

GraphVQA: Language-Guided Graph Neural Networks for Scene Graph Question Answering

Weixin Liang¹, Yanhao Jiang¹, Zixuan Liu¹

Stanford University, Stanford, CA 94305

{wxliang, jiangyh, zucks626}@stanford.edu

Abstract

Images are more than a collection of objects or attributes — they represent a web of relationships among interconnected objects. Scene Graph has emerged as a new modality as a structured graphical representation of images. Scene Graph encodes objects as nodes connected via pairwise relations as edges. To support question answering on scene graphs, we propose GraphVQA, a language-guided graph neural network framework that translates and executes a natural language question as multiple iterations of message passing among graph nodes. We explore the design space of GraphVQA framework, and discuss the trade-off of different design choices. Our experiments on GQA dataset show that GraphVQA outperforms the state-of-the-art model by a large margin (88.43% vs. 94.78%). Our code is available at <https://github.com/codexxxl/GraphVQA>

1 Introduction

Images are more than a collection of objects or attributes. Each image represents a web of relationships among interconnected objects. Towards formalizing a representation for images, Visual Genome (Krishna et al., 2017a) defined scene graphs, a structured formal graphical representation of an image that is similar to the form widely used in knowledge base representations. As shown in Figure 1, scene graph encodes objects (e.g., girl, burger) as nodes connected via pairwise relationships (e.g., holding) as edges. Scene graphs have been introduced for image retrieval (Johnson et al., 2015), image generation (Johnson et al., 2018), image captioning (Anderson et al., 2016), understanding instructional videos (Huang et al., 2018), and situational role classification (Li et al., 2017).

To support question answering on scene graphs, we propose GraphVQA, a language-guided graph

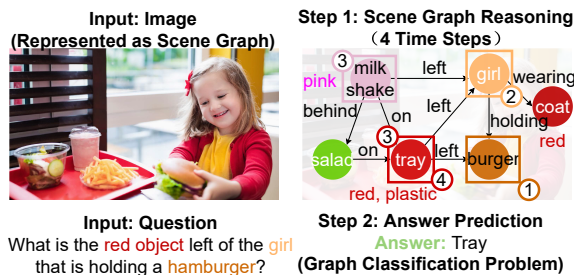


Figure 1: **Scene Graph:** Scene graph encodes objects (e.g., girl, burger) as nodes connected via pairwise relationships (e.g., holding) as edges. **GraphVQA Framework:** Our core insight is to translate and execute a natural language question as multiple iterations of message passing among graph nodes (e.g., hamburger → small girl → red tray). The final state after message passing represents the answer (e.g., tray).

neural network framework for Scene Graph Question Answering (Scene Graph QA). Our core insight is to translate a natural language question into multiple iterations of message passing among graph nodes. Figure 1 shows an example question “What is the red object left of the girl that is holding a hamburger”. This question can be naturally answered by the following iterations of message passing “hamburger → small girl → red tray”. The final state after message passing represents the answer (e.g., tray), and the intermediate states reflect the model’s reasoning. Each message passing iteration is accomplished by a graph neural network (GNN) layer. We explore various message passing designs in GraphVQA, and discuss the trade-off of different design choices.

Scene Graph QA is closely related to Visual Question Answering (VQA). Although there are many research efforts in scene graph generation, Scene Graph QA remains relatively under-explored. Sporadic attempts in scene graph based VQA (Hu et al., 2019; Li et al., 2019; Santoro et al., 2017) mostly propose various attention mechanisms designed primarily for fully-connected graphs, thereby failing to model and capture the important structural information of the scene graphs.

¹Equal Contribution. Authors listed in alphabetical order.

We evaluate GraphVQA on GQA dataset (Hudson and Manning, 2019a). We found that GraphVQA with de facto GNNs can outperform the state-of-the-art model by a large margin (88.43% vs. 94.78%). We discuss additional related work in appendix A. Our results suggest the importance of incorporating recent advances from graph machine learning into our community.

2 Machine Learning with Graphs

Modeling graphical data has historically been challenging for the machine learning community. Traditionally, methods have relied on Laplacian regularization through label propagation, manifold regularization or learning embeddings. Today’s de facto choice is graph neural network (GNN), which is an operator on local neighborhoods of nodes.

