

Reasoning about Fine-grained Attribute Phrases using Reference Games

Jong-Chyi Su* Chenyun Wu* Huaizu Jiang Subhransu Maji
University of Massachusetts, Amherst
{jcsu, chenyun, hzjiang, smaji}@cs.umass.edu

Abstract

We present a framework for learning to describe fine-grained visual differences between instances using attribute phrases. Attribute phrases capture distinguishing aspects of an object (e.g., “propeller on the nose” or “door near the wing” for airplanes) in a compositional manner. Instances within a category can be described by a set of these phrases and collectively they span the space of semantic attributes for a category. We collect a large dataset of such phrases by asking annotators to describe several visual differences between a pair of instances within a category. We then learn to describe and ground these phrases to images in the context of a reference game between a speaker and a listener. The goal of a speaker is to describe attributes of an image that allows the listener to correctly identify it within a pair. Data collected in a pairwise manner improves the ability of the speaker to generate, and the ability of the listener to interpret visual descriptions. Moreover, due to the compositionality of attribute phrases, the trained listeners can interpret descriptions not seen during training for image retrieval, and the speakers can generate attribute-based explanations for differences between previously unseen categories. We also show that embedding an image into the semantic space of attribute phrases derived from listeners offers 20% improvement in accuracy over existing attribute-based representations on the FGVC-aircraft dataset.

1. Introduction

Attribute-based representations have been used for describing instances within a basic-level category as they often share a set of high-level properties. These attributes serve as basis for human-centric tasks such as retrieval and categorization [47, 23, 36], and for generalization to new categories based on a description of their attributes [12, 11, 39, 25]. However, most prior work has relied on a fixed set of attributes designed by experts. This limits their scalability to new domains since collecting expert annotations are expensive, and results in models that are less robust to noisy open-ended descriptions provided by a non-expert user.



Figure 1. *Left*: Each annotation in our dataset consists of five pairs of attribute phrases. *Right*: A reference game played between a speaker who describes an attribute of an image within a pair and a listener whose goal is to pick the right one.

Instead of discrete attributes this work investigates the use of *attribute phrases* for describing instances. Attribute phrases are short sentences that describe a unique semantic visual property of an object (e.g., “red and white color”, “wing near the top”). Like captions, they can describe properties in a compositional manner, but are typically shorter and only capture a single aspect. Like attributes, they are modular, and can be combined in different ways to describe instances within a category. Their compositionality allows the expression of large number of properties in a compact manner. For example, colors of objects, or their parts, can be expressed by combining color terms (e.g., “red and white”, “green and blue”, etc.). A collection of these phrases constitutes the semantic space of describable attributes and can be used as a basis for communication between a human and computer for various tasks.

We begin by collecting a dataset of attribute phrases by asking annotators to describe five visual differences between random pairs of airplanes from the OID airplane dataset [44]. Each difference is of the form “ P_1 vs. P_2 ” with phrases P_1 and P_2 corresponding to the properties of the left and right image respectively (Figure 1). By collecting multiple properties at a time we obtain a diverse set of describable attributes. Moreover, phrases collected in a contrastive manner reveal attributes that are better suited for fine-grained discrimination. The two phrases in a comparison describe the same underlying attribute (e.g., *round nose* and *pointy nose* both describe the shape), and reflect an axis

* Authors contributed equally

of comparison in the underlying semantic space. We then analyze the ability of automatic methods to generate these attribute phrases using the collected dataset. In particular we learn to generate descriptions and ground them in images in the context of a *reference game* (RG) between a *speaker* S and a *listener* L (Figure 1). S is provided with a pair of images $\{I_1, I_2\}$ and produces a visual difference of the form P_1 (or “ P_1 vs. P_2 ”). L’s goal is to identify which of the two images corresponds to P_1 . Reference games have been widely used to collect datasets describing objects within a scene. This work employs the framework to generate and reason about compositional language-based attributes for fine-grained visual categorization.

Our experiments show that a speaker trained to describe visual differences displays remarkable pragmatic behavior allowing a neural listener to rank the correct image with **91.4% top-5** accuracy in the RG compared with **80.6%** of a speaker trained to generate captions non-contrastively. We also investigate a family of *pragmatic speakers* who generate descriptions by jointly reasoning about the listener’s ability to interpret them, based on the work of Andreas and Klein [3]. Contrastively trained pragmatic speakers offer significant benefits (on average **7%** higher top-5 accuracy in RG across listeners) over simple pragmatic speakers. The resulting speakers can be used to generate attribute-based explanations for differences between two categories. Moreover, given a set of attribute phrases, the score of an image with respect to each phrase according to a listener provides a natural embedding of the image into the space of semantic attributes. On the task of image classification on the *FGVC aircraft dataset* [30] this representation outperforms existing attribute-based representations by **20%** accuracy.

In summary, we show that reasoning about attribute phrases via reference games is a practical way of discovering and grounding describable attributes for fine-grained categories. We validate our approach on a dataset of 6,286 images with 9,400 pairs for a total of 47,000 phrases (Section 3). We systematically evaluate various speakers and listeners using the RG and human studies (Section 4.1-4.2), investigate the effectiveness of attribute phrases on various recognition tasks (Section 4.3-4.6), and conclude in Section 5.

2. Related Work

Attribute-based representations. Attributes have been widely used in the computer vision as an intermediate, interpretable representation for high-level recognition. They often represent properties that can be shared across categories, *e.g.*, both a car and bicycle have wheels, or within a subordinate category, *e.g.*, birds can be described by the shape of their beak. Due to their semantic nature they have been used for learning interpretable classifiers [11, 12], attribute-based retrieval systems [6], as high-level priors for unseen cate-

gories for zero-shot learning [25, 18], and as a means for communication in an interactive recognition system [23].

