COMPOSITIONAL ATTENTION NETWORKS FOR MACHINE REASONING

Anonymous authors
Paper under double-blind review

ABSTRACT

We present Compositional Attention Networks (CANs), a novel fully differentiable neural network architecture, designed to facilitate explicit and expressive reasoning. While many types of neural networks are effective at learning and generalizing from massive quantities of data, this model moves away from monolithic black-box architectures towards a design that provides a strong prior for iterative reasoning, enabling it to support explainable and structured learning, as well as generalization from a modest amount of data. The model builds on the great success of existing recurrent cells such as LSTMs: It sequences a single recurrent Memory, Attention, and Control (MAC) cell, and by careful design imposes structural constraints on the operation of each cell and the interactions between them, incorporating explicit control and soft attention mechanisms into their interfaces. We demonstrate the model's strength and robustness on the challenging CLEVR dataset for visual reasoning, achieving a new state-of-the-art 98.9% accuracy, halving the error rate of the previous best model. More importantly, we show that the new model is more computationally efficient, data-efficient, and requires an order of magnitude less time and/or data to achieve good results.

1 INTRODUCTION

This paper considers how best to design neural networks to perform the iterative reasoning necessary for complex problem solving. Putting facts and observations together to arrive at conclusions is a central necessary ability as we work to move neural networks beyond their current great success with sensory perception tasks (LeCun et al., 1998; Krizhevsky et al., 2012) towards displaying Artificial General Intelligence.

Concretely, we develop a novel model for the CLEVR dataset (Johnson et al., 2016) for visual question answering (VQA). VQA (Antol et al., 2015; Gupta, 2017) is a challenging multimodal task that requires responding to natural language questions about images. However, Agrawal et al. (2016) shows how the first generation of successful models on VQA tasks tend to acquire only superficial comprehension of both the image and the question, exploiting dataset biases rather than capturing a sound perception and reasoning process that would lead to the correct answer (Sturm, 2014). CLEVR was created to address this problem. As illustrated in figure 1, instances in the dataset consist of rendered images featuring 3D objects of several shapes, colors, materials and sizes, coupled with unbiased, compositional questions that require an array of challenging reasoning skills such as following transitive relations, counting objects and comparing their properties, without allowing any shortcuts around such reasoning. Importantly, each instance in CLEVR is also accompanied by a tree-structured functional program that was both used to construct the question and reflects a reasoning procedure – a series of predefined operations that can be composed together to answer it.

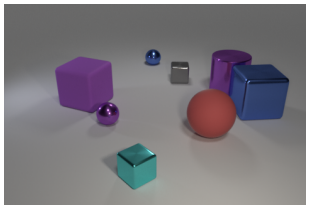

Figure 1: A sample instance from the CLEVR dataset, with a question: "There is a purple cube that is behind of a metal object right to a large ball; what material is it?"

Many neural networks are essentially very large correlation engines that will hone in on any statistical (or spurious) pattern that allows them to model the observed data more accurately. In contrast,

we seek to create a model structure that requires combining sound inference steps in solving a problem instance. At the other extreme, some approaches adopt symbolic structures that resemble the expression trees of programming languages to perform reasoning (Andreas et al., 2016b; Hu et al., 2017). In particular, some approaches to CLEVR use the supplied functional programs for supervised or semi-supervised training (Andreas et al., 2016a; Johnson et al., 2017). Not only do we wish to avoid using such supervision in our work, but we in general suspect that the rigidity of these structures and the use of an inventory of operation-specific neural modules undermines robustness and generalization, and at any rate requires the use of more complex reinforcement learning methods.

To address these weaknesses, while still seeking use of a sound and transparent underlying reasoning process, we propose Compositional Attention Networks (CANs), a novel, fully differentiable, non-modular architecture for reasoning tasks. A CAN is a straightforward recurrent neural network with attention; the novelty lies in the use of a new Memory, Attention and Composition (MAC) cell. The constrained and deliberate design of the MAC cell was developed as a kind of strong structural prior that encourages the network to solve problems by stringing together a sequence of transparent reasoning steps. MAC cells are versatile but constrained neural units. They explicitly separate out memory from control, both represented recurrently. The unit contains three subunits: The control unit updates the control representation based on outside instructions (for VQA, the question), learning to successively attend to different places in the instructions; the read unit gets stuff out of a knowledge base (for VQA, the image) based on the control signal and the previous memory; the write unit updates the memory based on soft self-attention to previous memories, controlled by what was retrieved and the control signal. A universal MAC unit with a single set of parameters is used throughout the reasoning process, but its behavior can vary widely based on the context in which it is applied – the input to the control unit (and the contents of the knowledge base). With attention, a CAN has the capacity to represent arbitrarily complex acyclic reasoning graphs (in a soft manner), while having physically sequential structure. The result is a continuous counterpart to module networks that can be trained end-to-end simply by backpropagation.

We test the behavior of our new network on CLEVR and its associated datasets. On the primary CLEVR reasoning task, we achieve an accuracy of 98.9%, halving the error rate compared to the previous state-of-the-art FiLM model (Perez et al., 2017). In particular we show that our architecture yields better performance on questions involving counting and aggregation. In supplementary studies, we show that a CAN learns more quickly (in terms of number of training epochs and training time) and more effectively from limited amounts of training data. Moreover, it also achieves a new state-of-the-art performance of 82.5% on the more varied and difficult human-authored questions of the CLEVR-Humans dataset. The careful design of our cell encourages compositionality, versatility and transparency. We achieve these properties by defining attention-based interfaces that constrict input and output spaces, and constrain the interactions both between and inside cells in order to guide them towards simple reasoning behaviors. Although each cell's functionality has only a limited range of possible continuous reasoning behaviors, when chained together in a CAN, the whole system becomes expressive and powerful. In the future, we believe that this architecture will also prove beneficial for other multi-step reasoning and inference tasks, for instance in machine comprehension and textual question answering.

2 Related Work

There have been several prominent models that address the CLEVR task. By and large they can be partitioned into two groups: module networks, which in practice have all used the strong supervision provided in the form of tree-structured functional programs that accompany each data instance, and large, relatively unstructured end-to-end differentiable networks that complement a fairly standard stack of CNNs with components that aid in performing reasoning tasks. In contrast to modular approaches (Andreas et al., 2016a;b; Hu et al., 2017; Johnson et al., 2017), our model does not require additional supervision and makes use of a single computational cell chained in sequence (like an LSTM) rather than a collection of custom modules deployed in a less flexible tree structure. In contrast to augmented CNN approaches (Santoro et al., 2017; Perez et al., 2017), we suggest that our approach provides an ability for relational reasoning with better generalization capacity and higher computational efficiency. These approaches and other related work are discussed and contrasted in more detail in the supplementary material in section C.

