

Abstract Visual Reasoning with Tangram Shapes

Anya Ji¹, Noriyuki Kojima^{1*}, Noah Rush^{1*}, Alane Suhr^{1,3*},
Wai Keen Vong², Robert D. Hawkins⁴, and Yoav Artzi¹

¹Cornell University ²New York University ³Allen Institute for AI ⁴Princeton University
{aj592, nk654}@cornell.edu noahjrush@gmail.com
waikeen.vong@nyu.edu suhr@cs.cornell.edu
rdhawkins@princeton.edu yoav@cs.cornell.edu

Abstract

We introduce KILOGRAM, a resource for studying abstract visual reasoning in humans and machines. Drawing on the history of tangram puzzles as stimuli in cognitive science, we build a richly annotated dataset that, with >1k distinct stimuli, is orders of magnitude larger and more diverse than prior resources. It is both visually and linguistically richer, moving beyond whole shape descriptions to include segmentation maps and part labels. We use this resource to evaluate the abstract visual reasoning capacities of recent multi-modal models. We observe that pre-trained weights demonstrate limited abstract reasoning, which dramatically improves with fine-tuning. We also observe that explicitly describing parts aids abstract reasoning for both humans and models, especially when jointly encoding the linguistic and visual inputs.

1 Introduction

Reference is a core function of natural language that relies on shared conventions and visual concepts. For example, in English, a speaker may use the term *dog* to refer to a particular animal of the species *canis familiaris*, or, through abstraction, to an object with a less strongly conventionalized name, such as the shape at the top of Figure 1. A speaker might refer to such a shape as looking like a *dog*, and even point to its parts, like its *head* and *tail*, despite having few visual features in common with the ordinary referent.

Comprehension and generation of references are critical for systems to engage in natural language interaction, and have been studied extensively with focus on ordinary references (e.g., Viethen and Dale, 2008; Mitchell et al., 2010; FitzGerald et al., 2013; Mao et al., 2016; Yu et al., 2016), in contrast to the visual abstraction illustrated in Figure 1.

*Equal contribution, alphabetically ordered.




Figure 1: Two example tangrams, each with two different annotations. Each annotation includes a whole-shape description (bold), segmentation to parts (in color), and naming of parts (linked to each part). The top example shows low variability with near-perfect agreement, while the bottom shows high variability with divergence of language and segmentation.

We address this gap by adopting an influential paradigm for probing human coordination in the cognitive science literature: reference games with abstract tangram shapes (e.g. Clark and Wilkes-Gibbs, 1986; Fox Tree, 1999; Hawkins et al., 2020).

Unlike photographs of natural objects, where there is often a single canonical label, tangrams are fundamentally ambiguous. While some shapes fall under strong existing conventions and elicit consensus about appropriate names (e.g., Figure 1, top), others are characterized by weaker conventions (e.g., Figure 1, bottom) and every speaker may arrive at a distinct but valid description (Zettersten and Lupyan, 2020; Hupet et al., 1991). While such diversity is a key consideration motivating their use as stimuli, existing behavioral studies have typically been limited to a relatively small set of 10–20 shapes, highly restricting the overall diversity of the stimulus class. It also limits their applicability for training and analyzing vision and language models, where significantly more data is necessary.

In this paper, we significantly expand this resource. We introduce KILOGRAM,¹ a large collec-

¹KILOGRAM is a portmanteau of *kilo* and *tangram*.

tion of tangrams with rich language annotations. KILOGRAM dramatically improves on existing resources along two dimensions. First, we curate and digitize 1,016 shapes, creating a set that is two orders of magnitude larger than collections used in existing work. This set dramatically increases coverage over the full range of naming variability, providing a more comprehensive view of human naming behavior. Second, rather than treating each tangram as a single whole shape, our images are vector graphics constructed from the original component puzzle pieces. This decomposition enables reasoning about both whole shapes and their parts.

We use this new collection of digitized tangram shapes to collect a large dataset of textual descriptions, reflecting a high diversity of naming behaviors. While existing work has focused on naming the complete shape, we also ask participants to segment and name semantically meaningful parts. We use crowdsourcing to scale our annotation process, collecting multiple annotations for each shape, thereby representing the distribution of annotations it elicits, rather than a single sample. In total, we collect 13,404 annotations, each describing a complete object and its segmented parts.

The potential of KILOGRAM is broad. For example, it enables the data-driven scaling of studies of human interactions and models of whole-part reasoning in language and vision models. In this paper, we use KILOGRAM to evaluate the visual reasoning capacities of recent pre-trained multi-modal models, focusing on generalizing concepts to abstract shapes. We observe limited generalization of this type in pre-trained models, but significant improvements following fine-tuning with our data. We also see how explicitly referring to and visualizing parts can help reference resolution. Data and code, as well as a data viewer are available at: <https://lil.nlp.cornell.edu/kilogram/>.

2 Background and Related Work

Abstract or ambiguous visual stimuli have been widely used to investigate how human partners coordinate when talking about things in the absence of strong naming conventions going back to Krauss and Weinheimer (1964). Tangrams as stimuli were introduced by Clark and Wilkes-Gibbs (1986). These shapes are all built from the same seven primitives, but elicit a wide range of figurative descriptions that conceptualize shapes in different ways (Schober and Clark, 1989; Hor-

ton and Gerrig, 2002; Duff et al., 2006; Holler and Wilkin, 2011; Horton and Slaten, 2012; Ibarra and Tanenhaus, 2016; Shore et al., 2018; Atkinson et al., 2019; Castillo et al., 2019; Bangerter et al., 2020). It has been observed that some shapes are easier or harder to describe (Hupet et al., 1991; Zettersten and Lupyan, 2020; Brashears and Minda, 2020), a property known as *nameability* or *codability*, which has also been studied with non-tangram shapes (e.g., line drawings; Snodgrass and Vanderwart, 1980; Cykowicz et al., 1997; Duñabeitia et al., 2018). Even though diversity is a key consideration in working with tangrams, existing stimuli sets are relatively small, limiting their usefulness as NLP benchmarks, where scale is critical. Even the largest studies of variability in naming (e.g., Murfitt and McAllister, 2001) have used a relatively small set of 60 tangrams. Fasquel et al. (2022) present a resource that is related and complementary to ours, including 332 PNG-formatted tangrams with whole-shape naming annotations in French.

Contemporary pre-trained vision and language approaches can be categorized along an axis characterizing how they encode the data, from jointly encoding the two inputs (Lu et al., 2019; Chen et al., 2020; Kim et al., 2021) to encoding them separately (Radford et al., 2021; Jia et al., 2021). Joint encoding aims to capture tighter interaction between the input modalities compared to separate encoding, but is generally more computationally expensive, and can only operate on multi-modal input. We study recent models on both ends: ViLT (Kim et al., 2021) for joint encoding and CLIP (Radford et al., 2021) for separate encoding.

These models are typically evaluated on image captioning (e.g., Chen et al., 2015) or visual question answering (e.g., Antol et al., 2015) benchmarks. Several benchmarks, such as NLVR (Suhr et al., 2017, 2019) and Winoground (Thrush et al., 2022), aim for more focused evaluations with a focus on compositionality. We build on these efforts, but target generalization through abstraction using visually ambiguous stimuli. This is inspired by the role of abstraction in human cognition. Abstraction is a key step in human perception (Biederman, 1987) that is critical for generalization (Gentner and Markman, 1997; Medin et al., 1993; Shepard, 1987), and forms the shared foundation on which human language communication is layered (Lupyan and Winter, 2018; McCarthy et al., 2021; Wong et al., 2022). Our focus on part de-

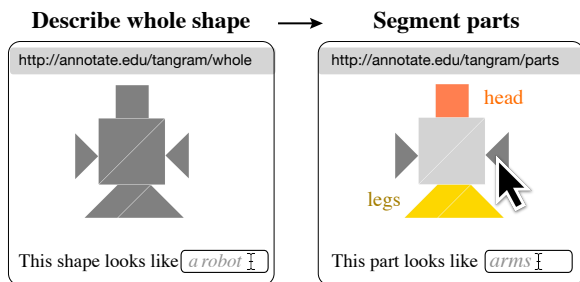


Figure 2: The two phases of our annotation task.

composition is aligned with how part identification plays an important role in human abstraction (Tversky and Hemenway, 1984).

