

Latent Compositional Representations Improve Systematic Generalization in Grounded Question Answering

Ben Bogin¹ Sanjay Subramanian² Matt Gardner² Jonathan Berant^{1,2}

¹Tel-Aviv University ²Allen Institute for AI

{ben.bogin, joberant}@cs.tau.ac.il, {sanjays, mattg}@allenai.org

Abstract

Answering questions that involve multi-step reasoning requires decomposing them and using the answers of intermediate steps to reach the final answer. However, state-of-the-art models in grounded question answering often do not explicitly perform decomposition, leading to difficulties in generalization to out-of-distribution examples. In this work, we propose a model that computes a representation and denotation for all question spans in a bottom-up, compositional manner using a CKY-style parser. Our model induces latent trees, driven by end-to-end (the answer) supervision only. We show that this inductive bias towards tree structures dramatically improves systematic generalization to out-of-distribution examples, compared to strong baselines on an arithmetic expressions benchmark as well as on CLOSURE, a dataset that focuses on systematic generalization for grounded question answering. On this challenging dataset, our model reaches an accuracy of 96.1%, significantly higher than prior models that almost perfectly solve the task on a random, in-distribution split.

1 Introduction

Humans can effortlessly interpret new language utterances, if they are composed of previously observed primitives and structure (Fodor and Pylyshyn, 1988). Neural networks, conversely, do not exhibit this *systematicity*: Although they generalize well to examples sampled from the distribution of the training set, they have been shown to struggle in generalizing to out-of-distribution (OOD) examples that contain new compositions in grounded question answering (Bahdanau et al., 2019a,b) and semantic parsing (Finegan-Dollak et al., 2018; Keyzers et al., 2020). Consider the

question in Figure 1. This question requires querying object sizes, comparing colors, identifying spatial relations, and computing intersections between sets of objects. Neural models succeed when these concepts are combined in ways that were seen at training time. However, they commonly fail when concepts are combined in new ways at test time.

A possible reason for this phenomenon is the expressivity of architectures such as LSTMs (Hochreiter and Schmidhuber, 1997) and Transformers (Vaswani et al., 2017), where “global” representations that depend on the entire input are computed. Contextualizing token representations using the entire utterance potentially lets the model avoid step-by-step reasoning, “collapse” reasoning steps, and rely on shortcuts (Jiang and Bansal, 2019; Subramanian et al., 2019, 2020). Such failures are revealed when evaluating models for systematic generalization on OOD examples. This stands in contrast to pre-neural log-linear models, where hierarchical representations were explicitly constructed over the input (Zettlemoyer and Collins, 2005; Liang et al., 2013).

In this work, we propose a model for visual question answering (QA) that, analogous to these classical pre-neural models, computes for every span in the input question a *representation* and a *denotation*, that is, the set of objects in the image that the span refers to (see Figure 1). Denotations for long spans are recursively computed from shorter spans using a bottom-up CKY-style parser *without* access to the entire input, leading to an inductive bias that encourages compositional computation. Because training is done from the final answer only, the model must effectively learn to induce latent trees that describe the compositional structure of the problem. We hypothesize that this explicit grounding of the meaning of sub-spans through hierarchical computation should result in better generalization to new compositions.

We evaluate our approach in two setups: (a) a synthetic arithmetic expressions dataset, and

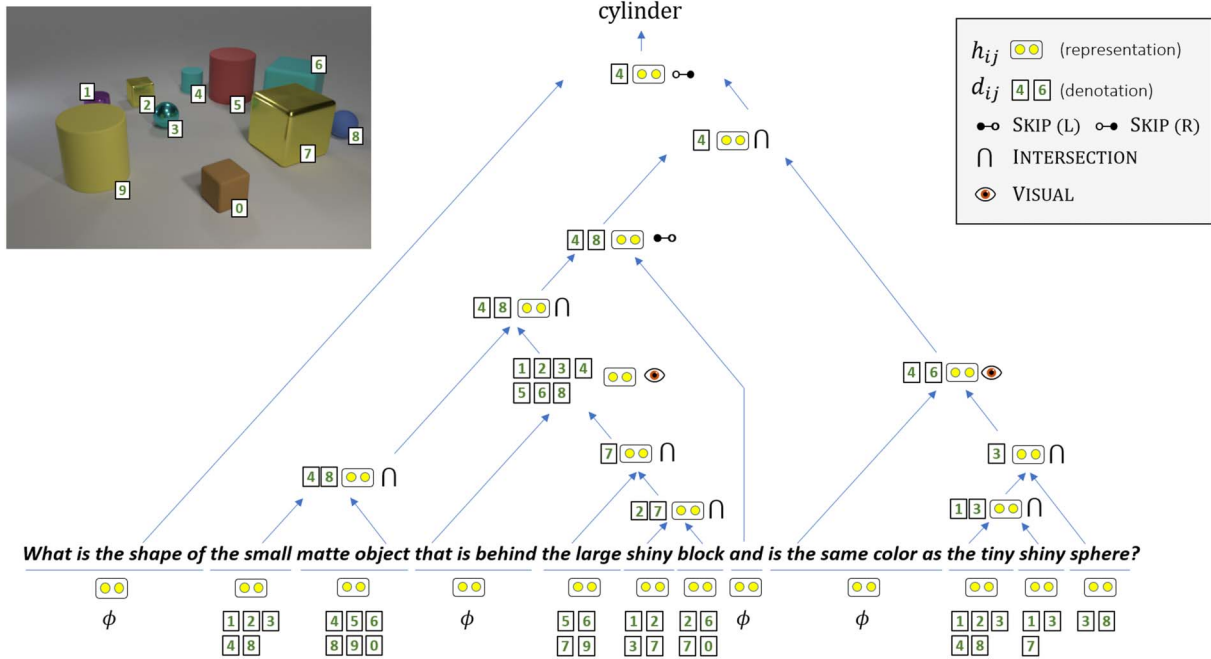


Figure 1: An example from CLOSURE illustrating how our model learns a latent structure over the input, where a representation and denotation is computed for every span (for denotation we show the set of objects with probability > 0.5). For brevity, some phrases were merged to a single node of the tree. For each phrase, we show the split point and module with the highest probability, although all possible split points and module outputs are softly computed. SKIP(L) and SKIP(R) refer to taking the denotation of the left or right sub-span, respectively.

(b) CLOSURE (Bahdanau et al., 2019b), a visual QA dataset focusing on systematic generalization. On a random train/test split (*i.i.d split*), both our model and prior baselines obtain near perfect performance. However, on splits that require generalization to new compositions (*compositional split*) our model dramatically improves performance: On the arithmetic dataset, a Transformer-based model fails to generalize and obtains 2.9% accuracy, whereas our model, Grounded Latent Trees (GLT), gets 98.4%. On CLOSURE, our model’s accuracy is 96.1%, 24 absolute points higher than strong baselines and even 19 points higher than models that use gold structures at training time or depend on domain-knowledge.

To conclude, we propose a model with an inherent inductive bias for compositional computation, which leads to gains in systematic generalization, and induces latent structures that are useful for understanding its inner workings. This suggests that despite the success of general-purpose architectures built on top of contextualized representations, restricting information flow inside the network can benefit compositional generalization. Our code and data can be found

at <https://github.com/benbogin/glt-grounded-latent-trees-qa>.

2 Compositional Generalization

Language is mostly compositional; humans can understand and produce a potentially infinite number of novel combinations from a finite set of components (Chomsky, 1957; Montague, 1970). For example, a person would know what a “winged giraffe” is even if she’s never seen one, assuming she knows the meaning of “winged” and “giraffe”. This ability, termed *compositional generalization*, is fundamental for building robust models that learn from limited data (Lake et al., 2018).

Neural networks have been shown to generalize well in many language understanding tasks in *i.i.d splits* (Devlin et al., 2019; Linzen, 2020). However, when evaluated on splits that require compositional generalization, a significant drop in performance is observed. For example, in SCAN (Lake and Baroni, 2018) and gSCAN (Ruis et al., 2020), synthetically generated commands are mapped into a sequence of actions. When

tested on unseen command combinations, models perform poorly. A similar trend was shown in text-to-SQL parsing Finegan-Dollak et al. (2018), where splitting examples by the template of the target SQL query resulted in a dramatic performance drop. SGOOP (Bahdanau et al., 2019a) shows the same phenomena on a synthetic visual task, which tests for generalization over unseen combinations of object properties and relations. This also led to methods that construct compositional splits automatically (Keysers et al., 2020).