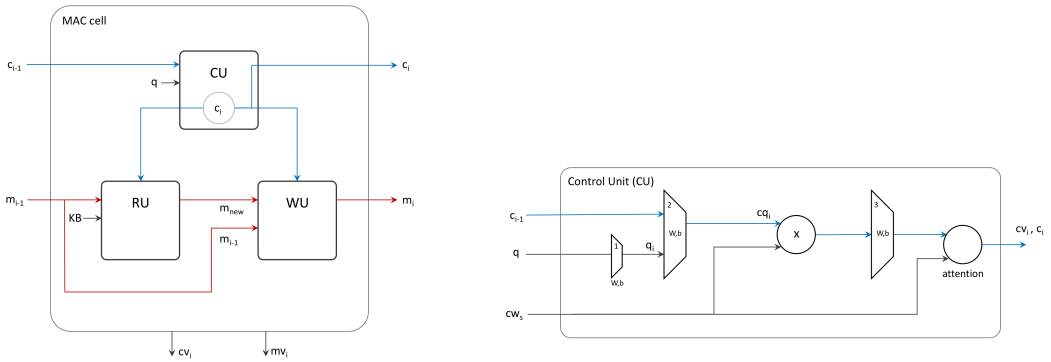

Figure 2: *Left:* The MAC cell, which is a recurrent unit comprised of a Control Unit, Read Unit, and Write Unit. Blue shows the control flow and red shows the memory flow. See section 3.2 for details. *Right:* The Control Unit (CU) of the MAC cell. See section 3.2.1 for details. Best viewed in color.

3 Compositional Attention Networks

Compositional Neural Networks is an end-to-end architecture for question-answering tasks that sequentially performs an explicit reasoning process by stringing together small building blocks, called MAC cells, each is responsible for performing one reasoning step.

We now provide an overview of the model, and a detailed discussion of the MAC unit. The model is composed of three components: an Input unit, the core MAC network, and an output unit. A TensorFlow implementation of the network, along with pretrained models will be made publicly available.

In this paper we explore the model in the context of VQA. However, it should be noted that while the input and output units are naturally domain-specific and should be designed to fit the task at hand, the MAC network has been designed to be generic and more broadly applicable, and may prove useful in contexts beyond those explored in the paper, such as machine comprehension or question answering over knowledge bases, which in our belief is a promising avenue for future work.

3.1 The Input Unit

The input unit processes the raw inputs given to the system into distributed vector representations. It receives a text question (or in general, a *query*), and an image (or in general, a *Knowledge Base* (KB)) and processes each of them with a matching sub-unit, for the query and the KB, here a biLSTM and CNNs. More details can be found in the supplementary material, section A.

At the end of this stage, we get from the query subunit a series of biLSTM output states, which we refer to as *contextual words*, $[cw_1, ..., cw_S]$, where S is the length of the question. In addition, we get $q = [\overleftarrow{cw_1}, \overrightarrow{cw_S}]$, the concatenation of the hidden states from the backward and forward LSTM passes. We refer to q as the *question representation*. Furthermore, we get from the Knowledge-Base subunit a static representation of the knowledge base. For the case of VQA, it will be represented by a continuous matrix KB_V of dimension H, W, d, where $H = W = 14$ are the height and width of the transformed image, corresponding to each of its regions.

3.2 The MAC cell

The MAC network, which is the heart of our model, chains a sequence of small building blocks, called MAC cells, each responsible for performing one reasoning step. The model is provided access to a Knowledge Base (KB), which is, for the specific case of VQA, the given image, and then upon receiving a query, i.e. a question, the model iteratively focuses, in p steps, on the query's various parts, each reflects in turn the current reasoning step, which we term the *control*. Consequently, guided by this control, it retrieves the relevant information from the KB, that is then passed to the next cell in a recurrent fashion.

Drawing inspiration from the Model-View-Controller paradigm used in software design and from the commonly exercised separation between control and data paths in computer architecture, the MAC cell is composed of three units: control unit, read unit and write unit. Each has a clearly defined role and an interface through which it interacts with the other units. See figure 2 (left).

The careful design and imposed interfaces that constrain the interaction between the units inside the MAC cell, as described below, serve as structural prior that limits the space of hypotheses it can learn, thereby guiding it towards acquiring the intended reasoning behaviors. As such, this prior facilitates the learning process and mitigate overfitting issues.

In particular, and similar in spirit to Perez et al. (2017), we allow the question to interact with the Knowledge Base – the image for the case of VQA, only through indirect means: by guiding the cell to attend to different elements in the KB, as well as controlling its operation through gating mechanisms. Thus, in both cases, the interaction between these mediums, visual and textual, or knowledge and query, is mediated through probability distributions, either in the form of attention map, or as a gate, further detailed below. This stands in stark contrast to many common approaches that fuse the question and image together into the same vector space through linear combinations, multiplication, or concatenation. Rather, our controlled interaction distills the influence that the query should have in processing the Knowledge Base, casting it onto discrete probability distributions instead.

The MAC cell has been designed to replace the discrete and predefined modules used in the modular approach (Andreas et al., 2016a;b; Hu et al., 2017; Johnson et al., 2017). Rather, we creates one uniformal and versatile cell that is applied across all the reasoning steps, sharing both its architecture as well as its parameters, across all of its instantiations. In contrast to the discrete modules, each trained to specialize to some specific elementary reasoning task, the MAC cell is capable of demonstrating a continuous range of possible reasoning behaviors conditioned on the context in which it is applied – namely, the inputs it receives from the prior cell.

Each cell MAC_i maintains two states: *control* c_i and *memory* m_i, both are continuous vectors of dimension d. The control c_i represents the reasoning operation the MAC cell should accomplish in the current step – focusing only on some aspect of the whole question. This is represented by a weighted-average summing only the attended question words. The memory m_i represents the current context information deemed relevant to respond to the query, or answer the question.This is represented practically by a weighted average over elements from the KB, or for the case of VQA, regions of the image. m_0 and c_0 are initialized to a random vector parameter of dimension d. The memory and control states are passed from one cell to the next in a recurrent fashion, and used in a way reminiscent of Key-Value memory networks (Miller et al., 2016), as discussed below.

3.2.1 The Control Unit

The control unit determines the reasoning operation that should be applied at this step. It receives the contextual words $[cw_1, ..., cw_S]$, the question representation q, and the control state from the previous MAC cell c_{i-1}, all of which are vectors of dimension d.