3 Data Collection

We scan a large set of tangram puzzles to vector graphics, and crowdsource annotations of natural language descriptions and part segmentations.

3.1 Collecting Tangram Puzzles

Tangram puzzles are made of seven primitive shapes (Elffers, 1977), which can be combined in a large variety of configurations evoking different concepts. We scan 1,004 tangrams depicting a broad set of concepts to vector graphic SVGs from Slocum (2003). Appendix A.1 shows example tangrams, Appendix A.2 details on our process.² We also manually add 12 tangrams commonly used in previous studies (Hawkins et al., 2020).

3.2 Whole-Part Annotation

We design a two-stage crowdsourcing task to elicit natural language English descriptions for each tangram, both of the whole shape and of its parts (Figure 2). First, in the *whole-shape description* stage, the worker is shown a tangram image in grayscale and asked to complete the prompt “This shape, as a whole, looks like ____.” In the *part annotation* stage, the worker is asked to select one or more puzzle pieces, and complete the prompt “The part(s) you selected look(s) like ____.” These pieces are then colored and the annotation appears in the corresponding color. The annotator can delete annotations, annotate a part as UNKNOWN when they are not sure about its semantics, and add pieces to existing

²The scanned documents are authorized for use for educational and research purposes (“fair use”) as per U.S. copyright law (Title 17, §108, United States Code). This use does not require permission or usage fees, including for publication. Copyright and use agreement are attached to the data.

Mean Description Length	
Whole-shape description	2.28±1.62
Part description	1.31±0.77
Vocabulary Size	
Whole-shape description	3,031
Part description	3,110
Overall	4,522
Part Segmentation	
Mean parts per shape	3.63±1.28
Mean pieces per part	1.93±1.20

Table 1: Data statistics for the complete dataset.

parts. All pieces must be annotated to submit the task, yielding a complete segmentation map.

We use Amazon Mechanical Turk for data collection. Workers are required to be located in the United States with at least a 98% HIT acceptance rate, must pass a qualification task, and complete a survey about their language proficiency (see Appendix A.3 for further details). To prevent a small group of workers from dominating the data, each annotator is only allowed to annotate each tangram once, and cannot annotate more than 200 distinct tangrams. Workers are paid 0.14 USD per task.³

We first collect 10,053 annotations for the 1,004 scanned tangrams, at least 10 annotations for each tangram (mean=10.01). Following this stage of annotation, we collect additional annotations for a subset of the tangrams to create a set with denser language and part segmentation annotation. We sample 62 tangrams to be representative of the different levels of diversity in annotations we observe in the initially collected data. Appendix A.4 describes the sampling procedure. We also add the 12 tangrams from previous studies for a total of 74 tangrams for dense annotation. We conduct additional annotation tasks to have at least 50 annotations for each of the 74 tangrams selected for dense annotation (mean=53.66).⁴ The dense annotation gives us a better estimate of the distribution of language for the 74 selected tangrams, for example to use as reference texts in generation tasks.

In total, we collect 13,404 annotations for 1,016 tangrams at a total cost of 2,172.94 USD. We lowercase and stem to compute vocabulary size, and tokenize on white spaces to compute description length. Table 1 shows basic data statistics. A total

³We set this rate aiming for an hourly rate of 12–15 USD for workers familiar with the task.

⁴This includes collecting data from scratch for the 12 commonly used tangrams that were added following the initial collection.

	FULL	DENSE	DENSE10
SND	0.91 ± 0.11	0.93 ± 0.06	0.90 ± 0.15
PND	0.76 ± 0.19	0.79 ± 0.15	0.73 ± 0.20
PSA	5.30 ± 0.62	5.09 ± 0.53	5.34 ± 0.77

Table 2: Mean and standard deviation of our analysis measures on the three sets.

of 297 MTurk workers participate in the annotation, with 98.0% of the workers speaking English as their first language. Those who do not speak English as their first language still rate their English proficiency level as native or close to native. 1.0% of the workers speak more than one language, among which the most common are Spanish, German, Japanese, and Chinese.

3.3 Standard Data Splits

We split the dataset for analysis and learning experiments. For analysis, we create two overlapping sets: FULL and DENSE. FULL includes 1,016 tangrams, each with 10–11 annotations (mean=10.11). It includes the 10,053 annotations initially collected for the scanned 1,004 tangrams. For the 12 commonly used tangrams, we sample 10 annotations from the later collection effort. DENSE includes all annotations for the 74 densely annotated tangrams, with at least 50, and 53.66 on average annotations per tangram. We also define the set DENSE10 to include only the annotations from the sparse set for the densely annotated tangrams. For learning experiments, we split according to tangrams to create training (692 tangrams), development (125), test (125), and test-dense sets (74). All densely annotated tangrams are in test-dense. The other three sets are split randomly.

4 Data Analysis

The language and concepts annotators use reflect varying degrees of consensus around conventions for describing the appearance of shapes and their parts. For analysis, we preprocess the annotations by lowercasing, tokenizing, lemmatizing, and removing stop words using NLTK (Bird, 2004). We use the larger FULL set for our analyses (Section 3.3), unless otherwise noted.

For a broad overview of the types of concepts evoked, we manually tag 250 randomly sampled annotations: 30.8% use human-like concepts (e.g., *dancer*), 31.2% animate but non-human concepts (e.g., *dog*), and 38.0% non-animate concepts (e.g.,

house). We examine how part words differ across whole-shape concepts by extracting head words from whole-shape and part descriptions. Figure 3 shows the distribution of part head words for each of 272 whole-shape head words with >10 occurrences, ranked in order of frequency. Figure A.2 in the appendix illustrates how the most common part word *head* is used in different tangrams.

A central problem of visual abstraction is the degree of ambiguity or subjectivity that a shape evokes across different people (Murthy et al., 2022): some descriptions have higher consensus than others. We define three measures of variability along different dimensions: shape naming divergence (SND), part naming divergence (PND), and part segmentation agreement (PSA). Table 2 lists the mean and standard deviation for these three measures over the sparsely and densely annotated data.

Shape Naming Divergence (SND) A tangram’s SND quantifies the variability among whole-shape annotations. SND is an operationalization of nameability, a criteria that is commonly used to measure how consistent is naming of an object across individuals (e.g., Zettersten and Lupyan, 2020).

Formally, a whole-shape annotation is a sequence of M tokens $\bar{x} = \langle x_1, \dots, x_M \rangle$. Given a tangram with N annotations $\bar{x}^{(j)}$, $j = 1, \dots, N$, each of length $M^{(j)}$, we define $w_i^{(j)}$ for each token $x_i^{(j)}$ in annotation $\bar{x}^{(j)}$ as the proportion of other annotations of that tangram that do not contain $x_i^{(j)}$:

$$w_i^{(j)} = \frac{1}{N-1} \sum_{j'=1}^N \mathbb{1}[x_i^{(j)} \notin \bar{x}^{j'}] , \quad (1)$$

where $\mathbb{1}$ is an indicator function. The divergence of annotation $\bar{x}^{(j)}$ is $W^{(j)} = \frac{1}{M^{(j)}} \sum_{i=0}^k w_i^{(j)}$. The divergence of a tangram is $W = \frac{1}{N} \sum_{j=0}^N W^{(j)}$. For example, the SNDs of the tangrams in Figure 1 computed only with the two annotations displayed are 0.00 (top) and 1.00 (bottom).

Mean SND is relatively high in our data, with 0.91 on FULL (Table 2). We observe relatively similar values for DENSE and DENSE10, albeit with lower standard deviation for DENSE, as expected with more annotations. Annotators often use words that are unique to their annotation. We observe perfect consensus for only one tangram, and mostly similar annotations with relatively few deviations for a few others. Figure 5 shows several examples.

