

# Cascaded Mutual Modulation for Visual Reasoning

Yiqun Yao<sup>123</sup>, Jiaming Xu<sup>12\*</sup>, Feng Wang<sup>1</sup> and Bo Xu<sup>1234</sup>

<sup>1</sup>Institute of Automation, Chinese Academy of Sciences (CASIA), Beijing, China

<sup>2</sup>Research Center for Brain-inspired Intelligence, CASIA

<sup>3</sup>University of Chinese Academy of Sciences

<sup>4</sup>Center for Excellence in Brain Science and Intelligence Technology, CAS, China

{yaoyiqun2014, jiaming.xu, feng.wang, xubo}@ia.ac.cn

## Abstract

Visual reasoning is a special visual question answering problem that is multi-step and compositional by nature, and also requires intensive text-vision interactions. We propose CMM: Cascaded Mutual Modulation as a novel end-to-end visual reasoning model. CMM includes a multi-step comprehension process for both question and image. In each step, we use a Feature-wise Linear Modulation (FiLM) technique to enable textual/visual pipeline to mutually control each other. Experiments show that CMM significantly outperforms most related models, and reach state-of-the-arts on two visual reasoning benchmarks: CLEVR and NLVR, collected from both synthetic and natural languages. Ablation studies confirm that both our multi-step framework and our visual-guided language modulation are critical to the task. Our code is available at <https://github.com/FlamingHorizon/CMM-VR>.

## 1 Introduction

It is a challenging task in artificial intelligence to perform reasoning with both textual and visual inputs. Visual reasoning task is designed for researches in this field. It is a special visual question answering (VQA) (Antol et al., 2015) problem, requiring a model to infer the relations between entities in both image and text, and generate a textual answer to the question correctly. Unlike other VQA tasks, questions in visual reasoning often contain extensive logical phenomena, and refer to multiple entities, specific attributes and complex relations. Visual reasoning datasets such as CLEVR (Johnson et al., 2017a) and NLVR (Suhr et al., 2017) are built on unbiased, synthetic images, with either complex synthetic questions or natural-language descriptions, facilitating in-depth analyses on reasoning ability itself.

\* Corresponding Author

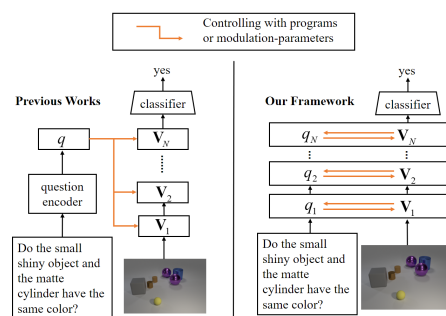


Figure 1: Connections and differences between previous “program-generating” works and our model: other models generate/control multi-step image-comprehension processes with single question representation, while we put more attention on language logics and let multi-modal information modulate each other in each step. The question and image are taken as a visual-reasoning example from CLEVR dataset.

Most previous visual reasoning models focus on using the question to guide the multi-step computing on visual features (which can be defined as a image-comprehension “program”). Neural Module Networks (NMN) (Andreas et al., 2016a,b) and Program Generator + Execution Engine (PG+EE) (Johnson et al., 2017b) learn to combine specific image-processing modules, guided by question semantics. Feature-modulating methods like FiLM (De Vries et al., 2017; Perez et al., 2018) control image-comprehension process using modulation-parameters generated from the question, allowing models to be trained end-to-end. However, the image-comprehension program in visual reasoning tasks can be extremely long and sophisticated. Using a single question representation to generate or control the whole image-comprehension process raises difficulties in learning. Moreover, since information comes from multiple modalities, it is not intuitive to assume

that one (language) is the “program generator”, and the other (image) is the “executor”. One way to avoid making this assumption is to perform multiple steps of reasoning with each modality being generator and executor alternately in each step. For these two reasons, we propose Cascaded Mutual Modulation (Figure 1), a novel visual reasoning model to solve the problem that previous “program-generating” models lack a method to use visual features to guide multi-step reasoning on language logics. CMM reaches state-of-the-arts on two benchmarks: CLEVR (complex synthetic questions) and NLVR (natural-language).