We would like to allow our MAC cell to perform continuously varied and adaptive range of behaviors, as demanded by the question. Therefore, we define the behavior of each cell to be a function of some of the contextual words $[cw_1, ..., cw_S]$ that the control unit chooses to attend to at this step. This will allow the cell to adapt its behavior – the reasoning operation it performs – to the question it receives, instead of having a fixed set of predefined behaviours as is the case in competing approaches Andreas et al. (2016a;b); Johnson et al. (2017).

The formal specification of the control unit is shown in figure 2. The question q is linearly transformed into a vector q_i of the same dimension, which in turn is concatenated with the previous control state c_{i-1} and linearly transformed again to a d-dimensional vector cq_i.

$$q_i = ({W_1}^{d,d})_i \cdot q + ({b_1}^d)_i$$
$$cq_i = {W_2}^{2d,d} [q_i, c_{i-1}] + {b_2}^d$$

Note that in contrast to all other parameters of the cell, which are shared across its instantiations at the different steps $i = 1, ..., p$, the parameters $\left({W_1}^{d,d}\right)_i$ and $\left({b_1}^d\right)_i$ are different for each iteration.

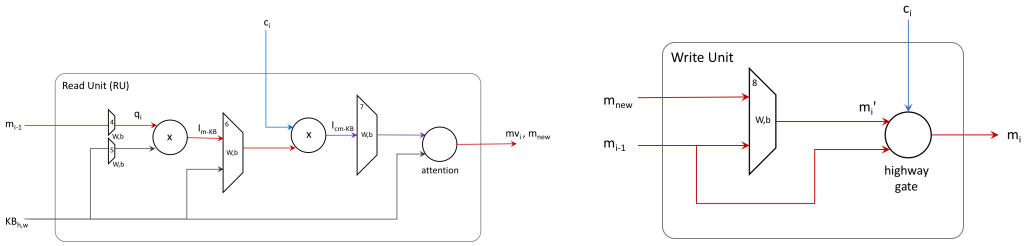

Figure 3: *Left:* The Read Unit (RU) diagram. Blue refers to control flow and red to memory flow. See section 3.2.2 for description. *Right:* The Write Unit (WU) diagram. Blue refers to control flow and red to memory flow. See section 3.2.3 for description.

This is done to allow each cell to attend more readily to different aspects of the questions, depending on the index of the current step – its relative position in the context of the whole reasoning process.

cq_i represents the current reasoning operation we would like to perform in a continuous way, taking into account both the overall meaning of the question q_i, as well as the words the model attended to in the previous step, c_{i-1}.

However, we would like to prevent the cell from diverging in the reasoning operations it tries to perform, and instead anchor it back in the question words, by using them to represent the reasoning operation of the current step. We can achieve that by computing an attention distribution cv_i over the contextual words $[cw_1, ..., cw_S]$ based on their similarity to a linear transformation of cq_i. Then, summing the contextual words according to the attention distribution cv_i will allow us to have a new control state, c_i, which is represented again in terms of words from the question. Intuitively, it is the gist of the question that is relevant to the reasoning operation we would like to perform in the current step.

$$cv_{i,s} = \text{softmax}({W_3}^{d,1}(cq_s \circ cw_s) + b_3)$$
$$c_i = \sum_{s=1}^{S} cv_s \cdot cw_s$$

Finally, the control unit returns the current control state c_i, along with an attention map cv_i over the contextual words.

3.2.2 The Read Unit

The Read Unit is provided with access to the knowledge base KB_V, along with the previous memory state m_{i-1} and the current control c_i. It is responsible for retrieving relevant content from the Knowledge Base KB_V for the reasoning task that the MAC cell should accomplish at this step, which is represented by the current control state c_i, as explained above. Figure 3 shows a diagram.

The relevance of the new information is judged in two stages by the relatedness of each element in the KB (or for the case of VQA, each region in the image) to either the memory m_{i-1} that has accumulated relevant information from previous iterations, or to the current control c_i, pointing towards the next piece of information that should be taken into account. Here, relatedness is measured by trained linear transformations comparing each element to the previous memory and the current control.

More formally, the interaction between each element $KB_{h,w}$, where $h = 1, ..., H, w = 1, ..., W$, and the previous memory m_{i-1} is computed by:

$$(I_{m-KB})_{h,w} = \left({W_4}^{d,d} m_{i-1} + {b_4}^d\right) \circ \left({W_5}^{d,d} KB_{h,w} + {b_5}^d\right)$$

These memory-KB interactions measure the relatedness of each element in the KB to the memory accumulated so far, which holds information that has been deemed relevant to handle previous

reasoning steps towards addressing the question. They allow the model to perform transitive inference, retrieving a new piece of information that now seems important in light of the recent memory retrieved in a prior iteration.

However, there are cases which necessitate the model to temporarily ignore current memories, when choosing the new information to retrieve. Logical OR is a classical example: when the model has to look at two different objects at the same time, and assuming it stored one of them at the first iteration, it should briefly *ignore* it, considering *new information that is relevant to the question but is unrelated to the memory*. In order to achieve such capability, the Read Unit concatenates the original KB elements to each corresponding memory-KB interaction, which are then projected back to d-dimensional space:

$$\left(I_{m-KB}{}'\right)_{h,w} = {W_6}^{2d,d}\left[I_{m-KB}, KB_{h,w}\right] + {b_6}^{d}$$

Finally, the read unit compares the current c_i with these memory-KB interactions, in order to focus on the information that is relevant to the current reasoning operation that the MAC cell seeks to accomplish. The result is then passed to a softmax layer yielding an attention map mv_i over the KB, which is used in turn to retrieve the relevant information to perform the current reasoning step.

$$\begin{aligned}
&I_{cm-KB} = c_i \circ \left(I_{m-KB}{}'\right)_{h,w} \\
&\left(mv_i\right)_{h,w} = \text{softmax}\left({W_7}^{d,d} I_{cm-KB} + {b_7}^{d}\right) \\
&m_i = \sum_{h,w=1,1}^{H,W} mv_i \cdot KB_{h,w}
\end{aligned}$$

Finally, the read unit returns the newly retrieved information m_{new}, along with an attention map mv_i over the Knowledge Base.

To give an example of the Read Unit operation, assume a given question q such as "What object is located left to the blue ball?", whose associated answer is "cube". Initially, no cue is provided to the model to attend to that cube, since no direct information about it presents in the question. Instead, based on its comprehension of the question, the model may start by focusing on the blue ball at the first iteration, such that the memory state m_1 will capture the blue ball. However, in the second iteration, the Control Unit, after re-examining the question, may realize it should now look left, storing the word "left" in c_2. Then, when considering *both* m_1 and c_2, the Read Unit will realize it should perform a reasoning operation corresponding to the word "left" (stored in c_2) given a memory representing the blue ball in m_1, thereby allowing it to look left to the blue ball and find the cube.