A different line of work has explored the question of discovering task-specific attributes. Berg *et al.* [5] discover attributes by mining frequent n-grams in captions. Parikh and Grauman [35] ask annotators to name directions that maximally separate the data according to some underlying features. Other approaches [38, 17, 2] have mined phrases from online text repositories to discover common-sense knowledge about properties of categories (*e.g.*, cars have doors). For a detailed description of the above methods see this recent survey [29].

The interface for collecting attribute phrases is based on our earlier work [28], which showed that annotations collected in a pairwise manner could be analyzed to discover a lexicon of parts and attributes. This work extends the prior work in a several ways. We (a) consider the complete problem of generating and interpreting attribute phrases on a significantly larger dataset, (b) systematically evaluate speaker and listener models on the data, and (c) show their utility in various recognition tasks.

Referring expression comprehension and generation.

Modern captioning systems [22, 10, 46] produce descriptions by using encoder-decoder architectures, typically consisting of a convolutional network for encoding an image and a recurrent network for decoding a sentence. A criticism of these tasks is that captions in existing datasets (*e.g.*, MS COCO dataset [26]) can be generated by identifying the dominant categories and relying on a language model. State-of-the-art systems are often matched by simple nearest-neighbor retrieval approaches [9, 50]. Visual question-answering systems [4] face a similar issue that most questions can be answered by relying on common-sense knowledge (*e.g.*, the sky is often blue). Some recent attempts have been made to address these issues [19].

Tasks where annotators are asked to describe an object in an image such that another can correctly identify it provides a way to collect context-sensitive captions [20]. These tasks have been widely studied in the linguistics community in an area called pragmatics (see Grice’s maxims [13]). Much prior work in computer vision has focused on generating referring expressions to distinguish an object within an image [32, 20, 31, 15, 33, 48, 49, 27]. More recently, referring expression generation have been extended to interactive dialogue systems [7, 8]. In contrast, our work aims to collect and generate referring expressions for fine-grained discrimination between instances.

For the task of fine-grained recognition, the work of Reed *et al.* [37] is the most related to ours. They ask annotators on Amazon Mechanical Turk to describe properties of birds

and flowers, and use the data to train models of images and text. They show the utility of such models for zero-shot recognition where a description of a novel category is provided as supervision, and for text-based image retrieval. Another recent work [45] showed that referring expressions for images within a set can be generated simply by enforcing separation of image probabilities during decoding using beam search. However, their model was trained on context agnostic captions. Our work takes a different approach. First, we collect attribute phrases in a contrastive manner that encourages pragmatic behavior among annotators. Second, we ask annotators to provide multiple attribute descriptions, which as we described earlier, allows modular reuse across instances, and serves as an intermediate representation for various tasks. Attribute phrases capture the spectrum between basic attributes and detailed captions. Like “visual phrases” [39] they capture frequently occurring relations between basic attributes.

3. Method

The framework used to collect our dataset is described in Section 3.1. Various speaker and listener models are described in Section 3.2 and Section 3.3 respectively.

3.1. A dataset of attribute phrases

We rely on human annotators to discover the space of descriptive attributes. Our annotations are collected on images from the OID aircraft dataset [44]. The annotations are organized into 4700 image pairs (1851 images) in training set, 2350 pairs (1730 images) in validation set, and 2350 pairs (2705 images) in test set. Each pair is chosen by picking two different images uniformly at random within the provided split in the OID aircraft dataset.

Annotators from Amazon Mechanical Turk are asked to describe five properties in the form “ P_1 vs. P_2 ”, each corresponding to a different aspect of the objects in the left and the right image respectively. We also provide some examples as guidance to the annotators. The interface shown in Figure 2 is lightweight and allows rapid deployment compared to existing approaches for collecting attribute annotations where an expert decides the set and semantics of attributes ahead of time. However, the resulting annotations are noisier and reflect the diversity of open-ended language-based descriptions. A second pass over the data is done to check for consistency, after which about 15% of the description pairs were discarded.

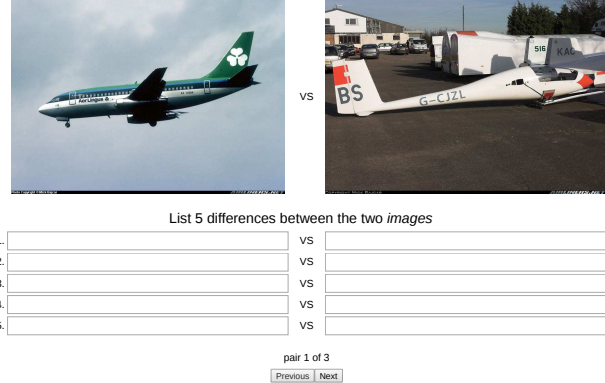
Figure 1 shows an example of our dataset (more examples are in the supplementary material). Annotations describe the shapes of parts (nose, wings and tail), relative sizes, orientation, colors, types of engines, *etc.* Most descriptions are short with an average length of 2.4 words on each side, although about 4.3% of them have more than 4 words. These

Describe differences between the two aeroplane images

Instructions:

- Annotate each one of the three tasks
- Press *Next* to move to the next pair and *Submit* once done.
- If the images do not display, your browser may not support this interface. Try the latest Chrome, Safari or Firefox browsers.

[Click here](#) to see example answers before you start.