Part Naming Divergence (PND) SND measures annotation divergence for part name annotations




Figure 3: Part distributions for different head words. Whole-shape head words (shown in descending order of frequency from left) elicit a variety of part head word distributions. Colors are randomly assigned to part head words, but are fixed across all bars. Grey indicates part head words with < 0.005 frequency.




Figure 4: Per tangram SND, PND, and PSA mean values and 95% confidence interval. Tangrams are ordered along the x -axis in ascending order according to the plotted measure. Values are calculated by bootstrapping with 1,000 resamplings. In the FULL plots, the 74 densely annotated tangrams are colored red.

collected in the second step of the annotation task. PND is computed identically to SND, but with the concatenation of all part names of an annotation as the input text \bar{x} . For example, the PNDs of the two tangrams in Figure 1 computed with only the two annotations displayed are 0.19 (top) and 1.00 (bottom). In general, part descriptions are more similar than whole-shape descriptions with mean PND of 0.76 (Table 2).

Part Segmentation Agreement (PSA) Annotators segment the tangrams into parts by grouping the tangram puzzle pieces. PSA quantifies the agreement between part segmentations as the maximum number of pieces that does not need to be moved to another group in order to edit one segmentation to another. We compute PSA as a linear

sum assignment problem with maximum weight matching. For each pair of segmentations, we create a cost matrix, where the number of rows is the number of parts in one annotation and the number of columns is the number of parts in the second annotation. The value of each matrix element is the number of matching puzzle pieces between the two corresponding parts in the two annotations. The tangram PSA is the mean of costs for all annotation pairs. For example, the PSAs of the two tangrams in Figure 1 computed with only the two annotations displayed are 6.00 (top) and 3.00 (bottom).

The mean PSA in our data is 5.30 (Table 2), with an approximately normal distribution of values. Some tangrams have strong segmentation cues, such that annotators reach perfect consensus, while others elicit significant segmentation disagreement.

Dense Annotations The comparison of FULL, DENSE, and DENSE10 illustrates how well our data approximates the real distribution of annotations for each tangram, and the advantage of DENSE. Figure 4 shows the complete distribution of values. Comparing DENSE10 and DENSE, the rankings of the tangrams are largely the same with the additional annotations: for SND, Spearman’s rank correlation coefficient is $r(72) = .78$, $p \ll .001$; for PND, $r(72) = .87$, $p \ll .001$; for PSA, $r(72) = .76$, $p \ll .001$. The tangrams sampled for DENSE represent well the distribution of tangrams along the different measures, as illustrated by the red highlights in Figure 4.

Inter-measure Correlations Figure 5 illustrates the correlations between the three measures. The divergences of the two types of language annotations, whole-shape and part descriptions, show moderate positive correlation $r(1014) = .531$, $p \ll .001$. This indicates that tangrams that are annotated with similar whole-shape descriptions are often annotated with similar part descriptions.

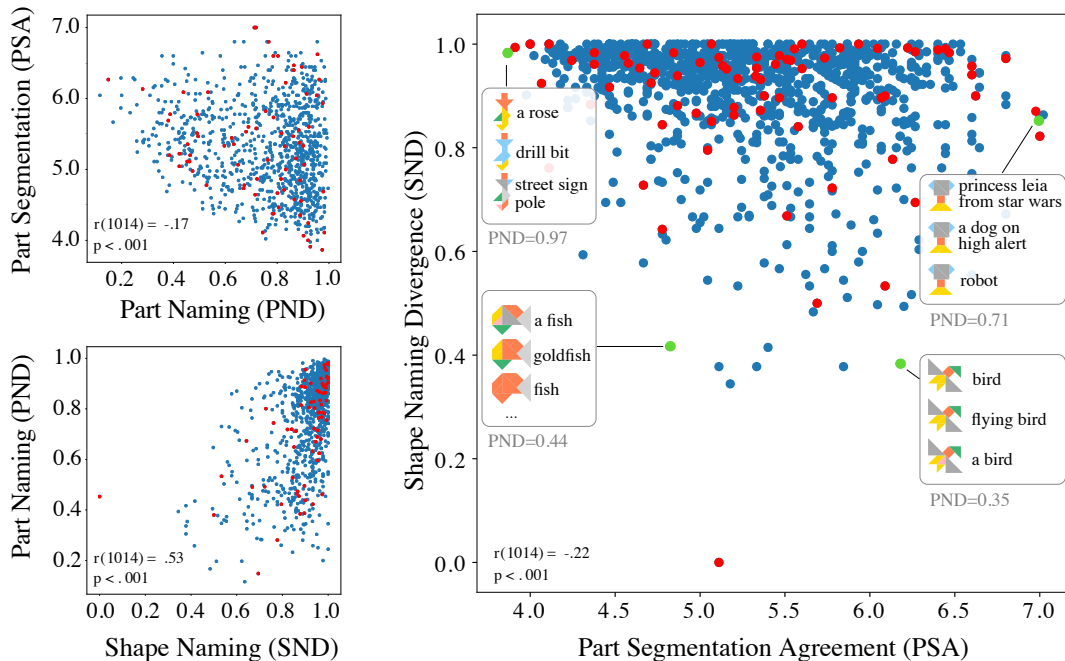


Figure 5: SND, PND, and PSA correlations computed over the FULL set. Representative examples of different SND and PSA values are illustrated on the right. Densely annotated examples are highlighted in red.

Nevertheless, many tangrams with similar whole-shape descriptions have diverse part descriptions. The correlations between language annotation divergence and PSA are lower, $r(1014) = -.216$, $p \ll .001$ for SND and PSA and $r(1014) = -.165$, $p \ll .001$ for PND and PSA.

5 Visual Reasoning with Tangrams

We use KILOGRAM to evaluate the reasoning of CLIP (Radford et al., 2021) and ViLT (Kim et al., 2021) through a reference game task, where the model is given a textual description and selects the corresponding image from a set of images. Formally, given a textual description \bar{x} and a set of k images $\mathcal{I} = \{I_1, \dots, I_k\}$, the task is to select the image $I_i \in \mathcal{I}$ corresponding to \bar{x} . We cast the task as computing a similarity score $f(\bar{x}, I_i)$ between the description \bar{x} and an image I_i . We select the corresponding image as $I^* = \arg \max_{I_i \in \mathcal{I}} f(\bar{x}, I_i)$.

5.1 Reference Game Generation

We randomly generate reference games for an annotated text-image pair (\bar{x}, I) by sampling additional $k - 1$ images from data under several constraints. We do not include repeating images in the set of k images or images that have identical whole-shape text annotations. This avoids obvious ambiguity that is impossible to resolve in the target selection. We also require all images to be annotated with the

same number of parts. This reduces the chance of the model relying on simple part counting to discriminate between target images when including parts in the text (condition PARTS below). Appendix A.8 shows the impact of these constraints through analyzing experiments not using them.

5.2 Models

We instantiate f using CLIP or ViLT, two models based on the Transformer architecture (Vaswani et al., 2017). We provide a brief review of the models, and refer the reader to the respective papers for further details.

CLIP uses two separate encoders to generate separate fixed-dimension representations of the text and images. It uses contrastive pre-training with a symmetric cross entropy loss on a large amount of aligned, but noisy web image-text data. We implement the scoring function f with CLIP by encoding the text \bar{x} and all images $I \in \mathcal{I}$ separately, and then computing the dot-product similarity score of the text with each image. This is identical to the CLIP pre-training objective, which potentially makes CLIP suitable for our task out of the box.

ViLT uses a single encoder that takes as input both the text and image inputs together. ViLT pre-training also uses aligned image-text data, but from existing benchmarks (Lin et al., 2014; Krishna et al., 2016; Ordonez et al., 2011; Sharma



Figure 6: Illustration of the language and vision modalities under the different experimental conditions.

et al., 2018). It is pre-trained using multiple self-supervised objectives, including image-text matching via a binary classification head, which is suitable for our task out of the box. We implement f using this classification head. Given a text \bar{x} and an image $I \in \mathcal{I}$, we compute their similarity using the matching classification head.

