



Graphhopper: Multi-hop Scene Graph Reasoning for Visual Question Answering

Rajat Koner^{1(✉)}, Hang Li^{1,2}, Marcel Hildebrandt^{1,2}, Deepan Das³,
Volker Tresp^{1,2}, and Stephan Günnemann³

¹ Ludwig Maximilian University of Munich, Munich, Germany
koner@dbs.ifi.lmu.de

² Siemens AG, Munich, Germany

³ Technical University of Munich, Munich, Germany

Abstract. Visual Question Answering (VQA) is concerned with answering free-form questions about an image. Since it requires a deep semantic and linguistic understanding of the question and the ability to associate it with various objects that are present in the image, it is an ambitious task and requires multi-modal reasoning from both computer vision and natural language processing. We propose Graphhopper, a novel method that approaches the task by integrating knowledge graph reasoning, computer vision, and natural language processing techniques. Concretely, our method is based on performing context-driven, sequential reasoning based on the scene entities and their semantic and spatial relationships. As a first step, we derive a scene graph that describes the objects in the image, as well as their attributes and their mutual relationships. Subsequently, a reinforcement learning agent is trained to autonomously navigate in a multi-hop manner over the extracted scene graph to generate reasoning paths, which are the basis for deriving answers. We conduct an experimental study on the challenging dataset GQA, based on both manually curated and automatically generated scene graphs. Our results show that we keep up with human performance on manually curated scene graphs. Moreover, we find that Graphhopper outperforms another state-of-the-art scene graph reasoning model on both manually curated and automatically generated scene graphs by a significant margin.

Keywords: Visual Question Answering (VQA) · Knowledge graph reasoning · Scene graph reasoning · Multi-modal reasoning · Reinforcement learning

1 Introduction

Visual Question Answering (VQA) is a challenging task that involves understanding and reasoning over two data modalities, i.e., images and natural lan-

R. Koner, H. Li and M. Hildebrandt—Equal contribution.

Electronic supplementary material The online version of this chapter (https://doi.org/10.1007/978-3-030-88361-4_7) contains supplementary material, which is available to authorized users.

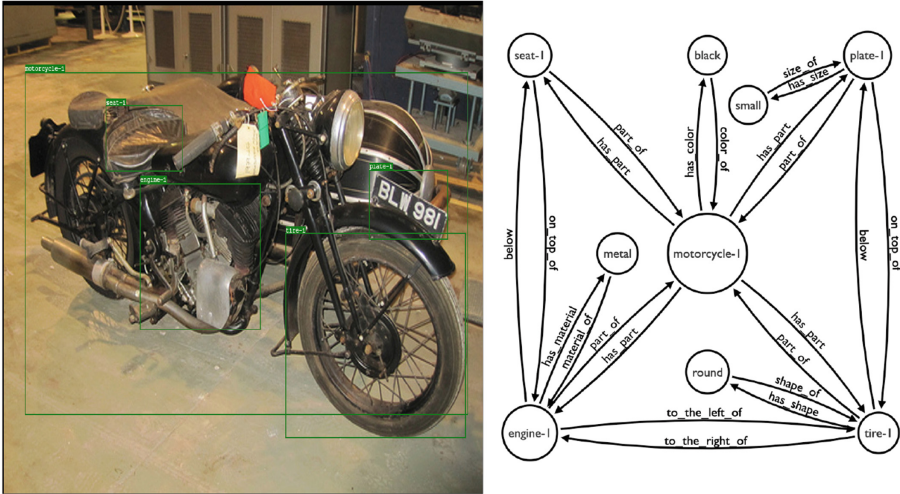


Fig. 1. Example of an image and the corresponding scene graph. Since the scene graph is a directed graph with typed edges, it resembles a knowledge graph and permits the application of knowledge-base completion techniques.

guage. Given an image and a free-form question which formulates a query about the presented scene—the issue is for the algorithm to find the correct answer.

VQA has been studied from the perspective of scene and knowledge graphs [6, 33], as well as vision-language reasoning [1, 10]. To study VQA, various real-world data sets, such as the VQA data set [4, 24], have been generated. It has been argued that, in the VQA data set, many of the apparently challenging reasoning tasks can be solved by an algorithm through exploiting trivial prior knowledge, and thus by shortcuts to proper reasoning (e.g., clouds are white or doors are made of wood). To address these shortcomings, the GQA dataset [17] has been developed. Compared to other real-world datasets, GQA is more suitable for evaluating reasoning abilities since the images and questions are carefully filtered to make the data less prone to biases.

