

Object Ordering with Bidirectional Matchings for Visual Reasoning

Hao Tan and Mohit Bansal
UNC Chapel Hill
{haotan, mbansal}@cs.unc.edu

Abstract

Visual reasoning with compositional natural language instructions, e.g., based on the newly-released Cornell Natural Language Visual Reasoning (NLVR) dataset, is a challenging task, where the model needs to have the ability to create an accurate mapping between the diverse phrases and the several objects placed in complex arrangements in the image. Further, this mapping needs to be processed to answer the question in the statement given the ordering and relationship of the objects across three similar images. In this paper, we propose a novel end-to-end neural model for the NLVR task, where we first use joint bidirectional attention to build a two-way conditioning between the visual information and the language phrases. Next, we use an RL-based pointer network to sort and process the varying number of unordered objects (so as to match the order of the statement phrases) in each of the three images and then pool over the three decisions. Our model achieves strong improvements (of 4-6% absolute) over the state-of-the-art on both the structured representation and raw image versions of the dataset.

1 Introduction

Visual Reasoning (Antol et al., 2015; Andreas et al., 2016; Bisk et al., 2016; Johnson et al., 2017) requires a sophisticated understanding of the compositional language instruction and its relationship with the corresponding image. Suhr et al. (2017) recently proposed a challenging new NLVR task and dataset in this direction with natural and complex language statements that have to be classified as true or false given a multi-image set (shown in Fig. 1). Specifically, each task instance consists of an image with three sub-images and a statement which describes the image. The model is asked to answer the question whether the given statement is consistent with the image or not.

To solve the task, the designed model needs to fuse the information from two different domains,

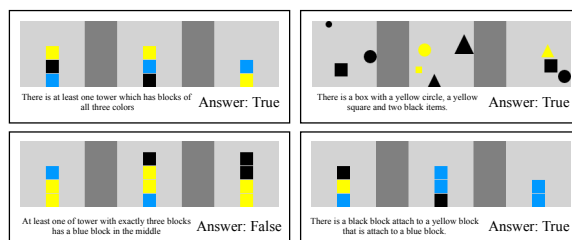


Figure 1: NLVR task: given an image with 3 sub-images and a statement, the model needs to predict whether the statement correctly describes the image or not. We show 4 such examples which our final BiATT-Pointer model correctly classifies but the strong baseline models do not (see Sec. 5).

the visual objects and the language, and learn accurate relationships between the two. Another difficulty is that the objects in the image do not have a fixed order and the number of objects also varies. Moreover, each statement reasons for truth over three sub-images (instead of the usual single image setup), which also breaks most of the existing models. In our paper, we introduce a novel end-to-end model to address these three problems, leading to strong gains over the previous best model. Our pointer network based LSTM-RNN sorts and learns recurrent representations of the objects in each sub-image, so as to match it better with the order of the phrases in the language statement. For this, it employs an RL-based policy gradient method with a reward extracted from the subsequent comprehension model. With these strong representations of the visual objects and the statement units, a joint-bidirectional attention flow model builds consistent, two-way matchings between the representations in different domains. Finally, since the scores computed by the bidirectional attention are about the three sub-images, a pooling combination layer over the three sub-image representations is required to give the final score of the whole image.

On the structured-object-representation version of the dataset, our pointer-based, end-to-end bidi-

rectional attention model achieves an accuracy of 73.9%, outperforming the previous (end-to-end) state-of-the-art method by 6.2% absolute, where both the pointer network and the bidirectional attention modules contribute significantly. We also contribute several other strong baselines for this new NLVR task based on Relation Networks (Santoro et al., 2017) and BiDAF (Seo et al., 2016). Furthermore, we also show the result of our joint bidirectional attention model on the raw-image version (with pixel-level, spatial-filter CNNs) of the NLVR dataset, where our model achieves an accuracy of 69.7% and outperforms the previous best result by 3.6%. On the unreleased leaderboard test set, our model achieves an accuracy of 71.8% and 66.1% on the structured and raw-image versions, respectively, leading to 4% absolute improvements on both tasks.