GNNs follow the message passing scheme. The high level idea is to update each node’s feature using its local neighborhoods of nodes. Specifically, node i ’s representation at l -th layer $\mathbf{h}_i^{(l)}$ can be calculated using previous layer’s node representations $\mathbf{h}_i^{(l-1)}$ and $\mathbf{h}_j^{(l-1)}$ as:

$$\mathbf{h}_{\mathcal{N}_i}^{(l)} = \text{AGG}_{j \in \mathcal{N}_i} \phi^{(l)}(\mathbf{h}_i^{(l-1)}, \mathbf{h}_j^{(l-1)}, e_{ji}) \quad (1)$$

$$\mathbf{h}_i^{(l)} = \gamma^{(l)}(\mathbf{h}_i^{(l-1)}, \mathbf{h}_{\mathcal{N}_i}^{(l)}) \quad (2)$$

where e_{ji} denotes the feature of edge from node j to node i , $\mathbf{h}_{\mathcal{N}_i}^{(l)}$ denotes aggregated neighborhood information, $\gamma^{(l)}$ and $\phi^{(l)}$ denotes differentiable functions such as MLPs, and AGG denotes aggregation functions such as mean or sum pooling.

3 GraphVQA Framework

Figure 2 shows an overview of four modules in GraphVQA: (1) Question Parsing Module translates the question to M instruction vectors. (2) Scene Graph Encoding Module initializes node features X and edge features E with word embeddings. (3) Graph Reasoning Module performs message passing with graph neural networks for each instruction vector. (4) Answering Module summarizes the final state after message passing and predicts the answer.

3.1 Question Parsing Module

Question Parsing Module uses a sequence-to-sequence transformer architecture to translate the question $[q_1, \dots, q_Q]$ into a sequence of instruction

vectors $[\mathbf{i}^{(1)}, \dots, \mathbf{i}^{(M)}]$ with a fixed M .

$$[\mathbf{i}^{(1)}, \dots, \mathbf{i}^{(M)}] = \text{Seq2Seq}(q_1, \dots, q_Q) \quad (3)$$

3.2 Scene Graph Encoding Module

Scene Graph Encoding Module first initializes node features $\hat{X} = [\hat{x}_1, \dots, \hat{x}_N]$ with the word embeddings of the object name and attributes, and edge features E with the word embedding of edge type. We then obtain contextualized node features X by:

$$\mathbf{x}_i = \sigma\left(\frac{1}{|\mathcal{N}_i|} \sum_{j \in \mathcal{N}_i} (W_{\text{enc}} [\hat{\mathbf{x}}_j; e_{ij}])\right) \quad (4)$$

where σ denotes the activation function, e_{ij} denotes the feature of the edge that connects node i and node j , and $X = [x_1, x_2, \dots, x_N]$ denotes the contextualized node features.

3.3 Graph Reasoning Module

Graph Reasoning Module is the core of GraphVQA framework. Graph Reasoning Module executes the M instruction vectors step-by-step, with N graph neural network layers. One major difference between our Graph Reasoning Module and standard GNN is that, we want the message passing in layer L conditioned on the L^{th} instruction vector. Inspired by language model type condition (Liang et al., 2020b), we adopt a general design that is compatible with *any* graph neural network design: Before running the L^{th} GNN layer, we concatenate the L^{th} instruction vector to every node and edge feature from the previous layer. Specifically,

$$\hat{\mathbf{h}}_i^{(L-1)} = [\mathbf{h}_i^{(L-1)}; \mathbf{i}^{(L)}] \quad (5)$$

$$\hat{\mathbf{e}}_{ij}^{(L-1)} = [e_{ij}^{(L-1)}; \mathbf{i}^{(L)}] \quad (6)$$

where $\hat{\mathbf{h}}_i^{(L-1)}$ and $\hat{\mathbf{e}}_{ij}^{(L-1)}$ denotes the node feature and edge feature as inputs to the L^{th} GNN layer. Next, we introduce three standard GNNs that we have explored, starting from the simplest one.