## 2 Related Work

Perez et al. (2018) proposed FiLM as an end-to-end feature-modulating method. The original ResBlock+GRU+FiLM structure uses single question representation, and conditions all image-modulation-parameters on it, without sufficiently handling multi-step language logics. In contrast, we modulate both image and language features alternately in each step, and condition the modulation-parameters on the representations from previous step. We design an image-guided language attention pipeline and use it in combination with FiLM in our CMM framework, and significantly outperform the original structure.

Other widely-cited works on CLEVR/NLVR include Stacked Attention Networks (SAN) (Yang et al., 2016), NMN (Andreas et al., 2016b), N2NMN (Hu et al., 2017), PG+EE (Johnson et al., 2017b) and Relation Networks (RN) (Santoro et al., 2017). The recent CAN model (Hudson and Manning, 2018) also uses multiple question representations and has strong performances on CLEVR. However, these representations are not modulated by the visual part as in our model.

In other VQA tasks, DAN (Nam et al., 2017) is the only multi-step dual framework related to ours. For comparison, in every time step, DAN computes textual and visual attention *in parallel* with the same key-vector, while we perform textual attention and visual modulation (instead of attention) in a *cascaded* manner.

## 3 Model

We review and extend FiLM in Section 3.1-3.2, and introduce CMM model in Section 3.3-3.4.

### 3.1 Visual Modulation

Perez et al. (2018) proposed Feature-wise Linear Modulation (FiLM), an affine transformation on intermediate outputs of a neural network ( $v$  stands for visual):

$$FiLM^v(\mathbf{F}_{i,c}|\gamma_{i,c},\beta_{i,c}) = \gamma_{i,c}\mathbf{F}_{i,c} + \beta_{i,c}, \quad (1)$$

where  $\mathbf{F}_{i,c}$  is the  $c$ -th feature map ( $C$  in total) generated by Convolutional Neural Networks (CNN) in the  $i$ -th image-comprehension step. Modulation-parameters  $\gamma_{i,c}$  and  $\beta_{i,c}$  can be conditioned on any other part of network (in their work the *single* question representation  $q$ ). If the output tensor  $\mathbf{V}_i$  of a CNN block is of size  $C \times H \times W$ , then  $\mathbf{F}_{i,c}$  is a single slice of size  $1 \times H \times W$ .  $H$  and  $W$  are the height and width of each feature map.

Unlike (Perez et al., 2018), in each step  $i$ , we compute a new question vector  $q_i$ . Modulation-parameters  $\gamma_i$  and  $\beta_i$  ( $C \times 1$  vectors,  $\gamma_i = [\gamma_{i,1}, \dots, \gamma_{i,C}]$ , etc.) are conditioned on the previous question vector  $q_{i-1}$  instead of a single  $q$ :

$$\gamma_i, \beta_i = MLP^i(q_{i-1}). \quad (2)$$

MLP stands for fully connected layers with linear activations. The weights and biases are not shared among all steps.

### 3.2 Language Modulation

In each step  $i$ , we also apply FiLM to modulate every language “feature map”. If the full question representation is a  $D \times T$  matrix, a question “feature map”  $\mathbf{f}_{i,d}$  is defined as a  $1 \times T$  vector gathering  $T$  features along a single dimension.  $D$  is the hidden-state dimension of language encoder, and  $T$  is a fixed maximum length of word sequences.

$$FiLM^l(\mathbf{f}_{i,d}|\gamma_{i,d},\beta_{i,d}) = \gamma_{i,d}\mathbf{f}_{i,d} + \beta_{i,d}, \quad (3)$$

where  $l$  stands for language. Concatenated modulation-parameters  $\gamma_i$  and  $\beta_i$  ( $D \times 1$ ) are conditioned on the visual features  $\mathbf{V}_i$  computed in the same step:

$$\gamma_i, \beta_i = g_m(\mathbf{V}_i), \quad (4)$$

where  $g_m$  (Section 3.4) is an interaction function that converts 3-d visual features to language-modulation-parameters. The weights in  $g_m$  are shared among all  $N$  steps.

