

# Yin and Yang: Balancing and Answering Binary Visual Questions

Peng Zhang<sup>\*†</sup> Yash Goyal<sup>\*†</sup> Douglas Summers-Stay<sup>‡</sup> Dhruv Batra<sup>†</sup> Devi Parikh<sup>†</sup>  
<sup>†</sup>Virginia Tech <sup>‡</sup>Army Research Laboratory

<sup>†</sup>{zhangp, ygoyal, dbatra, parikh}@vt.edu <sup>‡</sup>douglas.a.summers-stay.civ@mail.mil

## Abstract

The complex compositional structure of language makes problems at the intersection of vision and language challenging. But language also provides a strong prior that can result in good superficial performance, without the underlying models truly understanding the visual content. This can hinder progress in pushing state of art in the computer vision aspects of multi-modal AI.

In this paper, we address binary Visual Question Answering (VQA) on abstract scenes. We formulate this problem as visual verification of concepts inquired in the questions. Specifically, we convert the question to a tuple that concisely summarizes the visual concept to be detected in the image. If the concept can be found in the image, the answer to the question is “yes”, and otherwise “no”. Abstract scenes play two roles (1) They allow us to focus on the high-level semantics of the VQA task as opposed to the low-level recognition problems, and perhaps more importantly, (2) They provide us the modality to balance the dataset such that language priors are controlled, and the role of vision is essential. In particular, we collect fine-grained pairs of scenes for every question, such that the answer to the question is “yes” for one scene, and “no” for the other for the exact same question. Indeed, language priors alone do not perform better than chance on our balanced dataset. Moreover, our proposed approach matches the performance of a state-of-the-art VQA approach on the unbalanced dataset, and outperforms it on the balanced dataset.

## 1. Introduction

Problems at the intersection of vision and language are increasingly drawing more attention. We are witnessing a move beyond the classical “bucketed” recognition paradigm (e.g. label every image with categories) to rich compositional tasks involving natural language. Some of these problems concerning vision and language have proven surprisingly easy to take on with relatively simple techniques. Consider image captioning, which involves generating a

<sup>\*</sup>The first two authors contributed equally.

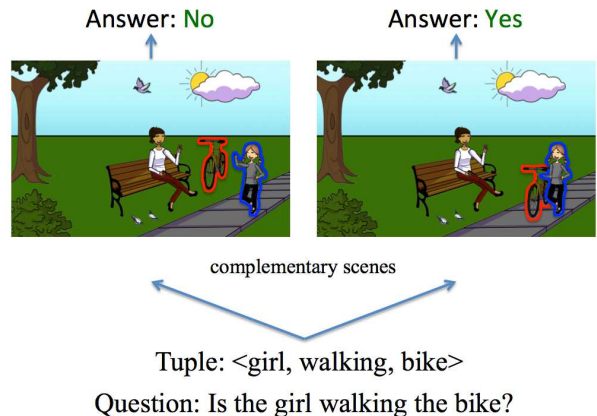


Figure 1: We address the problem of answering binary questions about images. To eliminate strong language priors that shadow the role of detailed visual understanding in visual question answering (VQA), we use abstract scenes to collect a *balanced* dataset containing pairs of *complementary* scenes: the two scenes have opposite answers to the same question, while being visually as similar as possible. We view the task of answering binary questions as a visual verification task: we convert the question into a tuple that concisely summarizes the visual concept, which if present, result in the answer of the question being “yes”, and otherwise “no”. Our approach attends to relevant portions of the image when verifying the presence of the visual concept.

sentence describing a given image [12, 6, 10, 26, 21, 19, 35]. It is possible to get state of the art results with a relatively coarse understanding of the image by exploiting the statistical biases (inherent in the world and in particular datasets) that are captured in standard language models.

For example, giraffes are usually found in grass next to a tree in the MS COCO dataset images [22]. Because of this, the generic caption “A giraffe is standing in grass next to a tree” is applicable to most images containing a giraffe in the dataset. The machine can confidently generate this caption just by recognizing a “giraffe”, without recognizing “grass”, or “tree”, or “standing”, or “next to”. In general, captions borrowed from nearest neighbor images result in a surprisingly high performance [8].

A more recent task involving vision and language is Visual Question Answering (VQA). A VQA system takes an image and a free-form natural language question about the image as input (e.g. “What is the color of the girl’s shoes?”, or “Is the boy jumping?”), and produces a natural language answer as its output (e.g. “blue”, or “yes”). Unlike image captioning, answering questions requires the ability to identify specific details in the image (e.g. color of an object, or activity of a person). There are several recently proposed VQA datasets on real images e.g. [2, 24, 25, 14, 29], as well as on abstract scenes [2]. The latter allows research on semantic reasoning without first requiring the development of highly accurate detectors.

Even in this task, however, a simple prior can give the right answer a surprisingly high percentage of the time. For example, in the VQA dataset (with images from MS COCO) [2], the most common sport answer “tennis” is the correct answer for 41% of the questions starting with “What sport is”. Similarly, “white” alone is the correct answer for 23% of the questions starting with “What color are the”. Almost half of all questions in the VQA dataset [2] can be answered correctly by a neural network that ignores the image completely and uses the question alone, relying on systematic regularities in the kinds of questions that are asked and what answers they tend to have.

This is true even for binary questions, where the answer is either “yes” or “no”, such as “Is the man asleep?” or “Is there a cat in the room?”. One would think that without considering the image evidence, both answers would be equally plausible. Turns out, one can answer 68% of binary questions correctly by simply answering “yes” to all binary questions. Moreover, a language-only neural network can correctly answer more than 78% of the binary questions, without even looking at the image.