3.3.1 Graph Convolution Networks (GCN)

GCN (Kipf and Welling, 2017) treats neighborhood nodes as equally important sources of information, and simply averages the transformed features of neighborhood nodes.

$$\mathbf{h}_i^{(L)} = \sigma\left(\frac{1}{|\mathcal{N}_i|} \sum_{j \in \mathcal{N}_i} (W_{\text{GCN}}^{(L)} \hat{\mathbf{h}}_j^{(L-1)})\right) \quad (7)$$

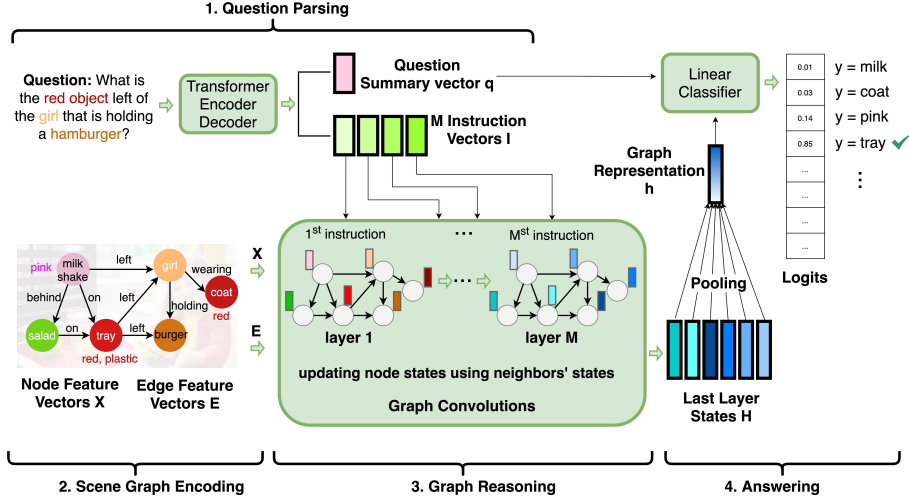


Figure 2: Semantics of the GraphVQA Framework. (1) Question Parsing Module translates the question to M instruction vectors. (2) Scene Graph Encoding Module initializes node features X and edge features E with word embeddings. (3) Graph Reasoning Module perform message passing with graph neural networks for each instruction vector. (4) Answering Module summarizes the final state after message passing and predicts the answer.

3.3.2 Graph Isomorphism Network (GINE)

GIN (Xu et al., 2019) is provably as powerful as the Weisfeiler-Lehman graph isomorphism test. GINE (Hu et al., 2020) augments GIN by also considering edge features during the message passing:

$$\mathbf{h}_i^{(L)} = \Theta((1 + \epsilon)\hat{\mathbf{h}}_i^{(L-1)} + \sum_{j \in \mathcal{N}(i)} \sigma(\hat{\mathbf{h}}_j^{(L-1)} + \hat{\mathbf{e}}_{j,i}^{(L-1)}))$$

where Θ denotes expressive functions such as MLPs, and ϵ is a scale factor for the emphasis of the central node.

3.3.3 Graph Attention Network (GAT)

Different from GIN and GINE, GAT (Veličković et al., 2018) learns to use attention mechanism to weight neighbour nodes differently. Intuitively, GAT fits more naturally with our Scene Graph QA task, since we want to emphasis different neighbor nodes given different instruction vectors. Specifically, the attention score $\alpha_{ij}^{(L)}$ for message passing from node j to node i at L^{th} layer is calculated as:

$$\alpha_{ij}^{(L)} = \text{Softmax}_{\mathcal{N}_i}(\text{MLP}(\hat{\mathbf{h}}_i^{(L-1)}, \hat{\mathbf{h}}_j^{(L-1)}, \hat{\mathbf{e}}_{ij}^{(L-1)})) \quad (8)$$

where $\text{Softmax}_{\mathcal{N}_i}$ is a normalization to ensure that the attention scores from one node to its neighbor nodes sum to 1. After calculating the attention scores, we calculate each node's new representation as a weighted average from its neighbour nodes.

$$\mathbf{h}_i^{(L)} = \sigma\left(\sum_{j \in \mathcal{N}_i} \alpha_{ij}^{(L)} \mathbf{W}_{\text{GAT}}^{(L)} \hat{\mathbf{h}}_j^{(L-1)}\right) \quad (9)$$

where σ denotes the activation function. Similar to transformer models, we use multiple attention heads in practice. In addition, many modern deep learning tool-kits can be incorporated into GNNs, such as batch normalization, dropout, gating mechanism, and residual connections.