Plenty of VQA approaches are agnostic towards the explicit relational structure of the objects in the presented scene and rely on monolithic neural network architectures that process regional features of the image separately [2, 39]. While these methods led to promising results on previous datasets, they lack explicit compositional reasoning abilities, which results in weaker performance on more challenging datasets such as GQA. Other works [15, 31, 34] perform reasoning on explicitly detected objects and interactive semantic and spatial relationships among them. These approaches are closely related to the scene graph representations [19] of an image, where detected objects are labeled as nodes and relationships between the objects are labeled as edges. In this work, we aim to combine VQA techniques with recent research advances in the area of statistical relation learning on knowledge graphs (KGs). KGs provide human-understandable, struc-

tured representations of knowledge about the real world via collections of factual statements. Inspired by multi-hop reasoning methods on KGs such as [8, 12, 38], we propose Graphhopper, a novel method that models the VQA task as a path-finding problem on scene graphs. The underlying idea can be summarized with the phrase: Learn to walk to the correct answer. More specifically, given an image, we consider a scene graph and train a reinforcement learning agent to conduct a policy-guided random walk on the scene graph until a conclusive inference path is obtained. In contrast to purely embedding-based approaches, our method provides explicit reasoning chains that lead to the derived answers. To sum up, our major contributions are as follows.

- Graphhopper is the first VQA method that employs reinforcement learning for multi-hop reasoning on scene graphs.
- We conduct a thorough experimental study on the challenging VQA dataset named QGA to show the compositional and *interpretable* nature of our model.
- To analyze the reasoning capabilities of our method, we consider manually curated (ground truth) scene graphs. This setting isolates the noise associated with the visual perception task and focuses solely on the language understanding and reasoning task. Thereby, we can show that our method achieves human-like performance.
- Based on both the manually curated scene graphs and our own automatically generated scene graphs, we show that Graphhopper outperforms the Neural State Machine (NMS), a state-of-the-art scene graph reasoning model that operates in a setting, similar to Graphhopper.

Moreover, we are the first group to conduct experiments and publish the code on generated scene graphs for the GQA dataset.¹ The remainder of this work is organized as follows. We review related literature in the next section. Section 3 introduces the notation and describes the methodology of Graphhopper. Section 4 and Sect. 5 detail an experimental study on the benchmark dataset GQA. Furthermore, through a rigorous study using both manually-curated ground-truth and generated scene graphs, we examine the reasoning capabilities of Graphhopper. We conclude in Sect. 6.

2 Related Work

Visual Question Answering: Various models have been proposed that perform VQA on both real-world [4, 17] and artificial datasets [18]. Currently, leading VQA approaches can be categorized into two different branches: First, monolithic neural networks, which perform implicit reasoning on latent representations obtained from fusing the two data modalities. Second, multi-hop methods that form explicit symbolic reasoning chains on a structured representation of the data. Monolithic network architectures obtain visual features from the image either in the form of individual detected objects or by processing the whole

¹ Code is available at: <https://github.com/rajatkoner08/Graphhopper>.

image directly via convolutional neural networks (CNNs). The derived embeddings are usually scored against a fixed answer set along with the embedding of the question obtained from a sequence model. Moreover, co-attention mechanisms are frequently employed to couple the vision and the language models allowing for interactions between objects from both modalities [2, 5, 20, 40, 41]. Monolithic networks are among the dominant methods on previous real-world VQA datasets such as [4]. However, they suffer from the black-box problem and possess limited reasoning capabilities with respect to complex questions that require long reasoning chains (see [7] for a detailed discussion).

Explicit reasoning methods combine the sub-symbolic representation learning paradigm with symbolic reasoning approaches over structured representations of the image. Most of the popular explicit reasoning approaches follow the idea of neural module networks (NMNs) [3] which perform a sequence of reasoning steps realized by forward passes through specialized neural networks that each correspond to predefined reasoning subtasks. Thereby, NMNs construct functional programs by dynamically assembling the modules resulting in a question-specific neural network architecture. In contrast to the monolithic neural network architectures described above, these methods contain a natural transparency mechanism via functional programs. However, while NMN-related methods (e.g., [14, 26]) exhibit good performance on synthetic datasets such as CLEVR [18], they require functional module layouts as additional supervision signals to obtain good results. Closely related to our method is the Neural State Machine (NSM) proposed by [16]. NSM’s underlying idea consists of first constructing a scene graph from an image and treating it as a state machine. Concretely, the nodes correspond to states and edges to transitions. Then, conditioned on the question, a sequence of instructions is derived that indicates how to traverse the scene graph and arrive at the answer. In contrast to NSM, we treat path-finding as a decision problem in a reinforcement learning setting. Concretely, we outline in the next section how extracting predictive paths from scene graphs can be naturally formulated in terms of a goal-oriented random walk induced by a stochastic policy that allows the approach to balance between exploration and exploitation. Moreover, our framework integrates state-of-the-art techniques from graph representation learning and NLP. This paper only considers basic policy gradient methods, but more sophisticated reinforcement learning techniques will be employed in future works.