In this work, we focus on answering complex grounded questions over images. The CLEVR benchmark (Johnson et al., 2017a) contains synthetic images and questions that require multi-step reasoning, for example, “*Are there any cyan spheres made of the same material as the green sphere?*”. While performance on this task is high, with an accuracy of 97%–99% (Perez et al., 2018; Hudson and Manning, 2018), recent work (Bahdanau et al., 2019b) introduced CLOSURE: a new set of questions with identical vocabulary but different structure than CLEVR, asked on the same set of images. They evaluated generalization of different models and showed that all fail on a large fraction.

The most common approach for grounded QA is based on end-to-end differentiable models such as FiLM (Perez et al., 2018), MAC (Hudson and Manning, 2018), LXMERT (Tan and Bansal, 2019), and UNITER (Chen et al., 2020). These models do not explicitly decompose the problem into smaller sub-tasks, and are thus prone to fail on compositional generalization. A different approach (Yi et al., 2018; Mao et al., 2019) is to parse the image into a symbolic or distributed knowledge graph with objects, attributes, and relations, and then parse the question into an executable logical form, which is deterministically executed. Last, Neural Module Networks (NMFs; Andreas et al. 2016) parse the question into an executable program, where execution is learned: Each program module is a neural network designed to perform an atomic task, and modules are composed to perform complex reasoning. The latter two model families construct compositional programs and have been shown to generalize better on compositional splits (Bahdanau et al., 2019a,b) compared with end-to-end models. However, programs are not explicitly tied to question spans, and search over the space of programs is not differentiable, leading to difficulties in training.

In this work, we learn a latent structure for the question and tie each question span to an executable module in a differentiable manner. Our model balances distributed and symbolic approaches: We learn from downstream supervision only and output an inferred tree, describing how the answer was computed. We base our work on latent tree parsers (Le and Zuidema, 2015; Liu et al., 2018; Maillard et al., 2019; Drozdov et al., 2019) that produce representations for all spans, and softly weight all possible trees. We extend these parsers to answer grounded questions, grounding sub-trees in image objects.

Closest to our work is Gupta and Lewis (2018) (GL), who have proposed to compute denotations for each span by constructing a CKY chart. Our model differs in several important aspects: First, while both models compute the denotations of all spans, we additionally compute dense representations for each span using a learned composition function, increasing the expressiveness of our model. Second, we compute the denotations using simpler and more generic composition operators, while in GL the operators are fine-grained and type-driven. Last, we work with images rather than a knowledge graph, and propose several modifications to the chart construction mechanism to overcome scalability challenges.

3 Model

We first give a high-level overview of our proposed GLT model (§3.1), and then explain our model in detail.

3.1 High-level Overview

Problem Setup Our task is visual QA, where given a question $q = (q_0, \dots, q_{n-1})$ and an image I , we aim to output an answer $a \in \mathcal{A}$ from a fixed set of natural language phrases. We train a model from a training set $\{(q^i, I^i, a^i)\}_{i=1}^N$. We assume we can extract from the image up to n_{obj} features vectors of objects, and represent them as a matrix $V \in \mathbb{R}^{n_{\text{obj}} \times h_{\text{dim}}}$ (details on object detection and representation are in §3.4).

Our goal is to compute for every question span $q_{ij} = (q_i, \dots, q_{j-1})$ a *representation* $\mathbf{h}_{ij} \in \mathbb{R}^{h_{\text{dim}}}$ and a *denotation* $\mathbf{d}_{ij} \in [0, 1]^{n_{\text{obj}}}$, interpreted as the probability that the question span refers to each object. We compute \mathbf{h}_{ij} and \mathbf{d}_{ij} in a bottom-up fashion, using CKY (Cocke, 1969; Kasami, 1965; Younger, 1967). Algorithm 1 provides a high-level description of the procedure. We compute

Algorithm 1

Require: question q , image I , word embedding matrix E , visual representations matrix V

- 1: H : tensor holding representations $\mathbf{h}_{ij}, \forall i, j$ s.t. $i < j$
- 2: D : tensor holding denotations $\mathbf{d}_{ij}, \forall i, j$ s.t. $i < j$
- 3: **for** $i = 1 \dots n$ **do** //Handle single-token spans
- 4: $\mathbf{h}_i \leftarrow E_{q_i}, \mathbf{d}_i \leftarrow f_{\text{ground}}(E_{q_i}, V)$ (see §3.4)
- 5: **for** $l = 1 \dots n$ **do** //Handle spans of length l
- 6: compute $\mathbf{h}_{ij}, \mathbf{d}_{ij}$ for all entries s.t. $j - i = l$
- 7: $p(a | q, I) \leftarrow \text{softmax}(W[\mathbf{h}_{0n}; \mathbf{d}_{0n}V])$
- 8: **return** $\arg \max_a p(a | q, I)$

representations and denotations for length-1 spans (we use $\mathbf{h}_i = \mathbf{h}_{i(i+1)}, \mathbf{d}_i = \mathbf{d}_{i(i+1)}$ for brevity) by setting the representation $\mathbf{h}_i = E_{q_i}$ to be the corresponding word representation in an embedding matrix E , and grounding each word in the image objects: $\mathbf{d}_i = f_{\text{ground}}(E_{q_i}, V)$ (lines 3–4; f_{ground} is described in §3.4). Then, we recursively compute representations and denotations of larger spans (lines 5–6). Last, we pass the question representation, \mathbf{h}_{0n} , and the weighted sum of the visual representations ($\mathbf{d}_{0n}V$) through a softmax layer to produce a final answer distribution (line 1), using a learned classification matrix $W \in \mathbb{R}^{A \times 2h_{\text{dim}}}$.

Computing $\mathbf{h}_{ij}, \mathbf{d}_{ij}$ for all spans requires overcoming some challenges. Each span representation \mathbf{h}_{ij} should be a function of two *sub-spans* $\mathbf{h}_{ik}, \mathbf{h}_{kj}$. We use the term sub-spans to refer to all adjacent pairs of spans that cover q_{ij} : $\{(q_{ik}, q_{kj})\}_{k=i+1}^{j-1}$. However, we have no supervision for the “correct” split point k . Our model considers all possible split points and learns to induce a latent tree structure from the final answer only (§3.2). We show that this leads to an interpretable compositional structure and denotations that can be inspected at test time.

In §3.3 we describe the form of the *composition functions*, which compute both span representations and denotations from two sub-spans. These functions must be expressive enough to accommodate a wide range of interactions between sub-spans, while avoiding reasoning shortcuts that might hinder compositional generalization.

3.2 Grounded Chart Parsing

We now describe how to compute $\mathbf{h}_{ij}, \mathbf{d}_{ij}$ from previously computed representations and denotations. In standard CKY, each constituent over a span q_{ij} is constructed by combining two

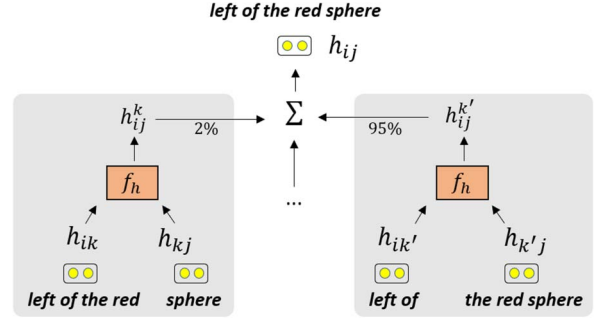


Figure 2: Computing \mathbf{h}_{ij} : we consider all split points and compose pairs of sub-spans using f_h . The output is a weighted sum of these representations.

sub-spans q_{ik}, q_{kj} that meet at a split point k . Similarly, we define a representation \mathbf{h}_{ij}^k conditioned on the split point and constructed from previously computed representations of two sub-spans:

$$\mathbf{h}_{ij}^k = f_h(\mathbf{h}_{ik}, \mathbf{h}_{kj}), \quad (1)$$

where $f_h(\cdot)$ is a *composition function* (§3.3).