5.3 Experimental Conditions

We study several input variants. Figure 6 illustrates the modalities under the different conditions, and Appendix A.5 shows complete example inputs. For the textual description \bar{x} , we experiment with including the whole-shape description only (WHOLE) or adding part names (PARTS) by combining with the whole-shape description using the template *<whole shape> with <part>, <part>, ..., and <part>*. This tests the ability of models to benefit from part names. We consider two image I conditions: coloring all parts with the same color (BLACK) or coloring parts differently (COLOR). The color choice in COLOR corresponds to the position of the part name in \bar{x} , when the text includes part names (PARTS).

We experiment with the original pre-trained model weights, and with contrastive fine-tuning on our data using a symmetric cross entropy loss (Radford et al., 2021). During fine-tuning only, we consider a data augmentation condition (AUG), where we augment the data by creating examples that include only a subset of the part names in the text and coloring only the parts corresponding to the included parts names in the image, while all other parts remain black. We generate partial part examples for all possible subsets of parts for each example. Appendix A.5 illustrates the generated examples. When generating reference games for the augmented data, we constrain all the examples within a reference game to have the same number of parts in their full annotation, otherwise the task could be solved by counting parts. Part names are shuffled when creating the augmented data, and part colors correspond to the sequential position of the part name in the templated text.

5.4 Implementation Details

We set the size of the reference game context to $k = 10$ throughout our experiments. During contrastive fine-tuning, we create a text-image matching matrix of size $k \times k$ for each generated reference game in our training data by randomly selecting a text description for each tangram distractor from its annotations. We compute matching loss in both directions, from text to images and vice versa. In practice, this is equivalent to creating $2k$ reference games in both directions, and provides more informative updates. For all experiments, we use an ensemble of three models combined by element-wise multiplication of their outputs. Appendix A.7 provides model-specific implementation details. Appendix A.9 provides a reproducibility list.

5.5 Estimating Human Performance

We conduct an initial estimation of expected human performance on the same evaluation task by recruiting an independent group of 217 human participants. Each participant is randomly assigned to one of the four conditions and shown a random sequence of 20 trials from that condition, preventing leakage across conditions. On each trial, we present an annotation from our development set along with the corresponding context of ten tangrams and ask the participant to click the tangram that was being described. We randomly sample one referential context per annotation, which provides coverage over all 125 tangrams and over 600 unique descriptions in each condition. Before the actual test trials, each participant is provided with a fixed set of 10 practice trials with feedback indicating whether they have selected the correct tangram, and if not, we highlight the correct answer. Performance in the practice trials is not considered in our analysis. Appendix A.6 provides further details.

5.6 Results and Analysis

Table 3 shows development and test reference game accuracies under different experimental setups, including for human studies. Figure 7 shows the accuracy distribution for human participants.

Condition	CLIP		ViLT		Human
	PT	FT	PT	FT	
Development Results					
WHOLE+BLACK	16.1	43.3	12.9	40.9	47.7
PARTS+BLACK	16.4	45.3	12.5	45.7	49.1
WHOLE+COLOR	15.9	40.8	11.7	41.0	49.5
PARTS+COLOR	15.0	45.4	10.7	75.2	63.0
PARTS+COLOR+AUG	-	47.6	-	72.2	
Held-out Test Results					
WHOLE+BLACK	17.9	42.5	13.1	44.5	
PARTS+BLACK	18.6	45.8	13.3	50.3	
WHOLE+COLOR	18.1	41.4	12.8	44.8	
PARTS+COLOR	17.0	46.5	11.7	77.3	
PARTS+COLOR+AUG	-	50.2	-	74.4	

Table 3: Reference game accuracies (%) for the different experimental conditions with pre-trained (PT) or fine-tuned (FT) models, as well as for human subjects.

While both models perform better than a random baseline (10%) out of the box, we generally observe poor performance with the pre-trained weights (PT). CLIP slightly outperforms ViLT throughout, potentially because it is trained with a contrastive objective similar to a reference game. Whereas ViLT’s matching loss is aligned with our goal, it is only one of several losses in its objective. We observe no reliable improvement from adding part information, either textual or visual. The low performance on WHOLE+BLACK indicates the models fail to generalize familiar concepts to abstract shapes and the lack of consistent improvement with part information indicates an inability to reason about the correspondence of text and colored parts.

Fine-tuning (FT) dramatically improves performance for both models. Adding part names to the text description improves both models (PARTS+BLACK). However, segmentation information in the form of part coloring without part names (WHOLE+COLOR) shows no benefit. Although ViLT does not benefit from color information alone, the combination with part names (PARTS+COLOR) shows significant added improvement in performance over having access to part information in one of the modalities. Overall, we observe small consistent differences in performance between the two models, except when having access to both part names and colors (PARTS+COLOR), which ViLT effectively uses following fine tuning. This may be because ViLT’s tight integration of the modalities in its single encoder allows it to take advantage of the part correspondence information provided




Figure 7: The distribution of each human participant’s mean accuracy in the four conditions. The white dashed lines are the estimated means of a two-component Gaussian mixture model.

when both part names and colors are given.

Human performance follows a similar trend to the fine-tuned models: adding part names and segmentation helps performance, and their benefit is most pronounced when both are provided. Human performance is significantly higher than pre-trained (PT) models across all four conditions. Fine-tuning (FT) closes this gap. Indeed, in the PARTS+COLOR condition, ViLT significantly outperforms mean human performance. To better analyze human results, we fit a two-component Gaussian mixture model to the distribution of individual participants’ accuracies (Figure 7). We observe two components for all conditions except WHOLE+BLACK, indicating two distinct sub-populations. For example, for PARTS+COLOR, the low-performing sub-population has a mean accuracy of 52.5%, while the high-performing has a mean of 83.8%, significantly outperforming the fine-tuned ViLT. It is possible that the lower-performance sub-population is not making full use of the additional information.

Data augmentation (AUG) improves performance for CLIP, but not for ViLT, which even shows a small decrease in performance, although still significantly outperforming CLIP. We hypothesize that the presence of training examples with partial part information complicates resolving the correspondence between parts and their name, resulting in overall lower ViLT performance. We leave further study of this hypothesis for future work.

The augmentation condition fine-tunes the models to handle examples with partial part information, and allows to study the impact of gradually




Figure 8: Mean probability assigned to the correct image using fine-tuned CLIP (left) or fine-tuned ViLT (right) on the development set, by number of parts included in text and colored in the images. Curves are separated by total number of parts in the annotation of the target example. Error bands are bootstrapped 95% confidence intervals.

adding part information. We apply the augmentation process to the development data to generate the data for this analysis. Figure 8 shows the effect of gradually adding part information on the probability of the correct prediction, separated by the total number of parts in the example. Overall, part information is beneficial, but with a diminishing return as more part information is added. We observe this for both models, but with a much faster rate for CLIP, which overall shows much lower performance. ViLT is able to benefit from increasing part information, with the benefit diminishing only after four parts are provided.

6 Discussion

KILOGRAM provides a new window into the visual abstraction capacity of grounded language models and their ability to generalize concepts beyond their photographic appearance, an integral component of human concept representations (Fan et al., 2015). Our experiments show that there is significant room to improve pre-trained models, which should be able to perform zero-shot reference game tasks without fine-tuning as well as humans do (Clark and Wilkes-Gibbs, 1986). The improved performance after fine-tuning indicates the multi-modal architecture itself has the potential for higher performance, which current pre-training regimes likely do not support. In particular, ViLT’s improved performance as a function of additional part information suggests that more structured concept alignment may play a role in this effort (e.g., between parts expressed as lexical items and the corresponding elements of the image).

While we focused on the task of reference resolution, KILOGRAM is also well-suited for production tasks (e.g., generating human-like distributions of

descriptions or coloring named parts on a blank tangram) as well as instruction-following tasks (e.g., placing pieces in the described configuration to reconstruct a tangram). More broadly, our data emphasizes the need for maintaining well-calibrated distributions over the many different possible ways that people may conceptualize or talk about things, rather than collapsing to a “best” prediction.