Statistical Relational Learning: Machine learning methods for KG reasoning aim at exploiting statistical regularities in observed connectivity patterns. These methods are studied under the umbrella of statistical relational learning (SRL) [27]. In recent years, KG embeddings have become the dominant approach in SRL. The underlying idea is that graph features that explain the connectivity pattern of KGs can be encoded in low-dimensional vector spaces. In the embedding spaces, the interactions among the embeddings for entities and relations can be efficiently modeled to produce scores that predict the validity of a triple. Despite achieving good results in KG reasoning tasks, most embedding-based methods have problems capturing the compositionality expressed by long rea-

soning chains. This often limits their applicability in complex reasoning tasks. Recently, multi-hop reasoning methods such as MINERVA [8] and DeepPath [38] were proposed. Both methods are based on the idea that a reinforcement learning agent is trained to perform a policy-guided random walk until the answer entity to a query is reached. Thereby, the path finding problem of the agent can be modeled in terms of a sequential decision making task framed as a Markov decision process (MDP). The method that we propose in this work follows a similar philosophy, in the sense that we train an RL agent to navigate on a scene graph to the correct answer node. However, a conceptual difference is that the agents in MINERVA and DeepPath perform walks on large-scale knowledge graphs exploiting repeating statistical patterns. Thereby, the policies implicitly incorporate approximate rules. In addition, instead of free-form processing questions, the query in the KG reasoning setting is structured as a pair of symbolic entities. That is why we propose a wide range of modifications to adjust our method to the challenging VQA setting.

3 Method

The task of VQA is framed as a scene graph traversal problem. Starting from a hub node that is connected to all other nodes, an agent sequentially samples transitions to neighboring nodes on the scene graph until the node corresponding to the answer is reached. In this way, by adding transitions to the current path, the reasoning chain is successively extended. Before describing the decision problem of the agent, we introduce the notation that we use throughout this work.

Notation: A scene graph is a directed multigraph where each node corresponds to a scene entity which is either an object associated with a bounding box or an attribute of an object. Each scene entity comes with a type that corresponds to the predicted object or attribute label. Typed edges specify how scene entities are related to each other. More formally, let \mathcal{E} denote the set of scene entities and consider the set of binary relations \mathcal{R} . Then a scene graph $\mathcal{SG} \subset \mathcal{E} \times \mathcal{R} \times \mathcal{E}$ is a collection of ordered triples (s, p, o) - subject, predicate, and object. For example, as shown in Fig. 1, the triple $(motorcycle-1, has_part, tire-1)$ indicates that both a motorcycle (subject) and a tire (object) are detected in the image. The predicate has_part indicates the relation between the entities. Moreover, we denote with p^{-1} the inverse relation corresponding to the predicate p . For the remainder of this work, we impose completeness with respect to inverse relations in the sense that for every $(s, p, o) \in \mathcal{SG}$ it is implied that $(o, p^{-1}, s) \in \mathcal{SG}$.

Environment. The state space of the agent \mathcal{S} is given by $\mathcal{E} \times \mathcal{Q}$ where \mathcal{E} are the nodes of a scene graph \mathcal{SG} and \mathcal{Q} denotes the set of all questions. The state at time t is the entity e_t at which the agent is currently located and the question Q . Thus, a state $S_t \in \mathcal{S}$ for time $t \in \mathbb{N}$ is represented by $S_t = (e_t, Q)$. The set of available actions from a state S_t is denoted by \mathcal{A}_{S_t} . It contains all outgoing

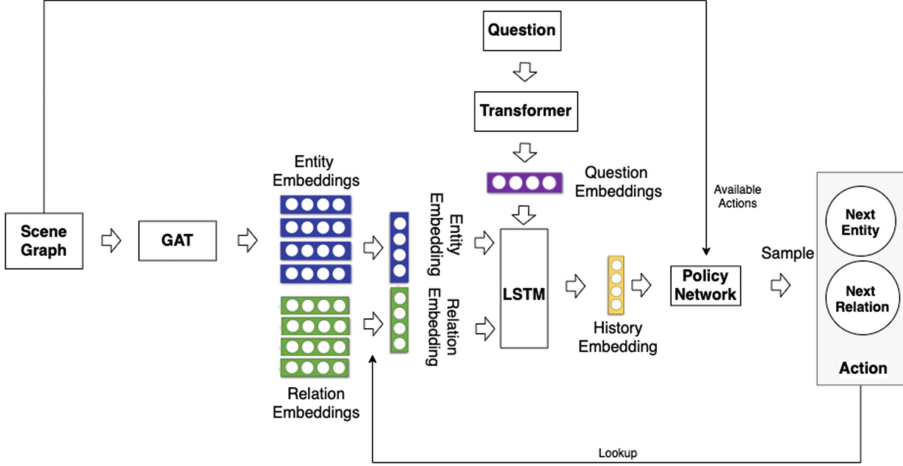


Fig. 2. The architecture of our scene graph reasoning module.