3.2.3 THE WRITE UNIT

The Write Unit is responsible for creating the new memory state m_i that will reflect all the informated considered to be important to answer the question so far, i.e. up to the current iteration in the reasoning process. It receives the last memory state m_{i-1} from the previous MAC cell, along with the newly retrieved information from the Read Unit in the current iteration, m_{new}. See figure 3 for a diagram.

In the most basic design we have explored, merging the new information with the previous memory state is done simply by a linear transformation.

$$m_i = {W_8}^{2d,d}[m_{new}, m_{i-1}] + {b_8}^{d}$$

However, there are couple of problems with this basic way to merge memories which we address by modifying the Write Unit, as presented below.

Self-Attention. The current architecture that we have presented allows the model to perform reasoning steps in a sequence, passing control and memory states from one cell to the following. However, we would like to grant the system with more flexibility. Particularly, we would like to allow it

to capture more complicated reasoning processes such as trees and graphs - Directed Acyclic Graph (DAG) in particular, where several branches of reasoning sub-processes are merged together in later stages. Indeed, the CLEVR dataset includes cases where the questions embody tree-like reasoning process, rather than just sequences, which we would like to address correctly in our model.

Modular networks approach this task by dynamically building a tailor-made tree layout to fit the given question. However, this results in non-differentiability of their overall model, since the layout-prediction task is discrete. In contrast to this approach, we would like to retain the end-to-end differentiability of our model, having universal physically static but versatile layout that will be able to capture any required reasoning layout.

We achieve this by adding self-attention connections between each MAC cell and all the prior cells. Since each cell can look on all the prior reasoning steps and their corresponding memories retrieved from the Knowledge Base, it can virtually capture any directed acyclic graph, while still having physically sequential layout.

More formally, the current MAC cell, of the ith iteration, is granted with access to $c_1, ..., c_{i-1}$ along with the corresponding $m_1, ..., m_{i-1}$, that have been computed by the prior MAC cells. It begins by computing the similarity between c_i and $c_1, ..., c_{i-1}$, and use it to derive an attention map over the prior MAC cells $SA_{i,j}$ for $j = 1, ..., i-1$. This represents the relevance of the jth prior reasoning step to the current one i.

$$SA_{ij} = \text{softmax}\left(W_9{}^{d,1}(c_i \circ c_j) + b_9\right)$$

Then, we average the previous *memories* according to this resulted attention map SA_{ij}. We obtain m_{sa}, representing the information from all the other reasoning steps that is relevant to the current one.

$$(m_{sa})_i = \sum_{j=1}^{i-1} sa_{i,j} \cdot m_j$$

This resembles the approach of Key-Value networks (Miller et al., 2016). The similarity between control states, corresponding to the reasoning operations that are performed in each prior step, allows the model to select which memories should be taken into account, when creating the new memory – namely, which branches of the reasoning process should be merged together at this point.

Finally, we use m_{sa} along with m_{i-1} and m_{new}, to compute m_i, similarly to what has been done in the basic Write Unit, presented before.

$$m_i = W_{10}{}^{3d,d}[m_{new}, m_{i-1}, m_{sa}] + b_{10}{}^d$$

Memory Gate. The currently presented MAC network has some fixed number N of concatenated MAC cells, representing the length of the overall reasoning process we perform. However, not all questions require reasoning process of the same length. Some questions are simpler while others more complex.

In order to let our network support questions with varied complexities, we add a gate over the new memory computed at each step, that may keep its previous value m_{i-1} unchanged. That way, the MAC network may skip steps when necessary – when the question demands short reasoning process – bypassing those memories intact further along the network. Overall, the gating mechanism confers the MAC network with the ability to adjust the length of the reasoning process to the complexity of the given question.

Practically, the gate functions in a similar way to a highway network (Srivastava et al., 2015), where the gate value is conditioned on the current reasoning operation, c_i.

$$m_i = W_8{}^{2d,d}[m_{new}, m_{i-1}] + b_8{}^d$$

$$m_i = \text{softmax}\,(c_i) \cdot m_{i-1} + (1 - \text{softmax}\,(c_i)) \cdot m_i{}'$$

The write unit returns the new memory state m_i, that will be passed along with c_i to the next MAC cell.

3.2.4 DISCUSSION

Overall, when designing the MAC cell, we have attempted to model and formulate the inner workings of an elementary, yet generic reasoning skills: the model decomposes the problem into steps, focusing on one at a time. At each such step, it takes into account:

- **The control** c_i: Some aspect of the task - pointing to the **future** work that has left to be done.
- **The previous memory** m_{i-1}: The partial solution or evidence the cell has acquired so far – pointing to the **past** work that has already been achieved.
- **The newly retrieved information** m_{new}: that is retrieved from the knowledge base KB and may or may not be transitively related to that partial solution or evidence - the **present**, or current work.

Considering these three sources of information together, the cell finally adds the new information up into its working memory, m_i, progressing one more step towards the final answer.

Indeed, section 4 shows clear evidence that our architecture serves as a strong prior allowing the model to learn much faster, and generalize from smaller amounts of data than competing approaches.

Furthermore, we would like to stress that our architecture can handle datasets more diverse than CLEVR, and does not assume that the data has the specific properties of this dataset. Indeed, in section 4 we demonstrate that our model is robust to linguistic variations and diverse vocabulary that may demand unconfined variety of reasoning processes. This is substantiated by its state-of-the-art performance on the CLEVR-humans dataset, which features crowd-sourced natural language questions on the given image.

3.3 THE OUTPUT UNIT

The output unit receives the question representation q, along with the memory state passed from the last MAC cell m_i for $i = p$, where p is the number of MAC cells in the network – representing the number of reasoning steps in the whole process. It inspects both and predicts an answer based on their concatenation. Intuitively, we would like our model to consider both the question as well as the relevant information that has been progressively retrieved from the KB, deemed the necessary information to answer it.

Note that considering both q and $m_{i=p}$ is critical to answer the question. While $m_{i=p}$ represents the information collected from KB, we still need to recall what has been asked about it to be able to answer accordingly. This is especially true in our case, when all other interactions between the question and the KB, are mediated through attention distributions, rather than being transformed into a shared continuous vector space.