7 Limitations

Although randomly constructed reference games provide an interpretable evaluation metric, they also pose several limitations. Performance is limited by the fact that descriptions were elicited for isolated images. These descriptions do not reflect the kind of pragmatic reasoning commonly deployed by human speakers in reference games to resolve ambiguities (Goodman and Frank, 2016). In other words, annotators were not able to anticipate the necessary level of detail to disambiguate the object from a specific context of distractors, hence the descriptions may be underinformative. Randomly generated reference games may include ambiguities that make them impossible to solve (e.g., two objects that could both plausibly be described as a *bird*). The possible performance ceiling on these games is likely below 100%. Extending the data through interactive reference games is an important direction for future work. Likewise, our studies of baseline human performance on this task are preliminary. We found that participants clustered into higher- and lower-performing groups, likely reflecting attentional and motivational factors (e.g., some participants may not have fully attended to the provided part information). A better understanding of human behavior is critical before making any clear conclusions comparing humans and

model performance. Ultimately, models only outperformed mean human performance significantly only after fine-tuning on approximately 6,600 example reference games.

Our resource contribution and analysis are focused on English. While the data collection design does not make language-specific assumptions, it depends on the availability of proficient speakers, which is limited in contemporary crowdsourcing services for certain languages. Our large collection of visual stimuli is well suited to extend our data collection to other languages and cultures, which may display different abstractions. This is an important direction for future work. Extending our analysis to other languages depends on the availability of pre-trained models in these languages, which may be limited by the availability of aligned language vision data and the computational resources required for pre-training.

Acknowledgements

This research was supported by ARO W911NF21-1-0106, NSF under grant No. 1750499, and a gift from Open Philanthropy. NK is supported by Masason Fellowship, AS by a Facebook PhD Fellowship and an NSF GRF under grant No. 1650441, and RDH by a CV Starr Fellowship. We thank Rob Goldstone, Judith Fan, Cathy Wong, and the anonymous reviewers for their helpful comments and suggestions. We are grateful for the contributions of the workers on Mechanical Turk.

References

- Stanislaw Antol, Aishwarya Agrawal, Jiasen Lu, Margaret Mitchell, Dhruv Batra, C. Lawrence Zitnick, and Devi Parikh. 2015. [VQA: Visual question answering](#). In *IEEE International Conference on Computer Vision*, pages 2425–2433.
- Mark Atkinson, Gregory J Mills, and Kenny Smith. 2019. Social group effects on the emergence of communicative conventions and language complexity. *Journal of Language Evolution*, 4(1):1–18.
- Adrian Bangerter, Eric Mayor, and Dominique Knutsen. 2020. Lexical entrainment without conceptual pacts? revisiting the matching task. *Journal of Memory and Language*, 114:104129.
- Irving Biederman. 1987. Recognition-by-components: a theory of human image understanding. *Psychological review*, 94 2:115–147.
- Steven Bird. 2004. Nltk: The natural language toolkit. *ArXiv*, cs.CL/0205028.
- G. Bradski. 2000. The OpenCV Library. *Dr. Dobb's Journal of Software Tools*.
- Bailey Brashears and John Paul Minda. 2020. The effects of feature verbalizability on category learning. In *Proceedings of the 42nd Conference of the Cognitive Science Society*.
- Lucía Castillo, Kenny Smith, and Holly P Branigan. 2019. Interaction promotes the adaptation of referential conventions to the communicative context. *Cognitive science*, 43(8):e12780.
- Xinlei Chen, Hao Fang, Tsung-Yi Lin, Ramakrishna Vedantam, Saurabh Gupta, Piotr Dollár, and C. Lawrence Zitnick. 2015. Microsoft COCO captions: Data collection and evaluation server. *CoRR*, abs/1504.00325.
- Yen-Chun Chen, Linjie Li, Licheng Yu, Ahmed El Kholy, Faisal Ahmed, Zhe Gan, Yu Cheng, and Jingjing Liu. 2020. Uniter: Universal image-text representation learning. In *European conference on computer vision*, pages 104–120.
- Herbert H Clark and Deanna Wilkes-Gibbs. 1986. Referring as a collaborative process. *Cognition*, 22(1):1–39.
- Yael M. Cychowicz, D Friedman, M Rothstein, and Joan Gay Snodgrass. 1997. Picture naming by young children: norms for name agreement, familiarity, and visual complexity. *Journal of experimental child psychology*, 65 2:171–237.
- Melissa C Duff, Julie Hengst, Daniel Tranel, and Neal J Cohen. 2006. Development of shared information in communication despite hippocampal amnesia. *Nature neuroscience*, 9(1):140–146.
- Jon Andoni Duñabeitia, Davide Crepaldi, Antje S. Meyer, Boris New, Christos Pliatsikas, Eva Smolka, and Marc Brysbaert. 2018. Multipic: A standardized set of 750 drawings with norms for six european languages. *Quarterly Journal of Experimental Psychology*, 71:808 – 816.
- Joost Elffers. 1977. *Tangram: The Ancient Chinese Puzzle*. Penguin Books.
- Judith E. Fan, Daniel Yamins, and Nicholas B. Turk-Browne. 2015. Common object representations for visual recognition and production. *Cognitive Science*.
- Alicia Fasquel, Angèle Brunellière, and Dominique Knutsen. 2022. A modified procedure for naming 332 pictures and collecting norms: Using tangram pictures in psycholinguistic studies. *Behavior research methods*.
- Nicholas FitzGerald, Yoav Artzi, and Luke Zettlemoyer. 2013. [Learning distributions over logical forms for referring expression generation](#). In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, pages 1914–1925.