3.4 Answering Module

After executing the Graph Reasoning module, we obtain the final states of all graph nodes after M iterations of message passing $[\mathbf{h}_1^{(M)}, \dots, \mathbf{h}_N^{(M)}]$. We first summarize the final states after message passing, and then predict the answer token with the question summary vector q :

$$\mathbf{h} = \text{Aggregate}([\mathbf{h}_1^{(M)}, \mathbf{h}_2^{(M)}, \dots, \mathbf{h}_N^{(M)}]) \quad (10)$$

$$\mathbf{y} = \text{Softmax}(\text{MLP}(\mathbf{h}, \mathbf{q})) \quad (11)$$

where \mathbf{y} is the predicted answer. We note that GraphVQA does not require any explicit supervision on how to solve the question step-by-step, and we only supervise on the final answer prediction.

4 Experiments

Setup We evaluate our GraphVQA framework on the GQA dataset (Hudson and Manning, 2019a) which contains 110K scene graphs, 1.5M questions, and over 1000 different answer tokens. We use the official train/validation split of GQA. Since the scene graphs of the test set are not publicly available, we use validation split as test set. We set the number of instructions $M = 5$. More dataset and training details are included in Appendix C.

	Method	Binary	Open	Consistency	Validity	Plausibility	Distribution	Accuracy
Baseline1	GCN	86.84	84.63	90.21	95.51	94.44	0.13	85.70
Baseline2	LCGN	90.57	88.43	93.88	95.40	93.89	0.16	88.43
Ablation1	Only Questions	61.90	22.69	68.68	96.39	87.30	0.17	41.07
Ablation2	Only Scene Graphs	21.86	17.54	46.98	36.89	32.63	7.22	19.63
Proposed1	GraphVQA-GCN	92.11	88.37	95.44	95.5	94.4	0.12	90.18
Proposed2	GraphVQA-GINE	92.36	88.56	94.79	95.44	94.39	0.13	90.38
Proposed3	GraphVQA-GAT	96.30	93.37	98.37	95.55	95.15	0.07	94.78

Table 1: Evaluation Results on GQA. All numbers are in percentages. The lower the better for distribution.

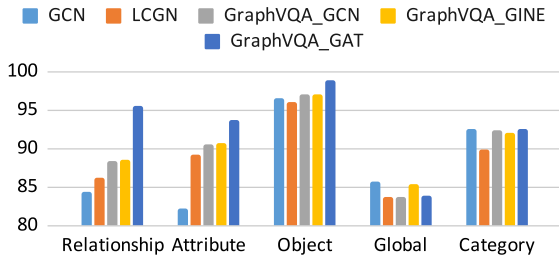


Figure 3: Accuracy breakdown on question semantic types. GraphVQA-GAT achieves significantly higher accuracy in relationship questions (95.53%).

Models and Metrics We evaluate three instantiations of GraphVQA: GraphVQA-GCN, GraphVQA-GINE, GraphVQA-GAT. We compare with the state-of-the-art model LCGN (Hu et al., 2019). We discuss LCGN in appendix B.3. We also compare with a simple GCN without instruction vector concatenation discussed in § 3.3 to study the importance of language guidance. We report the standard evaluation metrics defined in Hudson and Manning (2019a) such as accuracy and consistency.

Results The first take-away message is that GraphVQA outperforms the state-of-the-art approach LCGN, even with the simplest GraphVQA-GCN. Besides, GraphVQA-GAT outperforms LCGN by a large margin (88.43% vs. 94.78% accuracy), highlighting the benefits of incorporating recent advances from graph machine learning. The second take-away message is that conditioning on instruction vectors is important. Removing such conditioning drops performance (GCN vs. GraphVQA-GCN, 85.7% vs. 90.18%). The third take-away message is that attention mechanism is important for Scene Graph QA, as GraphVQA-GAT also outperforms both GraphVQA-GCN and GraphVQA-GINE by a large margin (94.78% vs. 90.38%), even though GINE is provably more expressive than GAT (Xu et al., 2019).

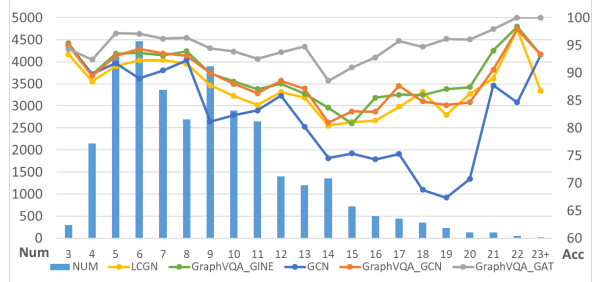


Figure 4: Accuracy breakdown on question word count. Num denotes the number of questions of each length. GraphVQA-GAT shows significant better performance for long question answering tasks.