As also discussed in [32], such dataset bias effects can give a false impression that a system is making progress towards the goal of understanding images correctly. Ideally, we want language to pose challenges involving the visual understanding of rich semantics while not allowing the systems to get away with ignoring the visual information. Similar to the ideas in [15], we propose to unbiase the dataset, which would force machine learning algorithms to exploit *image* information in order to improve their scores instead of simply learning to game the test. This involves not only having an equal number of “yes” and “no” answers on the test as a whole, but also ensuring that *each particular question is unbiased*, so that the system has no reason to believe, without bringing in visual information, that a question should be answered with “yes” or “no.”

In this paper, we focus on binary (yes/no) questions for two reasons. First, unlike open-ended questions (Q: “what is the man playing?” A: “tennis”), in binary questions (Q: “is the man playing tennis?”) all relevant semantic informa-

tion (including “tennis”) is available in the question alone. Thus, answering binary questions can be naturally viewed as *visual verification* of concepts inquired in the question (“man playing tennis”). Second, binary questions are easier to evaluate than open-ended questions.

Although our approach of visual verification is applicable to real images (more discussion in Sec. 6), we choose to use abstract images [2, 3, 39, 38, 40] as a test bed because abstract scene images allow us to focus on high-level semantic reasoning. They also allow us to balance the dataset by making changes to the images, something that would be difficult or impossible with real images.

Our main contributions are as follows: (1) We balance the existing abstract binary VQA dataset [2] by creating complementary scenes so that all questions<sup>1</sup> have an answer of “yes” for one scene *and* an answer of “no” for another closely related scene. We show that a language-only approach performs significantly worse on this balanced dataset. (2) We propose an approach that summarizes the content of the question in a tuple form which concisely describes the visual concept whose existence is to be verified in the scene. We answer the question by verifying if the tuple is depicted in the scene or not (See Fig. 1). We present results when training and testing on the balanced and unbalanced datasets.

## 2. Related work

**Visual question answering.** Recent work has proposed several datasets and methods to promote research on the task of visual question answering [15, 4, 33, 24, 2, 25, 14, 29], ranging from constrained settings [15, 24, 29] to free-form natural language questions and answers [4, 33, 2, 25, 14]. For example, [15] proposes a system to generate binary questions from templates using a fixed vocabulary of objects, attributes, and relationships between objects. [33] has studied joint parsing of videos and corresponding text to answer queries about videos. [24] studied VQA with synthetic (templated) and human-generated questions, both with the restriction of answers being limited to 16 colors and 894 object categories or sets of categories. A number of recent papers [2, 14, 25, 29] proposed neural network models for VQA composing LSTMs (for questions) and CNNs (for images). [2] introduced a large-scale dataset for free-form and open-ended VQA, along with several natural VQA models. [4] uses crowdsourced workers to answer questions about visual content asked by visually-impaired users.

**Data augmentation.** Classical data augmentation techniques (such as mirroring, cropping) have been widely used in past few years [18, 31] to provide high capacity models additional data to learn from. These transformations are designed to not change the label distribution in the training

<sup>1</sup>nearly all. About 6% of test questions do not lend themselves to this modification. See Sec. 3 for details.

data. In this work, we “augment” our dataset to explicitly change the label distribution. We use human subjects to collect additional scenes such that *every question* in our dataset has equal number of ‘yes’ and ‘no’ answers (to the extent possible). In that sense, our approach can be viewed as semantic data augmentation. Several classification datasets, such as ImageNet [7] try to be balanced. But this is infeasible for the VQA task on real images because of the heavy-tail of concepts captured by language. This motivates our use of abstract scenes.

**Visual abstraction + language.** A number of works have used abstract scenes to focus on high-level semantics and study its connection with other modalities such as language [23, 16, 38, 40, 39, 3, 13, 34], including automatically describing abstract scenes [16], generating abstract scenes that depict a description [39], capturing common sense [13, 23, 34], learning models of fine-grained interactions between people [3], and learning the semantic importance of visual features [38, 40]. Some of these works have also taken advantage of visual abstraction to “control” the distribution of data, for example, [3] collects equal number of examples for each verb/preposition combinations, and [38] have multiple scenes that depict the exact same sentence/caption. Similarly, we balance the dataset by making sure that each question in the dataset has a scene for “yes” and another scene for “no” to the extent possible.

**Visual verification.** [30, 34] reason about the plausibility of commonsense assertions (men, ride, elephants) by gathering visual evidence for them in real images [30] and abstract scenes [34]. In contrast, we focus on visually-grounded image-specific questions like “Is the man in the picture riding an elephant?”. [39] also reasons about relations between two objects, and maps these relations to visual features. They take as input a description and automatically generate a scene that is compatible with all tuples in the description and is a plausible scene. In our case, we have a single tuple (summary of the question) and we want to verify if it exists in a given image or not, for the goal of answering a free form “yes/no” question about the image.

**Visual attention** involves searching and attending to relevant image regions. [20, 36] uses alignment/attention for image caption generation. Input is just an image, and they try to describe the entire image and local regions with phrases and sentences. We address a different problem: visual question answering. We are given an image *and text* (a question) as input. We want to align parts of the question to regions in the image so as to extract detailed visual features of the regions of the image being referred to in the text.

### 3. Datasets

We first describe the VQA dataset for abstract scenes collected by [2]. We then describe how we balance this dataset by collecting more scenes.