- Jean E Fox Tree. 1999. Listening in on monologues and dialogues. *Discourse processes*, 27(1):35–53.
- Dedre Gentner and Arthur B. Markman. 1997. Structure mapping in analogy and similarity. *American Psychologist*, 52:45–56.
- Noah D. Goodman and Michael C. Frank. 2016. Pragmatic language interpretation as probabilistic inference. *Trends in Cognitive Sciences*, 20:818–829.
- Chris Harris, Mike Stephens, et al. 1988. A combined corner and edge detector. In *Alvey vision conference*, volume 15, pages 10–5244. Citeseer.
- Robert D. Hawkins, Michael C. Frank, and Noah D. Goodman. 2020. [Characterizing the dynamics of learning in repeated reference games](#). *Cognitive science*, 44(6):e12845.
- Judith Holler and Katie Wilkin. 2011. Co-speech gesture mimicry in the process of collaborative referring during face-to-face dialogue. *Journal of Nonverbal Behavior*, 35(2):133–153.
- William S Horton and Richard J Gerrig. 2002. Speakers’ experiences and audience design: Knowing when and knowing how to adjust utterances to addressees. *Journal of Memory and Language*, 47(4):589–606.
- William S Horton and Daniel G Slaten. 2012. Anticipating who will say what: The influence of speaker-specific memory associations on reference resolution. *Memory & cognition*, 40(1):113–126.
- Michel Hupet, Xavier Seron, and Yves Chantraine. 1991. The effects of the codability and discriminability of the referents on the collaborative referring procedure. *British Journal of Psychology*, 82(4):449–462.
- Alyssa Ibarra and Michael K Tanenhaus. 2016. The flexibility of conceptual pacts: Referring expressions dynamically shift to accommodate new conceptualizations. *Frontiers in psychology*, 7:561.
- Chao Jia, Yinfei Yang, Ye Xia, Yi-Ting Chen, Zarana Parekh, Hieu Pham, Quoc Le, Yun-Hsuan Sung, Zhen Li, and Tom Duerig. 2021. Scaling up visual and vision-language representation learning with noisy text supervision. In *International Conference on Machine Learning*, pages 4904–4916. PMLR.
- Wonjae Kim, Bokyoung Son, and Ildoo Kim. 2021. Vilt: Vision-and-language transformer without convolution or region supervision. In *International Conference on Machine Learning*, pages 5583–5594. PMLR.
- Robert M Krauss and Sidney Weinheimer. 1964. Changes in reference phrases as a function of frequency of usage in social interaction: A preliminary study. *Psychonomic Science*, 1(1):113–114.
- Ranjay Krishna, Yuke Zhu, Oliver Groth, Justin Johnson, Kenji Hata, Joshua Kravitz, Stephanie Chen, Yannis Kalantidis, Li-Jia Li, David A. Shamma, Michael S. Bernstein, and Li Fei-Fei. 2016. Visual genome: Connecting language and vision using crowdsourced dense image annotations. *International Journal of Computer Vision*, 123:32–73.
- Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. 2014. Microsoft coco: Common objects in context. In *European conference on computer vision*, pages 740–755.
- Jiasen Lu, Dhruv Batra, Devi Parikh, and Stefan Lee. 2019. Vilbert: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks. *Advances in neural information processing systems*, 32.
- Gary Lupyan and Bodo Winter. 2018. Language is more abstract than you think, or, why aren’t languages more iconic? *Philosophical Transactions of the Royal Society B: Biological Sciences*, 373.
- Junhua Mao, Jonathan Huang, Alexander Toshev, Oana Camburu, Alan L. Yuille, and Kevin Murphy. 2016. [Generation and comprehension of unambiguous object descriptions](#). In *Proceedings of the Conference on Computer Vision and Pattern Recognition*, pages 11–20. IEEE.
- William McCarthy, Robert X. D. Hawkins, Haoliang Wang, Cameron Holdaway, and Judith E. Fan. 2021. Learning to communicate about shared procedural abstractions. *ArXiv*, abs/2107.00077.
- Douglas L. Medin, Robert L. Goldstone, and Dedre Gentner. 1993. Respects for similarity. *Psychological Review*, 100:254–278.
- Margaret Mitchell, Kees van Deemter, and Ehud Reiter. 2010. [Natural reference to objects in a visual domain](#). In *Proceedings of the International Natural Language Generation Conference*.
- Tara Murfitt and Jan McAllister. 2001. The effect of production variables in monolog and dialog on comprehension by novel listeners. *Language and Speech*, 44(3):325–350.
- Sonia K Murthy, Thomas L Griffiths, and Robert D Hawkins. 2022. Shades of confusion: Lexical uncertainty modulates ad hoc coordination in an interactive communication task. *Cognition*, 225:105152.
- Vicente Ordonez, Girish Kulkarni, and Tamara Berg. 2011. Im2text: Describing images using 1 million captioned photographs. *Advances in neural information processing systems*, 24.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. 2021. Learning transferable visual models from natural language supervision. In *International Conference on Machine Learning*, pages 8748–8763. PMLR.

- Michael F Schober and Herbert H Clark. 1989. Understanding by addressees and overhearers. *Cognitive psychology*, 21(2):211–232.
- Piyush Sharma, Nan Ding, Sebastian Goodman, and Radu Soricut. 2018. Conceptual captions: A cleaned, hypernymed, image alt-text dataset for automatic image captioning. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2556–2565.
- Roger N. Shepard. 1987. Toward a universal law of generalization for psychological science. *Science*, 237 4820:1317–23.
- Todd Shore, Theofronia Androulakaki, and Gabriel Skantze. 2018. [KTH tangrams: A dataset for research on alignment and conceptual pacts in task-oriented dialogue](#). In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, Miyazaki, Japan. European Language Resources Association (ELRA).
- J Slocum. 2003. *The Tangram Book: The Story of the Chinese Puzzle with over 2000 Puzzles to Solve*. Sterling Publishing, New York.
- Joan Gay Snodgrass and Mary Vanderwart. 1980. A standardized set of 260 pictures: norms for name agreement, image agreement, familiarity, and visual complexity. *Journal of experimental psychology. Human learning and memory*, 6 2:174–215.
- Alane Suhr, Mike Lewis, James Yeh, and Yoav Artzi. 2017. [A corpus of natural language for visual reasoning](#). In *Proceedings of the Annual Meeting of the Association for Computational Linguistics*, pages 217–223.
- Alane Suhr, Stephanie Zhou, Ally Zhang, Iris Zhang, Huajun Bai, and Yoav Artzi. 2019. [A corpus for reasoning about natural language grounded in photographs](#). In *Proceedings of the Annual Meeting of the Association for Computational Linguistics*, pages 6418–6428.
- Tristan Thrush, Ryan Jiang, Max Bartolo, Amanpreet Singh, Adina Williams, Douwe Kiela, and Candace Ross. 2022. Winoground: Probing vision and language models for visio-linguistic compositionality. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5238–5248.
- Barbara Tversky and Kathleen Hemenway. 1984. Objects, parts, and categories. *Journal of Experimental Psychology: General*, 113:169–193.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in neural information processing systems*.
- Jette Viethen and Robert Dale. 2008. The use of spatial relations in referring expression generation. In *Proceedings of the International Conference on Natural Language Generation*.
- Catherine Wong, William McCarthy, Gabriel Grand, Yoni Friedman, Joshua B. Tenenbaum, Jacob Andreas, Robert D. Hawkins, and Judith E. Fan. 2022. Identifying concept libraries from language about object structure. *ArXiv*, abs/2205.05666.
- Licheng Yu, Patrick Poirson, Shan Yang, Alexander C Berg, and Tamara L Berg. 2016. Modeling context in referring expressions. In *The European Conference on Computer Vision*, pages 69–85.
- Martin Zettersten and Gary Lupyan. 2020. Finding categories through words: More nameable features improve category learning. *Cognition*, 196:104135.

A Appendix

A.1 Examples from KILOGRAM

Figure A.1 shows example tangrams from our data. Figure A.2 shows examples of the use of the part name *head*, the most common part head word in the data. All data can be browsed on the data visualization dashboard: <https://lil.nlp.cornell.edu/kilogram/>.

A.2 Collecting Tangrams

We scan all the pages of tangram solutions from Slocum (2003) into JPEG files to extract SVG files of individual tangrams. We use heuristics based on edge and corner detection (Harris et al., 1988) to extract individual tangrams into separate files by detecting the four corners of each puzzle and adding padding.⁵ We heuristically detect the individual standard pieces in each tangram using corner detection. Because the shapes are standard, we can test if an extracted shape is an expected puzzle’s piece and if we obtain the expected number of such shapes. We resize each tangram and all its pieces to a standard size, and label the ID of each puzzle piece consistently across all tangrams. We heuristically and manually validate the outputs, and prune solutions that fail to vectorize properly, for example if the process fails to recover exactly seven pieces.

A.3 Crowdsourcing Qualifications and Survey

The qualifier includes three multiple choice questions aimed to ensure that (a) the annotator describes the abstract shape meaningfully instead of simply describing its geometry; (b) each part description only contains one part (*body* and *arms* instead of *body with arms*); and (c) the part descriptions correspond to the description of the whole shape. We provide a short video tutorial of the task and examples of invalid annotations for workers to view before completing the qualifier. We also collect basic non-identifying demographic data from each worker, including the languages that they speak and their proficiency, if English is their first language, and where they learned English. We retain the correspondence of anonymized hashed worker IDs to the annotations and language information they provide.

A.4 Dense Annotation Sampling

The set DENSE is made of 62 tangrams sampled from FULL and 12 tangrams commonly used in

⁵We use OpenCV for this process (Bradski, 2000).

prior work. We sample the 62 tangrams from FULL to represent the diversity of tangrams using the first set of annotations we collect. We plot the annotated tangrams by average log perplexity of whole-shape descriptions with $\frac{1}{100}$ smoothing and PSA and apply a 5×5 grid to the plot (Figure A.3). Using perplexity and PSA allows us to sample a set of tangrams with diverse degrees of annotation and segmentation agreement. With a relatively high smoothing factor, we are able to spread out the data points, because the majority of the data set has high divergence in descriptions. We randomly pick 12 periphery points to collect more annotations for outliers, uniformly sample 25 from all the 1004 tangrams, and randomly sample 25, one from each grid, to represent the entire distribution.