edges from the node e_t together with their corresponding object nodes. More formally, $\mathcal{A}_{S_t} = \{(r, e) \in \mathcal{R} \times \mathcal{E} : S_t = (e_t, Q) \wedge (e_t, r, e) \in \mathcal{SG}\}$. Moreover, we denote with $A_t \in \mathcal{A}_{S_t}$ the action that the agent performed at time t . We include self-loops for each node in \mathcal{SG} that produce a *NO_OP*-label. These self-loops allow the agent to remain at the current location if it reaches the answer node. Furthermore, the introduction of inverse relations allows agent to transit freely in any direction between two nodes (Fig. 2).

The environments evolve deterministically by updating the state according to previous action. Formally, the transition function at time t is given by $\delta_t(S_t, A_t) := (e_{t+1}, Q)$ with $S_t = (e_t, Q)$ and $A_t = (r, e_{t+1})$.

Auxiliary Nodes: In addition to standard entity relation nodes present in a scene graph, we introduce a few auxiliary nodes (e.g. hub node). The underlying rationale for the inclusion of auxiliary nodes is that they facilitate the walk for the agent or help to frame the QA-task as a goal-oriented walk on the scene graph. These additional nodes are included during run-time graph traversal, but they are ignored during the compile time such as when computing node embedding. For example, we add a hub node (*hub*) to every scene graph which is connected to all other nodes. The agent then starts the scene graph traversal from a *hub* with global connectivity. Furthermore for a binary question, we add YES and NO nodes to the scene entities that correspond to the final location of the agent. The agent can then transition to either the YES or the NO node.

Question and Scene Graph Processing. We initialize words in Q with GloVe embeddings [29] with dimension $d = 300$. Similarly we initialize entities and relations in \mathcal{SG} with the embeddings of their type labels. In the scene graph, the node embeddings are passed through a multi-layered graph attention network (GAT) [36]. Extending the idea from graph convolutional networks [22] with

a self-attention mechanism, GATs mimic the convolution operator on regular grids where an entity embedding is formed by aggregating node features from its neighbors. Relations and inverse relations between nodes allows context to flow in both ways through GAT. Thus, the resulting embeddings are context-aware, which makes nodes with the same type, but different graph neighborhoods, distinguishable. To produce an embedding for the question Q , we first apply a Transformer [35], followed by a mean pooling operation.

Finally, since we added auxiliary *YES* and *NO* nodes to the scene graph for binary questions, we train a feedforward neural network to classify query-type (i.e., questions that query for an object in the depicted scene) and binary questions. This network consists of two fully connected layers with ReLU activation on the intermediate output. We find that it is easy to distinguish between query and binary questions (e.g., query questions usually begin with *What*, *Which*, *How*, etc., whereas binary questions usually begin with *Do*, *Is*, etc.). Since our classifier achieves 99.99% accuracy we will ignore the error in question classification in the following discussions.

Policy. We denote the agent’s history until time t with the tuple $H_t = (H_{t-1}, A_{t-1})$ for $t \geq 1$ and $H_0 = \text{hub}$ along with $A_0 = \emptyset$ for $t = 0$. The history is encoded via a multilayered LSTM [13]

$$\mathbf{h}_t = \text{LSTM}(\mathbf{a}_{t-1}), \quad (1)$$

where $\mathbf{a}_{t-1} = [\mathbf{r}_{t-1}, \mathbf{e}_t] \in \mathbb{R}^{2d}$ corresponds to the embedding of the previous action with \mathbf{r}_{t-1} and \mathbf{e}_t denoting the embeddings of the edge and the target node into \mathbb{R}^d , respectively. The history-dependent action distribution is given by

$$\mathbf{d}_t = \text{softmax}(\mathbf{A}_t (\mathbf{W}_2 \text{ReLU}(\mathbf{W}_1 [\mathbf{h}_t, \mathbf{Q}]))) , \quad (2)$$

where the rows of $\mathbf{A}_t \in \mathbb{R}^{|\mathcal{A}_{S_t}| \times d}$ contain latent representations of all admissible actions. Moreover, $\mathbf{Q} \in \mathbb{R}^d$ encodes the question Q . The action $A_t = (r, e) \in \mathcal{A}_{S_t}$ is drawn according to categorical (\mathbf{d}_t) . Equations (1) and (2) induce a stochastic policy π_θ , where θ denotes the set of trainable parameters.