2 Related work

Besides the NLVR corpus with a focus on complex and natural compositional language (Suhr et al., 2017), other useful visual reasoning datasets have been proposed for navigation and assembly tasks (MacMahon et al., 2006; Bisk et al., 2016), as well as for visual Q&A tasks which focus more on complex real-world images (Antol et al., 2015; Johnson et al., 2017). Specifically for the NLVR dataset, previous models have incorporated property- and count-based features of the objects and the language (Suhr et al., 2017), or extra semantic parsing (logical form) annotations (Goldman et al., 2017) – we focus on end-to-end models for this visual reasoning task.

Attention mechanism (Bahdanau et al., 2014; Luong et al., 2015; Xu et al., 2015) has been widely used for conditioned language generation tasks. It is further used to learn alignments between different modalities (Lu et al., 2016; Wang and Jiang, 2016; Seo et al., 2016; Andreas et al., 2016; Chaplot et al., 2017). In our work, a bidirectional attention mechanism is used to learn a joint representation of the visual objects and the words by building matchings between them.

Pointer network (Vinyals et al., 2015) was introduced to learn the conditional probability of an output sequence. Bello et al. (2016) extended this to near-optimal combinatorial optimization via reinforcement learning. In our work, a policy gradient based pointer network is used to “sort” the objects conditioned on the statement, such that the sequence of ordered objects is sent to the subse-

quent comprehension model for a reward.

3 Model

The training datum for this task consists of the statement s , the structured-representation objects o in the image I , and the ground truth label y (which is 1 for true and 0 for false). Our BiATT-Pointer model (shown in Fig. 2) for the structured-representation task uses the pointer network to sort the object sequence (optimized by policy gradient), and then uses the comprehension model to calculate the probability $P(s, o)$ of the statement s being consistent with the image. Our CNN-BiATT model for the raw-image I dataset version is similar but learns the structure directly via pixel-level, spatial-filter CNNs – details in Sec. 5 and the appendix. In the remainder of this section, we first describe our BiATT comprehension model and then the pointer network.

3.1 Comprehension Model with Joint Bidirectional Attention

We use one bidirectional LSTM-RNN (Hochreiter and Schmidhuber, 1997) (denoted by LANG-LSTM) to read the statement $s = w_1, w_2, \dots, w_T$, and output the hidden state representations $\{h_i\}$. A word embedding layer is added before the LSTM to project the words to high-dimension vectors $\{\tilde{w}_i\}$.

$$h_1, h_2, \dots, h_T = \text{LSTM}(\tilde{w}_1, \tilde{w}_2, \dots, \tilde{w}_T) \quad (1)$$

The raw features of the objects in the j -th sub-image are $\{o_k^j\}$ (since the NLVR dataset has 3 sub-images per task). A fully-connected (FC) layer without nonlinearity projects the raw features to object embeddings $\{e_k^j\}$. We then go through all the objects in random order (or some learnable order, e.g., via our pointer network, see Sec. 3.2) by another bidirectional LSTM-RNN (denoted by OBJ-LSTM), whose output is a sequence of vectors $\{g_k^j\}$ which is used as the (left plus right memory) representation of the objects (the objects in different sub-images are handled separately):

$$e_k^j = W o_k^j + b \quad (2)$$

$$g_1^j, g_2^j, \dots, g_{N_j}^j = \text{LSTM}(e_1^j, e_2^j, \dots, e_{N_j}^j) \quad (3)$$

where N_j is the number of the objects in j th sub-image. Now, we have two vector sequences for the representations of the words and the objects, using which the bidirectional attention then calculates the score measuring the correspondence be-

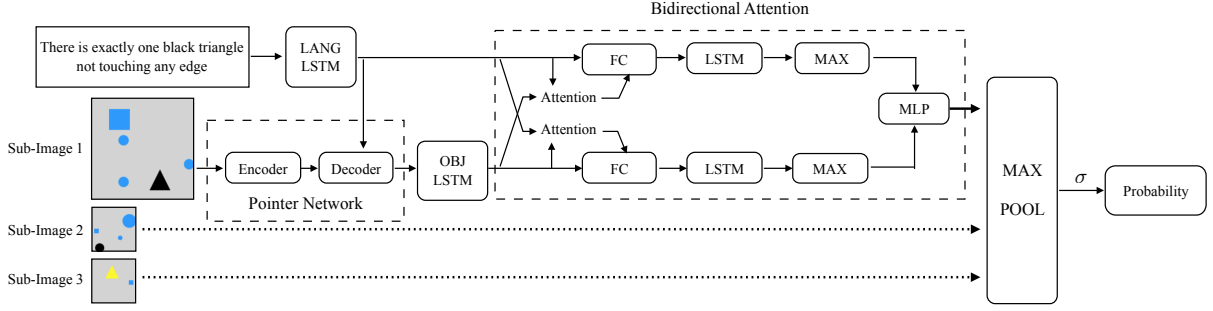


Figure 2: Our BiATT-Pointer model with a pointer network and a joint bidirectional attention module.