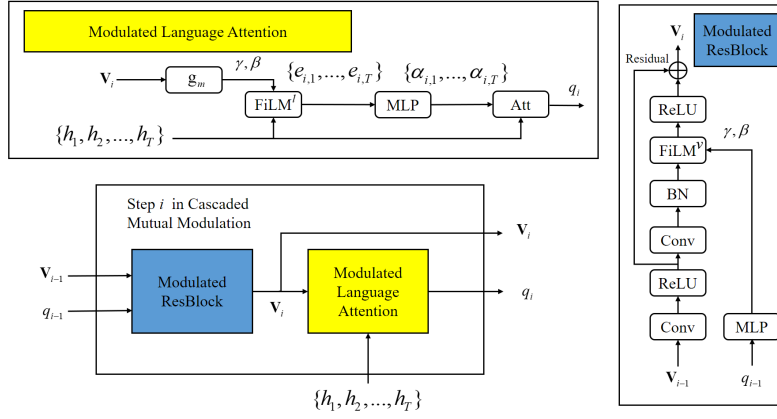


Figure 2: Details in CMM step  $i$  (middle), with a modulated ResBlock (right) and a modulated textual attention pipeline (top). Visual and textual features modulate each other in each step to compute new representations.

### 3.3 Cascaded Mutual Modulation

The whole pipeline of our model is built up with multiple *steps*. In each step  $i$  ( $N$  in total), previous question vector  $q_{i-1}$  and the visual features  $\mathbf{V}_{i-1}$  are taken as input;  $q_i$  and  $\mathbf{V}_i$  are computed as output. Preprocessed questions/images are encoded by language/visual encoders to form  $q_0$  and  $\mathbf{V}_0$ .

In each step, we cascade a FiLM-ed ResBlock with a modulated textual-attention. We feed  $\mathbf{V}_{i-1}$  into the ResBlock modulated by parameters from  $q_{i-1}$  to compute  $\mathbf{V}_i$ , and then control the textual attention process with modulation-parameters from  $\mathbf{V}_i$ , to compute the new question vector  $q_i$ . (Figure 2, middle).

Each ResBlock contains a  $1 \times 1$  convolution, a  $3 \times 3$  convolution, a batch-normalization (Ioffe and Szegedy, 2015) layer before FiLM modulation, followed by a residual connection (He et al., 2016). (Figure 2, right. We keep the same ResBlock structure as (Perez et al., 2018)). To be consistent with (Johnson et al., 2017b; Perez et al., 2018), we concatenate the input visual features  $\mathbf{V}_{i-1}$  of each ResBlock  $i$  with two “coordinate feature maps” scaled from  $-1$  to  $1$ , to enrich representations of spatial relations. All CNNs in our model use ReLU as activation functions; batch-normalization is applied before ReLU.

After the ResBlock pipeline, we apply language modulation on the full language features  $\{h_1, \dots, h_T\}$  ( $D \times T$ ) conditioned on  $\mathbf{V}_i$  and rewrite along the time dimension, yielding:

$$e_{i,t} = \text{FiLM}^l(h_t | g_m(\mathbf{V}_i)), \quad (5)$$

and compute visual-guided attention weights:

$$\alpha_{i,t} = \text{softmax}_t(\mathbf{W}_i^{\text{att}} e_{i,t} + b_i^{\text{att}}), \quad (6)$$

and weighted summation over time:

$$q_i = \sum_{t=1}^T \alpha_{i,t} h_t. \quad (7)$$

In equation (6),  $W_i^{\text{att}} \in \mathbb{R}^{1 \times D}$  and  $b_i^{\text{att}} \in \mathbb{R}^{1 \times 1}$  are network weights and bias;  $h_t$  is the  $t$ -th language state vector ( $D \times 1$ ), computed using a bi-directional GRU (Chung et al., 2014) from word embeddings  $\{w_1, \dots, w_T\}$ . In each step  $i$ , the language pipeline does not re-compute  $h_t$ , but re-modulate it as  $e_{i,t}$  instead. (Figure 2, top).

### 3.4 Feature Projections

We use a function  $g_p$  to project the last visual features  $\mathbf{V}_N$  into a final representation:

$$u_{\text{final}} = g_p(\mathbf{V}_N). \quad (8)$$

$g_p$  includes a convolution with  $K$   $1 \times 1$  kernels, a batch-normalization afterwards, followed by global max pooling over all pixels ( $K = 512$ ).

We also need a module  $g_m$  (equation (4)) to compute language-modulations with  $\mathbf{V}_i$ , since  $\mathbf{V}_i$  is 3-d features (not a weighted-summed vector as in traditional visual-attention). We choose  $g_m$  to have the same structure as  $g_p$ , except that  $K$  equals to the total number of modulation-parameters in each step. This design is critical (Section 4.3).