The prediction is built out of a standard 2-layers fully-connected softmax-based classifier with hidden dimension d and output dimension that matches the number of possible answers in the dataset. The classifier receives $[m_p, q]$ as input and returns a probability distribution over the answers.

4 EXPERIMENTS

We evaluate our model on the recent CLEVR dataset (Johnson et al., 2016). CLEVR is a synthetic dataset consisting of 700K tuples; each consists of a 3D-rendered image featuring objects of various shapes, colors, materials and sizes, coupled with compositional multi-step questions that measure performance on an array of challenging reasoning skills such as following transitive relations, counting objects and comparing their properties. In addition, each question is associated with a formal program, specifying the reasoning operations that should be performed to compute the answer, among 28 possibilities.

We first perform experiments on the original 700k CLEVR dataset (Johnson et al., 2016), comparing to prior work. As shown in table 1, our model matches or outperforms all existing models both in overall accuracy, as well as in each category, testing different reasoning skills. In particular, for the overall performance, we achieve 98.94% accuracy, more than halving the error rate of the prior best model, FiLM (Perez et al., 2017).

Table 1: CLEVR accuracy by baseline methods, competing methods, and our method (CAN). (*) denotes use of extra supervisory information through program labels. (†) denotes use of data augmentation. (‡) denotes training from raw pixels.

Model	Overall	Count	Exist	Compare Numbers	Query Attribute	Compare Attribute
Human (Johnson et al., 2017)	92.6	86.7	96.6	86.5	95.0	96.0
Q-type baseline (Johnson et al., 2017)	41.8	34.6	50.2	51.0	36.0	51.3
LSTM (Johnson et al., 2017)	46.8	41.7	61.1	69.8	36.8	51.8
CNN+LSTM (Johnson et al., 2017)	52.3	43.7	65.2	67.1	49.3	53.0
CNN+LSTM+SA (Johnson et al., 2016)	76.6	64.4	82.7	77.4	82.6	75.4
N2NMN* (Hu et al. 2017)	83.7	68.5	85.7	84.9	90.0	88.7
PG+EE (9K prog.)* (Johnson et al., 2017)	88.6	79.7	89.7	79.1	92.6	96.0
PG+EE (700K prog.)* (Johnson et al., 2017)	96.9	92.7	97.1	98.7	98.1	98.9
CNN+LSTM+RN†‡ (Santoro et al., 2017)	95.5	90.1	97.8	93.6	97.9	97.1
CNN+GRU+FiLM (Perez et al., 2017)	97.7	94.3	99.1	96.8	99.1	99.1
CNN+GRU+FiLM‡ (Perez et al., 2017)	97.6	94.3	99.3	93.4	**99.3**	99.3
CAN (this paper)	**98.9**	**97.2**	**99.5**	**99.4**	**99.3**	**99.5**

Counting and Numerical Comparison. Remarkably, our performance on questions testing counting and numerical comparisons is significantly higher than the competing models, which consistently struggle on this question type. Again, we nearly halve the corresponding error rate. These results demonstrate the aptitude of attention mechanisms to perform counting, reduction and aggregation, in contrast to alternative, CNN-based approaches.

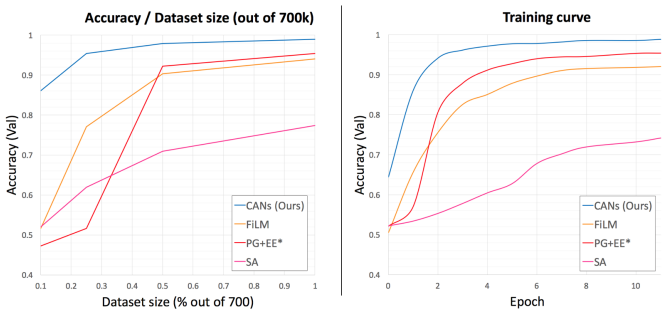

Figure 4: Training curves and accuracies for CANs (our model), FiLM (Perez et al., 2017), PG+EE (Johnson et al., 2017) and stacked-attention (Yang et al., 2016; Johnson et al., 2017). (Note: PG+EE uses the supported CLEVR programs as strong supervision.) *Left:* Training curve (accuracy/epoch). *Right:* Learning curve: Accuracy for 10%, 25%, 50% and 100% of the 700k CLEVR samples.

Training Length and Computational-Efficiency. We examine the learning curves of our and competing models. As shown in figure 4, our model learns significantly faster than the other leading methods, FiLM (Perez et al., 2017) and PG+EE (Johnson et al., 2017). While we do not have learning curves for the Relational Network model, Santoro et al. (2017) report approximately 1.4 million iterations to achieve 95.5% accuracy, which are equivalent to 125 epochs approximately, whereas our model achieves a comparable accuracy after 3 epochs only, yielding 50x reduction in the length of the training process.

Naturally, the smaller number of required training steps also translates to comparably shorter training time. Perez et al. (2017) report training time of 4 days, equivalent to 80 epochs, to reach accuracy of 97.7%. In contrast, we achieve higher accuracy in 6 epochs, taking 9.5 hours overall, leading to 10x reduction in training time.

4.1 Data Efficiency

We have explored the performance of our and other leading approaches on smaller subsets of the CLEVR dataset, in order to study the ability of models to generalize from smaller amount of data. We sampled at random subsets of CLEVR, with 10%, 25% and 50% of its original 700k size, and used them to train our and other 3 proposed models for the CLEVR task: FiLM

Table 2: Accuracy on CLEVR-Humans of previous methods and our method (CAN), before (left) and after (right) fine-tuning on the CLEVR-Humans training data. PG+EE uses supervised data.

Model	Train CLEVR	Train CLEVR + fine-tune HUMANS
LSTM (Johnson et al., 2017)	27.5	36.5
CNN+LSTM (Johnson et al., 2017)	37.7	43.2
CNN+LSTM+SA+MLP (Johnson et al., 2016)	50.4	57.6
PG+EE (18K prog.)* (Johnson et al., 2017)	54.0	66.6
CNN+GRU+FiLM (Perez et al., 2017)	56.6	75.9
CAN (this paper)	**58.6**	**82.5**

(Perez et al., 2017), the strongly-supervised PG+EE (Johnson et al., 2017), and stacked-attention networks (Johnson et al., 2017; Yang et al., 2016).

As shown in figure 4, our model outperforms the other models by a wide margin for all subsets of the CLEVR dataset. For 50% of the data, equivalent to 350k samples, other models obtain accuracies ranging between 70% and 92%, while our model achieves 97.9%. The gap becomes larger as the dataset size reduces: for 25% of the data, equivalent to 175k samples, performance of other models is between 50% and 77%, while our model maintains a high 95.4% accuracy.