Because we want the loss to be differentiable, we do not pick a particular value k , but instead use a continuous relaxation. Specifically, we compute α_{ij}^k that represents the probability of k being the split point for the span q_{ij} , given the tensor H of computed representations of shorter spans. We then define the representation \mathbf{h}_{ij} to be the expected representation over all possible split points:

$$\mathbf{h}_{ij} = \sum_k \alpha_{ij}^k \cdot \mathbf{h}_{ij}^k = \mathbb{E}_{\alpha_{ij}^k}[\mathbf{h}_{ij}^k]. \quad (2)$$

The split point distribution is defined as $\alpha_{ij}^k \propto \exp(\mathbf{s}^T \mathbf{h}_{ij}^k)$, where $\mathbf{s} \in \mathbb{R}^{h_{\text{dim}}}$ is a parameter vector that determines what split points are likely. Figure 2 illustrates computing \mathbf{h}_{ij} .

Next, we compute the denotation \mathbf{d}_{ij} of each span. Conceptually, computing \mathbf{d}_{ij} can be analogous to \mathbf{h}_{ij} ; that is, a function f_d will compute \mathbf{d}_{ij}^k for every split point k , and $\mathbf{d}_{ij} = \mathbb{E}_{\alpha_{ij}^k}[\mathbf{d}_{ij}^k]$. However, the function f_d (see §3.3) interacts with the visual representations of all objects and is thus computationally costly. Therefore, we propose a less expressive but more efficient approach, where $f_d(\cdot)$ is applied only once for each span q_{ij} .

Specifically, we compute the expected denotation of the left and right sub-spans of q_{ij} :

$$\bar{\mathbf{d}}_{ijL} = \mathbb{E}_{\alpha_{ij}^k}[\mathbf{d}_{ik}] \in \mathbb{R}^{n_{\text{obj}}}, \quad (3)$$

$$\bar{\mathbf{d}}_{ijR} = \mathbb{E}_{\alpha_{ij}^k}[\mathbf{d}_{kj}] \in \mathbb{R}^{n_{\text{obj}}}. \quad (4)$$

If α_{ij}^k puts most probability mass on a single split point k' , then the expected denotation will be similar to picking that particular split point.

Now we can compute \mathbf{d}_{ij} given the expected sub-span denotations and representations with a single application of $f_d(\cdot)$:

$$\mathbf{d}_{ij} = f_d(\bar{\mathbf{d}}_{ij_L}, \bar{\mathbf{d}}_{ij_R}, \mathbf{h}_{ij}), \quad (5)$$

which is more efficient than the alternative $\mathbb{E}_{\alpha_{ij}^k} [f_d(\mathbf{d}_{ik}, \mathbf{d}_{kj}, \mathbf{h}_{ij})]$: f_d is applied $O(n^2)$ times compared to $O(n^3)$ with the simpler solution. This makes training tractable in practice.

3.3 Composition Functions

We now describe the exact form of the composition functions f_h and f_d .

Composing Representations We describe the function $f_h(\mathbf{h}_{ik}, \mathbf{h}_{kj})$, used to compose the representations of two sub-spans (Equation 1). The goal of this function is to compose the ‘‘meanings’’ of two adjacent spans, *without* access to the rest of the question or the denotations of the sub-spans. For example, composing the representations of ‘‘same’’ and ‘‘size’’ to a representation for ‘‘same size’’. At a high-level, composition is based on an attention mechanism. Specifically, we use attention to form a convex combination of the representations of the two sub-spans (Equations 6, 7), and apply a non-linearity with a residual connection (Equation 8).

$$\beta_{ij}^k = \text{softmax}([\mathbf{a}_L \mathbf{h}_{ik}, \mathbf{a}_R \mathbf{h}_{kj}]) \in \mathbb{R}^2 \quad (6)$$

$$\hat{\mathbf{h}}_{ij}^k = \beta_{ij(1)}^k W_L \mathbf{h}_{ik} + \beta_{ij(2)}^k W_R \mathbf{h}_{kj} \in \mathbb{R}^{h_{\text{dim}}} \quad (7)$$

$$f_h(\mathbf{h}_{ik}, \mathbf{h}_{kj}) = \text{FF}_{\text{rep}}(\hat{\mathbf{h}}_{ij}^k) + \hat{\mathbf{h}}_{ij}^k \in \mathbb{R}^{h_{\text{dim}}} \quad (8)$$

where $\mathbf{a}_L, \mathbf{a}_R \in \mathbb{R}^{h_{\text{dim}}}$, $W_L, W_R \in \mathbb{R}^{h_{\text{dim}} \times h_{\text{dim}}}$, and $\text{FF}_{\text{rep}}(\cdot)$ is a linear layer of size $h_{\text{dim}} \times h_{\text{dim}}$ followed by a non-linear activation.¹

Composing Denotations Next, we describe the function $f_d(\bar{\mathbf{d}}_{ij_L}, \bar{\mathbf{d}}_{ij_R}, \mathbf{h}_{ij})$, used to compute the span denotation \mathbf{d}_{ij} (Equation 5). This function has access only to words in the span q_{ij} and not to the entire input utterance. We want $f_d(\cdot)$ to support both simple compositions that depend only on the denotations of sub-spans, as well as more complex

¹We also use Dropout and Layer-Norm (Ba et al., 2016) throughout the paper, omitted for simplicity.

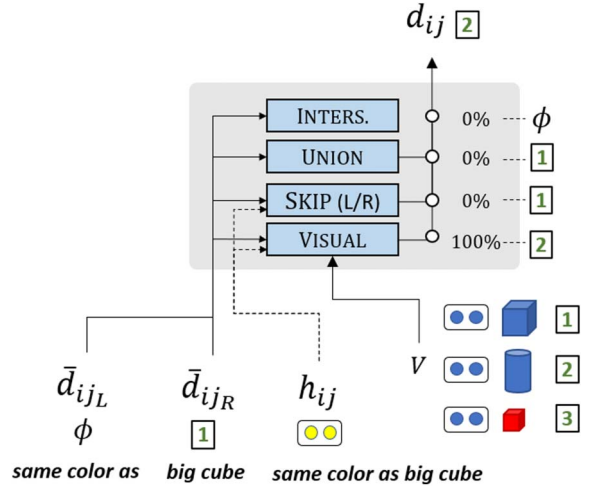


Figure 3: Illustration of how \mathbf{d}_{ij} is computed. We compute the denotations of all modules, and a weight for each one of the modules. The span denotation is then the weighted sum of the module outputs.

functions that consider the visual representations of different objects (spatial relations, colors, etc.).

We define four modules in $f_d(\cdot)$ for computing denotations and let the model learn when to use each module. The modules are: SKIP, INTERSECTION, UNION, and a general-purpose VISUAL function, where only VISUAL uses the visual representations V . As illustrated in Figure 3, each module m outputs a denotation vector $\mathbf{d}_{ij}^m \in [0, 1]^{n_{\text{obj}}}$, and the denotation \mathbf{d}_{ij} is a weighted average of the four modules:

$$p(m|i, j) \propto \exp(W_{\text{mod}} \mathbf{h}_{ij}) \in \mathbb{R}^4 \quad (9)$$

$$\mathbf{d}_{ij} = \sum_m p(m|i, j) \cdot \mathbf{d}_{ij}^m \in \mathbb{R}^{n_{\text{obj}}}, \quad (10)$$

where $W_{\text{mod}} \in \mathbb{R}^{h_{\text{dim}} \times 4}$. Next, we define the four modules (see Figure 4).

SKIP In many cases, only one of the left or right sub-spans have a meaningful denotation: for example, for ‘‘there is a’’ and ‘‘red cube’’, we should only keep the denotation of the right sub-span. To that end, the SKIP module weighs the two denotations and sums them ($W_{\text{sk}} \in \mathbb{R}^{h_{\text{dim}} \times 2}$):

$$(c_{ij(1)}, c_{ij(2)}) = \text{softmax}(W_{\text{sk}} \mathbf{h}_{ij}) \in \mathbb{R}^2 \quad (11)$$

$$\mathbf{d}_{ij}^{\text{sk}} = c_{ij(1)} \cdot \mathbf{d}_{ij_L} + c_{ij(2)} \cdot \mathbf{d}_{ij_R} \in \mathbb{R}^{n_{\text{obj}}}. \quad (12)$$

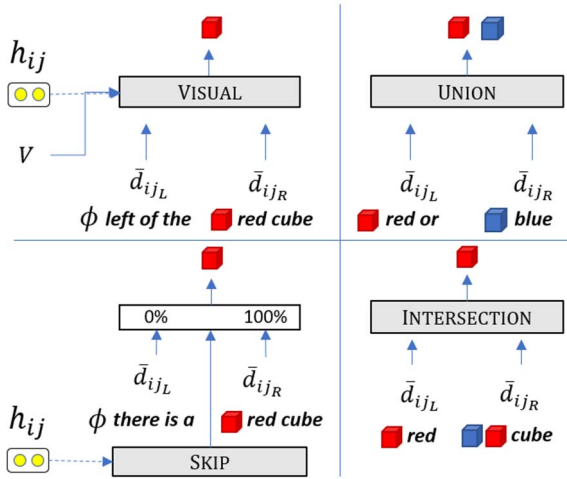


Figure 4: The different modules used with their inputs and expected output.