tween the statement and the image’s object structure. To simplify the notation, we will ignore the sub-image index j . We first merge the LANG-LSTM hidden outputs $\{h_i\}$ and the object-aware context vectors $\{c_i\}$ together to get the joint representation $\{\hat{h}_i\}$. The object-aware context vector c_i for a particular word w_i is calculated based on the bilinear attention between the word representation h_i and the representations of the objects $\{g_k\}$:

$$\alpha_{i,k} = \text{softmax}_k (h_i^\top B_1 g_k) \quad (4)$$

$$c_i = \sum_k \alpha_{i,k} \cdot g_k \quad (5)$$

$$\hat{h}_i = \text{relu} (W_{\text{LANG}} [h_i; c_i; h_i \circ c_i]) \quad (6)$$

where the symbol \circ denotes element-wise multiplication.

Improvement over BiDAF The BiDAF model of Seo et al. (2016) does not use a full object-to-words attention mechanism. The query-to-document attention module in BiDAF added the attended-context vector to the document representation instead of the query representation. However, the inverse attention from the objects to the words is important in our task because the representation of the object depends on its corresponding words. Therefore, different from the BiDAF model, we create an additional ‘symmetric’ attention to merge the OBJ-LSTM hidden outputs $\{g_k\}$ and the statement-aware context vectors $\{d_k\}$ together to get the joint representation $\{\hat{g}_k\}$. The improvement (6.1%) of our BiATT model over the BiDAF model is shown in Table 1.

$$\beta_{k,i} = \text{softmax}_i (g_k^\top B_2 h_i) \quad (7)$$

$$d_k = \sum_i \beta_{k,i} \cdot h_i \quad (8)$$

$$\hat{g}_k = \text{relu} (W_{\text{OBJ}} [g_k; d_k; g_k - d_k; g_k \circ d_k]) \quad (9)$$

These above vectors $\{\hat{h}_i\}$ and $\{\hat{g}_k\}$ are the representations of the words and the objects which

are aware of each other bidirectionally. To make the final decision, two additional bidirectional LSTM-RNNs are used to further process the above attention-based representations via an additional memory-based layer. Lastly, two max pooling layers over the hidden output states create two single-vector outputs for the statement and the sub-image, respectively:

$$\bar{h}_1, \bar{h}_2, \dots, \bar{h}_T = \text{LSTM}(\hat{h}_1, \hat{h}_2, \dots, \hat{h}_T) \quad (10)$$

$$\bar{g}_1, \bar{g}_2, \dots, \bar{g}_N = \text{LSTM}(\hat{g}_1, \hat{g}_2, \dots, \hat{g}_N) \quad (11)$$

$$\bar{h} = \text{ele max}_i \{\bar{h}_i\} \quad (12)$$

$$\bar{g} = \text{ele max}_k \{\bar{g}_k\} \quad (13)$$

where the operator ele max denotes the element-wise maximum over the vectors. The final scalar score for the sub-image is given by a 2-layer MLP over the concatenation of \bar{h} and \bar{g} as follows:

$$\text{score} = W_2 \tanh (W_1 [\bar{h}; \bar{g}] + b_1) \quad (14)$$

Max-Pooling over Sub-Images In order to address the 3 sub-images present in each NLVR task, a max-pooling layer is used to combine the above-defined scores of the sub-images. Given that the sub-images do not have any specific ordering among them (based on the data collection procedure (Suhr et al., 2017)), a pooling layer is suitable because it is permutation invariant. Moreover, many of the statements are about the existence of a special object or relationship in one sub-image (see Fig. 1) and hence the max-pooling layer effectively captures the meaning of these statements. We also tried other combination methods (mean-pooling, concatenation, LSTM, early pooling on the features/vectors, etc.); the max pooling (on scores) approach was the simplest and most effective method among these (based on the dev set).