We use a fully connected layer with 1024 ReLU hidden units as our answer generator. It takes  $u_{\text{final}}$  as input, and predicts the most probable answer in the answer vocabulary.

Model	Overall	Count	Exist	Compare Numbers	Query Attribute	Compare Attribute
Human	92.6	86.7	96.6	86.5	95.0	96.0
SAN (Yang et al., 2016)	76.7	64.4	82.7	77.4	82.6	75.4
N2NMN (Andreas et al., 2016a)	83.7	68.5	85.7	84.9	90.0	88.7
PG+EE-9K	88.6	79.7	89.7	79.1	92.6	96.0
PG+EE-700K (Johnson et al., 2017b)	96.9	92.7	97.1	98.7	98.1	98.9
RN (Santoro et al., 2017)	95.5	90.1	97.8	93.6	97.9	97.1
COG-model (Yang et al., 2018)	96.8	91.7	99.0	95.5	98.5	98.8
FiLM	97.7	94.3	99.1	96.8	99.1	99.1
FiLM-raw (Perez et al., 2018)	97.6	94.3	99.3	93.4	99.3	99.3
DDRprog (Suarez et al., 2018)	98.3	96.5	98.8	98.4	99.1	99.0
CAN (Hudson and Manning, 2018)	98.9	97.1	<b>99.5</b>	<b>99.1</b>	99.5	<b>99.5</b>
CMM-single (ours)	98.6	96.8	99.2	97.7	99.4	99.1
CMM-ensemble (ours)	<b>99.0</b>	<b>97.6</b>	<b>99.5</b>	98.5	<b>99.6</b>	99.4

Table 1: Accuracies on CLEVR test set. N2NMN and PG+EE need extra supervision to train with reinforcement learning. FiLM-raw uses raw image as input (others use pre-extracted features). Another work (Mascharka et al., 2018) gets 99.1% accuracy but uses strong program supervision, which is a totally different setting.

## 4 Experiments

We are the first to achieve top results on both datasets (CLEVR, NLVR) with one structure. See Appendix for more ablation and visualization results.

### 4.1 CLEVR

CLEVR (Johnson et al., 2017a) is a commonly-used visual reasoning benchmark containing 700,000 training samples, 150,000 for validation and test. Questions in CLEVR cover several typical elements of reasoning: counting, comparing, querying the memory, etc. Many well-designed models on VQA have failed on CLEVR, revealing the difficulty to handle the multi-step and compositional nature of logical questions.

On CLEVR dataset, we embed the question words into a 200-dim continuous space, and use a bi-directional GRU with 512 hidden units to generate 1024-dim question representations. Questions are padded with NULL token to a maximum length  $T = 46$ . As the first-step question vector in CMM,  $q_0$  can be arbitrary RNN hidden state in the set  $\{h_1, \dots, h_T\}$  (Section 3.3). We choose the one at the end of the unpadded question.

In each ResBlock, the feature map number  $C$  is set to 128. Images are pre-processed with a ResNet101 network pre-trained on ImageNet (Russakovsky et al., 2015) to extract  $1024 \times 14 \times 14$  visual features (this is also common practice on CLEVR). We use a trainable one-layer CNN with 128 kernels ( $3 \times 3$ ) to encode the extracted features into  $\mathbf{V}_0$  ( $128 \times 14 \times 14$ ). Convolutional paddings are used to keep the feature map size to be  $14 \times 14$  through the visual pipeline.

We train the model with an ADAM (Kingma and Ba, 2014) optimizer using a learning rate of  $2.5e-4$  and a batch-size of 64 for about 90 epoches, and switch to an SGD with the same learning rate and 0.9 momentum, fine-tuning for another 20 epoches. SGD generally brings around 0.3 points gains to CMM on CLEVR.

We achieve 98.6% accuracy with single model (4-step), significantly better than FiLM and other related work, only slightly lower than CAN, but CAN needs at least 8 model-blocks for  $>98\%$  (and 12 for best). We achieve state-of-the-art of 99.0% with ensemble of 4/5/6 step CMM models. Table 1 shows test accuracies on all types of questions. The main improvements over program-generating models come from ‘‘Counting’’ and ‘‘Compare Numbers’’, indicating that CMM framework significantly enhances language (especially numeric) reasoning without sophisticated memory design like CAN.