INTERSECTION and UNION We define two modules that only use the denotations $\bar{\mathbf{d}}_{ijL}$ and $\bar{\mathbf{d}}_{ijR}$. The first corresponds to intersection of two sets, and the second to union:

$$\mathbf{d}_{ij}^{\text{int}} = \min(\bar{\mathbf{d}}_{ijL}, \bar{\mathbf{d}}_{ijR}) \in \mathbb{R}^{n_{\text{obj}}} \quad (13)$$

$$\mathbf{d}_{ij}^{\text{uni}} = \max(\bar{\mathbf{d}}_{ijL}, \bar{\mathbf{d}}_{ijR}) \in \mathbb{R}^{n_{\text{obj}}}, \quad (14)$$

where $\min(\cdot)$ and $\max(\cdot)$ are computed element-wise.

VISUAL This module is responsible for visual computations, such as computing spatial relations (“*left of the red sphere*”) and comparing attributes of objects (“*has the same size as the red sphere*”). Unlike other modules, it also uses the visual representations of the objects, $V \in \mathbb{R}^{n_{\text{obj}} \times h_{\text{dim}}}$. For example, for the sub-spans “*left of*” and “*the red object*”, we expect the function to ignore $\bar{\mathbf{d}}_{ijL}$ (since the denotation of “*left to*” is irrelevant), and return a denotation with high probability for objects that are left to objects with high probability in $\bar{\mathbf{d}}_{ijR}$.

To determine if an object with index o should have high probability, we need to consider its relation to all other objects. A simple scoring function between object o and another arbitrary object o_z might be $(\mathbf{h}_{ij} + \mathbf{v}_o)^T (\mathbf{h}_{ij} + \mathbf{v}_{o_z})$, which will capture the relation between the objects conditioned on the span representation. However, computing this for all pairs is quadratic in n_{obj} . Instead, we propose a linear alternative that leverages expected denotations of sub-spans. Specifically, we com-

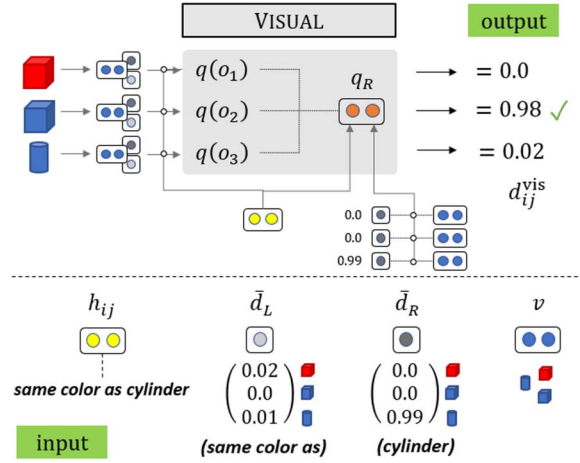


Figure 5: Overview of the VISUAL module. We compute a representation for each object, $q(o)$, augmented with the span denotation, representation, and visual embeddings, and then compute its dot-product with q_R , a representation conditioned on the weighted sum of the right sub-span objects and representations.

pute the expected visual representation of the right sub-span and pass it through a feed-forward layer:

$$\bar{\mathbf{v}}_R = \bar{\mathbf{d}}_R V \in \mathbb{R}^{h_{\text{dim}}} \quad (15)$$

$$\mathbf{q}_R = \text{FF}_R(W_h \mathbf{h}_{ij} + \bar{\mathbf{v}}_R) \in \mathbb{R}^{h_{\text{dim}}}. \quad (16)$$

We use the right sub-span because the syntax in CLEVR is mostly right-branching.² Then, we generate a representation $\mathbf{q}(o)$ for every object conditioned on the span representation \mathbf{h}_{ij} , the object probability under the sub-span denotations, and its visual representation. The final object probability is based on the interaction of $\mathbf{q}(o)$ and \mathbf{q}_R :

$$\mathbf{q}(o) = \text{FF}_{\text{vis}}(W_h \mathbf{h}_{ij} + \mathbf{v}_o + \bar{\mathbf{d}}_L(o) \mathbf{s}_L + \bar{\mathbf{d}}_R(o) \mathbf{s}_R)$$

$$\mathbf{d}_{ij}^{\text{vis}}(o) = \sigma(\mathbf{q}(o)^T \mathbf{q}_R),$$

where $W_h \in \mathbb{R}^{h_{\text{dim}} \times h_{\text{dim}}}$, $\mathbf{s}_L, \mathbf{s}_R \in \mathbb{R}^{h_{\text{dim}}}$ are learned embeddings and FF_{vis} is a feed-forward layer of size $h_{\text{dim}} \times h_{\text{dim}}$ with a non-linear activation. This is the most expressive module we propose. A visual overview is shown in Figure 5.

3.4 Grounding

In lines 3–4 of Algorithm 1, we handle length-1 spans. The span representation \mathbf{h}_i is initialized as the corresponding word embedding E_{q_i} , and the

²Using the left sub-span instead substantially reduces performance on CLEVR and CLOSURE.

denotation is computed with a *grounding function*. A simple implementation for f_{ground} would use the dot product between the word representation and the visual representations of all objects: $\sigma(\mathbf{h}_i^\top V)$.³ But, in the case of co-reference (“*it*”), we should ground the co-referring pronoun in the denotation of a previous span. We now describe this case.

Co-reference Sentences such as “*there is a red sphere; what is its material?*” are challenging given a CKY parser, because the denotation of “*its*” depends on a distant span. We propose a simple heuristic that addresses the case where the referenced object is the denotation of a previous sentence. This could be expanded in future work, to a wider array of coreference phenomena.

In every example that comprises two sentences: (a) We compute the denotation $\mathbf{d}^{\text{first}}$ for the entire first sentence as previously described; (b) We ground each word in the second sentence as proposed above: $\hat{\mathbf{d}}_i^{\text{second}} = \sigma(\mathbf{h}_i^{\text{second}\top} V)$; (c) For each word in the second sentence, we predict whether it co-refers to $\mathbf{d}^{\text{first}}$ using a learned gate $(r_1, r_2) = \text{softmax}(\text{FF}_{\text{coref}}(\mathbf{h}_i^{\text{second}}))$, where $\text{FF}_{\text{coref}} \in \mathbb{R}^{h_{\text{dim}} \times 2}$. (d) We define $\hat{\mathbf{d}}_i^{\text{second}} = r_1 \cdot \mathbf{d}^{\text{first}} + r_2 \cdot \hat{\mathbf{d}}_i^{\text{second}}$.

Visual Representations Next, we describe how we compute the visual embedding matrix V . Two common approaches to obtain visual features are (1) computing a feature map for the entire image and letting the model learn to attend to the correct feature position (Hudson and Manning, 2018; Perez et al., 2018); and (2) predicting the locations of objects in the image, and extracting features just for these objects (Anderson et al., 2018; Tan and Bansal, 2019; Chen et al., 2020). We use the latter approach, because it simplifies learning over discrete sets, and has better memory efficiency – the model only attends to a small set of objects rather than the entire image feature map.

We train an object detector, Faster R-CNN (Ren et al., 2015), to predict the location $\text{bb}_{\text{pred}} \in \mathbb{R}^{n_{\text{obj}} \times 4}$ of all objects, in the format of bounding boxes (horizontal and vertical positions).

We use gold scene data of 5,000 images from CLEVR for training (and 1,000 images for validation). In order to extract features for each predicted bounding box, we use the bottom-up top-down

³To improve runtime efficiency in the case of a large knowledge-graph, one could introduce a fast filtering function to reduce the number of proposed objects.

attention method of Anderson et al. (2018),⁴ which produces the features matrix $V_{\text{pred}} \in \mathbb{R}^{n_{\text{obj}} \times D}$, where $D = 2048$. Bounding boxes and features are extracted and fixed as a pre-processing step.