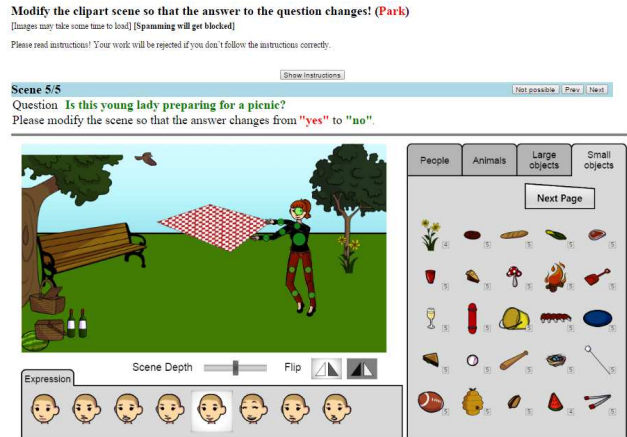


Figure 2: A snapshot of our Amazon Mechanical Turk (AMT) interface to collect complementary scenes.

### 3.1. VQA dataset on abstract scenes

**Abstract library.** The clipart library contains 20 “paperdoll” human models [3] spanning genders, races, and ages with 8 different expressions. The limbs are adjustable to allow for continuous pose variations. In addition to humans, the library contains 99 objects and 31 animals in various poses. The library contains two different scene types – “indoor” scenes, containing only indoor objects, e.g. desk, table, etc., and “outdoor” scenes, which contain outdoor objects, e.g. pond, tree, etc. The two different scene types are indicated by different background in the scenes.

**VQA abstract dataset** consists of 50K abstract scenes, with 3 questions for each scene, with train/val/test splits of 20K/10K/20K scenes respectively. This results in total 60K train, 30K validation and 60K test questions. Each question has 10 human-provided ground-truth answers. Questions are categorized into 3 types – ‘yes/no’, ‘number’, and ‘other’. In this paper, we focus on ‘yes/no’ questions, which gives us a dataset of 36,717 questions- 24,396 train and 12,321 val questions. Since test annotations are not publicly available, it is not possible to find the number of ‘yes/no’ type questions in test set. We use the binary val questions as our unbalanced test set, a random subset of 2,439 training questions as our unbalanced validation set, and rest of the training questions as our unbalanced train set.

### 3.2. Balancing abstract binary VQA dataset

We balance the abstract VQA dataset by posing a counterfactual task – given an abstract scene and a binary question, what would the scene have looked like *if the answer to the binary question was different?* While posing such counterfactual questions and obtaining corresponding scenes is nearly impossible in real images, abstract scenes allow us to perform such reasoning.

We conducted the following Mechanical Turk study – given

an abstract scene, and an associated question from the VQA dataset, we ask subjects to *modify the clipart scene such that the answer changes* from ‘yes’ to ‘no’ (or ‘no’ to ‘yes’). For example, for the question “Is a cloud covering the sun?”, a worker can move the ‘sun’ into open space in the scene to change the answer from ‘yes’ to ‘no’. A snapshot of the interface is shown in Fig. 2.

We ask the workers to modify the scene as little as possible. We encourage minimal changes because these complementary scenes can be thought of as hard-negatives/positives to learn subtle differences in the visual signal that are relevant to answering questions. This signal can be used as additional supervision for training models such as [37, 9, 28, 5] that can leverage explanations provided by the annotator in addition to labels. Our complementary scenes can also be thought of as analogous to good pedagogical techniques where a learner is taught concepts by changing one thing at a time via contrasting (e.g., one fish vs. two fish, red ball vs. blue ball, etc.). Full instructions on our interface can be found in supp. material.

Note that there are some (scene, question) pairs that do not lend themselves to easy creation of complementary scenes with the existing clipart library. For instance, if the question is “Is it raining?”, and the answer needs to be changed from ‘no’ to ‘yes’, it is not possible to create ‘rain’ in the current clipart library. Fortunately, these scenes make up a small minority of the dataset (e.g., 6% of the test set).

To keep the balanced train and test set comparable to unbalanced ones in terms of size, we collect complementary scenes for  $\sim$ half of the respective splits – 11,760 from train and 6,000 from test set. Since Turkers indicated that 2,137 scenes could not be modified to change the answer because of limited clipart library, we do not have complementary scenes for them. In total, we have 10,295 complementary scenes for the train set and 5,328 complementary scenes for test, resulting in balanced train set containing 22,055 samples and balanced test set containing 11,328 samples. We further split a balanced set of 2,202 samples from balanced train set for validation purposes. Examples from our balanced dataset are shown in Fig. 1 and Fig. 4.

We use the publicly released VQA evaluation script in our experiments. The evaluation metric uses 10 ground-truth answers for each question to compute performance. To be consistent with the VQA dataset, we collected 10 answers from human subjects using AMT for all complementary scenes in the balanced test set.

We compare the degree of balance in our unbalanced and balanced datasets. We find that 92.65% of the (scene, question) pairs in the unbalanced test set do not have a corresponding complementary scene (where the answer to the same question is the opposite). Only 20.48% of our balanced test set does not have corresponding complementary scenes. Note that our dataset is not 100% balanced either

because there are some scenes which could not be modified to flip the answers to the questions (5.93%) or because the most common answer out of 10 human annotated answers for some questions does not match with the intended answer of the person creating the complementary scene (14.55%) either due to inter-human disagreement, or if the worker did not succeed in creating a good scene.

## 4. Approach

We present an overview of our approach before describing each step in detail in the following subsections. To answer binary questions about images, we propose a two-step approach: (1) Language Parsing: where the question is parsed into a tuple, and (2) Visual Verification: where we verify whether that tuple is present in the image or not.

