pennington, Richard Socher, and Christopher D. Manning. Glove: Global vectors for word representation. In *EMNLP*, 2014. 12
- [46] Sainandan Ramakrishnan, Aishwarya Agrawal, and Stefan Lee. Overcoming language priors in visual question answering with adversarial regularization. In *NeurIPS*, 2018. 2
- [47] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *NeurIPS*, 2015. 8
- [48] Ramprasaath R. Selvaraju, Stefan Lee, Yilin Shen, Hongxia Jin, Shalini Ghosh, Larry Heck, Dhruv Batra, and Devi Parikh. Taking a hint: Leveraging explanations to make vision and language models more grounded. In *ICCV*, 2019. 2, 3
- [49] Meet Shah, Xinlei Chen, Marcus Rohrbach, and Devi Parikh. Cycle-consistency for robust visual question answering. In *CVPR*, 2019. 6
- [50] Jiaxin Shi, Hanwang Zhang, and Juanzi Li. Explainable and explicit visual reasoning over scene graphs. In *CVPR*, 2019. 2, 6, 7, 12
- [51] Amanpreet Singh, Vedanuj Goswami, and Devi Parikh. Are we pretraining it right? digging deeper into visio-linguistic pretraining. *arXiv preprint arXiv:2004.08744*, 2020. 2, 7
- [52] Amanpreet Singh, Vivek Natarajan, Meet Shah, Yu Jiang, Xinlei Chen, Dhruv Batra, Devi Parikh, and Marcus Rohrbach. Towards vqa models that can read. In *CVPR*, 2019. 3
- [53] 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. 2
- [54] Hao Tan and Mohit Bansal. Lxmert: Learning cross-modality encoder representations from transformers. In *EMNLP*, 2019. 2, 7, 12
- [55] Damien Teney, Peter Anderson, Xiaodong He, and Anton Van Den Hengel. Tips and tricks for visual question answering: Learnings from the 2017 challenge. In *CVPR*, 2018. 4, 12
- [56] Damien Teney and Anton van den Hengel. Zero-shot visual question answering. *arXiv preprint arXiv:1611.05546*, 2016. 3
- [57] 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. 2, 7, 11, 13
- [58] Oriol Vinyals, Meire Fortunato, and Navdeep Jaitly. Pointer networks. In *NeurIPS*, 2015. 11
- [59] Jason Weston, Antoine Bordes, Sumit Chopra, Alexander M Rush, Bart van Merriënboer, Armand Joulin, and Tomas Mikolov. Towards ai-complete question answering: A set of prerequisite toy tasks. In *ICLR*, 2016. 1
- [60] Jialin Wu and Raymond Mooney. Self-critical reasoning for robust visual question answering. In *NeurIPS*, 2019. 2, 3
- [61] Zhou Yu, Jun Yu, Yuhao Cui, Dacheng Tao, and Qi Tian. Deep modular co-attention networks for visual question answering. In *CVPR*, 2019. 1, 2, 7
- [62] Xiaoyu Zeng, Yanan Wang, Tai-Yin Chiu, Nilavra Bhattacharya, and Danna Gurari. Vision skills needed to answer visual questions. *Proc. ACM Hum.-Comput. Interact.*, 2020. 1, 11

Appendix

A. Skill and Concept Details

To construct a comprehensive list of common skills required to answer a VQA question, we draw information from three sources: (1) our own annotation on 400 randomly selected VQA questions; (2) user study from [62]; and (3) previous work on question types [27, 31]. The user study in [62] only provides four types of vision skills. Existing work on question types have relevant information, however, the question types are not always directly translatable to our paradigm of skill and concept composition. For example, concept recognition is considered as a question type in [31] (object presence), but in our framework, it is considered as *concept grounding* rather than as a separate skill. Besides, existing question types are sometimes incomplete [31], or not representative of natural questions typically asked about images [27]. For instance, skills that require comparison or text reading form $\sim 6\%$ of the questions according to our labeling results, but they are not covered in [31]. We consolidate our annotations with groupings in existing work, which results in the following set of skills:

- Color recognition: *What color hair does the woman have? What color is his shirt?*
- Attribute recognition (non-color attributes): *Is the bed made? Is this desk messy?*
- Subcategory recognition: *What kind of car is parked? What kind of animals are shown?*
- Action recognition: *What is the man doing in the street? Are they comparing their phones?*
- Scene recognition: *Is this on a farm? Are they outside?*
- Counting: *How many lights are there? How many zebras are in this picture?*
- Commonsense knowledge: *Is the sun going down? Is this in America?*
- Positional reasoning: *What is on top of the toaster? What is the zebra standing on?*
- Text Recognition: *What number bus is it? What is the store called?*
- Comparison: *Is the tank the same color as the toilet? Are they facing the same direction?*