Analysis Figure 3 shows the accuracy breakdown on question semantic types. We found that GraphVQA-GAT achieves significantly higher accuracy in relationship questions (95.53%). This shows the strength in the attention mechanism in modeling the relationships in scene graphs.

Figure 4 shows the accuracy breakdown on question word count. As expected, longer questions are harder to answer by all models. In addition, we found that as questions become longer, the accuracy GraphVQA-GAT deteriorates drops than other methods, showing that GraphVQA-GAT is better at answering long questions.

5 Conclusion

In this paper, we present GraphVQA to support question answering on scene graphs. GraphVQA translates and executes a natural language question as multiple iterations of message using graph neural networks. We explore the design space of GraphVQA framework, and found that GraphVQA-GAT (Graph Attention Network) is the best design. GraphVQA-GAT outperforms the state-of-the-art model by a large margin (88.43% vs. 94.78%). Our results suggest the potential benefits of revisiting existed Vision + Language multimodal models from the perspective of graph machine learning.

Acknowledgments

We would like to sincerely thank NAACL-HLT 2021 MAI-Workshop Program Committee for their review efforts and helpful feedback. We would also like to extend our gratitude to Jingjing Tian and the teaching staff of Stanford CS 224W for their valuable feedback.

References

- Aishwarya Agrawal, Dhruv Batra, and Devi Parikh. 2016. Analyzing the behavior of visual question answering models. In *EMNLP*, pages 1955–1960. The Association for Computational Linguistics.
- Saeed Amizadeh, Hamid Palangi, Alex Polozov, Yichen Huang, and Kazuhito Koishida. 2020. Neuro-symbolic visual reasoning: Disentangling "visual" from "reasoning". In *ICML*, volume 119 of *Proceedings of Machine Learning Research*, pages 279–290. PMLR.
- Peter Anderson, Basura Fernando, Mark Johnson, and Stephen Gould. 2016. SPICE: semantic propositional image caption evaluation. In *ECCV (5)*, volume 9909 of *Lecture Notes in Computer Science*, pages 382–398. Springer.
- Abhishek Das, Harsh Agrawal, Larry Zitnick, Devi Parikh, and Dhruv Batra. 2016. Human attention in visual question answering: Do humans and deep networks look at the same regions? In *EMNLP*, pages 932–937. The Association for Computational Linguistics.
- Ronghang Hu, Anna Rohrbach, Trevor Darrell, and Kate Saenko. 2019. Language-conditioned graph networks for relational reasoning. In *ICCV*, pages 10293–10302. IEEE.
- Weihua Hu, Bowen Liu, Joseph Gomes, Marinka Zitnik, Percy Liang, Vijay S. Pande, and Jure Leskovec. 2020. Strategies for pre-training graph neural networks. In *ICLR*. OpenReview.net.
- De-An Huang, Shyamal Buch, Lucio M. Dery, Animesh Garg, Li Fei-Fei, and Juan Carlos Niebles. 2018. Finding "it": Weakly-supervised reference-aware visual grounding in instructional videos. In *CVPR*, pages 5948–5957. IEEE Computer Society.
- Drew A. Hudson and Christopher D. Manning. 2019a. GQA: A new dataset for real-world visual reasoning and compositional question answering. In *CVPR*, pages 6700–6709. Computer Vision Foundation / IEEE.
- Drew A. Hudson and Christopher D. Manning. 2019b. Learning by abstraction: The neural state machine. In *NeurIPS*, pages 5901–5914.
- Justin Johnson, Agrim Gupta, and Li Fei-Fei. 2018. Image generation from scene graphs. In *CVPR*, pages 1219–1228. IEEE Computer Society.
- Justin Johnson, Ranjay Krishna, Michael Stark, Li-Jia Li, David A. Shamma, Michael S. Bernstein, and Fei-Fei Li. 2015. Image retrieval using scene graphs. In *CVPR*, pages 3668–3678. IEEE Computer Society.
- Thomas N. Kipf and Max Welling. 2017. [Semi-supervised classification with graph convolutional networks](#).
- Ranjay Krishna, Yuke Zhu, Oliver Groth, Justin Johnson, Kenji Hata, Joshua Kravitz, Stephanie Chen, Yannis Kalantidis, Li-Jia Li, David A. Shamma, Michael S. Bernstein, and Li Fei-Fei. 2017a. Visual genome: Connecting language and vision using crowdsourced dense image annotations. *Int. J. Comput. Vis.*, 123(1):32–73.
- Ranjay Krishna, Yuke Zhu, Oliver Groth, Justin Johnson, Kenji Hata, Joshua Kravitz, Stephanie Chen, Yannis Kalantidis, Li-Jia Li, David A. Shamma, Michael S. Bernstein, and Li Fei-Fei. 2017b. [Visual genome: Connecting language and vision using crowdsourced dense image annotations](#). *International Journal of Computer Vision*, 123(1):32–73.
- Linjie Li, Zhe Gan, Yu Cheng, and Jingjing Liu. 2019. Relation-aware graph attention network for visual question answering. In *ICCV*, pages 10312–10321. IEEE.
- Ruiyu Li, Makarand Tapaswi, Renjie Liao, Jiaya Jia, Raquel Urtasun, and Sanja Fidler. 2017. Situation recognition with graph neural networks. In *ICCV*, pages 4183–4192. IEEE Computer Society.
- Weixin Liang, Feiyang Niu, Aishwarya N. Reganti, Govind Thattai, and Gökhan Tür. 2020a. [LRTA: A transparent neural-symbolic reasoning framework with modular supervision for visual question answering](#). *CoRR*, abs/2011.10731.
- Weixin Liang, Youzhi Tian, Chengcai Chen, and Zhou Yu. 2020b. [MOSS: end-to-end dialog system framework with modular supervision](#). In *AAAI*, pages 8327–8335. AAAI Press.
- Weixin Liang and James Zou. 2020. [Neural group testing to accelerate deep learning](#). *CoRR*, abs/2011.10704.
- Weixin Liang, James Zou, and Zhou Yu. 2020c. [AL-ICE: active learning with contrastive natural language explanations](#). In *EMNLP (1)*, pages 4380–4391. Association for Computational Linguistics.
- Weixin Liang, James Zou, and Zhou Yu. 2020d. [Beyond user self-reported likert scale ratings: A comparison model for automatic dialog evaluation](#). In *ACL*, pages 1363–1374. Association for Computational Linguistics.