Finally, to compute V , similar to LXMERT and UNITER we augment object representations in V_{pred} with their position embeddings, and pass them through a single Transformer self-attention layer to add context about other objects: $V = \text{TransformerLayer}(V_{\text{pred}}W_{\text{feat}} + \text{bb}_{\text{pred}}W_{\text{pos}})$, where $W_{\text{feat}} \in \mathbb{R}^{D \times h_{\text{dim}}}$ and $W_{\text{pos}} \in \mathbb{R}^{4 \times h_{\text{dim}}}$.

Complexity Similar to CKY, we go over all $O(n^2)$ spans in a sentence, and for each span compute \mathbf{h}_{ij}^k for each of the possible $O(n)$ splits (there is no grammar constant since the grammar has effectively one rule). To compute denotations \mathbf{d}_{ij} , for all $O(n^2)$ spans, we perform a linear computation over all n_{obj} objects. Thus, the algorithm runs in time $O(n^3 + n^2n_{\text{obj}})$, with similar memory consumption. This is higher than end-to-end models that do not compute explicit span representations.

3.5 Training

The model is fully differentiable, and we simply maximize the log probability $\log p(a^* | q, I)$ of the correct answer a^* (see Algorithm 1).

4 Experiments

In this section, we evaluate our model on both in-distribution and OOD splits.

4.1 Arithmetic Expressions

It has been shown that neural networks can be trained to perform numerical reasoning (Zaremba and Sutskever, 2014; Kaiser and Sutskever, 2016; Trask et al., 2018; Geva et al., 2020). However, models are often evaluated on expressions that are similar to the ones they were trained on, where only the numbers change. To test generalization, we create a simple dataset and evaluate on two splits that require learning the correct operator precedence. In the first split, sequences of operators that appear at test time do not appear at training time. In the second split, the test set contains longer sequences compared to the training set.

⁴Specifically, we use this version: <https://github.com/airsplay/py-bottom-up-attention>.

	Easy split	Operation split
Training	$8*2*5+3+9 = 92$ $6+3*4+8+4 = 30$	$8*2*5+3+9 = 92$ $6+3*4+8+4 = 30$
Test	$5*3*6+2+3 = 95$ $9+4*7+1+2 = 40$	$7*3+4*5*2 = 61$ $1+2+2*7+5 = 22$

Figure 6: Arithmetic expressions: Unlike the easy setup, we evaluate models on expressions with operations ordered in ways unobserved at training time. Flipped operator positions are in red.

We define an arithmetic expression as a sequence containing n numbers with $n - 1$ arithmetic operators between each pair. The answer a is the result of evaluating the expression.

Evaluation Setups The sampled operators are addition and multiplication, and we use only expressions where $a \in \{0, 1, \dots, 100\}$ to train as a multi-class problem. At training, we randomly pick the length n to be up to n_{train} , and at test time we choose a fixed length n_{test} . We evaluate on three setups, (a) *Easy split*: we choose $n_{\text{train}} = n_{\text{test}} = 8$, and randomly sample operators from a uniform distribution for both training and test examples. In this setup, we only check that the exact same expression is not shared between the training and test set. (b) *Operation split*: We randomly pick 3 positions, and for each one randomly assign exactly one operator that will appear at training time. On the test set, the operators in all three positions are flipped, so that they contain the unseen operator (Figure 6). The same lengths are used as in the easy split. (c) *Length split*: we train with $n_{\text{train}} = 8$ and test with $n_{\text{test}} = 10$. Examples for all setups are generated on-the-fly for 3 million steps (batch size is 100).

Models We compare GLT to a standard Transformer, where the input is the expression, and the output is predicted using a classification layer over the [CLS] token. All models are trained with cross-entropy loss given the correct answer.

For both models, we use an in-distribution validation set for hyper-parameter tuning. Since here we do not have an image, we only compute \mathbf{h}_{ij} for all spans, and define $p(a | q) = \text{softmax}(W\mathbf{h}_{0n})$.

GLT layers are almost entirely recurrent, that is, the same parameters are used to compute representations for spans of all lengths. The only exception are layer-normalization parameters,

	Easy split	Op. split	Len. split
Transformer	100.0 \pm 0.0	2.9 \pm 1.1	10.4 \pm 2.4
GLT	99.9 \pm 0.2	98.4 \pm 0.7	94.9 \pm 1.1

Table 1: Arithmetic expressions results for easy split, operation split, and length split.

which are not shared across layers. Thus, at test time when processing an expression longer than observed at training time, we use the layer-normalization parameters (total of $2 \cdot h_{\text{dim}}$ parameters per layer) from the longest span seen at training time.⁵

Results Results are reported in Table 1. We see that both models almost completely solve the in-distribution setup, but on OOD splits the Transformer performs poorly, while GLT shows only a small drop in accuracy.

4.2 CLEVR and CLOSURE

We evaluate performance on grounded QA using CLEVR (Johnson et al., 2017a), consisting of 100,000 synthetic images with multiple objects of different shapes, colors, materials and sizes. 864,968 questions were generated using 80 different templates, including simple questions (“*what is the size of red cube?*”) and questions requiring multi-step reasoning (Figure 1). The split in this dataset is i.i.d: The same templates are used for the training, validation, and test sets.

To test compositional generalization when training on CLEVR, we use the CLOSURE dataset (Bahdanau et al., 2019b), which includes seven new question templates, with a total of 25,200 questions, asked on the CLEVR validation set images. The templates are created by taking referring expressions of various types from CLEVR and combining them in novel ways.

A problem found in CLOSURE is that sentences from the template `embed_mat_spa` are ambiguous. For example, in the question “*Is there a sphere on the left side of the cyan object that is the same size as purple cube?*”, the phrase “*that is the same size as purple cube*” can modify either “*the sphere*” or “*the cyan object*”, but the answer in CLOSURE is always the latter. Therefore, we deterministically compute both of the two

⁵Removing layer normalization leads to improved accuracy of 99% on the arithmetic expressions length split, but training convergence on CLEVR becomes too slow.

possible answers and keep two sets of question-answer pairs of this template for the entire dataset. We evaluate models⁶ on this template by taking the maximum score over these two sets (such that models must be consistent and choose a single interpretation for the template to get a perfect score).

Baselines We evaluate against the baselines from Bahdanau et al. (2019b). The most comparable baselines are MAC (Hudson and Manning, 2018) and FiLM (Perez et al., 2018), which are differentiable and do not use program annotations. We also compare to NMNs that require a few hundred program examples for training. We show results for PG+EE (Johnson et al., 2017b) and an improved version, PG-Vector-NMN (Bahdanau et al., 2019b). Last, we compare to NS-VQA, which in addition to parsing the question, also parses the scene into a knowledge graph. NS-VQA requires additional gold data to parse the image into a knowledge graph based on data from CLEVR (colors, shapes, locations, etc.).

Setup Baseline results are taken from previous papers (Bahdanau et al., 2019b; Hudson and Manning, 2018; Yi et al., 2018; Johnson et al., 2017b), except for MAC and FiLM on CLOSURE, which we re-executed due to the aforementioned evaluation change. For GLT, we use CLEVR’s validation set for hyper-parameter tuning and early-stopping, for a maximum of 40 epochs. We run 4 experiments to compute mean and standard deviation on CLOSURE test set.

Because of our model’s run-time and memory demands (see §3.4), we found that running on CLEVR and CLOSURE, where question length goes up to 43 tokens, can be slow. Thus, we delete function words that typically have empty denotations and can be safely skipped,⁷ reducing the maximum length to 25. We run a single experiment where stop words were not removed and report results in Table 3 (with stop-words), showing that performance only mildly change.

CLEVR and CLOSURE In this experiment we compare results on i.i.d and compositional splits. Results are in Table 2. We see that GLT performs

⁶We update the scores on CLOSURE for MAC, FiLM, and GLT due to this change in evaluation. The scores for the rest of the models were not affected.

⁷The removed tokens are punctuations, ‘the’, ‘there’, ‘is’, ‘a’, ‘as’, ‘of’, ‘are’, ‘other’, ‘on’, ‘that’.

	CLEVR	CLOSURE
MAC	98.5	72.4
FiLM	97.0	60.1
GLT (our model)	99.1	96.1 ± 2.5
NS-VQA † ‡	100	77.2
PG+EE (18K prog.) †	95.4	–
PG-Vector-NMN †	98.0	71.3
GT-Vector-NMN † ‡	98.0	94.4

Table 2: Test results for all models on CLEVR and CLOSURE. † stands for models trained with gold programs, ‡ for oracle models evaluated using gold programs, and ‡ for deterministically executed models that depend on domain-knowledge for execution (execution is not learned).