The overall probability that the statement correctly describes the full image (with three sub-images) is the sigmoid of the final max-pooled

score. The loss of the comprehension model is the negative log probability (i.e., the cross entropy):

$$P(s, o) = \sigma \left(\max_j \text{score}^j \right) \quad (15)$$

$$L(s, o, y) = -y \log P(s, o) - (1 - y) \log(1 - P(s, o)) \quad (16)$$

where y is the ground truth label.

3.2 Pointer Network

Instead of randomly ordering the objects, humans look at the objects in an appropriate order w.r.t. their reading of the given statement and after the first glance of the image. Following this idea, we use an additional pointer network (Vinyals et al., 2015) to find the best object ordering for the subsequent language comprehension model. The pointer network contains two RNNs, the encoder and the decoder. The encoder reads all the objects in a random order. The decoder then learns a permutation π of the objects’ indices, by recurrently outputting a distribution over the objects based on the attention over the encoder hidden outputs. At each time step, an object is sampled without replacement following this distribution. Thus, the pointer network models a distribution $p(\pi | s, o)$ over all the permutations:

$$p(\pi | s, o) = \prod_i p(\pi(i) | \pi(< i), s, o) \quad (17)$$

Furthermore, the appropriate order of the objects depends on the language statement, and hence the decoder importantly attends to the hidden outputs of the LANG-LSTM (see Eqn. 1).

The pointer network is trained via reinforcement learning (RL) based policy gradient optimization. The RL loss $L_{\text{RL}}(s, o, y)$ is defined as the expected comprehension loss (expectation over the distribution of permutations):

$$L_{\text{RL}}(s, o, y) = E_{\pi \sim p(\cdot | s, o)} L(s, o[\pi], y) \quad (18)$$

where $o[\pi]$ denotes the permuted input objects for permutation π , and L is the loss function defined in Eqn. 16. Suppose that we sampled a permutation π^* from the distribution $p(\pi | s, o)$; then the above RL loss could be optimized via policy gradient methods (Williams, 1992). The reward R is the negative loss of the subsequent comprehension model $L(s, o[\pi^*], y)$. A baseline b is subtracted from the reward to reduce the variance (we use the

self-critical baseline of Rennie et al. (2016)). The gradient of the loss L_{RL} could then be approximated as:

$$R = -L(s, o[\pi^*], y) \quad (19)$$

$$\nabla_{\theta} L_{\text{RL}}(s, o, y) \approx -(R - b) \nabla_{\theta} \log p(\pi^* | s, o) + \nabla_{\theta} L(s, o[\pi^*], y) \quad (20)$$

This overall BiATT-Pointer model (for the structured-representation task) is shown in Fig. 2.

4 Experimental Setup

We evaluate our model on the NLVR dataset (Suhr et al., 2017), for both the structured and raw-image versions. All model tuning was performed on the dev set. Given the fact that the dataset is balanced (the number of true labels and false labels are roughly the same), the accuracy of the whole corpus is used as the metric. We only use the raw features of the statement and the objects with minimal standard preprocessing (e.g., tokenization and UNK replacement; see appendix for reproducibility training details).

5 Results and Analysis

Results on Structured Representations Dataset:

Table 1 shows our primary model results. In terms of previous work, the state-of-the-art result for end-to-end models is ‘MAXENT’, shown in Suhr et al. (2017).¹ Our proposed BiATT-Pointer model (Fig. 2) achieves a 6.2% improvement on the public test set and a 4.0% improvement on the unreleased test set over this SotA model. To show the individual effectiveness of our BiATT and Pointer components, we also provide two ablation results: (1) the bidirectional attention BiATT model without the pointer network; and (2) our BiENC baseline model without any attention or the pointer mechanisms. The BiENC model uses the similarity between the last hidden outputs of the LANG-LSTM and the OBJ-LSTM as the score (Eqn. 14).

Finally, we also reproduce some recent popular frameworks, i.e., Relationship Network (Santoro et al., 2017) and BiDAF model (Seo et al., 2016), which have been proven to be successful in other machine comprehension and visual reasoning tasks. The results of these models are weaker than our proposed model. Reimplementation details are shown in the appendix.