The screenshot shows the annotation interface. At the top, there's a title 'Describe differences between the two aeroplane images' and instructions. Below the instructions is a link to see example answers. The main part of the interface displays two side-by-side images of airplanes. The left image is a green and white Airbus A320neo with the tail number G-CJZL. The right image is a white and red Cessna 441 with the tail number G-CJZL. Below the images, there's a prompt 'List 5 differences between the two images' and a table with 5 rows for input. Each row has a number (1-5) on the left, a 'vs' separator in the middle, and a text input field on the right. At the bottom, there's a 'pair 1 of 3' indicator and 'Previous' and 'Next' buttons.

Figure 2. The annotation interface used to collect five different *attribute phrase* pairs adapted from [28]. Amazon mechanical turkers were paid \$0.12 for annotating three pairs.

are qualitatively different from image captions which are typically longer and more grammatical. However, each annotation provides five different attribute pairs.

The OID dataset also comes with a set of expert-designed attributes. A comparison with OID attributes shows that attribute phrases capture novel properties that describe the relative arrangement of parts (*e.g.*, “door above the wing”, “wing on top”), color combinations, relative sizes, shape, and number of parts (*e.g.*, “big nose”, “more windows”, *etc.*) Section 4.4 shows a visualization of the space of attribute phrases. Section 4.3 provides a direct comparison of OID attributes and those derived from our data on the task of FGVC-aircraft variant classification [30].

3.2. Speaker models

A speaker maps visual inputs to attribute phrases. We consider two speakers; a *simple speaker* (SS) that takes a single image as input and produces a description, and a *discerning speaker* (DS) that takes two images as input and produces a single (or a pair of) description(s).

Both our speaker models are based on the show-and-tell model [46] developed for image captioning. Images are encoded using a convolutional network and decoded into a sentence using a recurrent network over words. We use one-hot encoding for 730 words with frequency greater than 5 in the training set. We consider fc7 layer outputs of the VGG-16 network [41] plus two fully-connected layers with ReLU units [34] on top as the image feature, and a LSTM model [14] with 2048 hidden units to generate the

sentences. The image feature is fed into the LSTM not only as the initial input, but also in each state input together with word embeddings. This led to an improved speaker in our experiments. For the *discerning speaker*, we concatenate two image features as input to the LSTM. At test time we apply beam search with beam size 10 and get 10 output descriptions from each image (pair). Although the *discerning speaker* is trained to generate phrase pairs, we can simply take the first (or second) half of the pair and evaluate it in the same way as a *simple speaker*.

We also consider a *pragmatic speaker* that generates contrastive captions by reasoning about the listener’s ability to pick the correct image based on the description. Andreas and Klein [3] proposed a simple strategy to do so by reranking descriptions of an image based on a weighted combination of (a) *fluency* – the score assigned by the speaker, and (b) *accuracy* – the score assigned by the listener on the referred image. Various pragmatic speakers are possible based on the choice of speakers and listeners. The details are described in Section 4.2.

Optimization details: Our implementation is based on Tensorflow [1]. The descriptions are truncated at length 14 when training the LSTM. The VGG-16 network is initialized with weights pre-trained on ImageNet dataset [24]. We first fix the VGG-16 weights and train the rest of the network, using Adam optimizer [21] with initial learning rate 0.001, $\beta_1 = 0.7$, $\beta_2 = 0.999$ and $\epsilon = 1.0 \times 10^{-8}$. We have batch normalization [16] in fully connected layers after VGG-16, and drop out with rate 0.7 in LSTM. We use batch size 64 for 40000 steps (~ 28 epochs). Second, we fine tune the whole network with initial learning rate modified to 5×10^{-6} , batch size 32 for another 40000 steps.

3.3. Listener models

A listener interprets a single (or a pair of) attribute phrase(s), and picks an image within a pair by measuring the similarity between the phrase(s) and images in a common embedded space. Once again we consider two listeners: a *simple listener* (SL) that interprets a single phrase, and a *discerning listener* (DL) that interprets a phrase pair. The *simple listener* models the score of the image I_1 within a pair (I_1, I_2) for a phrase P as:

$$p(I_1|P) = \sigma(\phi(I_1)^T \theta(P), \phi(I_2)^T \theta(P)).$$

Here ϕ and θ are embeddings of the image and the phrase respectively, and σ is the softmax function $\sigma(x, y) = \exp(x) / (\exp(x) + \exp(y))$. Similarly, a *discerning listener* models the score of an image by comparing it with an embedding of the phrase pair $\theta([P_1 \text{ vs. } P_2])$. A simple way to construct a discerning listener from a simple listener is by averaging the predictions from the left and right phrases, i.e., $p(I|[P_1 \text{ vs. } P_2]) = (p(I|P_1) + p(I|P_2)) / 2$.

We follow the setup of the speaker to embed phrases and use the final state of a LSTM with 1024 hidden nodes as the phrase embedding. The vocabulary of words is kept identical. For image features, once again we use the fc7 layer of the VGG-16 network and add a fully-connected layer with 1024 units and ReLU activation. The parameters are learned by minimizing the cross-entropy loss.

We also evaluate two variants of the *simple listener*, SL_r and SL , based on whether it is trained on non-contrastive data (I_1, I_2, P_1) where I_2 is a *random image* within the training set, or the contrastive data where I_2 is the other image in the annotation pair.

Optimization details: We first fix the VGG-16 network and use Adam optimizer with initial learning rate = 0.001, $\beta_1 = 0.7$, batch size = 32 for 2000 steps (4000 steps for SL_r model), then fine-tune the entire model with a learning rate 1×10^{-5} for another 7000-10000 steps.