Finally, for 10% of the data – 70k samples, still a sizeable amount – our model is the only one that manages to generalize, with performance of 86%, whereas the other three models completely fail, achieving 51.6% at best. Note that as pointed out by (Johnson et al., 2016) a simple baseline that predicts the most frequent answer for each of the question types achieves already 42.1%, suggesting that answering half of the questions correctly means that the competing models barely learn to generalize from the smaller dataset. These results demonstrate the robustness of our architecture and its key role as a structural prior guiding our network to learn the intended reasoning skills.

4.2 CLEVR Humans - Natural Language Questions

We analyze our model performance on the CLEVR-Humans dataset (Johnson et al., 2017), consisting of natural language questions collected through crowdsourcing. As such, the dataset has diverse vocabulary and linguistic variations, and it also demands more varied reasoning skills.

Since the training set is relatively small, consisting of 18k samples, we use it to finetune a model pretrained on the standard CLEVR dataset. However, since most of the vocabulary in CLEVR-Humans is not covered by CLEVR, we do not train the word vectors during the pre-training stage, so to prevent drift in their meaning compared to other uncovered words in CLEVR-Humans that may be semantically related.

As shown in table 2, our model achieves state-of-the-art performance on CLEVR-Humans both before and after fine-tuning. It surpasses the next-best FiLM model, (Perez et al., 2017) by 6.6% percent, achieving 82.5%.

The results substantiate the model's robustness against linguistic variations and noise, as well as its ability to adapt to diverse vocabulary and varied reasoning skills. Arguably, the soft attention performed over the question words allows the model to focus on the words that are most critical to answer the question and translate them to corresponding reasoning operations, giving less attention to irrelevant linguistic variations.

5 Conclusion

We have given a first demonstration of how a sequence of Memory, Attention and Control (MAC) cells combined into a Compositional Attention Network (CAN) provides a very effective tool for neural reasoning. In future work, we wish to explore this promising architecture for other tasks and domains, including real-world VQA, machine comprehension and textual question answering.

REFERENCES

Aishwarya Agrawal, Dhruv Batra, and Devi Parikh. Analyzing the behavior of visual question answering models. *arXiv preprint arXiv:1606.07356*, 2016.

Jacob Andreas, Marcus Rohrbach, Trevor Darrell, and Dan Klein. Neural module networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 39–48, 2016a.

Jacob Andreas, Marcus Rohrbach, Trevor Darrell, and Dan Klein. Learning to compose neural networks for question answering. *arXiv preprint arXiv:1601.01705*, 2016b.

Stanislaw Antol, Aishwarya Agrawal, Jiasen Lu, Margaret Mitchell, Dhruv Batra, C Lawrence Zitnick, and Devi Parikh. Vqa: Visual question answering. In *Proceedings of the IEEE International Conference on Computer Vision*, pp. 2425–2433, 2015.

Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E Hinton. Layer normalization. *arXiv preprint arXiv:1607.06450*, 2016.

Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473*, 2014.

Djork-Arné Clevert, Thomas Unterthiner, and Sepp Hochreiter. Fast and accurate deep network learning by exponential linear units (elus). *arXiv preprint arXiv:1511.07289*, 2015.

Alex Graves, Navdeep Jaitly, and Abdel-rahman Mohamed. Hybrid speech recognition with deep bidirectional lstm. In *Automatic Speech Recognition and Understanding (ASRU), 2013 IEEE Workshop on*, pp. 273–278. IEEE, 2013.

Alex Graves, Greg Wayne, and Ivo Danihelka. Neural turing machines. *arXiv preprint arXiv:1410.5401*, 2014.

Alex Graves, Greg Wayne, Malcolm Reynolds, Tim Harley, Ivo Danihelka, Agnieszka Grabska-Barwińska, Sergio Gómez Colmenarejo, Edward Grefenstette, Tiago Ramalho, John Agapiou, et al. Hybrid computing using a neural network with dynamic external memory. *Nature*, 538 (7626):471–476, 2016.

Akshay Kumar Gupta. Survey of visual question answering: Datasets and techniques. *arXiv preprint arXiv:1705.03865*, 2017.

Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778, 2016.

Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8): 1735–1780, 1997.

Ronghang Hu, Jacob Andreas, Marcus Rohrbach, Trevor Darrell, and Kate Saenko. Learning to reason: End-to-end module networks for visual question answering. *arXiv preprint arXiv:1704.05526*, 2017.

Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *International Conference on Machine Learning*, pp. 448–456, 2015.

Justin Johnson, Bharath Hariharan, Laurens van der Maaten, Li Fei-Fei, C Lawrence Zitnick, and Ross Girshick. Clevr: A diagnostic dataset for compositional language and elementary visual reasoning. *arXiv preprint arXiv:1612.06890*, 2016.

Justin Johnson, Bharath Hariharan, Laurens van der Maaten, Judy Hoffman, Li Fei-Fei, C Lawrence Zitnick, and Ross Girshick. Inferring and executing programs for visual reasoning. *arXiv preprint arXiv:1705.03633*, 2017.

Diederik Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.

Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pp. 1097–1105, 2012.

Ankit Kumar, Ozan Irsoy, Peter Ondruska, Mohit Iyyer, James Bradbury, Ishaan Gulrajani, Victor Zhong, Romain Paulus, and Richard Socher. Ask me anything: Dynamic memory networks for natural language processing. In *International Conference on Machine Learning*, pp. 1378–1387, 2016.

Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.

Jiasen Lu, Jianwei Yang, Dhruv Batra, and Devi Parikh. Hierarchical question-image co-attention for visual question answering. In *Advances In Neural Information Processing Systems*, pp. 289–297, 2016.

Alexander Miller, Adam Fisch, Jesse Dodge, Amir-Hossein Karimi, Antoine Bordes, and Jason Weston. Key-value memory networks for directly reading documents. *arXiv preprint arXiv:1606.03126*, 2016.

Jeffrey Pennington, Richard Socher, and Christopher Manning. Glove: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, pp. 1532–1543, 2014.

Ethan Perez, Florian Strub, Harm de Vries, Vincent Dumoulin, and Aaron Courville. Film: Visual reasoning with a general conditioning layer. *arXiv preprint arXiv:1709.07871*, 2017.

Adam Santoro, David Raposo, David GT Barrett, Mateusz Malinowski, Razvan Pascanu, Peter Battaglia, and Timothy Lillicrap. A simple neural network module for relational reasoning. *arXiv preprint arXiv:1706.01427*, 2017.