Rewards and Optimization. After sampling T transitions, a terminal reward is assigned according to

$$R = \begin{cases} 1 & \text{if } e_T \text{ is the answer to } Q, \\ 0 & \text{otherwise.} \end{cases} \quad (3)$$

We employ REINFORCE [37] to maximize the expected rewards. Thus, the agent’s maximization problem is given by

$$\arg \max_{\theta} \mathbb{E}_{Q \sim \mathcal{T}} \mathbb{E}_{A_1, A_2, \dots, A_N \sim \pi_\theta} \left[R \mid e_c \right], \quad (4)$$

where \mathcal{T} denote the set of training questions. During training the first expectation in Eq. (4) is substituted with the empirical average over the training set. The

second expectation is approximated by the empirical average over multiple roll-outs. We also employ a moving average baseline to reduce the variance. Further, we use entropy regularization with parameter $\lambda \in \mathbb{R}_{\geq 0}$ to enforce exploration. During inference, we do not sample paths but perform a beam search with width 20 based on the transition probabilities given by Eq. (2).

Additional details on the model, the training and the inference procedure along with sketches of the algorithms, and a complexity analysis can be found in the supplementary material.

4 Dataset and Experimental Setup

In this section we introduce the dataset and detail the experimental protocol.

4.1 Dataset

The *GQA* dataset [17] has been introduced with the goal of addressing key shortcomings of previous VQA datasets, such as *CLEVR* [18] or the *VQA* dataset [4]. *GQA* is more suitable for evaluating the reasoning and compositional abilities of a model in a realistic setting. It contains 113K images, and around 1.2M questions split into roughly 80%/10%/10% for the training, validation, and testing. The overall vocabulary size consists of 3097 words, including 1702 object classes, 310 relationships, and 610 object attributes.

Due to the large number of objects and relationships present in *GQA*, we used a pruned version of the dataset (see Sect. 5) for our generated scene graph. In this work, we have conducted two primary experiments. First, we report the results on manually curated scene graphs provided in the *GQA* dataset. In this setting, the true reasoning and language understanding capabilities of our model can be analyzed. Afterward, we evaluate the performance of our model with the generated scene graphs on pruned *GQA* dataset. It shows the performance of our model on noisy generated data. We have used state of the art Relation Transformer Network (RTN) [23] for the scene graph generation and DetectorS [30] for object detection. We have conducted all the experiments on “test-dev” split of the *GQA*.

Question Types: The questions are designed to evaluate the reasoning abilities such as visual verification, relational reasoning, spatial reasoning, comparison, and logical reasoning. These questions can be categorized either according to structural or semantic criteria. An overview of the different question types is given in supplementary (see Table 4).

4.2 Experimental Setup

Scene Graph Reasoning: Regarding the model parameters, we apply 300 dimensional GloVe embeddings to both the questions and the graphs (i.e., edges and nodes). Moreover, we employ a two-layer GAT [36] model. The dropout [32] probability of each layer is set to 0.1. The first layer has eight attention heads.

Each head has eight latent features which are concatenated to form the output features of that layer. The output layer has eight attention heads with mean aggregation, so that the output also has 300-dimensional features. We apply dropout with $p = 0.1$ to the attention coefficients at each layer. This essentially means that each node is exposed to a stochastically sampled neighborhood during training. Moreover, we employ a two-layer Transformer [35] decoder model. The model dimension is set to 300, and the key and query dimensions are both set to 64 with dropout $p = 0.1$. The LSTM of the policy networks consists of a uni-directional layer with hidden size 300. Finally, the agent performs a fixed number of transitions. In question answering, most questions concern one subject to be explored within one reasoning path originated from the start node. Hence, we set the maximum number of steps to 4, without resetting. By contrast, the binary questions have 8 steps and a reset frequency of 4. In other words, the agent is prompted to the hub node after the fourth step.

Training the Graphhopper: In terms of the training procedure, the GAT, the Transformer, and the policy networks are initialized with Glorot [11] initialization. We train our model with data from the *val_balanced_questions* tier. We use a batch size of 64 and sample a batch of questions along with their associated graphs. We collect 20 stochastic rollouts for each question performed in a vectorized form to utilize parallel computation. For each batch, we collect the rewards when a complete forward pass is done. Then the gradients are approximated from the rewards and applied to update the weights. We employ the Adam optimizer [21] with a learning rate of 10^{-4} for all trainable weights. The coefficient for the action entropy, which balances exploration and exploitation, starts from 0.2 and decreases exponentially at each step with a factor 0.99.