Human listener. We also consider human annotators to perform the task of the listener in the RG. For each generated phrase that describes one image out of an image pair, we let three users to pick which image out of the pair the phrase is referring to. However, unlike (most) human speakers, neural speakers can produce irrelevant descriptions. Thus, in addition to the choice of left and right image, users have the option to say “*not sure*” when the description is ambiguous. If two or more users out of three picked the same image, we say the human listener is certain about the choice, otherwise we say the human listener is uncertain. The interface is shown in the supplementary material.

4. Results

We evaluate various listeners and speakers on the dataset we collected in terms of their accuracy in the RG in Section 4.1 and Section 4.2 respectively. We then evaluate their effectiveness on a fine-grained classification task in Section 4.3, visualize the space of attribute phrases discovered from the data in Section 4.4, for text-based image retrieval in Section 4.5, and for generating visual explanations for differences between categories in Section 4.6.

4.1. Evaluating listeners

We first evaluate various listeners on human-generated phrases. For simple listeners, each annotation provides ten different reference tasks $(I_1, I_2, P) \rightarrow \{0,1\}$ corresponding to five different left and right attribute phrases. Each task is evaluated independently and accuracy is measured as the fraction of correct references made by the listener. Similarly, discerning listeners are evaluated by replacing P with “ $P_1 \text{ vs. } P_2$ ” or “ $P_2 \text{ vs. } P_1$ ”.

Accuracy using human speakers. The results are shown in Table 1. Training on contrastive data improves the accuracy of the simple listener slightly from 84.2% (SL_r) to 86.3% (SL) on the test set. Discerning listeners see both phrases at once and naturally perform better. There is almost no difference between a discerning listener that combines two simple listeners by averaging their predictions ($2 \times SL$), and one that interprets the two phrases at once (DL). The results indicate that on our dataset the listener’s task is relatively easy and contrastive data does not provide any significant benefits. As a reference the accuracy of a human listener is close to 100% on human-generated phrases.

Input	Speaker	Listener	Val	Test
P_1	Human	SL_r	82.7	84.2
		SL	85.3	86.3
P_1 vs. P_2	Human	DL	88.7	88.9
		$2 \times SL$	89.6	89.3

Table 1. Accuracy (%) of various listeners in the RG using attribute phrases provided by a human speaker.

Are the top attributes more salient? As annotators are asked to describe five different attributes they might pick ones that are more salient first. We evaluate this hypothesis by measuring the accuracy of the listener (SL) on phrases as a function of the position of the annotation in the interface ranging from one for the top attribute to five for the last one. The results are shown in Table 2. The accuracy decreases monotonically from one to five suggesting that the first attribute phrase is easier for the listener to discriminate. We are uncertain if this is because the attributes near the top are more discriminative, or because the listener is better at interpreting these as they are likely to be more frequent in the training data. Nevertheless, attribute saliency is a signal we did not model explicitly and may be used to train better speakers and listeners (*e.g.*, see Turakhia and Parikh [42]).

	1	2	3	4	5
Val	91.3	86.6	84.1	82.5	82.3
Test	92.3	87.4	85.9	84.0	81.6

Table 2. Accuracy (%) of the simple listener (SL) on RG using human-generated attribute phrases at positions one through five across the validation and test set. The accuracy decreases monotonically from one to five suggesting that the top attribute phrases are easier to discriminate.

4.2. Evaluating speakers

We use simple listeners, SL and SL_r , and the human listener to evaluate speakers. As described in Section 3.2 we use beam search to generate 10 descriptions for each image pair and evaluate them individually using various listeners. The discerning speaker generates phrase pairs but we simply take the first and second half separated by “vs.”, a special word in the vocabulary, and evaluate it using a simple

listener (that sees only one phrase). If the word “vs.” is missing in the generated output we simply consider the entire sentence as the P_1 . Only 1 out of 23500 phrase pairs did not contain the “vs.” token.

For evaluation with humans we collect three independent annotations on a subset of 100 image pairs (with 10 descriptions each) out of the full test set. The listeners are considered to be correct when the probability of the correct image is greater than half. For human listener, we report the accuracy of when there is a majority agreement on the correct image, *i.e.*, when two or more users picked the correct image. For direct comparison with the simple speaker models, we also report the human listener accuracy when they are allowed to guess. This is the sum of earlier accuracy, and half of the cases when there is no majority agreement. Human annotators are uncertain when the generated descriptions are not fluent or when they are not discriminative. Therefore, a better human accuracy reflects speaker quality both in terms of fluency and discriminativeness. Some examples of the generated attribute phrases using various speakers are shown in Figure 3.

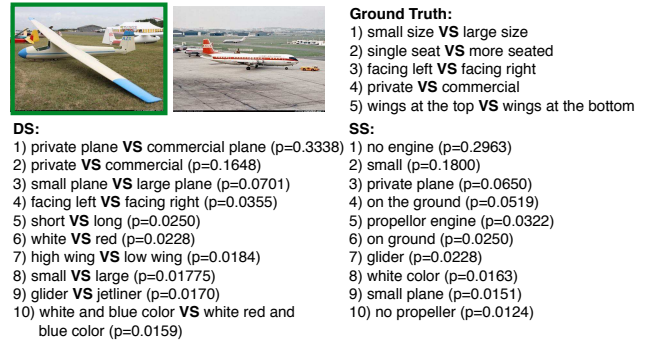


Figure 3. Example output of simple speaker SS and discerning speaker DS. Simple speaker takes the left image in the green box as input, while the discerning speaker takes both images as input. In brackets are the probabilities according to the speaker.