Our language parsing step summarizes a binary question into a tuple of the form  $\langle P, R, S \rangle$ , where P refers to primary object, R to relation and S to secondary object, e.g. for a binary question “Is there a cat in the room?”, our goal is to extract a tuple of the form:  $\langle \text{cat}, \text{in}, \text{room} \rangle$ . Tuples need not have all the arguments present. For instance, “Is the dog asleep”  $\rightarrow \langle \text{dog}, \text{asleep}, \rangle$ . The primary argument P is always present. Since we only focus on binary questions, this extracted tuple captures the entire visual concept to be verified in the image. If the concept is depicted in the image, the answer is “yes”, otherwise the answer is “no”.

Once we extract  $\langle P, R, S \rangle$  tuples from questions (details in Sec. 4.1), we align the P and S arguments to objects in the image (Sec. 4.2). We then extract text and image features (Sec. 4.4), and finally learn a model to reason about the consistency of the tuple with the image (Sec. 4.3).

### 4.1. Tuple extraction

In this section, we describe how we extract  $\langle P, R, S \rangle$  tuples from raw questions. Existing NLP work such as [11] has studied this problem, however, these approaches are catered towards statements, and are not directly applicable to questions. We only give an overview of our method, more details can be found in supp. material.

**Parsing:** We use the Stanford parser to parse the question. Each word is assigned an entity, e.g. nominal subject (“nsubj”), direct object (“dobj”), etc. We remove all characters other than letters and digits before parsing.

**Summarizing:** As an intermediate step, we first convert a question into a “summary”, before converting that into a tuple. First, we remove a set of “stop words” such as determiners (“some”, “the”, etc.) and auxiliary verbs (“is”, “do”, etc.). Our full list of stop words is provided in supp. material. Next, following common NLP practice, we remove all words before a nominal subject (“nsubj”) or a passive nominal subject (“nsubjpass”). For example, “Is the woman on couch petting the dog?” is parsed as “Is(*aux*) the(*det*) woman(*nsubj*) on(*case*) couch(*nmod*) petting(*root*)

the(*det*) dog(*obj*)?”. The summary of this question can be expressed as (woman, on, couch, petting, dog).

**Extracting tuple:** Now that we have extracted a summary of each question, next we split it into PRS arguments. Ideally, we would like P and S to be noun phrases (“woman on couch”, “dog”) and the relation R to be a verb phrase (“petting”) or a preposition (“in”) when the verb is a form of “to be”. For example, <dog, in, room>, or <woman on couch, petting, dog>. Thus, we apply the Hunpos Part of Speech (POS) tagger [17] to assign words to appropriate arguments of the tuple. See supp. material for details.

### 4.2. Aligning objects to primary (P) and secondary (S) arguments

In order to extract visual features that describe the objects in the scene being referred to by P and S, we need to align each of them with the image.

We extract PRS tuples from all binary questions in the training data. Among the three arguments, P and S contain noun phrases. To determine which objects are being referred to by the P and S arguments, we follow the idea in [39] and compute the mutual information<sup>2</sup> between word occurrence (e.g. “dog”), and object occurrence (e.g. clipart piece #32). We only consider P and S arguments that occur at least twice in the training set. At test time, given an image and a PRS tuple corresponding to a binary question, the object in the image with the highest mutual information with P is considered to be referred by the primary object, and similarly for S. If there is more than one instance of the object category in the image, we assign P/S to a random instance. Note that for some questions with ground-truth answer ‘no’, it is possible that P or S actually refers to an object that is not present in the image (e.g. Question: “Is there a cat in the image?” Answer: “no”). In such cases, some other object from images (say clipart #23, which is a table) will be aligned with P/S. However, since the category label (‘table’) of the aligned object is a feature, the model can learn to handle such cases, i.e., learn that when the question mentions ‘cat’ and the aligned clipart object category is ‘table’, the answer should be ‘no’.

We found that this simple mutual information based alignment approach does surprisingly well. This was also found in [39]. Fig. 3 shows examples of clipart objects and three words/phrases that have the highest mutual information.

### 4.3. Visual verification

We have extracted PRS tuples and aligned PS to the clipart objects in the image, we can now compute a score indicating the strength of visual evidence for the concept inquired in the question. Our scoring function measures compatibility between image and text features (described in Sec. 4.4).

<sup>2</sup>We compute MI separately for indoor and outdoor scenes. More details about scene types can be found in Sec. 3.1.

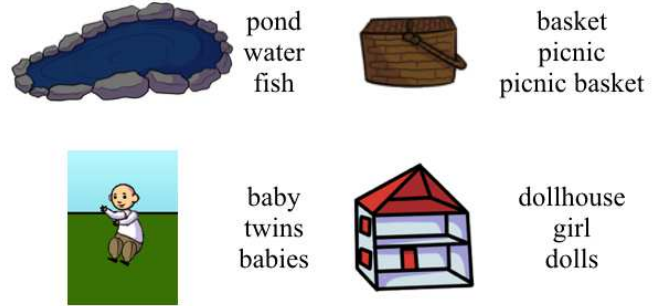


Figure 3: Most plausible words for an object determined using mutual information.