### 4.2 NLVR

NLVR (Suhr et al., 2017) is a visual reasoning dataset proposed by researchers in NLP field. NLVR has 74,460 samples for training, 5,940 for validation and 5,934 for public test. In each sample, there is a human-posed natural language description on an image with 3 sub-images, and requires a false/true response.

We use different preprocessing methods on NLVR. Before training, we reshape NLVR images into  $14 \times 56$  raw pixels and use them directly as visual inputs  $\mathbf{V}_0$ . For language part, we correct some obvious typos among the rare words (frequency  $< 5$ ) in the training set, and pad the sentences to a maximum length of 26. Different from



CLEVR, LSTM works better than GRU on the real-world questions. For training, we use ADAM with a learning rate of  $3.5e-4$  and a batch-size of 128 for about 200 epoches, without SGD fine-tuning.

Our model (3-step, 69.9%) outperforms all proposed models on both validation and public test set, showing that CMM is also suitable for real-world languages (Table 2).

Model	Dev	Test-P
Text only	56.6	57.2
Image only	55.4	56.1
CNN+RNN (Suhr et al., 2017)	56.6	58.0
NMN (Andreas et al., 2016b)	63.1	66.1
FiLM (our run)	59.0	61.3
CNN-BiAtt (Tan and Bansal, 2018)	66.9	69.7
CMM-3-steps (ours)	<b>68.0</b>	<b>69.9</b>

Table 2: Accuracies on valid and test set of NLVR.

### 4.3 Ablation Studies

We list CMM ablation results in Table 3. Ablations on CLEVR show that CMM is robust to step number but sensitive to  $g_m$  structure because it’s the key to multi-modal interaction. Section 3.4 is temporarily a best choice. CMM performs over 7-point higher than FiLM on NLVR in a setting of same hyper-parameters and ResBlocks, showing the importance of handling *language* logics (see also difficult CLEVR subtasks in Table 1).

CLEVR		NLVR		
Model	Val	Model	Dev	Test-P
5-step	98.4	FiLM-hyp	59.0	61.3
6-step	98.4	1-step	65.3	66.9
$g_m$ -CNN	93.3	2-step	67.7	66.8
$g_m$ w/o BN	94.1	3-step	68.5	68.4
$g_m$ -NS	97.0	4-step	67.2	66.5
4-step	98.6	3-step-LSTM	68.0	69.9

Table 3: Ablation studies on CLEVR/NLVR.  $g_m$ -CNN means using 2-layer-CNN with  $3 \times 3$  kernels, followed by concatenation and MLP, as  $g_m$ . BN means batch-normalization in  $g_m$ . NS means not sharing weights. “FiLM-hyp” uses all the same hyper-parameters as the 3-step CMM (both use GRU as question encoder).

### 4.4 A Case Study

We select an image-question pair from the validation set of CLEVR for visualization. In Table 4, we visualize the multi-step attention weights on question words, and the distribution of argmax position in the global max-pooling layer of  $g_p$  (equiv-

According to NLVR rules, we will run on the unreleased test set (Test-U) in the near future.

Words	Block 1	Block 2	Block 3	Visual Attention Map
<START>	0.017	0.030	0.052	
is	0.006	0.048	0.501	
there	0.010	0.041	0.299	
a	0.006	0.031	0.015	
big	0.008	0.031	0.001	
brown	0.049	0.040	0.002	
object	0.149	0.265	0.003	
of	0.032	0.063	0.003	
the	0.030	0.029	0.004	
same	0.062	0.032	0.002	
shape	0.345	0.152	0.007	
as	0.233	0.031	0.001	
the	0.000	0.005	0.004	
green	0.003	0.014	0.004	
thing	0.012	0.075	0.011	
<END>	0.006	0.036	0.025	

Table 4: Visualization of CMM intermediate outputs on a sample from CLEVR validation set. We colour the largest attention weight with dark gray, and top four attention weights in the rest with light gray.

alent to the last visual “attention map” although there isn’t explicit visual attention in our image-comprehension pipeline). On the bottom right is the original image, and on the top right is the distribution of argmax positions in the global max-pooling, multiplied by the original image.

Our model attends to phrases “same shape as” and “brown object” in the first two reasoning steps. These phrases are meaningful because “same shape as” is the core logic in the question, and “brown object” is the key entity to generating the correct answer. In the last step, the model attends to the phrase “is there”. This implicitly classifies the question into question-type “exist”, and directs the answer generator to answer “no” or “yes”. The visual map, guided by question-based modulation parameters, concentrates on the green and brown object correctly.