Accuracy of various speakers and listeners. Results on the full test set (Test) and the human-evaluated subset (Test*) are shown in Table 3. The accuracy of discerning speaker exceeds that of simple speaker by more than 10% no matter which listener to use. This result suggests that data collected contrastively using our annotation task allows direct training of speaker models that show remarkable context-sensitive behavior. Somewhat surprisingly we also see that the simple listeners are more accurate than the human listener when evaluated on descriptions generated by our speaker models. This is because humans tend to be more cautious in the reference game. For example, simple listeners will accept yellowish grass being referred to as “concrete” compared to green grass, but humans tend to view it as an unclear reference.

Accuracy (%)						
Top	SL _r		SL		Human	
	Test*	Test	Test*	Test	Test*	
SS	1	84.0	79.8	83.0	81.7	68.0 (77.0)
	5	80.0	79.2	78.0	80.6	64.2 (74.1)
	10	78.0	78.9	76.6	80.0	61.6 (72.4)
DS	1	94.0	92.8	92.0	92.8	82.0 (88.5)
	5	91.2	90.3	91.2	91.4	80.2 (86.7)
	10	88.6	88.8	90.0	90.5	77.9 (85.0)

Table 3. Accuracy in the RG using different speakers and listeners. Test represents the full test set consisting of 2350 image pairs. Test* represents a subset of 100 test set image pairs for which we collected human listener results. For the human listener, we report the accuracy when there is a majority agreement, and accuracy with guessing (in brackets). DS is significantly better at generating discriminative attribute phrases than SS.

Does pragmatics help? Given that our discerning speaker can generate highly accurate contrastive descriptions, we investigate if additional benefits can be achieved if the speaker jointly reasons about the listener’s ability to interpret the descriptions. We employ the *pragmatic speaker* model of Andreas and Klein [3] where a simple speaker generates descriptions that are reranked by a simple listener using a weighted combination of speaker and listener scores. In particular, we rerank the output 10 sentences from speakers by the probabilities from simple listeners. We combine the listener probability p_l and speaker beam-search probability p_s as $p = p_s^\lambda \cdot p_l^{(1-\lambda)}$, and pick the optimal λ on a validation set annotated by a human listener. We found that the optimal λ is close to 0, so we decided to use p_l only for reranking on test set.

In Table 4, we report the accuracy of top k sentences ($k = 1, 5, 7$) of the human listener and the results after reranking on the Test* set. When using the listener score from SL_r, the average accuracy of the top five generated descriptions after reranking improves dramatically from 64.2% to 82.6% for the simple speaker. The accuracy of the discerning speaker also improves to 90%. This suggests that better pragmatics can be achieved if both the speaker and listener are trained in a contrastive manner. Surprisingly the contrastively-trained simple listener SL is less effective at reranking than SL_r. We believe this is because the SL overfits on the human speaker descriptions and is less effective when used with neural speakers.

Figure 4 shows an example pair and the output of different speakers. Simple speaker suffers from generating descriptions that are true to the target image, but fail to differentiate two images. Discerning speaker can mostly avoid this mistake. Reranking by listeners can move better sentences to the top and improves the quality of top sentences.

Human listener accuracy (%)				
Top	Reranker listener			
	None	SL _r	SL	
SS	1	68.0 (77.0)	94.0 (96.0)	87.0 (92.0)
	5	64.2 (74.1)	82.6 (88.3)	80.8 (87.1)
	7	63.1 (72.8)	74.3 (82.0)	74.3 (82.4)
DS	1	82.0 (88.5)	95.0 (96.5)	95.0 (97.0)
	5	80.2 (86.7)	90.0 (93.3)	88.6 (92.8)
	7	79.1 (85.6)	86.7 (91.5)	86.1 (91.1)

Table 4. Accuracy of pragmatic speakers with human listeners on the Test* set. After generating the descriptions by the speaker model (either SS or DS), we use the listener model (SL_r or SL) to rerank them. We report the accuracy based on human listener from the user study. We report both the accuracy when there is majority agreement, and accuracy with guessing (in brackets). Pragmatic speakers are strictly better than non-pragmatic ones.

4.3. Fine-grained classification with attributes

We compare the effectiveness of attribute phrases to existing attributes in the OID dataset on the task of fine-grained classification on the FGVC aircraft dataset [30]. The OID dataset is designed with attributes in mind and has long-tail distribution over aircraft variants with 2728 models, while the FGVC dataset is designed for fine-grained classification task with 100 variants each with 100 images. Both datasets are based on the images from the airliners.net website and have a few overlapping images. We exclude the 169 images from the FGVC test set that appear in the OID training+validation set in our evaluation.

There are 49 attributes in the OID dataset organized into 14 categories. We exclude three attributes – two referring to the airline label and model, most of which have only one training examples per category, and another that is rare. We then trained linear classifiers to predict each attribute using the fc7 layer feature of the VGG-16 network. Using the same features and trained classifiers, we construct a 46 dimensional embedding of the FGVC images into the space of OID attributes. The attribute classifiers based on the VGG-16 network features are fairly accurate (66% mean AP across attributes) and outperforms the Fisher vector baseline included in the OID dataset paper.

For the attribute phrase embeddings, we first obtain the K most frequent ones in our training set. Given an image I, we compute the score $\phi(I)^T \theta(P)$ for each phrase P from a listener as the embedding. For a fair comparison the image features are kept identical to the OID attribute classifiers. We also explore an *opponent attribute space*, where instead of top phrases we consider the top phrase pairs. Phrase pairs represent an axis of comparison, e.g., “small vs. medium”, or “red and blue vs. red and white”, and are better suited for describing relative attributes. We use the discerning listener for the embedding on the opponent attribute space.