Our model is an ensemble of two similar models– Q-model and Tuple-model, whose common architecture is inspired from a recently proposed VQA approach [2]. Specifically, each model takes two inputs (image and question), each along a different branch. The two models (Q-model and Tuple-model) use the same image features, but different language features. Q-model encodes the sequential nature of the question by feeding it to an LSTM and using its 256-dim hidden representation as a language embedding, while Tuple-model focuses on the important words in the question and uses concatenation of word2vec [27] embeddings (300-dim) of P, R and S as the language features. If P, R or S consist of more than one word, we use the average of the corresponding word2vec embeddings. This 900-dimensional feature vector is passed through a fully-connected layer followed by a *tanh* non-linearity layer to create a dense 256-dim language embedding.

The image is represented by rich semantic features, described in Sec. 4.4. Our binary VQA model converts these image features into 256-dim with an inner-product layer, followed by a *tanh* layer. This inner-product layer learns to map visual features onto the space of text features.

Now that both image and text features are in a common space, they are point-wise multiplied resulting in a 256-dim *fused* language+image representation. This fused vector is then passed through two more fully-connected layers in a Multi-Layered Perceptron (MLP), which finally outputs a 2-way softmax score for the answers ‘yes’ and ‘no’. These predictions from the Q-model and Tuple-model are multiplied to obtain the final prediction. Both the models are learned separately and end-to-end (including LSTM) with a cross-entropy loss. Our implementation uses Keras [1]. Learning is performed via SGD with a batch-size of 32, dropout probability 0.5, and the model is trained till the validation loss plateaus.

At test time, given the question and image features, we can perform visual verification simply by performing forward pass through our network.

## 4.4. Visual Features

We use the same features as [23] for our approach. These visual features describe the objects in the image that are being referred to by the P and S arguments, their interactions, and the context of the scene within which these objects are present. In particular, the feature vector for each scene has 1432 dimensions, which are composed of 563 dimensions for each primary object and secondary object, encoding object category (e.g., cat vs. dog vs. tree), instance (e.g., which particular tree), flip (i.e., facing left or right), absolute location modeled via GMMs, pose (for humans and animals), expression, age, gender and skin color (for humans), 48 dimensions for relative location between primary and secondary objects (modeled via GMMs), and 258 dimensions encoding which other object categories and instances are present in the scene around P and S.

## 5. Experiments

### 5.1. Baselines

We compare our model with several strong baselines including language-only models as well as a state-of-the-art VQA method.

**Prior:** Predicting the most common answer in the training set, for all test questions. The most common answer is “yes” in the unbalanced set, and “no” in the balanced set.

**Blind-Q+Tuple:** A language-only baseline which has a similar architecture as our approach except that each model only accepts language input and does not utilize any visual information. Comparing our approach to Blind-Q+Tuple quantifies to what extent our model has succeeded in leveraging the image to answer questions correctly.

**SOTA Q+Tuple+H-IMG:** This VQA model has a similar architecture as our approach, except that it uses holistic image features (H-IMG) that describe the entire scene layout, instead of focusing on specific regions in the scene as determined by P and S. This model is analogous to the state-of-the-art models presented in [2, 25, 29, 14], except applied to abstract scenes.

These holistic features include a bag-of-words for clipart objects occurrence (150-dim), human expressions (8-dim), and human poses (7-dim). The 7 human poses refer to 7 clusters obtained by clustering all the human pose vectors (concatenation of (x, y) locations and global angles of all 15 deformable parts of human body) in the training set. We extract these 165-dim holistic features for the complete scene and for four quadrants, and concatenate them together to create a 825-dim vector. These holistic image features are similar to decaf features for real images, which are good at capturing what is present where, but (1) do not attend to different parts of the image based on the questions, and (2) may not be capturing intricate interactions between objects.

Comparing our model to SOTA Q+Tuple+H-IMG quantifies the improvement in performance by attending to specific regions in the image as dictated by the question being asked, and explicitly capturing the interactions between the relevant objects in the scene. In other words, we quantify the improvement in performance obtained by pushing for a deeper understanding of the image than generic global image descriptors. Thus, we name our model Q+Tuple+A-IMG, where A is for attention.

### 5.2. Evaluation on the original (unbalanced) dataset

In this subsection, we train all models on the train splits of both the unbalanced and balanced datasets, and test on our unbalanced test set. The results are shown in Table 1.

	Training set	
	Unbalanced	Balanced
Prior (“yes”)	68.67	68.67
Blind-Q+Tuple	78.90	60.80
SOTA Q+Tuple+H-IMG	78.49	69.19
Ours Q+Tuple+A-IMG	<b>79.20</b>	<b>72.80</b>

Table 1: Evaluation on unbalanced test set. All accuracies are calculated using the VQA [2] evaluation metric.

We draw the following key inferences:

**Vision helps.** We observe that models that utilize visual information tend to perform better than “blind” model when trained on the balanced dataset. This is because the lack of strong language priors in the balanced dataset forces the models to focus on the visual understanding.

**Attending to specific regions is important.** When trained on the balanced set where visual understanding is critical, our proposed model Q+Tuple+A-IMG, which focuses only on a specific region in the scene, outperforms all the baselines by a large margin.

**Bias is exploited.** As expected, the performance of all models trained on unbalanced dataset is better than the balanced dataset, because these models learn the language biases while training on unbalanced dataset, which are also present in the unbalanced test set.

### 5.3. Evaluation on the balanced dataset

We also evaluate all models trained on the train splits of both the unbalanced and balanced datasets, by testing on the balanced test set. The results are summarized in Table 2.