We also provide additional information and statistics of the novel compositions. To facilitate further research on novel-VQA evaluation, we will provide concept and skill annotations, and the respective data indices for each set of novel composition. The list of concepts within each concept group is:

- {animals}: *giraffe, zebra, bird, sheep, horse, elephant, cow, dog, cat*

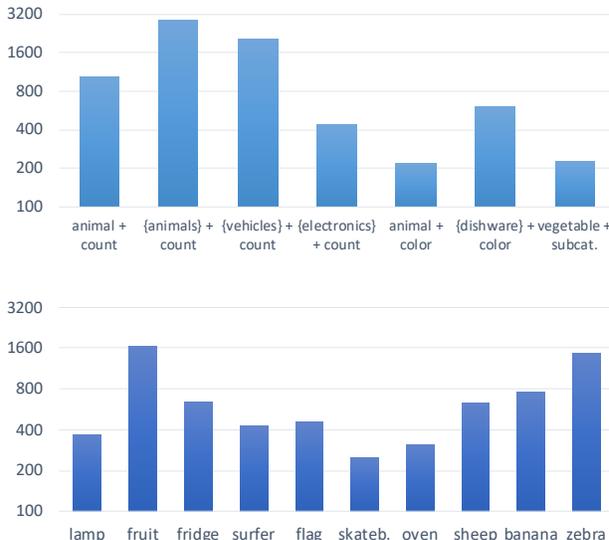


Figure 6. Novel skill-concept composition (top) and novel concept question statistics (bottom).

- {vehicles}: *motorcycle, airplane, plane, jet, bus, car, truck, bike, bicycle*
- {electronics}: *computer, monitor, laptop, phone, cellphone*
- {dishware}: *plate, bowl*

The list of sizes for each novel testing split is shown in Fig. 6. To determine the compositions/concepts that we use, we employ a few criteria: 1) each skill-concept composition (or concept) must have a minimum of 400 training questions and 200 testing questions; 2) for compositions, to increase coverage and ensure the minimally required size, we use concept groups where the concepts in a group all fall under a broader category (e.g., {animals} = {giraffe, zebra,...}). We then sample from these compositions/concepts to conduct experiments on.

B. Approach Details

Here, we detail the projection functions, similarity functions, and other settings for our approach. In the following equations, all W and b are learned parameters.

Concept Grounding. For our concept grounding loss, we want to maximize the similarity of the masked target concept token to the correct concept token in the positive reference example. Since we are directly comparing tokens between examples, we model the similarity computation as an attention [7, 43, 57] with which the model must point [58] to the correct concept token. Specifically, our projection func-

tion, $\phi_g(\cdot)$, and similarity function, $\text{sim}(\cdot, \cdot)$, are defined as

$$\phi_g(x) = W_g x + b_g \quad (4)$$

$$\text{sim}(x, y) = \frac{x^\top y}{\sqrt{d}}, \quad (5)$$

where d is the dimension of x and y , $W_g \in \mathbb{R}^{d \times d}$, and $b_g \in \mathbb{R}^d$. Though this is similar to an attention, our formulation matches more traditional contrastive learning objectives [8, 44], where \sqrt{d} is the temperature and we use a dot product as our similarity measure.

Skill Matching. Our skill matching loss seeks to maximize the similarity of the summary representation of the target question with the summary representations of other questions with the same skill. To obtain summary representations of questions, we simply use mean pooling over the question token representations. We define our projection function, $\phi_s(\cdot, \cdot)$, and similarity function, $\text{sim}(\cdot, \cdot)$, as

$$\phi_s(x) = W_s^{(2)} \psi(W_s^{(1)} x + b_s^{(1)}) + b_s^{(2)} \quad (6)$$

$$\text{sim}(x, y) = \frac{\cos(x, y)}{\tau_s}, \quad (7)$$

where ψ is a ReLU nonlinearity and τ_s is a temperature ($\tau_s = 0.5$ in our experiments). Since we are not directly comparing token representations, we use the more standard contrastive objective [8] as opposed to the attention-based formulation used for concept grounding.

Reference Sets and Training Procedure. When forming our CCC candidate references from which we sample our reference sets, we use $N^+ = 20$ and $N^- = 40$ (since we have two settings for negative examples), so there are N^+ positive and N^- negative examples that can be selected from to form a reference set for a given target example. Meanwhile, we use $N^+ = 200$ and $N^- = 200$ for our skill matching candidate references. Then, in our multi-task training procedure, we use $p_{\text{sep}} = 0.1$ as the probability of applying our framework at each training step. Additionally, we simply use $N_r^+ = 1$ and $N_r^- = 2$ for both concept grounding and skill matching, so the model will contrast between a single positive example and two distractor negative examples. For our concept grounding loss, we sample one negative example from both of our settings as our negative examples.