¹There is also recent work by Goldman et al. (2017), who use extra, manually-labeled semantic parsing data to achieve a released/unreleased test accuracy of 80.4%/83.5%, resp.

Model	Dev	Test-P	Test-U
STRUCTURED REPRESENTATIONS DATASET			
MAXENT (Suhr et al., 2017)	68.0%	67.7%	67.8%
MLP (Suhr et al., 2017)	67.5%	66.3%	65.3%
ImageFeat+RNN (Suhr et al., 2017)	57.7%	57.6%	56.3%
RelationNet (Santoro et al., 2017)	65.1%	62.7%	-
BiDAF (Seo et al., 2016)	66.5%	68.4%	-
BiENC Model	65.1%	63.4%	-
BiATT Model	72.6%	72.3%	-
BiATT-Pointer Model	74.6%	73.9%	71.8%
RAW IMAGE DATASET			
CNN+RNN (Suhr et al., 2017)	56.6%	58.0%	56.3%
NMN (Suhr et al., 2017)	63.1%	66.1%	62.0%
CNN-BiENC Model	58.7%	58.7%	-
CNN-BiATT Model	66.9%	69.7%	66.1%

Table 1: Dev, Test-P (public), and Test-U (unreleased) results of our model on the structured-representation and raw-image datasets, compared to the previous SotA results and other reimplemented baselines.

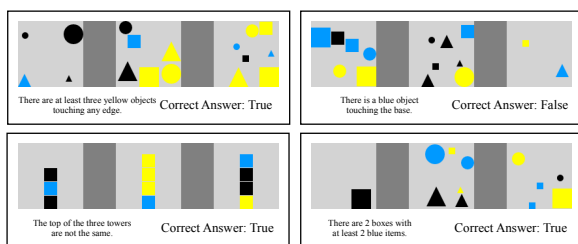


Figure 3: Incorrectly-classified examples.

Results on Raw Images Dataset: To further show the effectiveness of our BiATT model, we apply this model to the raw image version of the NLVR dataset, with minimal modification. We simply replace each object-related LSTM with a visual feature CNN that directly learns the structure via pixel-level, spatial filters (instead of a pointer network which addresses an unordered sequence of structured object representations). As shown in Table 1, this CNN-BiATT model outperforms the neural module networks (NMN) (Andreas et al., 2016) previous-best result by 3.6% on the public test set and 4.1% on the unreleased test set. More details and the model figure are in the appendix.

Output Example Analysis: Finally, in Fig. 1, we show some output examples which were successfully solved by our BiATT-Pointer model but failed in our strong baselines. The left two examples in Fig. 1 could not be handled by the BiENC model. The right two examples are incorrect for the BiATT model without the ordering-based pointer network. Our model can quite successfully understand the complex meanings of the attributes and their relationships with the diverse objects, as well as count the occurrence of and reason over objects without any specialized features.

Next, in Fig. 3, we also show some negative examples on which our model fails to predict the correct answer. The top two examples involve com-

plex high-level phrases e.g., “touching any edge” or “touching the base”, which are hard for an end-to-end model to capture, given that such statements are rare in the training data. Based on the result of the validation set, the max-pooling layer is selected as the combination method in our model. The max-pooling layer will choose the highest score from the sub-images as the final score. Thus, the layer could easily handle statements about single-subimage-existence based reasoning (e.g., the 4 positively-classified examples in Fig. 1). However, the bottom two negatively-classified examples in Fig. 3 could not be resolved because of the limitation of the max-pooling layer on scenarios that consider multiple-subimage-existence. We did try multiple other pooling and combination methods, as mentioned in Sec. 3.1. Among these methods, the concatenation, early pooling and LSTM-fusion approaches might have the ability to solve these particular bottom-two failed statements. In our future work, we are addressing multiple types of pooling methods jointly.

6 Conclusion

We presented a novel end-to-end model with joint bidirectional attention and object-ordering pointer networks for visual reasoning. We evaluate our model on both the structured-representation and raw-image versions of the NLVR dataset and achieve substantial improvements over the previous end-to-end state-of-the-art results.