	SS:	SS + SLr:	DS:	DS + SLr:
	<ul style="list-style-type: none"> ✓ passenger plane ? white ✓ jet engine ? facing right ✓ commercial plane ? _UNK ? on the ground ✓ large ✓ large size ✓ on runway 	<ul style="list-style-type: none"> ✓ commercial plane ✓ large ✓ large size ✓ jet engine ✓ on runway ✓ passenger plane ? on the ground ? _UNK ? white ? facing right 	<ul style="list-style-type: none"> ✓ commercial plane ? facing right ✓ turbofan engine ✓ on concrete ✓ t tail ✓ jet engine ✓ twin engine ✓ multi seater ✓ white and red ✓ white colour with red stripes 	<ul style="list-style-type: none"> ✓ commercial plane ✓ jet engine ✓ turbofan engine ✓ twin engine ✓ on concrete ✓ multi seater ✓ t tail ✓ white and red ? facing right ✓ white colour with red stripes

Figure 4. An example output of various speakers. Given the image pair, we use SS and DS to generate descriptions of the top left image. Outputs from SS and DS are listed in the order of probabilities from speaker beam search. Outputs of SS+SL_r and DS+SL_r are reranked by SL_r. Green checks mean human listener picks correct image with certain, while question marks mean human listener is uncertain which image is referred to. The results indicate that DS is better than SS, and reranking using listeners improves the quality of top sentences.




Figure 5. Classification accuracy on FGVC aircraft dataset using the 46 dimensional OID attributes and varying number of attribute phrases. See Section 4.3 for details.

Figure 5 shows a comparison of OID attributes and attribute phrases for various listeners and number of attributes. For the same number of attributes as the OID dataset, attribute phrases are **12%** better. With 300 attributes the accuracy improves to **32%**, about **20%** better than OID. These results indicate that attribute phrases provide a better coverage of the space of discriminative directions. The two simple listeners perform equally well and the opponent attribute space does not offer any additional benefits.

4.4. Visualizing the space of descriptive attributes

We visualize the space of the 500 most frequent phrases in the training set using the embedding of the simple listener model projected from 1024 dimensions to 2 using t-SNE [43] in Figure 6. Various semantically related phrases are clustered into groups. The cluster on the top right reflects color combinations; Phrases such as “less windows” and “small plane” are nearby (bottom right). Visualizations of the learned embeddings of images $\phi(I)$ and opponent attribute phrases $\theta([P_1 \text{ vs. } P_2])$ are provided in the supplementary material.

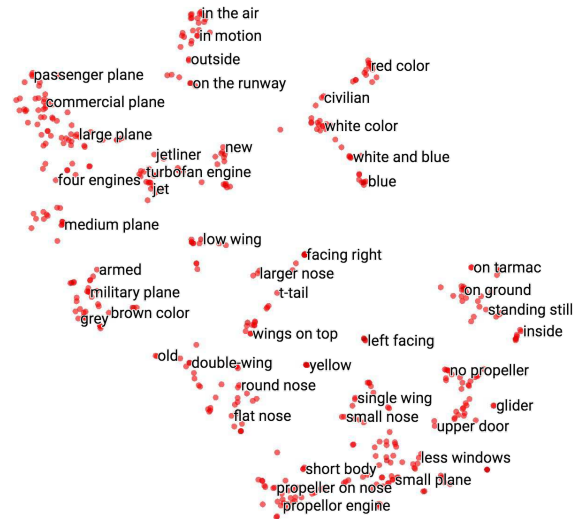


Figure 6. Visualization of the 500 most frequent descriptions. Each attribute is embedded into a 1024 dimensional space using the simple listener SL and projected into two dimensions using t-SNE [43]. (Best viewed digitally with zoom.)

4.5. Image retrieval with descriptive attributes

The listeners also allows us to retrieve an image given one or more attribute phrases. Given a phrase P we rank the images in the test set by the listener scores $\phi(I)^T \theta(P)$. Figure 7 shows some query phrases and the 18 most similar images retrieved from the test set. These results were obtained by simply concatenating all the query phrases to obtain a single phrase. More sophisticated schemes for combining scores from individual phrase predictions are likely to improve results [40]. Our model can retrieve images with multiple attribute phrases well even though the composition of phrases does not appear in the training set. For example, “red and blue” only shows five times in total of 47,000 phrases in the training set, “pointy nose” and “on the runway” are never seen in a single phrase together.



Figure 7. Top 18 images ranked by the listener for various attribute phrases as queries (shown on top). We rank the images by the scores from the simple listener on the concatenation of the attribute phrases. The images are ordered from top to bottom, left to right.



Figure 8. Top 10 discriminative attribute phrases for pairs of categories from FGVC aircraft dataset. Descriptions are generated by the discerning speaker for each pair of images in the first and second category. The phrases sorted by the occurrence frequency provides an attribute-based explanation of the visual difference between two categories.

4.6. Generating attribute explanations

The pairwise reasoning of a speaker can be extended to analyze an instance within a *set* by aggregating speaker utterances across all pairs that include the target. Similarly one can describe differences between two sets by considering all pairs of instances across the two sets. We use this to generate attribute-based explanations for visual differences between two categories. We select two categories A, B from FGVC aircraft dataset and randomly choose ten images from each category. For each image pair ($I_1 \in A, I_2 \in B$), we generate ten phrase pairs using our discerning speaker. We then sort unique phrases primarily by their image frequency (number of images from target category described by the given description minus that from the opposite category), and when tied secondarily by their phrase frequency (number of occurrences of the phrase in target category minus that in the opposite category.) The top ten attribute phrases for the two categories for an example pair of categories are shown in Figure 8. The algorithm reveals several discriminative attributes between two such as “engine under wings” for 747-400, and “stabilizer on top of tail” for ATR-42.