C. Experimental Details

C.1. Dataset Information

We use VQA v2 [17] for our main experiments. For training, we only use the training split. Since the testing data of VQA v2 [17] does not have public groundtruth information, we use the validation split of VQA v2 as the testing set for novel-VQA. To form our novel skill-concept and novel-concept VQA test splits, we automatically label

questions with the skills and concepts using different NLP-based rules. For labeling skills, we use question template matching (e.g., “How many ...”) as well as verifying that the answers fit the matched templates. For labeling concepts, we utilize lemmatization and POS tagging and collect the frequent nouns. We then create different training splits that have a specific skill-concept composition or concept removed.

We also run experiments on the test-dev, test-std, and VQA-CP [2] splits of VQA v2. When evaluating on test-dev and test-std, we train on the validation set and additional Visual Genome data [55].

C.2. Model Configurations

All models use the same visual features [4].⁴ We also use GloVe word embeddings [45].⁵ Our baselines from prior work follow the recommended settings provided by the authors, whenever possible.

For XNM [50], we use the implementation provided by the authors as well as the recommended settings.⁶ To ensure consistency between the two compositional models, we implement StackNMN [23] within the same code base as XNM. Specifically, we match the controller and the modules of StackNMN to the original paper. We use hidden dimension sizes of 512 for StackNMN and 1024 for XNM. We use the recommended number of reasoning steps, $T = 3$, for XNM and use the same for StackNMN. Both these models are trained with the Adam optimizer [34] and have the same learning rate of 0.0008 and batch size of 256.

For both X-Att [54]⁷ and X-BERT [9]⁸, we use the original model source code. For fair comparison, we do not use large-scale pre-training, same as our model. For X-Att, we use the recommended settings with a hidden size of 768, 12 layers, and 12 attention heads. X-Att uses the recommended learning rate of 0.0001, batch size of 64, 20 training epochs, and the Adam optimizer [34]. Due to their similarities in architecture, we use the same settings for Base, X-Bert and our framework for a more head-to-head comparison. Specifically, we use a hidden size of 512, 6 layers, and 8 attention heads. We match the training settings as well: a learning rate of 0.0001, batch size of 64, 13 training epochs, step learning rate decay with a rate of 0.2, and the Adam optimizer [34].

⁴<https://github.com/peteanderson80/bottom-up-attention>

⁵Common Crawl 840B: <https://nlp.stanford.edu/projects/glove/>

⁶<https://github.com/shijx12/XNM-Net>

⁷<https://github.com/airsplay/lxmert>

⁸<https://github.com/ChenRocks/UNITER>

D. Base Model Architecture

The base model (Base), to which we apply our framework, is based on the standard transformer encoder [9, 12] with a few modifications. As is standard with transformers, we input visual regions, question tokens, and a special CLS that is appended to the beginning of the inputs, which we use to predict answers via a softmax output layer. There are two minor differences between a standard transformer and our base model: First, before inputting the question into the transformer layers, we encode sequential information in the question tokens using an bi-directional LSTM, yielding a slight improvement than positional embeddings [57]. Second, in each layer, the CLS token and visual regions can attend to all inputs, including themselves, and the question tokens only directly attend to themselves and the CLS token. The change allows the CLS token to act as a bottleneck through which textual information interacts with the visual information.

E. Qualitative Examples

We show VQA output examples in Fig. 7 that compare the performance of our approach versus Base, where the first two rows show predictions on novel skill-concept compositions and the last row shows predictions on novel concept VQA. As a reminder, the models tested here never see labeled image-question pairs with the specific compositions/concepts during training. Our approach allows the model to adapt to these unseen compositions. We see that, for unseen compositions of *counting* and different concepts, the base model struggles to recognize and count these concepts. For example, we observe that despite the clear appearance of the animals in the images, the Base model is unable to transfer the skill of counting, whereas the model trained with our framework is able to handle these cases. Similarly, in the third and fourth examples of the first row, we see an interesting effect where our approach is able to more precisely locate the specific “*plate*” being referred to. Another interesting example of the improvements that our grounding framework can offer is shown in the first three examples of the last row, where our model is able to locate the specific object and produce the correct answer. The last two examples of the third row show some intriguing failure cases, where our model produces plausible yet somewhat generic answers compared to the baseline.