Here are the observations from this experiment:

**Training on balanced is better.** It is clear from Table 2 that both language+vision models trained on balanced data perform better than the models trained on unbalanced data. This may be because the models trained on balanced data *have* to learn to extract visual information to answer the question correctly, since they are no longer able to exploit

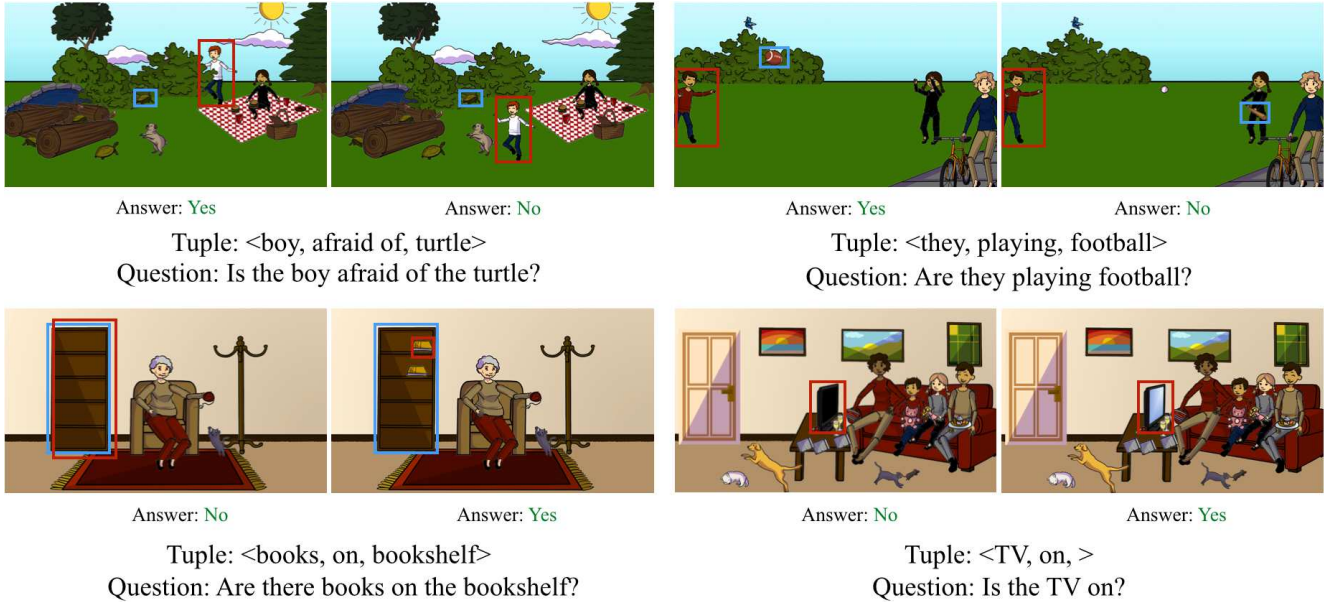


Figure 4: Qualitative results of our approach. We show input questions, complementary scenes that are subtle (semantic) perturbations of each other, along with tuples extracted by our approach, and objects in the scenes that our model chooses to attend to while answering the question. Primary object is shown in red and secondary object is in blue.

	Training set	
	Unbalanced	Balanced
Prior (“no”)	63.85	63.85
Blind-Q+Tuple	65.98	63.33
<i>SOTA</i> Q+Tuple+H-IMG	65.89	71.03
<i>Ours</i> Q+Tuple+A-IMG	<b>68.08</b>	<b>74.65</b>

Table 2: Evaluation on balanced test set. All accuracies are calculated using the VQA [2] evaluation metric.

language biases in the training set. Where as models trained on the unbalanced set are blindsided into learning strong language priors, which are then not available at test time.

**Blind models perform close to chance.** As expected, when trained on unbalanced dataset, the “blind” model’s performance is significantly lower on the balanced dataset (66%) than on unbalanced (79%). Note that the accuracy is higher than 50% because this is not binary classification accuracy but the VQA accuracy [2], which provides partial credit when there is inter-human disagreement in the ground-truth answers.

**Attention helps.** When trained on balanced dataset (where language biases are absent), our model Q+Tuple+A-IMG is able to outperform all baselines by a significant margin. Specifically, our model gives improvement in performance relative to the state-of-the-art VQA model from [2] (Q+Tuple+H-IMG), showing that attending to relevant regions and describing them in detail helps, as also seen in Sec. 5.2.

**Role of balancing.** We see clear improvements by reason-

ing about vision in addition to language. Note that in addition to the lack of language bias, the visual reasoning is also harder on the balanced dataset because now there are pairs of scenes with fine-grained differences but with opposite answers to the same question. So the model really needs to understand the subtle details of the scene to answer questions correctly. Clearly, there is a lot of room for improvement and we hope our balanced dataset will encourage more future work on detailed understanding of visual semantics towards the goal of accurately answering questions about images.

**Classifying a pair of complementary scenes.** We experiment with an even harder setting – a test point consists of a pair of complementary scenes and the associated question. Recall, that by construction, the answer to the question is “yes” for one image in the pair, and “no” for the other. This test point is considered to be correct only when the model is able to predict both its answers correctly.

Since language-only models only utilize the textual information in the question ignoring the image, and therefore, predict the same answer for both scenes, their accuracy is zero in this setting<sup>3</sup>. The results of the baselines and our model, trained on balanced and unbalanced datasets, are shown in Table 3. We observe that our model trained on the balanced dataset performs the best. And again, our

<sup>3</sup>Note that to create this pair-level test set, we only consider those pairs where the answers were opposites. We removed all scenes that workers were unable to create complementary scenes for due to a finite clipart library, as well as those scenes for which the majority answer from 10 workers did not agree with the intended answer of the creator of the scene.

model that focuses on relevant regions in the image to answer the question outperforms the state-of-the-art approach of [2] (Q+Tuple+H-IMG) that does not model attention.