Pramod Kaushik Mudrakarta, Ankur Taly, Mukund Sundararajan, and Kedar Dhamdhere. 2018. Did the model understand the question? In *ACL (1)*, pages 1896–1906. Association for Computational Linguistics.

Jeffrey Pennington, Richard Socher, and Christopher D. Manning. 2014. Glove: Global vectors for word representation. In *EMNLP*, pages 1532–1543. ACL.

Adam Santoro, David Raposo, David G. T. Barrett, Mateusz Malinowski, Razvan Pascanu, Peter W. Battaglia, and Tim Lillicrap. 2017. A simple neural network module for relational reasoning. In *NIPS*, pages 4967–4976.

Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua Bengio. 2018. *Graph attention networks*.

Keyulu Xu, Weihua Hu, Jure Leskovec, and Stefanie Jegelka. 2019. How powerful are graph neural networks? In *ICLR*. OpenReview.net.

A Related Work

A.1 Visual Question Answering

VQA requires an interplay of visual perception with reasoning about the question semantics grounded in perception. The predominant approach to visual question answering (VQA) relies on encoding the image and question with a “black-box” neural encoder, where each image is usually represented as a bag of object features, where each feature describes the local appearance within a bounding box detected by the object detection backbone. However, representing images as collections of objects fails to capture relationships which are crucial for visual question answering. Recent study has further demonstrated some unsettling behaviours of those models: they tend to ignore important question terms (Mudrakarta et al., 2018), look at wrong image regions (Das et al., 2016), or undesirably adhere to superficial or even potentially misleading statistical associations (Agrawal et al., 2016). In addition, it has been shown that recent advances are primarily driven by perception improvements (e.g. object detection) rather than reasoning (Amizadeh et al., 2020).