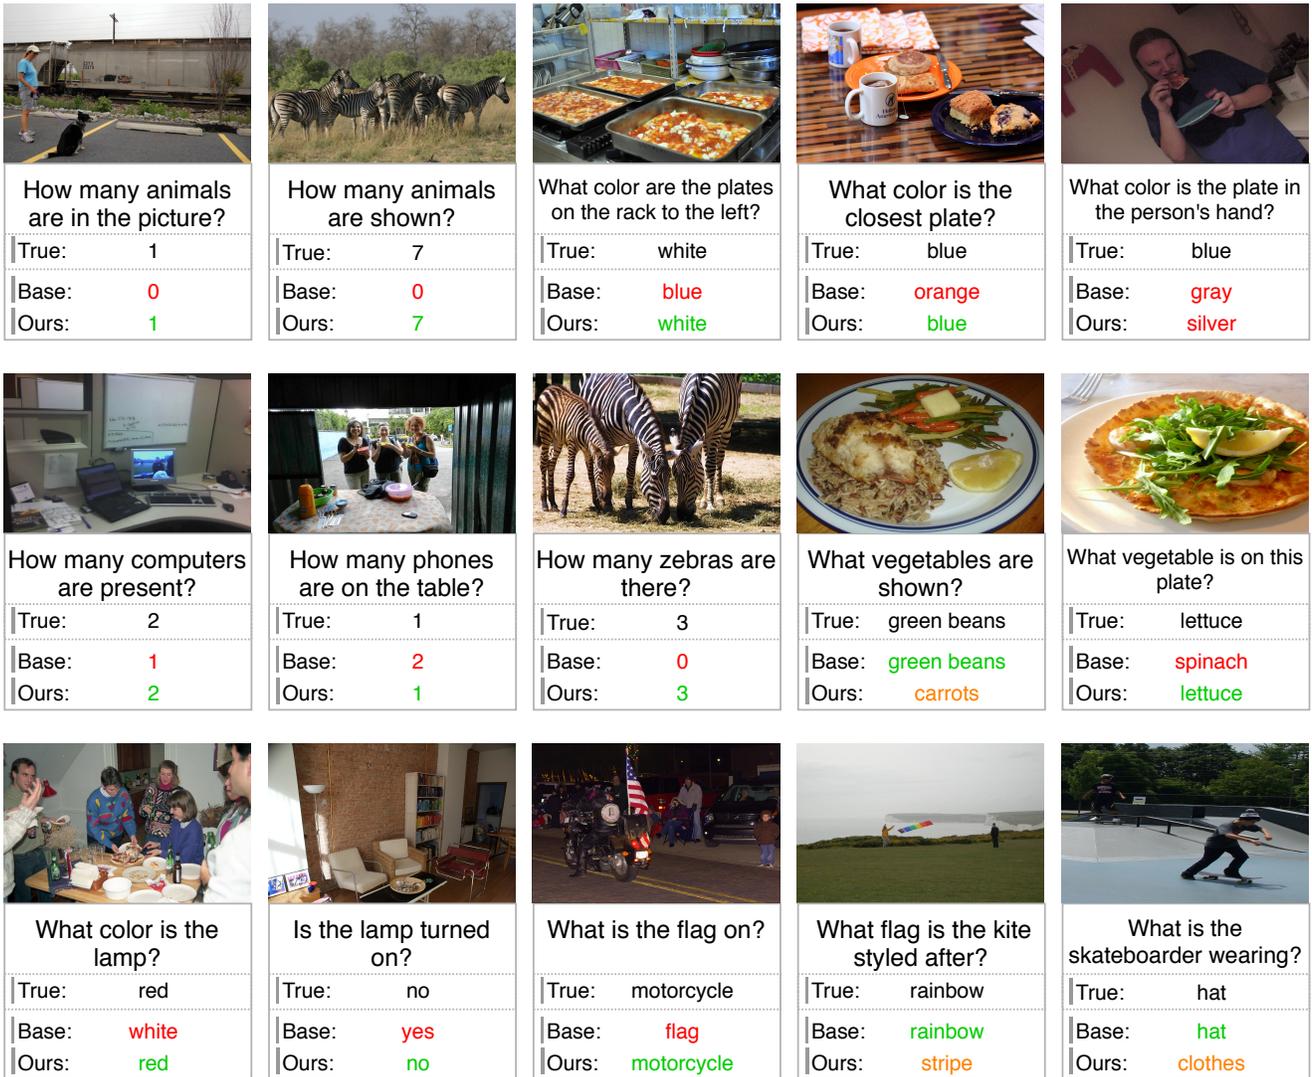


Figure 7. Correct, incorrect, and plausible VQA output examples for novel skill-concept composition VQA (rows 1 and 2) and novel concept VQA (row 3), comparing the predictions of our approach (Ours) and the Base model.