	Training set	
	Unbalanced	Balanced
Blind-Q+Tuple	0	0
Q+Tuple+H-IMG	03.20	23.13
Q+Tuple+A-IMG	<b>09.84</b>	<b>34.73</b>

Table 3: Classifying a pair of complementary scenes. All accuracies are percentage of test pairs that have been predicted correctly.

#### 5.4. Analysis

Our work involves three steps: tuple extraction, tuple and object alignment, and question answering. We conduct analyses of these three stages to determine the importance of each of the three stages. We manually inspected a random subset of questions, and found the tuple extraction to be accurate 86.3% of the time. Given perfect tuple extraction, the alignment step is correct 95% of the time. Given perfect tuple extraction and alignment, our approach achieves VQA accuracy of 81.06% as compared to 79.2% with imperfect tuple extraction and alignment. Thus,  $\sim 2\%$  in VQA accuracy is lost due to imperfect tuple extraction and alignment.

#### 5.5. Ablation Study

We conducted an ablation study to analyze the importance of the two kinds of language features— LSTM for question vs. word2vec for tuple. For “blind” (language only) models trained and tested on unbalanced datasets, we found that the combination (Q+Tuple) performs better than each of the individual methods. Specifically, Q+Tuple achieves a VQA accuracy of 78.9% as compared to 77.87% (Q-only) and 77.54% (Tuple-only).

#### 5.6. Qualitative results

Fig. 4 shows qualitative results for our approach. We show a question and two complementary scenes with opposite answers. We find that even though pairs of scenes with opposite ground truth answers to the same questions are visually similar, our model successfully predicts the correct answers for both scenes. Further, we see that our model has learned to attend to the regions of the scene that seem to correspond to the regions that are most relevant to answering the question at hand. The ability to (correctly) predict different answers to scenes that are subtle (semantic) perturbations of each other demonstrates visual understanding.

## 6. Discussion

The idea of balancing a dataset can be generalized to real images. For instance, we can ask MTurk workers to find images with different answers for a given question. The advantage with clipart is that it lets us make the complementary scenes very fine-grained forcing the models to learn subtle differences in visual information. The differences in complementary real images will be coarser and therefore easier for visual models. Overall, there is a trade-off between clipart and real images. Clipart is easier (trivial) for low-level recognition tasks, but is more difficult balanced dataset because it can introduce fine-grained semantic differences. Real is more difficult for low-level recognition tasks, but may be an easier balanced dataset because it will have coarse semantic differences.

## 7. Conclusion

In this paper, we take a step towards the AI-complete task of Visual Question Answering. Specifically, we tackle the problem of answering binary questions about images. We balance the existing abstract binary VQA dataset by augmenting the dataset with complementary scenes, so that nearly all questions in the balanced dataset have an answer “yes” for one scene and an answer “no” for another closely related scene. For an approach to perform well on this balanced dataset, it *must* understand the image. We will make our balanced dataset publicly available.

We propose an approach that extracts a concise summary of the question in a tuple form, identifies the region in the scene it should focus on, and verifies the existence of the visual concept described in the question tuple to answer the question. Our approach outperforms the language prior baseline and a state-of-the-art VQA approach by a large margin on the balanced dataset. We also present qualitative results showing that our approach attends to relevant parts of the scene in order to answer the question.

**Acknowledgements.** This work was supported in part by the The Paul G. Allen Family Foundation via an award to D.P., ICTAS at Virginia Tech via awards to D.B. and D.P., Google Faculty Research Awards to D.P. and D.B., the Army Research Lab via awards to D.P. and D.B., the National Science Foundation CAREER award to D.B., the Army Research Office YIP Award to D.B., and an Office of Naval Research grant to D.B. The authors would like to thank Larry Zitnick from Facebook AI Research, Lucy Vanderwende from Microsoft Research and Claire Bonial from Army Research Lab for useful discussions.