Acknowledgments

We thank the anonymous reviewers for their helpful comments. This work was supported by a Google Faculty Research Award, a Bloomberg Data Science Research Grant, an IBM Faculty Award, and NVidia GPU awards.

References

- Jacob Andreas, Marcus Rohrbach, Trevor Darrell, and Dan Klein. 2016. 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.
- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2014. Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473*.
- Irwan Bello, Hieu Pham, Quoc V Le, Mohammad Norouzi, and Samy Bengio. 2016. Neural combinatorial optimization with reinforcement learning. *arXiv preprint arXiv:1611.09940*.
- Yonatan Bisk, Deniz Yuret, and Daniel Marcu. 2016. Natural language communication with robots. In *HLT-NAACL*, pages 751–761.
- Devendra Singh Chaplot, Kanthashree Mysore Sathyendra, Rama Kumar Pasumarthi, Dheeraj Rajagopal, and Ruslan Salakhutdinov. 2017. Gated-attention architectures for task-oriented language grounding. *arXiv preprint arXiv:1706.07230*.
- Omer Goldman, Veronica Latcinnik, Udi Naveh, Amir Globerson, and Jonathan Berant. 2017. Weakly-supervised semantic parsing with abstract examples. *arXiv preprint arXiv:1711.05240*.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Identity mappings in deep residual networks. In *European Conference on Computer Vision*, pages 630–645. Springer.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation*, 9(8):1735–1780.
- Justin Johnson, Bharath Hariharan, Laurens van der Maaten, Li Fei-Fei, C Lawrence Zitnick, and Ross Girshick. 2017. 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. IEEE.
- Jiasen Lu, Jianwei Yang, Dhruv Batra, and Devi Parikh. 2016. Hierarchical question-image co-attention for visual question answering. In *Advances In Neural Information Processing Systems*, pages 289–297.
- Minh-Thang Luong, Hieu Pham, and Christopher D Manning. 2015. Effective approaches to attention-based neural machine translation. *arXiv preprint arXiv:1508.04025*.
- Matt MacMahon, Brian Stankiewicz, and Benjamin Kuipers. 2006. Walk the talk: Connecting language, knowledge, and action in route instructions. *Def*, 2(6):4.
- Steven J Rennie, Etienne Marcheret, Youssef Mroueh, Jarret Ross, and Vaibhava Goel. 2016. Self-critical sequence training for image captioning. *arXiv preprint arXiv:1612.00563*.
- Adam Santoro, David Raposo, David GT Barrett, Mateusz Malinowski, Razvan Pascanu, Peter Battaglia, and Timothy Lillicrap. 2017. A simple neural network module for relational reasoning. *arXiv preprint arXiv:1706.01427*.
- Minjoon Seo, Aniruddha Kembhavi, Ali Farhadi, and Hannaneh Hajishirzi. 2016. Bidirectional attention flow for machine comprehension. *arXiv preprint arXiv:1611.01603*.
- Alane Suhr, Mike Lewis, James Yeh, and Yoav Artzi. 2017. A corpus of natural language for visual reasoning. In *55th Annual Meeting of the Association for Computational Linguistics, ACL*.
- Oriol Vinyals, Meire Fortunato, and Navdeep Jaitly. 2015. Pointer networks. In *Advances in Neural Information Processing Systems*, pages 2692–2700.
- Shuohang Wang and Jing Jiang. 2016. Machine comprehension using match-1stm and answer pointer. *arXiv preprint arXiv:1608.07905*.
- Ronald J Williams. 1992. Simple statistical gradient-following algorithms for connectionist reinforcement learning. In *Reinforcement Learning*, pages 5–32. Springer.
- Kelvin Xu, Jimmy Ba, Ryan Kiros, Kyunghyun Cho, Aaron Courville, Ruslan Salakhutdinov, Rich Zemel, and Yoshua Bengio. 2015. Show, attend and tell: Neural image caption generation with visual attention. In *International Conference on Machine Learning*, pages 2048–2057.

A Supplementary Material

A.1 CNN-BiATT Model Details

As shown in Fig. 4, we apply our BiATT model to the raw image dataset with minimal modification. The visual input of the model for this task is changed from the unordered structured representation set of objects o to the raw image pixels I . Hence, we replace all object-related LSTMs (e.g., the OBJ-LSTM and the LSTM-RNN in the bidirectional attention in Fig. 2) with visual feature convolutional neural networks (CNNs) that directly learn the structure via pixel-level, spatial filters (instead of a pointer network which addresses an unordered sequence of structured object representations).