5. Conclusion

We analyzed attribute phrases that emerge when annotators describe visual differences between instances within a subordinate category (airplanes), and showed that speakers and listeners trained on this data can be used for various human-centric tasks such as text-based retrieval and attribute-based explanations of visual differences between unseen categories. Our experiments indicate that pragmatic speakers that combine listeners and speakers are effective on the reference game [3], and speakers trained on contrastive data offers significant additional benefits. We also showed that attribute phrases are modular and can be used to embed images into an interpretable semantic space. The resulting attribute phrases are highly discriminative and outperform existing attributes on FGVC aircraft dataset on the fine-grained classification task.

Acknowledgement: This research was supported in part by the NSF grants 1617917 and 1661259, and a faculty gift from Facebook. The experiments were performed using equipment obtained under a grant from the Collaborative R&D Fund managed by the Massachusetts Tech Collaborative and GPUs donated by NVIDIA.

References

- [1] M. Abadi, A. Agarwal, P. Barham, E. Brevdo, Z. Chen, C. Citro, G. S. Corrado, A. Davis, J. Dean, M. Devin, et al. Tensorflow: Large-scale machine learning on heterogeneous distributed systems. *arXiv preprint arXiv:1603.04467*, 2016. 4
- [2] Z. Akata, S. Reed, D. Walter, H. Lee, and B. Schiele. Evaluation of output embeddings for fine-grained image classification. In *Computer Vision and Pattern Recognition (CVPR)*, 2015. 2
- [3] J. Andreas and D. Klein. Reasoning About Pragmatics with Neural Listeners and Speakers. *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2016. 2, 4, 6, 8
- [4] S. Antol, A. Agrawal, J. Lu, M. Mitchell, D. Batra, C. L. Zitnick, and D. Parikh. VQA: Visual Question Answering. In *International Conference on Computer Vision (ICCV)*, 2015. 2
- [5] T. L. Berg, A. C. Berg, and J. Shih. Automatic attribute discovery and characterization from noisy web data. In *European Conference on Computer Vision (ECCV)*, 2010. 2
- [6] H. Chen, A. Gallagher, and B. Girod. Describing clothing by semantic attributes. In *European Conference on Computer Vision (ECCV)*, 2012. 2
- [7] A. Das, S. Kottur, K. Gupta, A. Singh, D. Yadav, J. M. F. Moura, D. Parikh, and D. Batra. Visual dialog. In *Computer Vision and Pattern Recognition (CVPR)*, 2017. 2
- [8] H. de Vries, F. Strub, S. Chandar, O. Pietquin, H. Larochelle, and A. Courville. Guesswhat?! visual object discovery through multi-modal dialogu. In *Computer Vision and Pattern Recognition (CVPR)*, 2017. 2
- [9] J. Devlin, S. Gupta, R. Girshick, M. Mitchell, and C. L. Zitnick. Exploring nearest neighbor approaches for image captioning. *arXiv preprint arXiv:1505.04467*, 2015. 2
- [10] J. Donahue, L. Anne Hendricks, S. Guadarrama, M. Rohrbach, S. Venugopalan, K. Saenko, and T. Darrell. Long-term recurrent convolutional networks for visual recognition and description. In *Computer Vision and Pattern Recognition (CVPR)*, 2015. 2
- [11] A. Farhadi, I. Endres, and D. Hoiem. Attribute-centric recognition for cross-category generalization. In *Computer Vision and Pattern Recognition (CVPR)*, 2010. 1, 2
- [12] A. Farhadi, I. Endres, D. Hoiem, and D. Forsyth. Describing objects by their attributes. In *Computer Vision and Pattern Recognition (CVPR)*, 2009. 1, 2
- [13] H. P. Grice. Logic and conversation. 1975. 2
- [14] S. Hochreiter and J. Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997. 3
- [15] R. Hu, H. Xu, M. Rohrbach, J. Feng, K. Saenko, and T. Darrell. Natural language object retrieval. *Computer Vision and Pattern Recognition (CVPR)*, 2016. 2
- [16] S. Ioffe and C. Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *International Conference on Machine Learning (ICML)*, 2015. 4
- [17] H. Izadinia, F. Sadeghi, S. K. Divvala, H. Hajishirzi, Y. Choi, and A. Farhadi. Segment-phrase table for semantic segmentation, visual entailment and paraphrasing. In *International Conference on Computer Vision (ICCV)*, 2015. 2
- [18] D. Jayaraman and K. Grauman. Zero-shot recognition with unreliable attributes. In *Neural Information Processing Systems (NIPS)*, 2014. 2
- [19] J. Johnson, B. Hariharan, L. van der Maaten, L. Fei-Fei, C. L. Zitnick, and R. Girshick. CLEVR: A diagnostic dataset for compositional language and elementary visual reasoning. In *Computer Vision and Pattern Recognition (CVPR)*, 2017. 2
- [20] S. Kazemzadeh, V. Ordonez, M. Matten, and T. L. Berg. ReferItGame: Referring to objects in photographs of natural scenes. In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2014. 2
- [21] D. Kingma and J. Ba. ADAM: A method for stochastic optimization. In *International Conference on Learning Representations (ICLR)*, 2015. 4
- [22] R. Kiros, R. Salakhutdinov, and R. S. Zemel. Unifying visual-semantic embeddings with multimodal neural language models. *Transactions of the Association for Computational Linguistics (TACL)*, 2015. 2
- [23] A. Kovashka, D. Parikh, and K. Grauman. WhittleSearch: Image search with relative attribute feedback. In *Computer Vision and Pattern Recognition (CVPR)*, 2012. 1, 2
- [24] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *Neural Information Processing Systems (NIPS)*, 2012. 4
- [25] C. H. Lampert, H. Nickisch, and S. Harmeling. Attribute-based classification for zero-shot visual object categorization. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 36(3):453–465, 2014. 1, 2
- [26] 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 *European Conference on Computer Vision (ECCV)*, 2014. 2
- [27] R. Luo and G. Shakhnarovich. Comprehension-guided referring expressions. In *Computer Vision and Pattern Recognition (CVPR)*, 2017. 2
- [28] S. Maji. Discovering a lexicon of parts and attributes. In *Workshop on Parts and Attributes, European Conference on Computer Vision (ECCV)*, 2012. 2, 3
- [29] S. Maji. A taxonomy of part and attribute discovery techniques. In *Visual Attributes*, pages 247–268. Springer International Publishing, 2017. 2
- [30] S. Maji, E. Rahtu, J. Kannala, M. Blaschko, and A. Vedaldi. Fine-grained visual classification of aircraft. *arXiv preprint arXiv:1306.5151*, 2013. 2, 3, 6
- [31] J. Mao, J. Huang, A. Toshev, O. Camburu, A. Yuille, and K. Murphy. Generation and comprehension of unambiguous object descriptions. *Computer Vision and Pattern Recognition (CVPR)*, 2016. 2
- [32] M. Mitchell, K. Van Deemter, and E. Reiter. Generating expressions that refer to visible objects. In *HLT-NAACL*, 2013. 2
- [33] V. K. Nagaraja, V. I. Morariu, and L. S. Davis. Modeling context between objects for referring expression understand-