Rupesh Kumar Srivastava, Klaus Greff, and Jürgen Schmidhuber. Highway networks. *arXiv preprint arXiv:1505.00387*, 2015.

Bob L Sturm. A simple method to determine if a music information retrieval system is a horse. *IEEE Transactions on Multimedia*, 16(6):1636–1644, 2014.

Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. *arXiv preprint arXiv:1706.03762*, 2017.

Caiming Xiong, Stephen Merity, and Richard Socher. Dynamic memory networks for visual and textual question answering. In *International Conference on Machine Learning*, pp. 2397–2406, 2016.

Zichao Yang, Xiaodong He, Jianfeng Gao, Li Deng, and Alex Smola. Stacked attention networks for image question answering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 21–29, 2016.

SUPPLEMENTARY MATERIAL

A DETAILS OF INPUT UNIT

The input unit processes the raw inputs given to the system into distributed vector representations. It receives a text question (or in general, a *query*), and an image (or in general, a *Knowledge Base* (KB)) and processes each of them with a matching sub-unit. Here we provide details of the Query Unit and the Image Unit used in this work.

A.0.1 THE QUERY UNIT

We encode a query of S words into a continuous representation using a bidirectional LSTM (Hochreiter & Schmidhuber, 1997; Graves et al., 2013). Each word is associated with a word embedding w_s, where $s = 1, ..., S$. In our case, we use GloVE words embeddings (Pennington et al., 2014). Then, these embeddings are processed by a bidirectional $LSTM$ of dimension d that outputs:

- a matching sequence of d-dimensional output states, which we refer to as *contextual words*, $[cw_1, ..., cw_S]$
- d-dimensional hidden state $q = [\overleftarrow{cw_1}, \overrightarrow{cw_S}]$, the concatenation of the hidden states from the backward and forward passes. We refer to q as the *question representation*.

Intuitively, each contextual word cw_s represents the meaning of sth word, in the context of the question, while the hidden state q represents the overall (compositional) meaning of the question.

A.0.2 THE IMAGE UNIT

Given an image, and following prior work on CLEVR (**?**Santoro et al., 2017; Perez et al., 2017), we extract conv4 features from ResNet101 (He et al., 2016) pretrained on ImageNet (Krizhevsky et al., 2012) which we treat as a fixed initial representation of the image, x of dimension H, W, C where $H = W = 14$ are the height and width of the transformed image and $C = 1024$ is the number of channels. Each feature $x_{h,w}$ represents one region in the original image.

Similar to prior work (Hu et al., 2017; Santoro et al., 2017; Perez et al., 2017), we would like to allow our model to reason explicitly about spatial locations, as required by many of the questions in CLEVR, and therefore we concatenate to this representation a spatial map that represents each of the positions in the image. However, in contrast to prior work that uses a linear meshgrid feature map with 2 features h and w ranging from -1 to 1, and to allow better representation of the positions, we use the positional encoding scheme proposed by Vaswani et al. (2017):

$$p_{(h,2i)} = \sin\left(h/10000^{2i/p_d}\right)$$
$$p_{(h,2i+1)} = \cos\left(h/10000^{2i/p_d}\right)$$

And similarly for w, where p_d is a hyperparameter. Overall, the positional encoding of a feature at position (h, w) is $[p_h, p_w]$, the concatenation of the positional encodings for h and w.

This positional encoding scheme allows better correspondence between the distance of 2 positions (x, y) and (x, y) in the image and a vector similarity of their positional encodings, even when p_d is larger than two.

We then concatenate the obtained spatial map with x, receiving a spatially-aware image representation x_p. Then, we pass this representation through two CNN layers with d output channels and obtain a final representation of the image, which we refer to as our *Visual Knowledge Base* (KB_V that is used in further components of the model.

B IMPLEMENTATION AND TRAINING DETAILS

For the question processing, we use GloVE (Pennington et al., 2014) word-vectors with dimension 300. For the image processing, we extract conv4 features from ResNet101 (He et al., 2016) pre-

trained on ImageNet (Krizhevsky et al., 2012), with dimension H, W, C where $H = W = 14$ and $C = 1024$, followed by 2 CNN layers with kernel size 2. We use MAC network with $p = 12$ cells, and train it using Adam (Kingma & Ba, 2014), with learning rate $(10)^{(} - 4)$. We train our model for $10 - 20$ epochs, with batch size 64, and use early stopping based on validation accuracies. During training, the moving averages of all weights of the model are maintained with the exponential decay rate of 0.999. At test time, the moving averages instead of the raw weights are used. We use dropout 0.85, and ELU (Clevert et al., 2015) which in our experience has reduce the training process compared to RELU.The training process takes roughly 10-20 hours on a single Titan X GPU.

C FURTHER DISCUSSION OF RELATED WORK

In this section we provide detailed discussion of related work. Several models have been applied to the CLEVR task. These can be partitioned into two groups, module networks that use the strong supervision provided as a tree-structured functional program associated with each instance, and end-to-end, fully differentiable networks that combine a fairly standard stack of CNNs with components that aid them in performing reasoning tasks. We also discuss the relation of CANs to other approaches, such as memory networks and neural computers.

C.1 MODULE NETWORKS

The modular approach (Andreas et al., 2016a;b; Hu et al., 2017; Johnson et al., 2017) first translates the given question into a tree-structured action plan, aiming to imitate the ground-truth programs provided as a form of strong-supervision. Then, it constructs a tailor-made network that executes the plan on the image in multiple steps. This network is composed of discrete units selected out of a collection of predefined modules, each responsible for an elementary reasoning operation, such as identifying an objects color, filtering them for their shape, or comparing two amounts. Each module has its own set of learned parameters (Johnson et al., 2017), or even hand-crafted design (Andreas et al., 2016a) to guide it towards its intended behavior.

Overall, this approach makes discrete choices at two levels: the identity of each module – the behavior it should learn among a fixed set of possible types of behaviors, and the network layout – the way in which these modules are wired together to compute the answer progressively. Hence, their differentiability is confined to the boundaries of a single module, disallowing end-to-end training.

Several key differences exist between our approaches. First, our model replaces the fixed modules collection with one versatile and universal cell that shares both its architecture and parameters across all of its instantiations, and is applied across all the reasoning steps. Second, it replaces the dynamic recursive tree structures with a sequential topology, augmented by soft attention mechanisms, as done in Bahdanau et al. (2014). This confers our network with a virtual capacity to represent arbitrarily complex Directed Acyclic Graphs (DAGs) while still having efficient and readily deployed physical sequential structure. Together, both of these relaxations allow us to effectively train our model end-to-end by backpropagation alone, whereas module networks demand a more involved training scheme that relies on the strongly-supervised programs at the first stage, and on various Reinforcement Learning (RL) techniques at the second. Furthermore, while our model can be train without the strong supervisory programs, developing adaptive reasoning skills to address the task is it trained for, the modular approach reliance on questions structured and formal representation hinder its applicability to real-world tasks.