## References

- [1] Keras: Theano-based deep learning library. <https://github.com/fchollet/keras.git>, 2015. 5
- [2] S. Antol, A. Agrawal, J. Lu, M. Mitchell, D. Batra, C. L. Zitnick, and D. Parikh. Vqa: Visual question answering. In *ICCV*, 2015. 2, 3, 5, 6, 7, 8
- [3] S. Antol, C. L. Zitnick, and D. Parikh. Zero-Shot Learning via Visual Abstraction. In *ECCV*, 2014. 2, 3
- [4] J. P. Bigham, C. Jayant, H. Ji, G. Little, A. Miller, R. C. Miller, R. Miller, A. Tatarowicz, B. White, S. White, and T. Yeh. VizWiz: Nearly Real-time Answers to Visual Questions. In *User Interface Software and Technology*, 2010. 2
- [5] A. Biswas and D. Parikh. Simultaneous active learning of classifiers & attributes via relative feedback. In *CVPR*, 2013. 4
- [6] X. Chen and C. L. Zitnick. Mind’s Eye: A Recurrent Visual Representation for Image Caption Generation. In *CVPR*, 2015. 1
- [7] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. ImageNet: A Large-Scale Hierarchical Image Database. In *CVPR*, 2009. 3
- [8] J. Devlin, S. Gupta, R. B. Girshick, M. Mitchell, and C. L. Zitnick. Exploring nearest neighbor approaches for image captioning. *CoRR*, abs/1505.04467, 2015. 1
- [9] J. Donahue and K. Grauman. Annotator rationales for visual recognition. In *ICCV*, 2011. 4
- [10] J. Donahue, L. A. Hendricks, S. Guadarrama, M. Rohrbach, S. Venugopalan, K. Saenko, and T. Darrell. Long-term Recurrent Convolutional Networks for Visual Recognition and Description. In *CVPR*, 2015. 1
- [11] A. Fader, S. Soderland, and O. Etzioni. Identifying relations for open information extraction. In *EMNLP*, 2011. 4
- [12] H. Fang, S. Gupta, F. N. Iandola, R. Srivastava, L. Deng, P. Dollár, J. Gao, X. He, M. Mitchell, J. C. Platt, C. L. Zitnick, and G. Zweig. From Captions to Visual Concepts and Back. In *CVPR*, 2015. 1
- [13] D. F. Fouhey and C. Zitnick. Predicting object dynamics in scenes. In *CVPR*, 2014. 3
- [14] H. Gao, J. Mao, J. Zhou, Z. Huang, and A. Yuille. Are you talking to a machine? dataset and methods for multilingual image question answering. In *NIPS*, 2015. 2, 6
- [15] D. Geman, S. Geman, N. Hallonquist, and L. Younes. A Visual Turing Test for Computer Vision Systems. In *PNAS*, 2014. 2
- [16] L. Gilberto Mateos Ortiz, C. Wolff, and M. Lapata. Learning to interpret and describe abstract scenes. In *Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 2015. 3
- [17] P. Halcsy, A. Kornai, and C. Oravecz. Hunpos - an open source trigram tagger. In *ACL*, 2007. 5
- [18] A. G. Howard. Some improvements on deep convolutional neural network based image classification. *CoRR*, abs/1312.5402, 2013. 2
- [19] A. Karpathy and L. Fei-Fei. Deep Visual-Semantic Alignments for Generating Image Descriptions. In *CVPR*, 2015. 1
- [20] A. Karpathy, A. Joulin, and F.-F. Li. Identifying relations for open information extraction. In *NIPS*, 2014. 3
- [21] R. Kiros, R. Salakhutdinov, and R. S. Zemel. Unifying Visual-Semantic Embeddings with Multimodal Neural Language Models. *TACL*, 2015. 1
- [22] T.-Y. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick. Microsoft COCO: Common Objects in Context. In *ECCV*, 2014. 1
- [23] X. Lin and D. Parikh. Don’t just listen, use your imagination: Leveraging visual common sense for non-visual tasks. In *CVPR*, 2015. 3, 6
- [24] M. Malinowski and M. Fritz. A Multi-World Approach to Question Answering about Real-World Scenes based on Uncertain Input. In *NIPS*, 2014. 2
- [25] M. Malinowski, M. Rohrbach, and M. Fritz. Ask your neurons: A neural-based approach to answering questions about images. In *ICCV*, 2015. 2, 6
- [26] J. Mao, W. Xu, Y. Yang, J. Wang, and A. L. Yuille. Explain Images with Multimodal Recurrent Neural Networks. *CoRR*, abs/1410.1090, 2014. 1
- [27] T. Mikolov, I. Sutskever, K. Chen, G. S. Corrado, and J. Dean. Distributed Representations of Words and Phrases and their Compositionality. In *NIPS*, 2013. 5
- [28] A. Parkash and D. Parikh. Attributes for classifier feedback. In *ECCV*, 2012. 4
- [29] M. Ren, R. Kiros, and R. Zemel. Exploring models and data for image question answering. In *NIPS*, 2015. 2, 6
- [30] F. Sadeghi, S. K. Kumar Divvala, and A. Farhadi. Viske: Visual knowledge extraction and question answering by visual verification of relation phrases. In *CVPR*, 2015. 3
- [31] P. Y. Simard, D. Steinkraus, and J. C. Platt. Best practices for convolutional neural networks applied to visual document analysis. Institute of Electrical and Electronics Engineers, Inc., August 2003. 2
- [32] A. Torralba and A. Efros. Unbiased look at the bias. In *CVPR*, 2011. 2
- [33] K. Tu, M. Meng, M. W. Lee, T. E. Choe, and S. C. Zhu. Joint Video and Text Parsing for Understanding Events and Answering Queries. *IEEE MultiMedia*, 2014. 2
- [34] R. Vedantam, X. Lin, T. Batra, C. L. Zitnick, and D. Parikh. Learning common sense through visual abstraction. In *ICCV*, 2015. 3
- [35] O. Vinyals, A. Toshev, S. Bengio, and D. Erhan. Show and Tell: A Neural Image Caption Generator. In *CVPR*, 2015. 1
- [36] K. Xu, J. Ba, R. Kiros, K. Cho, A. Courville, R. Salakhutdinov, R. Zemel, and Y. Bengio. Show, attend and tell: Neural image caption generation with visual attention. In *ICML*, 2015. 3
- [37] O. Zaidan, J. Eisner, and C. Piatko. Using annotator rationales to improve machine learning for text categorization. In *NAACL - HLT*, 2007. 4
- [38] C. L. Zitnick and D. Parikh. Bringing Semantics Into Focus Using Visual Abstraction. In *CVPR*, 2013. 2, 3
- [39] C. L. Zitnick, D. Parikh, and L. Vanderwende. Learning the Visual Interpretation of Sentences. In *ICCV*, 2013. 2, 3, 5
- [40] C. L. Zitnick, R. Vedantam, and D. Parikh. Adopting Abstract Images for Semantic Scene Understanding. *PAMI*, 2015. 2, 3