Next to other standard Python libraries, we mainly employed PyTorch [28]. All experiments were conducted on a machine with one NVIDIA RTX 2080 Ti GPU and 64 GB RAM. Training the scene graph reasoner of Graphhopper for 40 epochs on *GQA* takes around 10 h, testing about 1 h.

4.3 Performance Metrics

Along with the accuracy (i.e., Hits@1) on open questions (“Open”), binary questions (yes/no) (“Binary”), and the overall accuracy (“Accuracy”), we also report the additional metric “Consistency” (answers should not contradict themselves), “Validity” (answers are in the range of a question; e.g., *red* is a valid answer when asked for the color of an object), “Plausibility” (answers should be reasonable; e.g., red is a reasonable color of an apple reasonable, blue is not), as proposed in [17].

5 Results and Discussion

As outlined before, VQA is a challenging task, and there is still a significant performance gap between state-of-the-art VQA methods and human performance on

challenging, real-world datasets such as GQA (see [17]). Similar to other existing methods, our architecture involves multiple components, and it is important to be able to analyse the performance of the different modules and processing steps in isolation. Therefore we first present the results of our experiments on manually curated, ground-truth scene graphs provided in the GQA dataset and compare the performance of Graphhopper against NSM and humans. This setting allows us to isolate the noise from the visual perception component and quantify our methods’ reasoning capabilities. Subsequently, we present the results with our own generated scene graphs.

In addition, we also observed that the inclusion of auxiliary nodes helps the agent to achieve efficient performance. *Hub* node performs better compared to starting from any random nodes, as it facilitates easier forward and backtracking from a node. For binary question instead of YES or NO node, we experimented where the path of the agent was processed by another classifier (e.g., a logistic regression) and the classification logits were assigned as rewards. However, this led to inferior results; most likely due to the absence of a weight-sharing mechanism and due to the noisy reward signal produced by the classifier. These observations support our assumption on the role of auxiliary nodes we have used in scene graph.

Reproducing NSM: [15] proposed the state of the art method named NSM for VQA. NSM is the conceptually most similar method, as it also exploits the scene graph reasoning for VQA. We consider NSM to be our baseline method for comparison. However, their approach to reasoning is different from ours. To compare the reasoning ability of our method with the same generated scene graph, we tried to reproduce NSM, as the code for NSM is not open-sourced. We have used the available parameters from [15] and the implementation from [9].

Table 1. A comparison of Graphhopper with human performance and NSM based on manually curated scene graphs.

Method	Binary	Open	Consistency	Validity	Plausibility	Accuracy
Human [17]	91.2	87.4	98.4	98.9	97.2	89.3
NSM [15]	51.03	18.79	81.36	83.69	79.12	34.5
Graphhopper	92.18	92.40	91.92	93.68	93.13	92.30

5.1 Results on Manually Curated Scene Graphs

In this section, we report on an experimental study with Graphhopper on the manually curated scene graphs provided along with the *GQA* dataset. Table 1 shows the performance of Graphhopper and compares it with the human performance reported in [17] and with the performance of NSM on the same underlying manually curated scene graphs. We find that Graphhopper strictly outperforms

NSM with respect to all performance measures. In particular, on the open questions, the performance gap is significant. Moreover, Graphhopper also slightly outperforms humans with respect to the accuracy on both types of questions. On the other hand, concerning the supplementary performance measures consistency, validity, and plausibility, Graphhopper is outperformed by humans but nevertheless consistently reaches high values. Overall, these results can be seen as a testament of the reasoning capabilities and establish an upper bound to the performance of Graphhopper.

5.2 Results on Automatically Generated Graph

The process of generating a graph representation for visual data is a costly and complex procedure. Although the scene graph generation is not the main focus of this work, it constituted one of the major challenges to create good scene graph for GQA due to the following facts:

- There is no open source code for GQA scene graph generation or object detection.
- A large number of instances and an uneven class distribution in GQA leads to a significant drop in the accuracy compared to existing scene graph datasets (see [24]).
- There is a lack of attribute prediction models in modern object detection frameworks.

In this work, we address all of these challenges as our model’s performance is directly dependent on the quality of the scene graph. We will also open-source our code base for transparency and accelerate the development scene graph-based reasoning for VQA.

Generation of Scene Graph: To address these problems, first, we choose two state-of-the-art network, RTN [23] for scene graph generation, and DetectoRS [30] for object detection. The transformer [35] based architecture of RTN and its contextual scene graph embedding is most closely related to our architecture and for our future expansion. To make Graphhopper generic to any scene graph generator, we haven’t use contextualized embedding from RTN, instead we rely on GAT for contextualization.