The result shows that visual features can guide the comprehension of question logics with textual modulation. On the other hand, question-based modulation parameters enable the ResBlocks to filter out irrelative objects.

## 5 Conclusion

We propose CMM as a novel visual reasoning model cascading visual and textual modulation in each step. CMM reaches state-of-the-arts on visual reasoning benchmarks with both synthetic and real-world languages.

## Acknowledgements

We thank the reviewers for their insightful comments. This work is supported by the National Natural Science Foundation of China (61602479) and the Strategic Priority Research Program of the Chinese Academy of Sciences (XDBS01070000).

## References

- Jacob Andreas, Marcus Rohrbach, Trevor Darrell, and Dan Klein. 2016a. Learning to compose neural networks for question answering. In *Proceedings of NAACL-HLT*, pages 1545–1554.
- Jacob Andreas, Marcus Rohrbach, Trevor Darrell, and Dan Klein. 2016b. Neural module networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 39–48.
- Stanislaw Antol, Aishwarya Agrawal, Jiasen Lu, Margaret Mitchell, Dhruv Batra, C Lawrence Zitnick, and Devi Parikh. 2015. Vqa: Visual question answering. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2425–2433.
- Junyoung Chung, Caglar Gulcehre, KyungHyun Cho, and Yoshua Bengio. 2014. Empirical evaluation of gated recurrent neural networks on sequence modeling. *arXiv preprint arXiv:1412.3555*.
- Harm De Vries, Florian Strub, Jérémie Mary, Hugo Larochelle, Olivier Pietquin, and Aaron C Courville. 2017. Modulating early visual processing by language. In *Advances in Neural Information Processing Systems*, pages 6597–6607.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778.
- Ronghang Hu, Jacob Andreas, Marcus Rohrbach, Trevor Darrell, and Kate Saenko. 2017. Learning to reason: End-to-end module networks for visual question answering. *CoRR, abs/1704.05526*, 3.
- Drew A Hudson and Christopher D Manning. 2018. Compositional attention networks for machine reasoning. In *International Conference on Learning Representations*.
- Sergey Ioffe and Christian Szegedy. 2015. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *International conference on machine learning*, pages 448–456.
- 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 *IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2017*, pages 1988–1997. IEEE.
- 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 *Proceedings of the IEEE International Conference on Computer Vision*.
- Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- David Mascharka, Philip Tran, Ryan Soklaski, and Arjun Majumdar. 2018. Transparency by design: Closing the gap between performance and interpretability in visual reasoning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*.
- Hyeonseob Nam, Jung-Woo Ha, and Jeonghee Kim. 2017. Dual attention networks for multimodal reasoning and matching. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 299–307.
- Ethan Perez, Florian Strub, Harm De Vries, Vincent Dumoulin, and Aaron Courville. 2018. Film: Visual reasoning with a general conditioning layer. In *Proceedings of the 32nd AAAI Conference on Artificial Intelligence*.
- Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. 2015. Imagenet large scale visual recognition challenge. *International Journal of Computer Vision*, 115(3):211–252.
- Adam Santoro, David Raposo, David G Barrett, Mateusz Malinowski, Razvan Pascanu, Peter Battaglia, and Tim Lillicrap. 2017. A simple neural network module for relational reasoning. In *Advances in neural information processing systems*, pages 4974–4983.
- Joseph Suarez, Justin Johnson, and Fei-Fei Li. 2018. Ddrprog: A clevr differentiable dynamic reasoning programmer. *arXiv preprint arXiv:1803.11361*.
- Alane Suhr, Mike Lewis, James Yeh, and Yoav Artzi. 2017. A corpus of natural language for visual reasoning. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, volume 2, pages 217–223.
- Hao Tan and Mohit Bansal. 2018. Object ordering with bidirectional matchings for visual reasoning. In *Proceedings of NAACL-HLT*.
- Guangyu Robert Yang, Igor Ganichev, Xiao-Jing Wang, Jonathon Shlens, and David Sussillo. 2018. A dataset and architecture for visual reasoning with a working memory. *arXiv preprint arXiv:1803.06092*.
- Zichao Yang, Xiaodong He, Jianfeng Gao, Li Deng, and Alex Smola. 2016. Stacked attention networks for image question answering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 21–29.