	CLEVR	CLOSURE	C.Humans
GLT	99.1 ± 0.0	96.1 ± 2.7	75.8 ± 0.7
{VISUAL, SKIP}	98.9 ± 0.1	91.8 ± 10.1	76.3 ± 0.6
{VISUAL}	99.0 ± 0.1	83.8 ± 1.8	72.0 ± 1.6
with stop-words	99.3	93.2	75.4
contextualized input	98.9 ± 0.1	87.5 ± 12.1	76.1 ± 0.9
non-compositional	97.7 ± 2.0	83.5 ± 0.5	75.8 ± 2.3

Table 3: Ablations on the development sets of CLEVR, CLOSURE, and CLEVR-Humans.

well on CLEVR and gets the highest score on CLOSURE, improving by almost 24 points over comparable models, and by 19 points over NS-VQA, outperforming even the oracle GT-Vector-NMN which uses gold programs at **test time**.

Removing Modules We defined (§3.3) two modules specifically for CLEVR, INTERSECTION and UNION. We evaluate performance without them, keeping only VISUAL and SKIP, and show results in Table 3, observing only a moderate loss in accuracy and generalization. Removing these modules leads to more cases where the VISUAL function is used, effectively performing intersection and union as well. While the drop in accuracy is small, this model is harder to interpret since VISUAL now performs multiple functions.

Next, we evaluate performance when training only with the VISUAL module. Results show that even a single high-capacity module is enough for in-distribution performance, but generalization to CLOSURE now substantially drops to 83.8.

Contextualizing Question Tokens We hypothesized that the fact that question spans do not

	CLOSURE FS	C.Humans
MAC	90.2	81.5
FiLM	–	75.9
GLT (our model)	97.4 ± 0.3	76.2
NS-VQA	92.9	67.0
PG-Vector-NMN	88.0	–
PG+EE (18K prog.)	–	66.6

Table 4: Test results in the few-shot setup and for CLEVR-Humans.

observe the entire input might aid composition generalization. To check this, we pass the question tokens q through a Transformer self-attention layer, to get contextualized tokens embeddings E_{q_i} . We show in Table 3 that performance on CLEVR remains the same, but on CLOSURE it drops to 87.5.

Non-compositional Representations To assess the importance of *compositional* representations, we replace \mathbf{h}_{ij} (Equation 2) with a sequential encoder. Concretely, we set the representation of the span q_{ij} to be $\mathbf{h}_{ij} = \text{BiLSTM}(q_{ij})$, the output of a single layer bi-directional LSTM which is given the tokens of the span as input. We see a drop in both CLEVR and CLOSURE (experiments were more prone to overfitting), showing the importance of compositional representations.

Few-shot In the few-shot (FS) setup, we train with additional 252 OOD examples: 36 questions for each CLOSURE template. Similar to Bahdanau et al. (2019b), we take a model that was trained on CLEVR and fine-tune adding oversampled CLOSURE examples (300x) to the original training set. To make results comparable to Bahdanau et al. (2019b), we perform model selection based on CLOSURE validation set, and evaluate on the test set. As we see in Table 4, GLT gets the best accuracy. If we perform model selection based on CLEVR (the preferred way to evaluate in the OOD setup, Teney et al. 2020), accuracy remains similar at 97.0 ± 2.0 (and is still highest).

CLEVR-Humans To test performance on real language, we use CLEVR-Humans (Johnson et al., 2017b), which includes 32,164 questions over images from CLEVR. These questions, asked by humans, contain words and reasoning steps that

were unseen in CLEVR. We train on CLEVR and fine-tune on CLEVR-Human, similar to prior work. We use GloVe embeddings (Pennington et al., 2014) for words unseen in CLEVR. Table 4 shows that GLT gets better results than models that use programs, showing its flexibility to learn new concepts and phrasings. It is also comparable to FiLM, but gets lower results than MAC (see error analysis below).

4.3 Error analysis

We sampled 25 errors from each dataset (CLEVR, CLOSURE, and CLEVR-Humans) for analysis. On CLEVR, most errors (84%) are due to problems in *visual processing* such as grounding the word “*rubber*” to a metal object, problems in bounding box prediction or questions that require subtle spatial reasoning, such as identifying if an object is left to another object of different size, when they are at an almost identical x-position. The remaining errors (16%) are due to failed comparisons of numbers or attributes (“*does the red object have the same material as the cube*”).

On CLOSURE, 60% of the errors were similar to CLEVR, for example, problematic visual processing or failed comparisons. In 4% of cases, the execution of the VISUAL module was wrong, for example, it collapsed two reasoning steps (both intersection and finding objects of same shape), but did not output the correct denotation. Other errors (36%) are in the predicted latent tree, where the model was uncertain about the split point and softly predicted more than one tree, resulting in a wrong answer. In some cases (16%) this was due to question ambiguity (see §4.2), and in others cases the cause was unclear (e.g., for the phrase “*same color as the cube*” the model gave a similar probability for the split after “*same*” and after “*color*”, leading to a wrong denotation of that span).

On CLEVR-Humans, the model successfully learned certain new “concepts” such as colors (“*gold*”), superlatives (“*smallest*”, “*furthest*”), relations (“*next to*”, “*between*”), negation (see Figure 7), and the “*all*” quantifier. It also answered correctly questions that had different style from CLEVR (“*Are there more blocks or balls?*”). However, the model is not always robust to these new concepts, and fails in other cases that require counting of abstract concepts (such as colors or shapes) or that use CLEVR’s operators




Figure 7: An example from CLEVR-Humans. The model learned to negate (“not”) using the VISUAL module (negation is not part of CLEVR).

differently than the original uses. See Table 5 for our manually performed analysis and examples over 150 CLEVR-Humans validation questions.

4.4 Limitations

While some of the errors shown in Table 5 can be improved with careful design of modules architectures, improvement of the image features, or by introducing a pre-training mechanism, some question structures are inherently more difficult to correctly answer using our model. We describe two main limitations.

One key issue is discrepancy between the text of the question and the actual reasoning steps that are required to correctly answer it. For example, in the question “What is the most common shape?”, the phrase “most common” entails grouping and counting shapes, and then taking an argmax as an answer. This complex reasoning cannot be constructed compositionally from question spans.

The second issue is the supported types of phrase denotations. GLT only grounds *objects* in the image to a phrase, however, in some cases denotations should be a number or an attribute. For example, in comparison questions (“is the number of cubes higher than the number of spheres?”) a model would ideally have a numerical denotation for each group of objects.

Importantly, these limitations do not prevent the model from answering questions correctly, since

even when denotations do not accurately represent reasoning, the model can still “fall back” to the flexible answer function, which can take the question and image representations and, in theory, perform any needed computation. In such cases, the model will not be interpretable, similar to other high-capacity models such as Transformers. In future work, we will explore combining the compositional generalization abilities of GLT with the advantages of high-capacity architectures.

4.5 Interpretability

A key advantage of latent trees is interpretability—one can analyze the structure of a given question. Next, we quantitatively inspect the quality of the intermediate outputs produced by the model by (a) comparing these outputs to ground truth and (b) assessing how helpful they are for humans.

Evaluating Intermediate Outputs We assess how accurate our intermediate outputs are compared to ground truth trees. Evaluating against ground truth trees is not trivial, since for a given question, there could be many “correct” trees—that is, trees that represent a valid way to answer the question. For example, the two possible ways to split the phrase “large shiny block” are both potentially valid. Thus, even if we had one ground-truth tree for each question, comparing predicted trees to ground-truth trees might not be reliable.

Instead, we exploit the fact that for the synthetic questions in CLEVR, it is possible to define a set of certain constituents that *must* appear in any valid tree. For example, phrases starting with “the”, followed by adjectives and then a noun (“the shiny tiny sphere”), should always be considered a constituent. We use such manually defined rules to extract the set c_q of obligatory constituents for any question q in CLEVR and CLOSURE validation sets.