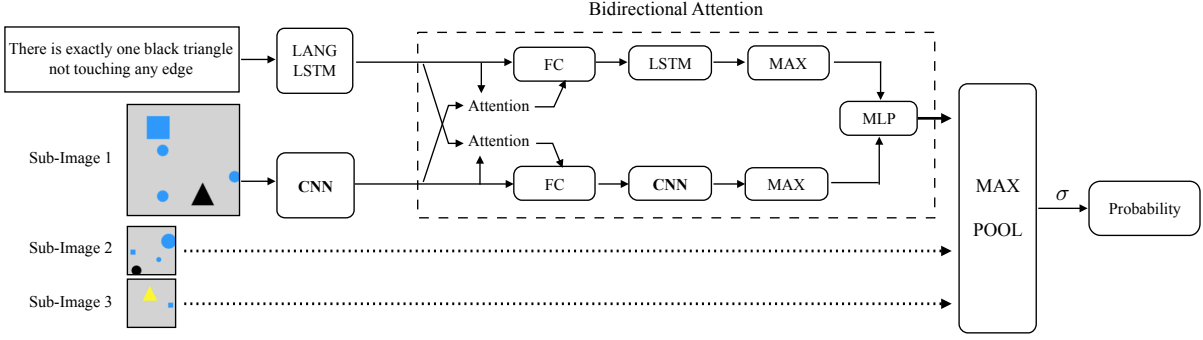


Figure 4: Our CNN-BiATT model for the raw-image dataset version replaces every object-related LSTM-RNN with a spatial-filter convolutional neural network (CNN). The CNN for the raw image-pixels is a pretrained ResNet-v2-101. A 3-layers CNN with relu activation is used in the bidirectional attention.

The training datum for the NLVR raw-image version consists of the statement s , the image I and the ground truth label y . The image I contains three sub-images x^1 , x^2 and x^3 . We will use x to indicate any sub-image. The superscript which indicates the index of the sub-image is ignored to simplify the notation. The representation of the statement $\{h_i\}$ is calculated by the LANG-LSTM as before. For the image representation, we project the sub-image to a sequence of feature vectors (i.e., the feature map) $\{a_l : l = 1, \dots, L\}$ corresponding to the different image locations. $L = m \times m$ is the size of the features and m is the width and height of the feature map. The projection consists of ResNet-V2-101 (He et al., 2016) and a following fully-connected (FC) layer. We only use the blocks in the ResNet before the average pooling layer and the output of the ResNet is a feature map of size $m \times m \times 2048$.

$$f_1, \dots, f_L = \text{ResNet}(x) \quad (21)$$

$$a_l = \text{relu}(W_x f_l + b_x) \quad (22)$$

The joint-representation of the statement $\{\hat{h}_i\}$ is the combination of the LANG-LSTM hidden output states $\{h_i\}$ and the image-aware context vectors $\{c_i\}$:

$$\alpha_{i,l} = \text{softmax}_l(h_i^\top B_1 a_l) \quad (23)$$

$$c_i = \sum_l \alpha_{i,l} \cdot a_l \quad (24)$$

$$\hat{h}_i = \text{relu}(W_{\text{LANG}}[h_i; c_i; h_i - c_i; h_i \circ c_i]) \quad (25)$$

The joint-representation of the image $\{\hat{a}_l\}$ is cal-

culated in the same way:

$$\beta_{l,i} = \text{softmax}_i(a_l^\top B_2 h_i) \quad (26)$$

$$d_l = \sum_i \beta_{l,i} \cdot h_i \quad (27)$$

$$\hat{a}_l = \text{relu}(W_{\text{IMG}}[a_l; d_l; a_l - d_l; a_l \circ d_l]) \quad (28)$$