We calculate average log perplexity of whole-shape annotations for each tangram. Let $\bar{x}^{(1)}, \dots, \bar{x}^{(N)}$ be annotations for a tangram, where each annotation is a sequence of tokens $\bar{x}^{(j)} = \langle x_1, \dots, x_{M^{(j)}} \rangle$ of length $M^{(j)}$. We create a language model $p^{(j)}$ for every annotation $\bar{x}^{(j)}$ using all other $N - 1$ annotations for the tangram:

$$p^{(j)}(x) = \frac{C_{x \in \bar{x}^{(j') \neq j}} + k}{\text{total}_{j' \neq j} + kV}, \quad (2)$$

where $C_{x \in \bar{x}^{(j') \neq j}}$ is the number of occurrences of x in the other annotations for the tangram, k is the smoothing factor, $\text{total}_{j' \neq j}$ is the total number of words used in the other annotations for the tangram and V is the vocabulary size of all whole-shape annotations across all tangrams. The log perplexity for annotation $\bar{x}^{(j)}$ is $\log PP^{(j)} = -\frac{1}{M^{(j)}} \sum_{i=1}^{M^{(j)}} \log_2 p(x_i^{(j)})$. The log perplexity for the tangram is the average of perplexity values for all its annotations $\log PP = \frac{1}{N} \sum_{j=1}^N \log PP^{(j)}$. We lowercase, stem, and remove stop words before computing the log perplexity.

A.5 Example Inputs for Experimental Conditions

Figure A.4 shows how one annotation, including both text and image, appears under the different experimental conditions. For conditions with PARTS annotations, we generate simple English sentences combining the whole shape description with part descriptions using the template *<whole shape> with <part>, <part>, ..., and <part>*. We add an indefinite article to each singular part description. BLACK images are tangrams with all pieces colored black with white borders. COLOR images are tangrams with each part colored with one of

the CSS preset colors in the order of coral, gold, lightskyblue, lightpink, mediumseagreen, darkgrey, lightgrey that correspond to the parts in the annotation. For the augmented condition (AUG), text inputs are whole annotations combined with each possible subset of the part descriptions. Image inputs are tangrams colored in the same way as colored images, but the parts excluded from the subset of part descriptions are colored black instead. All part descriptions in the annotations are randomly shuffled and not consistently associated with any particular color in the images, so that the coloring solely serves as an indication of the ordering of parts in the combined text.

A.6 Human Performance Baseline Details

We recruited an independent group of 233 human participants from the Prolific crowdsourcing platform (<https://www.prolific.co/>), and asked them to perform the same reference game task we used for model evaluation. Each participant was randomly assigned to one of the four conditions and shown a random sequence of 20 trials from that condition. On each trial, we showed a text annotation from the development set along with the corresponding context of ten tangrams and asked the participant to click the tangram that was being described. The information that was available varied across condition, just as in the model evaluations. The tangrams were either presented to participants in black-and-white (BLACK) or colored according to their segmentation map (COLOR), and the language was either the whole-shape description alone (WHOLE) or with the parts included (PARTS). In the PARTS+COLOR condition, the parts text was colored to match the image to facilitate visual comparison, providing the same alignment information available to the models.

We took several steps to ensure high-quality responses. First, participants began with a fixed set of 10 practice trials to familiarize with the task. For these practice trials, we provided feedback indicating whether they have selected the correct tangram, and if not, we highlight the correct answer. To assess whether participants were paying attention as opposed to responding randomly, we inserted an unambiguous “catch trial” where the target was the square tangram and the description was *square*. We excluded 16 participants who failed to select the correct target on this trial, yielding a final sample size of 217 participants out of the 233 recruited.

Because our aim was to obtain overall accuracy estimates for each condition, we did not require judgements for every individual annotation and context in the test set. However, we were able to ensure good coverage of the dataset, including annotations from all 125 tangrams and over 600 unique descriptions in each condition.

A.7 Model-specific Implementation Details

For experiments with CLIP, we use the ViT-B/32 variant. We fine-tune using an Adam optimizer with learning rate $5e-8$ and weight decay $1e-6$. At the end of each epoch, the training data is shuffled and rebatched. We train the models up to 200 epochs and use patience of 50 epochs to select the model with the highest image prediction accuracy on a non-augmented validation set taken from the training data. All images are resized to CLIP’s default input resolution of 224×224 , with white padding to make to rectangle images square. The total number of trainable parameters in CLIP is 151.2M. CLIP models are fine-tuned with either a single GeForce RTX 2080 Ti GPU with 11GB memory or a single Titan RTX GPU with 24GB memory. Fine-tuning takes approximately 40 minutes per epoch for augmented setups (AUG) and roughly 3 minutes for other setups.

For ViLT experiments, we fine-tune with an AdamW optimizer with learning rate $1e-4$ and weight decay $1e-2$. We use a cosine learning rate schedule with warm-up over the first epoch. We train the models up to 30 epochs with a patience of 10 epochs and follow the same model selection criterion as for CLIP. All images are resized to 384×384 . The total number of trainable parameters in ViLT is 87.4M. ViLT models are fine-tuned with a single Titan RTX GPU with 24 GB memory. Fine-tuning takes up to 5.5 hours per epoch for augmented setups (AUG) and roughly 15 minutes for other setups.

A.8 Random Generation of Reference Games

In our main experiments (Section 5), we randomly generate reference games subject to constraints (Section 5.1). In particular, we ensure that distractors contained the same total number of parts. We explore the impact of these constraints by repeating our experiments on reference games generated without the constraints. Without the constraints, part counting can help the model disqualify distractors and significantly narrow down the set of likely referents. This is because images with a different

Condition	CLIP		ViLT	
	PT	FT	PT	FT
WHOLE+BLACK	17.3	46.2	13.2	41.3
PARTS+BLACK	16.8	47.4	12.6	47.0
WHOLE+COLOR	15.9	48	12.4	46.2
PARTS+COLOR	15.9	71.3	12.1	89.0
PARTS+COLOR+AUG	–	74	–	86.0

Table A.1: Reference game development accuracies (%) for the different experimental conditions with pre-trained (PT) or fine-tuned (FT) models for games generated without constraints.

number of parts colored compared to the number of parts in the text description can be easily ignored without considering the semantics of the text or images. Table A.1 shows development accuracies for games generated without constraints, both for training and testing. Generally, the success rate achieved on unconstrained contexts is much higher compared to contexts generated with constraints (Figure 3). However, when analyzing the performance of this model on part-controlled contexts (Figure A.5), we observe roughly similar performance to the games generated with constraints (Figure 8), even though we would expect a significant performance increase given the results in Table A.1. We even observe a more pronounced decrease in performance when more parts are added, illustrating further difficulty generalizing. We conclude that the model trained on games generated without constraints (Table A.1) likely learns to rely on part-counting heuristics and may be less reliable in other settings.

A.9 Reproducibility Checklist

For all reported experimental results:

- A clear description of the mathematical setting, algorithm, and/or model: yes; see Section 5.
- Submission of a zip file containing source code, with specification of all dependencies, including external libraries, or a link to such resources: yes; attached to our submission.
- Description of computing infrastructure used: yes; see Appendix A.7.
- The average runtime for each model or algorithm (e.g., training, inference, etc.) or estimated energy cost: yes; see Appendix A.7.

- Number of parameters in each model: yes; see Appendix A.7.
- Corresponding validation performance for each reported test result: yes; see Appendix 3 and Appendix A.1 for results on the development set.
- Explanation of evaluation metrics, with links to code used: yes; see Section 5 for an explanation of the reference game metric. An implementation is included in the attached code zipfile.

For all experiments with hyperparameter search:

- We performed a minimal manual search for learning rate and weight decay, and used the same values for all experiments (described in Section A.7).

For all datasets used:

- Relevant details such as languages, and number of examples and label distributions: yes; see Section 3.
- Details of train/test/validation splits: yes; see Section 3.3.
- Explanation of any data that were excluded, and all pre-processing steps: yes; see Section 3 and Section A.2.
- A zip file containing data or link to a downloadable version of the data: yes; attached to our submission.
- For new data collected, a complete description of the data collection process, such as instructions to annotators and methods for quality control: yes; see Section 3.2 and Section A.3.