We compute a recall score for each question q , by producing predicted trees (we greedily take the maximum probability split in each node), extracting a set \hat{c}_q of all constituents, and then computing the proportion of constituents in c_q that are also in \hat{c}_q . We show results in Table 6 for both CLEVR and CLOSURE. We see that 81.6%–83.1% of the expected constituents are correctly output by

Category	Examples	Correct
Negation	– How many objects are not shiny? ✓ – What color is the cylinder that is not the same <u>colr</u> as the sphere? ✗	1/2 (50%)
Spelling mistakes	– Are there two green <u>cumes</u> ? ✓ – How many rubber <u>spehres</u> are there here? ✗	1/5 (20%)
Superlatives	– What color is the object furthest to the right? ✓ – What shape is the smallest matte object? ✓	8/10 (80%)
Visual Concepts: obscuring, between	– Is the sphere the same color as the object that is obscuring it? ✓ – What color is the object in between the two large cubes? ✓	5/7 (71%)
Visual Concepts: reflection, shadow	– Which shape shows the largest reflection ✗ – Are all of the objects casting a shadow? ✗	0/2 (0%)
Visual Concepts: relations	– What color is the small ball near the brown cube? ✓ – What is behind and right of the cyan cylinder? ✗	3/8 (38%)
All quantifier	– Are all the spheres the same size? ✓ – Are all the cylinders brown? ✗	10/12 (83%)
Counting abstract concepts	– How many different shapes are there? ✓ – How many differently colored cubes are there? ✗	1/4 (25%)
Complex logic	– if these objects were lined up biggest to smallest, what would be in the middle? ✓ – if most of the items shown are shiny and most of the items shown are blue, would it be fair to say most of the items are shiny and blue? ✗	3/4 (75%)
Different question structure	– Are more objects metallic or matte? ✓ – Each shape is present 3 times except for the ✗	1/2 (50%)
Uniqueness	– What color object is a different material from the rest? ✓ – What color is the object that does not match the others? ✗	1/2 (50%)
Long-tail concepts	– Can you roll all the purple objects? ✓ – How many of these things could be stacked on top of each other? ✗	2/5 (40%)
Operators used differently than CLEVR	– Are the large cylinders the same color? ✓ – Are there more rubber objects than matte cylinders and green cubes? ✗	2/5 (40%)
Same operators as CLEVR, possibly different phrasing	– What color is the cube directly in front of the blue cylinder? ✓ – What color is the little cylinder? ✗	1/80 (89%)

Table 5: Categorization of 150 CLEVR-Humans examples, together with the accuracy of the GLT model (for many categories, sample is too small to make generalizations on model performance). Spelling mistakes are underlined. Some examples were labeled with multiple categories. Six examples with gold-label mistakes were ignored.

GLT, for CLEVR and CLOSURE respectively. Next, we sample and analyze 25 cases from each dataset where the expected constituents were not output. We observe that almost all missed constituents actually *were* output as a constituent by the model, but with an additional prefix that caused a wrong split. For example, instead of the expected constituent “*tiny green objects*” we found the constituent “*any tiny green objects*” (split after “*tiny*”). These mistakes did not cause any error in module execution, except for one case in CLOSURE.

Additionally, we obtain the gold set of objects for each constituent in c_q by deterministically mapping each constituent to a program that we execute on the scene. For each constituent in $c_q \cap \hat{c}_q$, we compare the predicted set of objects (objects with probability above 50%) with the gold set of objects, and report the average F_1 . As

	CLEVR	CLOSURE
Constituents (%)	83.1	81.6
Denotations (F_1)	95.9	94.7

Table 6: Evaluation of the intermediate outputs, showing both the recall of expected constituents (%) and average F_1 of predicted objects.

can be seen in Table 6, reported score for CLEVR is 95.9% and for CLOSURE it is 94.7%.

Human Interpretability Next, we want to assess how useful are the intermediate outputs produced by the model for humans. We perform the “forward prediction” test (Hu et al., 2018) that evaluates interpretability by showing humans questions accompanied with intermediate outputs,

	CLEVR	CLOSURE	C.Humans
A - w/o tree	61.4%	58.1%	52.3%
B - w/ tree	62.5%	82.5%	77.5%

Table 7: Results for the “forward prediction” test. Group A sees only the question, image and gold answer, group B additionally sees the denotation tree.

and asking them to predict if the model will be correct. This test is based on the assumption that humans can more easily predict the accuracy of an interpretable model versus a non-interpretable one. Thus, we also ask a separate group of participants to predict the accuracy *without* showing them the intermediate outputs for comparison.

We show each of the 20 participants 12 questions evenly split into CLEVR, CLOSURE, and CLEVR-Humans, where for each dataset we show two correct and two incorrect examples per person. We assign 11 participants to the baseline group (group A) that sees only the question, image, and gold answer, and 9 participants to the tested group (group B) which are also given the predicted denotation tree of our model, similar to Figure 1. Both groups are given the same 8 “training” examples of questions, with the predicted answer of the model along with the gold answer, to improve their understanding of the task. In these training examples, group B participants also see example denotation trees, along with a basic explanation of how to read those.

We show results in Table 7. We see that for CLOSURE and CLEVR-Humans, accuracy for group B is significantly higher ($p < 0.05$), but for CLEVR results are similar. We observe that in CLEVR, where accuracy is 99.1%, wrong predictions are mostly due to errors in the less interpretable VISUAL module, while for CLOSURE and CLEVR-Humans errors are more often due to wrong selection of constituents and modules, which can be spotted by humans.

5 Conclusion

Developing models that generalize to compositions that are unobserved at training time has recently sparked substantial research interest. In this work we propose a model with a strong

inductive bias towards compositional computation, which leads to large gains in systematic generalization and provides an interpretable structure that can be inspected by humans. Moreover, our model also obtains high performance on real language questions (CLEVR-humans). In future work, we will investigate the structures revealed by our model in other grounded QA setups, and will allow the model freedom to incorporate non-compositional signals, which go hand in hand with compositional computation in natural language.

Acknowledgments

This research was partially supported by The Yandex Initiative for Machine Learning, and the European Research Council (ERC) under the European Union Horizons 2020 research and innovation programme (grant ERC DELPHI 802800). We thank Jonathan Herzig and the anonymous reviewers for their useful comments. This work was completed in partial fulfillment for the Ph.D. degree of Ben Bogin.

References

- Peter Anderson, Xiaodong He, Chris Buehler, Damien Teney, Mark Johnson, Stephen Gould, and Lei Zhang. 2018. Bottom-up and top-down attention for image captioning and visual question answering. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6077–6086. DOI: <https://doi.org/10.1109/CVPR.2018.00636>
- Jacob Andreas, Marcus Rohrbach, Trevor Darrell, and Dan Klein. 2016. Neural module networks. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 39–48. DOI: <https://doi.org/10.1109/CVPR.2016.12>
- Lei Jimmy Ba, Jamie Ryan Kiros, and Geoffrey E. Hinton. 2016. Layer normalization. *CoRR*, abs/1607.06450.
- Dzmitry Bahdanau, Shikhar Murty, Michael Noukhovitch, Thien Huu Nguyen, Harm de Vries, and Aaron Courville. 2019a. Systematic generalization: What is required and can it be learned? In *International Conference on Learning Representations*.

- Dzmitry Bahdanau, Harm de Vries, Timothy J. O'Donnell, Shikhar Murty, Philippe Beaudoin, Yoshua Bengio, and Aaron C. Courville. 2019b. CLOSURE: Assessing systematic generalization of CLEVR models. *CoRR*, abs/1912.05783.
- Yen-Chun Chen, Linjie Li, Licheng Yu, A. E. Kholy, Faisal Ahmed, Zhe Gan, Y. Cheng, and Jing jing Liu. 2020. UNITER: Universal image-text representation learning. In *ECCV*. DOI: https://doi.org/10.1007/978-3-030-58577-8_7
- Noam Chomsky. 1957. *Syntactic Structures*, Mouton. DOI: <https://doi.org/10.1515/9783112316009>
- John Cocke. 1969. *Programming Languages and Their Compilers: Preliminary Notes*, New York University.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Andrew Drozdov, Patrick Verga, Mohit Yadav, Mohit Iyyer, and Andrew McCallum. 2019. Unsupervised latent tree induction with deep inside-outside recursive auto-encoders. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 1129–1141, Minneapolis, Minnesota. Association for Computational Linguistics. DOI: <https://doi.org/10.18653/v1/N19-1116>
- Catherine Finegan-Dollak, Jonathan K. Kummerfeld, Li Zhang, Karthik Ramanathan, Sesh Sadasivam, Rui Zhang, and Dragomir Radev. 2018. Improving text-to-SQL evaluation methodology. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 351–360, Melbourne, Australia. Association for Computational Linguistics. DOI: <https://doi.org/10.18653/v1/P18-1033>
- Jerry A. Fodor and Zenon W. Pylyshyn. 1988. *Connectionism and Cognitive Architecture: A Critical Analysis*. MIT Press, Cambridge, MA, USA. ISBN 0262660644, pp. 3–71. DOI: [https://doi.org/10.1016/0010-0277\(88\)90031-5](https://doi.org/10.1016/0010-0277(88)90031-5)
- Mor Geva, Ankit Gupta, and Jonathan Berant. 2020. Injecting numerical reasoning skills into language models. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 946–958, Online. Association for Computational Linguistics. DOI: <https://doi.org/10.18653/v1/2020.acl-main.89>
- Nitish Gupta and Mike Lewis. 2018. Neural compositional denotational semantics for question answering. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2152–2161, Brussels, Belgium. Association for Computational Linguistics. DOI: <https://doi.org/10.18653/v1/D18-1239>, PMID: 30485774
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural Computation*, 9:1735–1780. DOI: <https://doi.org/10.1162/neco.1997.9.8.1735>, PMID: 9377276
- Ronghang Hu, Jacob Andreas, Trevor Darrell, and Kate Saenko. 2018. Explainable neural computation via stack neural module networks. In *Proceedings of the European Conference on Computer Vision (ECCV)*.
- Drew Arad Hudson and Christopher D. Manning. 2018. Compositional attention networks for machine reasoning. In *International Conference on Learning Representations*.
- Yichen Jiang and Mohit Bansal. 2019. Avoiding reasoning shortcuts: Adversarial evaluation, training, and model development for multi-hop QA. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2726–2736, Florence, Italy.