The joint-representation of the statement is further processed by a LSTM-RNN. Different from our BiATT model, a 3-layers CNN is used for modeling the joint-representation of the image $\{\hat{a}_l\}$. The output of the CNN layer is another feature map $\{\bar{a}_l\}$. Each CNN layer has kernel size 3×3 and uses relu as the activation function, and then we finally use element-wise max operator similar to Sec. 3.1:

$$\bar{h}_1, \bar{h}_2, \dots, \bar{h}_T = \text{LSTM}(\hat{h}_1, \hat{h}_2, \dots, \hat{h}_T) \quad (29)$$

$$\bar{a}_1, \bar{a}_2, \dots, \bar{a}_L = \text{CNN}(\hat{a}_1, \hat{a}_2, \dots, \hat{a}_L) \quad (30)$$

$$\bar{h} = \text{ele max}_i \{\bar{h}_i\} \quad (31)$$

$$\bar{a} = \text{ele max}_l \{\bar{a}_l\} \quad (32)$$

At last, we use the same method as our BiATT model to calculate the score and the loss function:

$$\text{score}(s, x) = W_2 \tanh(W_1[\bar{h}; \bar{a}] + b_1) \quad (33)$$

$$P(s, I) = \sigma \left(\max_j \text{score}(s, x^j) \right) \quad (34)$$

$$L(s, I, y) = -y \log P(s, I) - (1 - y) \log(1 - P(s, I)) \quad (35)$$

A.2 Reimplementation Details for Relationship Network and BiDAF Models

We reimplement a Relationship Network (Santoro et al., 2017), using a three-layer MLP with

256 units per layer in the G-net and a three-layer MLP consisting of 256, 256 (with 0.3 dropout), and 1 units with ReLU nonlinearities for F-net. We also reimplement a BiDAF model (Seo et al., 2016) using 128-dimensional word embedding, 256-dimensional LSTM-RNN and 0.3 dropout rate. A max pooling layer on top of the modeling layer of BiDAF is used to merge the hidden outputs to a single vector.

A.3 Experimental Setup and Training Details for Our BiATT-Pointer, BiENC, and CNN-BiATT Models

A.3.1 BiATT-Pointer

For preprocessing, we replace the words whose occurrence is less than 3 with the “UNK” token. We create a 9 dimension vector as the feature of each object. This feature contains the location (x, y) in 2D coordinate, the size of the object and two 3-dimensional hot vectors for the shape and the color. The (x, y) coordinates are normalized to the range $[-1, 1]$.

For the model hyperparameters (all lightly tuned on dev set), the dimension of the word embedding is 128, and the number of units in an LSTM cell is 256. The word embedding is trained from scratch. The object feature is projected to a 64-dimensional vector. The dimensions of joint representation \hat{h}_i and \hat{g}_k are both 512. The first fully-connected layer in calculating the sub-images score has 512 units. All the trainable variables are initialized with the Xavier initializer. To regularize the training process, we add a dropout rate 0.3 to the hidden output of the LSTM-RNNs and before the last MLP layer which calculates the score for sub-images. We also clip the gradients by their norm to avoid gradient exploding. The losses are optimized by a single Adam optimizer and the learning rate is fixed at $1e-4$.

For the pointer network, we sample the objects following the distribution of the objects at each decoder step during training. In inference, we select the object with maximum probability. We use the self-critical baseline (Rennie et al., 2016) to stabilize the RL training, where the final score in inference (choosing object with maximum probability) is subtracted from the reward. To reduce the number of parameters, we share the weight of the fully-connected layer which projects the raw object feature to the high dimensional vector in the pointer encoder, the pointer decoder, and the OBJ-

LSTM. The pointer decoder attends to the hidden outputs of the LANG-LSTM using bilinear attention (Luong et al., 2015).

A.3.2 CNN-BiATT

We initialize our model with weights of the public pretrained ResNet-V2-101 (based on the ImageNet dataset) and freeze it during training. The ResNet projects the sub-image to a feature map of $10 \times 10 \times 2048$. The feature map is normalized to a mean of 0 and a standard deviation of 1 before feeding into the FC layer. The fully connected layer after the ResNet has 512 units. Each layer of the 3-layers CNN in the bidirectional attention has kernel size 3×3 with 512 filters and no padding.

A.3.3 BiENC

The BiENC model uses LANG-LSTM and OBJ-LSTM to read the statement and the objects. A bilinear form calculates the similarity between the last hidden outputs of the two LSTM-RNNs. The similarity is directly used as the score of the sub-image. The CNN-BiENC model replaces the OBJ-LSTM with a CNN.