Figure A.1: Example tangrams from our dataset.

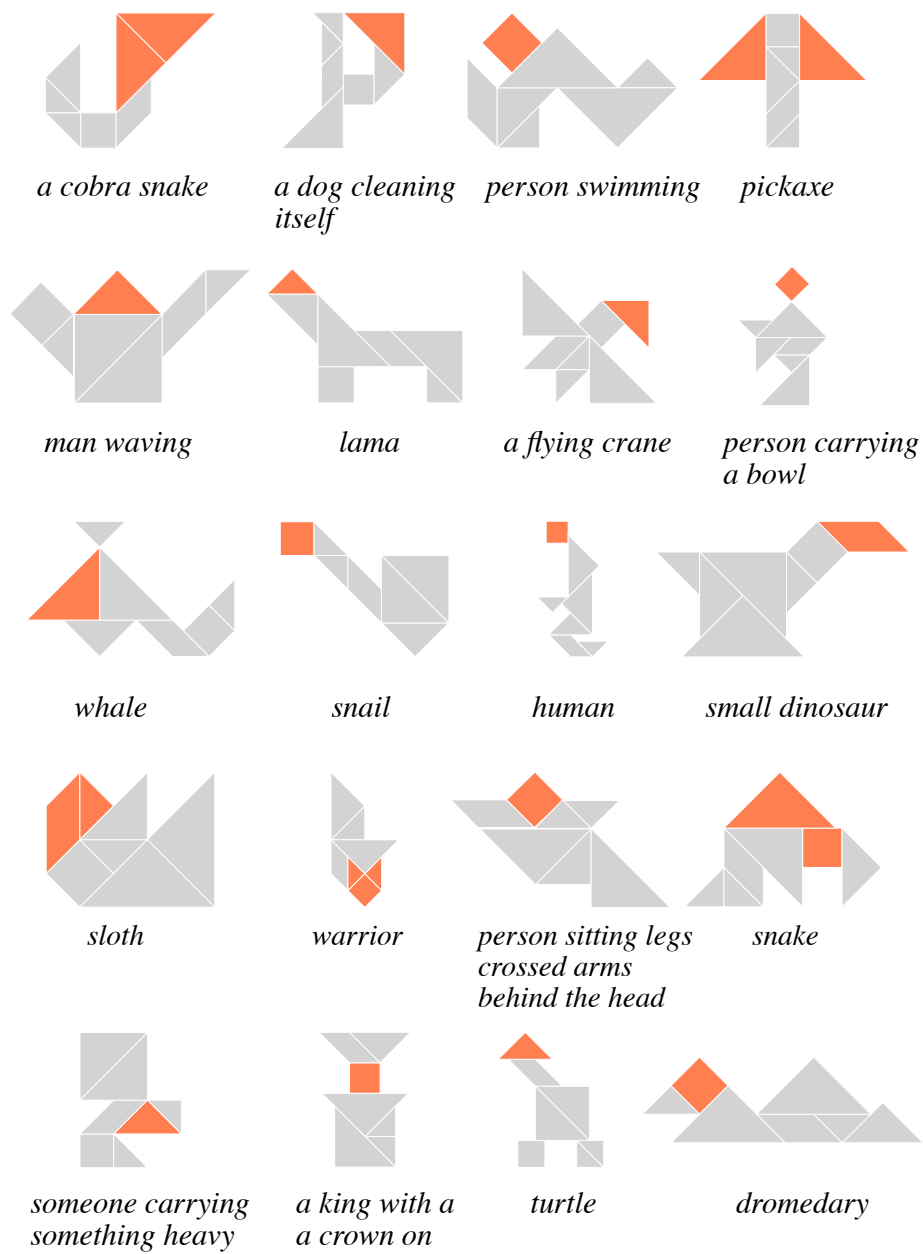


Figure A.2: Example tangrams containing the part description *head*. Each example includes a tangram and its whole-shape description. We highlight the segmentation corresponding to *head* in each tangram.




Figure A.3: Sampled tangrams for dense annotation collection: 12 purple points picked from the periphery, 25 red points randomly sampled from each grid, and 25 green points uniformly sampled from all points.

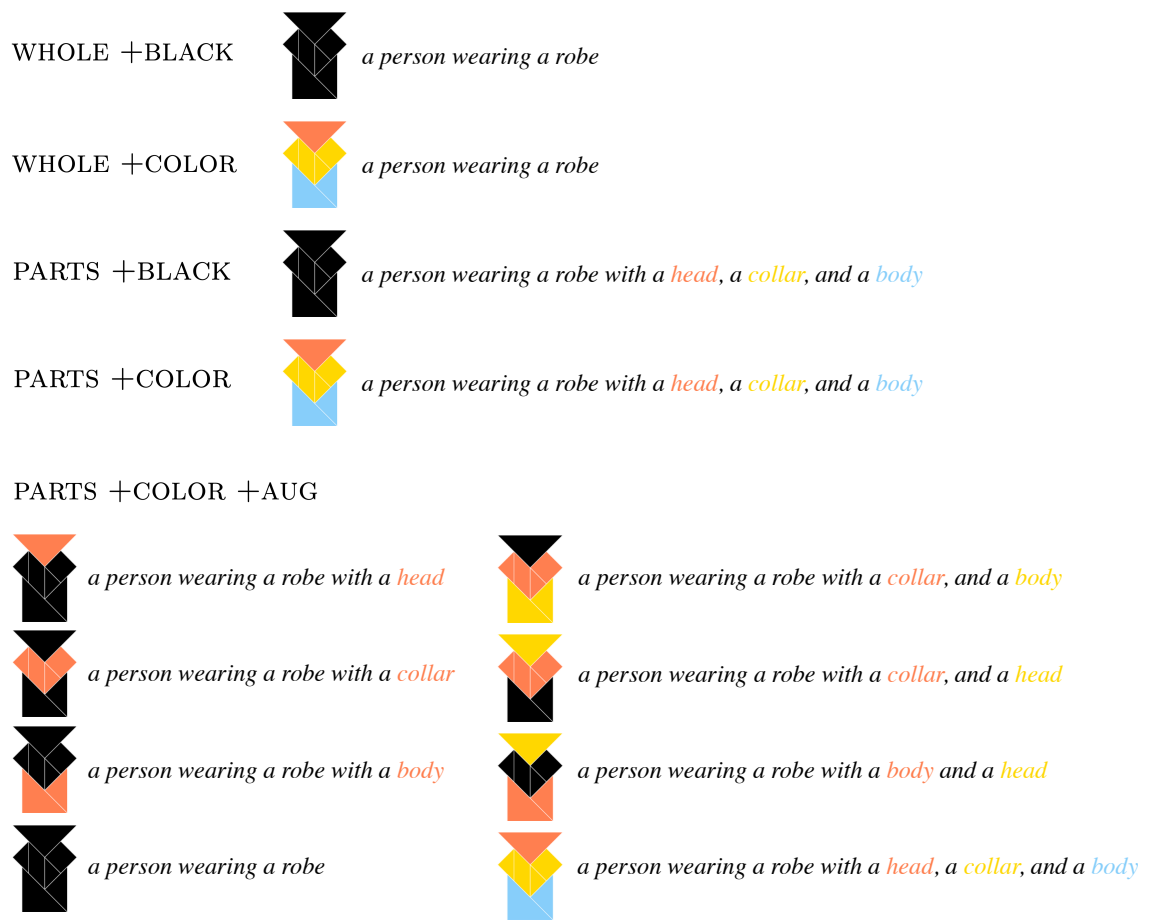


Figure A.4: An example of one annotation across the different experimental conditions. The augmentation condition (AUG) creates multiple examples from the same annotation.




Figure A.5: Mean development probabilities of predicting the correct image in reference games generated without constraints using fine-tuned CLIP (top) or fine-tuned ViLT (bottom) by number of parts included in text and colored in the images. We separate the curves by the total number of parts in the annotation of the target example. The error bands show the 95% confidence interval of the expected mean at each point by bootstrapping with 1000 resamplings.