- Association for Computational Linguistics. **DOI:** <https://doi.org/10.18653/v1/P19-1262>, **PMID:** 31353678
- Justin Johnson, Bharath Hariharan, Laurens van der Maaten, Li Fei-Fei, C. Lawrence Zitnick, and Ross Girshick. 2017a. CLEVR: A diagnostic dataset for compositional language and elementary visual reasoning. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1988–1997. **DOI:** <https://doi.org/10.1109/CVPR.2017.215>
- Justin Johnson, Bharath Hariharan, Laurens van der Maaten, Judy Hoffman, Li Fei-Fei, C. Lawrence Zitnick, and Ross Girshick. 2017b. Inferring and executing programs for visual reasoning. In *2017 IEEE International Conference on Computer Vision (ICCV)*, pages 3008–3017. **DOI:** <https://doi.org/10.1109/ICCV.2017.325>, **PMID:** 28556096
- Lukasz Kaiser and Ilya Sutskever. 2016. Neural GPUs learn algorithms. *International Conference on Learning Representations*.
- Tadao Kasami. 1965. An efficient recognition and syntax analysis algorithm for context-free languages, Air Force Cambridge Research Laboratory, Bedford, MA.
- Daniel Keysers, Nathanael Schärli, Nathan Scales, Hylke Buisman, Daniel Furrer, Sergii Kashubin, Nikola Momchev, Danila Sinopalnikov, Lukasz Stafiniak, Tibor Tihon, Dmitry Tsarkov, Xiao Wang, Marc van Zee, and Olivier Bousquet. 2020. Measuring compositional generalization: A comprehensive method on realistic data. In *International Conference on Learning Representations*.
- Brenden M. Lake and Marco Baroni. 2018. Generalization without systematicity: On the compositional skills of sequence-to-sequence recurrent networks. In *ICML*.
- Brenden M. Lake, Tomer D. Ullman, Joshua B. Tenenbaum, and Samuel Gershman. 2018. Building machines that learn and think like people. *Behavioral and Brain Sciences*, 40:e253.
- Phong Le and Willem Zuidema. 2015. The forest convolutional network: Compositional distributional semantics with a neural chart and without binarization. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 1155–1164, Lisbon, Portugal. Association for Computational Linguistics.
- Percy Liang, Michael I. Jordan, and Dan Klein. 2013. Learning dependency-based compositional semantics. *Computational Linguistics*, 39(2):389–446. **DOI:** <https://doi.org/10.1162/COLI-a-00127>
- Tal Linzen. 2020. How can we accelerate progress towards human-like linguistic generalization? In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5210–5217, Online. Association for Computational Linguistics. **DOI:** <https://doi.org/10.18653/v1/2020.acl-main.465>
- Yang Liu, Matt Gardner, and Mirella Lapata. 2018. Structured alignment networks for matching sentences. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 1554–1564, Brussels, Belgium. Association for Computational Linguistics. **DOI:** <https://doi.org/10.18653/v1/D18-1184>
- Jean Maillard, Stephen Clark, and Dani Yogatama. 2019. Jointly learning sentence embeddings and syntax with unsupervised tree-LSTMs. *Natural Language Engineering*, 25(4):433–449. **DOI:** <https://doi.org/10.1017/S1351324919000184>
- Jiayuan Mao, Chuang Gan, Pushmeet Kohli, Joshua B. Tenenbaum, and Jiajun Wu. 2019. The neuro-symbolic concept learner: Interpreting scenes, words, and sentences from natural supervision. In *International Conference on Learning Representations*.
- Richard Montague. 1970. Universal grammar. *Theoria*, 36(3):373–398. **DOI:** <https://doi.org/10.1111/j.1755-2567.1970.tb00434.x>
- Jeffrey Pennington, Richard Socher, and Christopher Manning. 2014. GloVe: Global

- vectors for word representation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1532–1543, Doha, Qatar. Association for Computational Linguistics. **DOI:** <https://doi.org/10.3115/v1/D14-1162>
- Ethan Perez, Florian Strub, Harm de Vries, Vincent Dumoulin, and Aaron Courville. 2018. FiLM: Visual reasoning with a general conditioning layer. In *AAAI Conference on Artificial Intelligence*.
- Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. 2015. Faster R-CNN: Towards real-time object detection with region proposal networks. In *Proceedings of the 28th International Conference on Neural Information Processing Systems - Volume 1*. NIPS’15, pages 91–99, Cambridge, MA, USA. MIT Press.
- Laura Ruis, Jacob Andreas, Marco Baroni, Diane Bouchacourt, and Brenden M. Lake. 2020. A benchmark for systematic generalization in grounded language understanding. *CoRR*, abs/2003.05161.
- Sanjay Subramanian, Ben Bogin, Nitish Gupta, Tomer Wolfson, Sameer Singh, Jonathan Berant, and Matt Gardner. 2020. Obtaining faithful interpretations from compositional neural networks. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5594–5608, Online. Association for Computational Linguistics. **DOI:** <https://doi.org/10.18653/v1/2020.acl-main.495>
- Sanjay Subramanian, Sameer Singh, and Matt Gardner. 2019. Analyzing compositionality in visual question answering. In *NeurIPS Workshop on Visually Grounded Interaction and Language*.
- Hao Tan and Mohit Bansal. 2019. LXMERT: Learning cross-modality encoder representations from transformers. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 5100–5111, Hong Kong, China. Association for Computational Linguistics. **DOI:** <https://doi.org/10.18653/v1/D19-1514>
- Damien Teney, Kushal Kafle, Robik Shrestha, Ehsan Abbasnejad, Christopher Kanan, and Anton van den Hengel. 2020. On the value of out-of-distribution testing: An example of goodhart’s law. *CoRR*, abs/2005.09241.
- Andrew Trask, Felix Hill, Scott E. Reed, Jack Rae, Chris Dyer, and Phil Blunsom. 2018. Neural arithmetic logic units. In *Advances in Neural Information Processing Systems*, pages 8035–8044.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *Advances in Neural Information Processing Systems 30*, pages 5998–6008. Curran Associates, Inc.
- Kexin Yi, Jiajun Wu, Chuang Gan, Antonio Torralba, Pushmeet Kohli, and Josh Tenenbaum. 2018. Neural-symbolic VQA: Disentangling reasoning from vision and language understanding. In *Advances in Neural Information Processing Systems*, volume 31, pages 1031–1042. Curran Associates, Inc.
- Daniel H. Younger. 1967. Recognition and parsing of context-free languages in time n^3 . *Information and Control*, 10(2):189–208. **DOI:** [https://doi.org/10.1016/S0019-9958\(67\)80007-X](https://doi.org/10.1016/S0019-9958(67)80007-X)
- Wojciech Zaremba and Ilya Sutskever. 2014. Learning to execute. *CoRR*, abs/1410.4615.
- Luke Zettlemoyer and Michael Collins. 2005. Learning to map sentences to logical form: Structured classification with probabilistic categorial grammars. In *UAI*.