Pruning of GQA: GQA has more than 6 times the number of relationships compared to Visual Genome [24], which is the most used scene graph generation dataset, and contains more than 18 times the number of objects compared to the most common object detection dataset COCO [25]. Also, the class distribution is highly skewed which causes a significant drop in the accuracy for both the object detection and the scene graph generation task. To efficiently prune the number of instances, we take the first 800 classes, 170 relationships, and 200 attributes based on their frequency of occurrence in the training questions and answers. This pruning allows us to reduce more than 60% of the words while covering more than 96% of the combined answers in the training set.

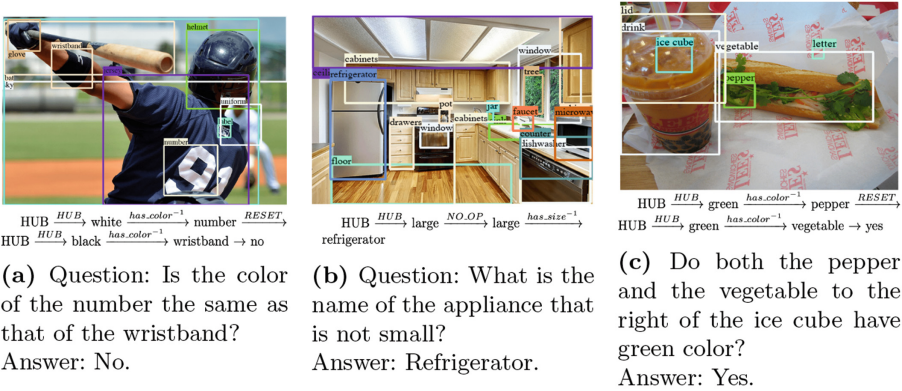


Fig. 3. Three examples question and the corresponding images and paths.

Attribute Prediction: One of the shortcomings of existing scene graph generation and object detection networks is that they do not predict the attributes (e.g., the color or size of an object) of a detected object. Therefore, we have incorporated the attribute prediction for answering the question on GQA. Contextualized object embedding from RTN [23] is used for attribute prediction as

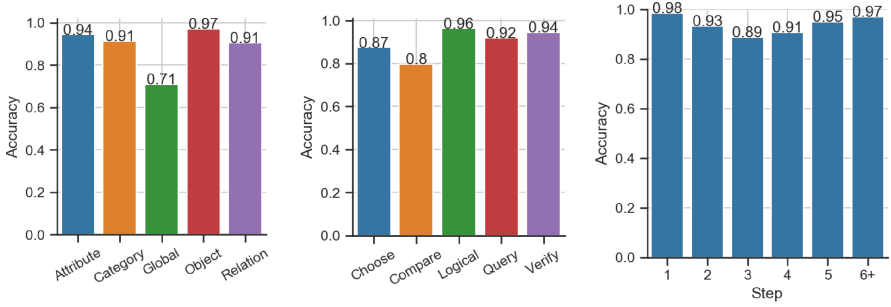
$$P_{attribute} = \sigma(W(Obj_{context}, P_{obj})), \tag{5}$$

where W , $Obj_{context}$, P_{obj} , $P_{attribute}$ are the weight matrices of a linear layer, the contextual embedding of an object, the probability distribution over all objects and the probability distribution over the attributes. σ denotes the sigmoid function.

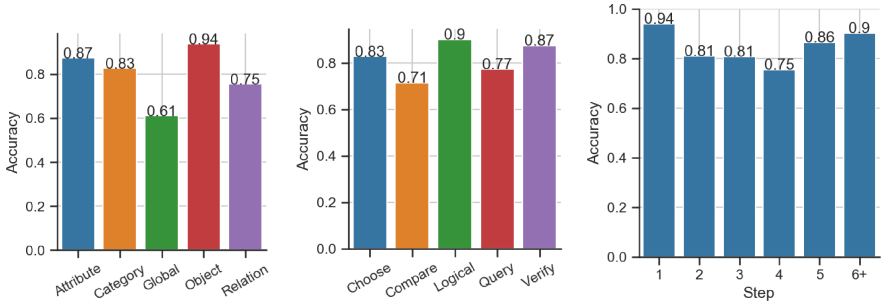
We have trained both the object detector and the scene graph generator on a pruned version of GQA with their respective default parameters after the pre-processing. This helps to increase the coverage of all the instances (e.g., objects, attributes, relationships) on training questions from 52% to 77% implying that our generated scene graph now covers 77% of all instances that represent answers to the training questions.

Table 2. A comparison of our method with NSM, based on generated scene graphs. Graphhopper (pr) indicates that we employed predicted relations from RTN [23].