- ing. In *European Conference on Computer Vision (ECCV)*, 2016. 2
- [34] V. Nair and G. E. Hinton. Rectified linear units improve restricted boltzmann machines. In *International Conference on Machine Learning (ICML)*, 2010. 3
- [35] D. Parikh and K. Grauman. Interactively building a discriminative vocabulary of nameable attributes. In *Computer Vision and Pattern Recognition (CVPR)*, 2011. 2
- [36] D. Parikh and K. Grauman. Relative attributes. In *International Conference on Computer Vision (ICCV)*, 2011. 1
- [37] S. Reed, Z. Akata, H. Lee, and B. Schiele. Learning deep representations of fine-grained visual descriptions. In *Computer Vision and Pattern Recognition (CVPR)*, 2016. 2
- [38] F. Sadeghi, S. K. Kumar Divvala, and A. Farhadi. VISKE: Visual knowledge extraction and question answering by visual verification of relation phrases. In *Computer Vision and Pattern Recognition (CVPR)*, 2015. 2
- [39] M. A. Sadeghi and A. Farhadi. Recognition using visual phrases. In *Computer Vision and Pattern Recognition (CVPR)*, 2011. 1, 3
- [40] W. J. Scheirer, N. Kumar, P. N. Belhumeur, and T. E. Boult. Multi-attribute spaces: Calibration for attribute fusion and similarity search. In *Computer Vision and Pattern Recognition (CVPR)*, 2012. 7
- [41] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. *CoRR*, abs/1409.1556, 2014. 3
- [42] N. Turakhia and D. Parikh. Attribute dominance: What pops out? In *Computer Vision and Pattern Recognition (CVPR)*, 2013. 5
- [43] L. van der Maaten and G. Hinton. Visualizing data using t-sne. *Journal of Machine Learning Research (JMLR)*, 9(Nov), 2008. 7
- [44] A. Vedaldi, S. Mahendran, S. Tsogkas, S. Maji, B. Girshick, J. Kannala, E. Rahtu, I. Kokkinos, M. B. Blaschko, D. Weiss, B. Taskar, K. Simonyan, N. Saphra, and S. Mohamed. Understanding objects in detail with fine-grained attributes. In *Computer Vision and Pattern Recognition (CVPR)*, 2014. 1, 3
- [45] R. Vedantam, S. Bengio, K. Murphy, D. Parikh, and G. Chechik. Context-aware captions from context-agnostic supervision. In *Computer Vision and Pattern Recognition (CVPR)*, 2017. 3
- [46] O. Vinyals, A. Toshev, S. Bengio, and D. Erhan. Show and tell: A neural image caption generator. In *Computer Vision and Pattern Recognition (CVPR)*, 2015. 2, 3
- [47] C. Wah, G. Van Horn, S. Branson, S. Maji, P. Perona, and S. Belongie. Similarity comparisons for interactive fine-grained categorization. In *Computer Vision and Pattern Recognition (CVPR)*, 2014. 1
- [48] L. Yu, P. Poirson, S. Yang, A. C. Berg, and T. L. Berg. Modeling context in referring expressions. In *European Conference on Computer Vision (ECCV)*, 2016. 2
- [49] L. Yu, H. Tan, M. Bansal, and T. L. Berg. A joint speaker-listener-reinforcer model for referring expressions. In *Computer Vision and Pattern Recognition (CVPR)*, 2017. 2
- [50] P. Zhang, Y. Goyal, D. Summers-Stay, D. Batra, and D. Parikh. Yin and Yang: Balancing and answering binary visual questions. In *Computer Vision and Pattern Recognition (CVPR)*, 2016. 2