A.2 Scene Graph Question Answering

Although there are many research efforts in scene graph generation, using scene graphs for visual question answering remains relatively under-explored (Hudson and Manning, 2019b; Hu et al.,

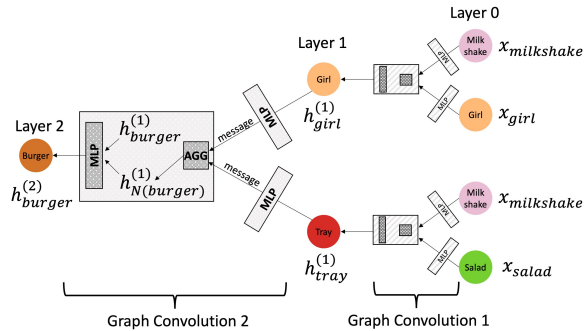


Figure 5: Structure of 2 Layer Graph Neural Network

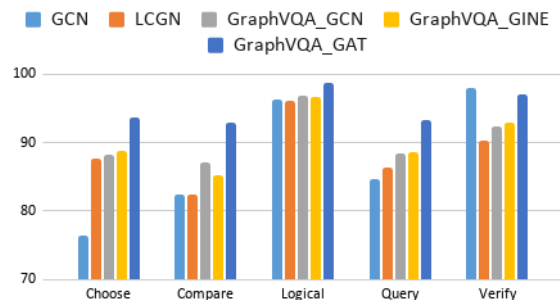


Figure 6: Accuracy breakdown on question structural types. GraphVQA-GAT achieves significantly higher accuracy in all types except for verify.

2019; Li et al., 2019; Liang et al., 2020a). Hudson and Manning (2019b) propose a task-specific graph traversal framework with neural networks. The framework requires specifying the detail ontology of the dataset (e.g., color: red, blue,...; material: wooden, metallic), and thus is not directly generalizable. Other attempts in graph based VQA (Hu et al., 2019; Li et al., 2019) mostly explore attention mechanism on fully-connected graphs, thereby failing to capture the important structural information of the scene graphs.

B Additional Results

B.1 Additional Performances Analysis

Figure 6 provides another set of accuracy breakdown result on question structural types. We found that GraphVQA-GAT achieves the best for all types of questions except for the verify types. Specifically, GraphVQA-GAT outperforms significantly than other methods on answering queries, comparing among objects and making choices. This intuitively matches the principle of attention mechanism and again shows its advantages in modeling structural information in scene graphs.

Scene Graphs Statistics	Validation data	Train data	All
Total Number of Graphs	10,696	74,942	85,638
Total Number of Nodes	174,331	1,231,134	1,405,465
Total Number of Edges	534,889	3,795,907	4,330,796
Average Number of Nodes per Graph	16	16	16
Average Number of Edges per Graph	50	51	51
Total Number of Node Types	1,536	1,702	1,703
Total Number of Edge Types	295	310	310
Total Number of Attributes Types	603	617	617

Table 2: Scene Graphs Statistics of the GQA Dataset

	Level of Classification	Structural	Semantic
Is there apples in the picture?	node	verify	object
What color is the apple?	node	query	attribute
Is the cat to the left or right of the flower?	edge type	choose	relation
Is it sunny or cloudy?	graph	query	global

Table 3: Typical types of questions

B.2 Expressive Ability Analysis of GraphVQA-GINE

As mentioned in Section 3.3.1 and Section 3.3.2, a expressive function Θ is used in GINE layer. When Θ is just a single layer MLP, the corresponding GIN/GINE structure will be very similar to the GCN structure. Since in Section 4 we implemented Θ as a single layer MLP, the performance of GraphVQA-GCN and GraphVQA-GINE stays at very similar stage. As GIN and GINE are now very popular as basic components for large-scale graph neural network design, one may ask if using Θ with more powerful expression ability will help the performance. The short answer is no. We provide a simple ablation study on different choice of Θ , using a two layer MLP-style network with (FC, ReLU, FC, ReLU, BN) structure. Table 4 shows that the result of GraphVQA-GINE-2 degrades to the worst. One possible reason is that the scale for each scen graph is generally small, therefore the expression ability might already be enough for a single layer MLP, and use a more complex Θ may leads to harder optimization problems, and thus leads to a downgrade of the performance. Such guess could possibly be further investigated and evaluated in our future work. In addition, the scene graph-based VQA as in this work might offer an opportunity for further accelerating the real world image-based applications (Liang and Zou, 2020). Exploring such deployment benefits is another direction of future work.