C.2 AUGMENTED CONVOLUTIONAL NEURAL NETWORKS

Alternative approaches for the CLEVR task that do not rely on the provided programs as a strong supervision signal are Santoro et al. (2017) and Perez et al. (2017). Both complement standard multi-layer Convolutional Neural Networks (CNNs) with components that aid them in handling compositional and relational questions.

Relational Networks. Santoro et al. (2017) appends a Relation Network (RN) layer to the CNN. This layer inspects all pairs of pixels in the image, thereby enhancing the network capacity to reason over binary relations between objects. While this approach is very simple and elegant conceptually, it suffers from quadratic computational complexity, in contrast to our and other leading approaches.

But beyond that, closer inspection reveals that this direct pairwise comparison might be unnecessary. Based on the analogy suggested by Santoro et al. (2017), according to which pixels are equivalent to objects and their pairwise interactions to relations, a RN layer attempts to grasp the induced graph between objects all at once in one shallow and broad layer. Conversely, our attention-based model proceeds in steps. It basically compares the image to its current memory and control for this step, aggregates the attended regions into the new memory, and repeats the process. By the same analogy, it traverses a narrow and deep path, progressively following transitive relations. Consequently, our model exhibits a relational capacity while circumventing the computational inefficiency.

FiLM. FiLM (Perez et al., 2017) is a recently proposed method that interleaves standard CNN layers that process the given image with linear layers, reminiscent of layer normalization techniques (Ba et al., 2016; Ioffe & Szegedy, 2015). Each of these layers, called FiLM, is conditioned on the question: the question words are processed by a GRU, and its output is linearly transformed into matching biases and variances for each of the CNN layers, tilting its activations to reflect the specifics of the given question and affect the computation done over the image.

Similarly to our model, this approach features distant modulation between the question and the image, where rather than being fused together into the same vector space, the question can affect the image processing only through constrained means – for the case of FiLM – linear transformations. However, since the same transformation is applied to all the activations homogeneously, agnostic to both their spatial location as well as the features values, this approach does not allow the question to differentiate between regions in the image based on the objects or concepts they represent – on the content of the image. This stands in stark contrast to our attention-based model, which readily allows and actually encourages the question to inform the model about relevant regions to focus on. We speculate that this still distant, yet more direct interaction between the question and the data, or image, for the case of VQA, facilitates learning and increases generalizability. It may be more suitable to VQA tasks, and CLEVR in particular, where the questions demand the responder to focus on specific objects, and reason about their properties or relations, rather than respond based only on a holistic view of the image that may lead to sub-optimal results (Yang et al., 2016), as is the case of FiLM. Indeed, as demonstrated in 4, there is significant evidence showing our models better generalization capacity, allowing it to achieve high accuracies much faster, and from less data than FiLM and other competing methods.

C.3 Memory and Attention

Our architecture draws inspiration from recent research on memory and attention (Kumar et al., 2016; Xiong et al., 2016; Graves et al., 2014; 2016). Kumar et al. (2016); Xiong et al. (2016) propose the Dynamic Memory Network model that proceeds in an iterative process, applying soft attention to retrieve relevant information from a visual or textual KB, which is in turn accumulated into memory passed from one iteration to the next. However, in contrast to our model, it views the question as an atomic unit, whereas our model decomposes it into a multi-step action plan informing each cell in our sequential network about its current objective. Another key difference is the distant interaction between the question and the KB that characterizes our model. Conversely, DMN fuses their corresponding representations together into the same vector space.

Graves et al. (2016; 2014) complements a neural network with a memory array it can interact with, through the means of soft attention. Analogously to our model, it partitions the model into a core neural network, called controller, as well as reading and writing heads that interact with external memory array. However, a main point distinguishing our model from this approach, is the use of dynamic memory, as in Kumar et al. (2016), instead of a fixed-array memory. Each MAC cell is associated with a memory state, our reading unit inspects only the latest memory passed from the previous state, and our writing unit creates a new memory state rather than writing to multiple slots in a fixed shared external memory. Notably, our approach is much more reminiscent of the widely successful RNN structure, rather than to Graves et al. (2016; 2014) .

Finally, our approach has potential ties to the VQA models Hu et al. (2017); Lu et al. (2016) which also attend both the to question words and the image while progressively addressing the given question. However, both of these models have distinct specialized designs for each of their attention layers or modules, and have a discrete or fixed layout in which they are composed together. In con-

trast, our approach relax both of these limitations, having one universal cell design and one universal self-attending sequential network layout.

C.4 ATTENTION VS. CONVOLUTION

Compared to other leading methods, our model stands out by being heavily based on soft attention, whereas most competing approaches are CNN-based, surprisingly lack any attention mechanism. Since attention is commonly used in models designed for standard VQA (Antol et al., 2015; Gupta, 2017; Lu et al., 2016; Yang et al., 2016), it is reasonable to assume that it would be beneficial to incorporate such methods into visual reasoning systems for the CLEVR task as well. In fact, attention mechanisms should be especially useful for multi-step reasoning questions such as those present in CLEVR. Such questions refer to several relations between different objects in the image and feature compositional structure that may be approached one step at a time. Thus, it should be beneficial for a cogent responder to have the capacity to selectively focus on on one or some objects at each step, traversing the relevant relational links one after the other, both at the image level, and at the question level.

Moreover, attention mechanisms enhance our model's ability to perform reasoning skills that pertain to aggregation of information across different regions, such as counting, finding maximum value, or performing other reduction operations over information that is spread across the image. Indeed, as discussed in 4, all existing models for visual reasoning, most of which lacking any attention mechanism, struggle with the counting and numerical comparisons questions present in CLEVR. Conversely, our model proves much more capable of performing these reasoning skills, outperforming the other approaches by a wide margin. Noticeably, incorporating soft attention into our model makes it much more adept at performing such aggregation reasoning skills, successfully addressing the this type of questions.

Finally, as pointed out by Lu et al. (2016); Yang et al. (2016), soft attention confers the model with robustness to noise introduced from irrelevant information presents in the image, and higher capacity for handling larger and more diverse vocabulary, the latter being demonstrated in 4. It allows the model to separate the wheat from the chaff, selectively attending to the relevant information only, and arguably, being more resilient to both visual and linguistic variations.