B.3 Brief Introduction of LCGN

Language-Conditioned Graph Networks (LCGN) (Hu et al., 2019) updates node representations recurrently using the same single layer graph neural network. Given a set of instruction vectors $[i_1, \dots, i_M]$, LCGN uses a single layer attention to convert them into context representations $[c_1, \dots, c_M]$. Then, given a set of node representations $[x_{loc,1}, \dots, x_{loc,n}]$, LCGN first randomly initialize another set of context representations $[x_{ctx,1}, \dots, x_{ctx,n}]$, and then use them to concatenate with node representations to form initial local features, i.e,

$$\tilde{x}_{t,i} = [x_{loc,i}, x_{ctx,i,t-1}, W_1 x_{loc,i} \circ W_2 x_{ctx,i,t-1}] \quad (12)$$

With the assumption that all nodes are connected, LCGN computes the edge weights $w_{j,i}^{(t)}$ for each node pair (i,j), i.e,

$$w_{j,i}^{(t)} = \text{Softmax}((W_3 \tilde{x}_{t,i})^T ((W_4 \tilde{x}_{t,j}) \circ (W_5 c_t))) \quad (13)$$

The messages $m_{i,j}^{(t)}$, are then computed as:

$$m_{j,i}^{(t)} = w_{j,i}^{(t)} ((W_6 \tilde{x}_{t,j}) \circ (W_7 c_t)) \quad (14)$$

Finally, LCGN aggregates the neighborhood message information to update the context local representation $x_{ctx,i,t}$.

$$x_{ctx,i,t} = W_8 [x_{ctx,i,t-1}; \sum_{j=1}^N m_{j,i}^{(t)}] \quad (15)$$

Method	Binary	Open	Consistency	Validity	Plausibility	Distribution	Accuracy
GraphGQA-GINE-2	86.83	83.85	89.8	95.54	94.25	0.16	85.04

Table 4: Ablation Study Results for 2 layer GraphVQA-GINE. All numbers are in percentages. The lower the better for distribution.

Note that the graph neural structure of LCGN can be regarded as a variant of recurrently-used single standard GAT layer, but with more self-designed learnable parameters. The main difference between LCGN’s and other proposed graph neural structure is that the output node and edge features will be recurrently fed into the same layer again for each reasoning step, leading to a RNN-style network structure, instead of a sequential-style network. Moreover, our LCGN implementation is a variant of original LCGN, including a few improvements. Firstly, we use a transformer encoder and decoder to obtain instruction vectors instead of Bi-LSTM (Liang et al., 2020d). Secondly, we incorporate the true scene graph relations as edges instead of densely connected edges. Thirdly, edge attributes are also used in the generation of initial node features.

C Implementation Details

C.1 Data Pre-processing

The edges in the original scene graphs are directed. This means in most of the cases where we only have one directed edge connecting two nodes in the graph, the messages can only flow through one direction. However, this does not make sense in the natural way of human reasoning. For example, an relation of "A is to the left of B" should obviously entail an opposite relation of "B is to the right of A". Therefore, in order to enhance the connectivity of our graphs, we introduce a synthetic symmetric edge for every non-paired edge, making it pointing reversely to the source node. And in order to encode this reversed relationship, we negate the original edge’s feature vector and use it as the representation of our synthetic symmetric edge.

C.2 Additional Dataset Information

These scene graphs are generated from 113k images on COCO and Flickr using the Visual Genome Scene Graph (Krishna et al., 2017b) annotations. Specifically, each node in the GQA scene graph is representing an object, such as a person, a window, or an apple. Along with the positional

information of bounding box, each object is also annotated with 1-3 different attributes. These attributes are the adjectives used to describe associated objects. For examples, there can be color attributes like "white", size attributes like "large", and action attributes like "standing". Attributes are important sources of information beyond the coarse-grained object classes (Liang et al., 2020c). Each edge in the scene graph denotes relation between two connected objects. These relations can be action verbs, spatial prepositions, and comparatives, such as "wearing", "below", and "taller".

We use the official split of the GQA dataset. We use two files "val_sceneGraphs.json" and "train_sceneGraphs.json" directly obtained on the GQA website as our raw dataset. Since each image (graph) is independent, GQA splits the dataset by individual graphs with rough split percentages of **train/validation: 88%/12%**. In the table 2, we summarize the statistics that we collected from the dataset. We did not report the statistics of the test set since the scene graphs in the test set is not publicly available.

C.3 Training details

We train the models using the Adam optimization method, with a learning rate of 10^{-4} , a batch size of 256, and a learning rate drop(divide by 10) each 90 epochs. We train all models for 100 epochs. Both hidden states and word embedding vectors have a dimension size of 300, the latter being initialized using GloVe (Pennington et al., 2014). The instruction vectors have a dimension size of 512. All results reported are for a single-model settings (i.e., without ensembling). We use cross validation for hyper-parameter tuning.