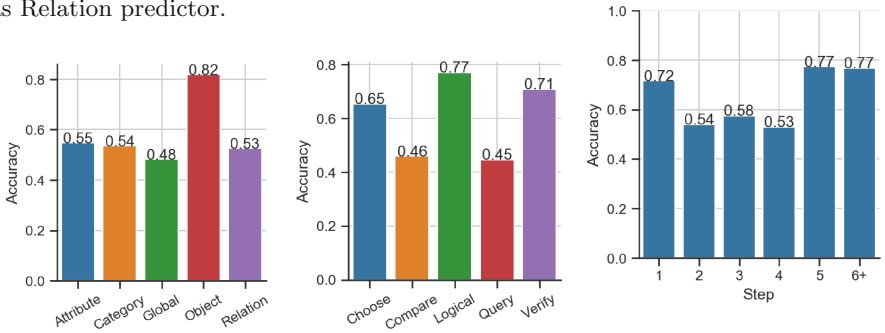
Method	Binary	Open	Consistency	Validity	Plausibility	Accuracy
NSM [15]	51.88	19.83	82.01	86.28	81.75	35.34
Graphhopper	69.48	44.69	83.64	89.42	85.13	56.69
Graphhopper (pr)	85.84	77.27	92.98	92.26	89.50	81.41



(a) Experiments on Manually Curated Scene Graph



(b) Experiments on Ground Truth objects and predicted relation from RTN [23] as Relation predictor.



(c) Experiments on Generated Scene Graph using DetectorRS [30] object detector and RTN [23] as Scene Graph generator.

Fig. 4. Comparison of the performance of our model on various Scene Graph generation settings, (left) accuracy across various semantic instances (“Attribute”, “Global”, “Relation” etc.) required to answer a question (middle) accuracy on multiple types of question category (“Choose”, “Logical”, “Verify” etc.) and (right) accuracy on minimum number of steps needed to reach the answer node.

Table 2, shows the performance of Graphhopper in two settings: First, with a generated graph where we predict the classes, the attributes, and relationships using our own pipeline. Second, where we only use the predicted relationships from RTN [23] (with ground truth objects and attributes). We find that Graphhopper consistently outperforms NSM [15] based on the generated graph. Moreover, in the “pr” or predicted relations setting, it achieves an even higher score as the graphs do not contain any misprediction from the object detector. These encouraging results show superior reasoning abilities both on the generated graph and generated relationships between objects.

5.3 Discussion on the Reasoning Ability

To further analyze the reasoning abilities of Graphhopper, Fig. 4 disentangles the results according to different types of questions: 5 semantic types (left) and 5 structural types (middle). Moreover, we report the performance of Graphhopper according to the length of the reasoning path (right) (see the supplementary material for additional information). Moreover, we show the performance of Graphhopper separately for each of the three scene graph settings that we considered in this work. Figure 4a shows performance on a manually curated scene graph that depicts the actual performance in an ideal environment. Figure 4b illustrates the performance based on only the predicted relationships between objects. This setting shows the performance of Graphhopper along with a scene graph generator. Finally, Fig. 4c depicts the performance based on the object detector, the scene graph generator, and Graphhopper. First and foremost, we find that Graphhopper consistently achieves high accuracy on all types of questions in every setting. Moreover, we find that the performance of Graphhopper does not suffer if answering the questions requires many reasoning steps. We conjecture that this is because high-complexity questions are harder to answer, but due to proper contextualization of the embeddings (e.g., via the GAT and the Transformer), the agent can extract the specific information that identifies the correct target node. The good performance on these high-complexity questions can be seen as evidence that Graphhopper can efficiently translate the question into a transition on the scene graph hopping until the correct answer is reached.

Examples of Reasoning Path: Figure 3 shows three examples of scene graph traversals of Graphhopper that lead to the correct answer. One can see in these examples that the sequential reasoning process over explicit scene graph entities makes the reasoning process more comprehensible. In the case of wrong predictions, the extracted path may offer insights into the mechanics of Graphhopper and facilitate debugging.

6 Conclusion

We have proposed Graphhopper, a novel method for visual question answering that integrates existing KG reasoning, computer vision, and natural language

processing techniques. Concretely, an agent is trained to extract conclusive reasoning paths from scene graphs. To analyze the reasoning abilities of our method, we conducted a rigorous experimental study on both manually curated and generated scene graphs. Based on the manually curated scene graphs we showed that Graphhopper reaches human performance. Moreover, we find that, on our own automatically generated scene graph, Graphhopper outperform another state-of-the-art scene graph reasoning model with respect to all considered performance metrics. In future works, we plan to combine scene graphs with common sense knowledge graphs to further enhance the reasoning abilities of Graphhopper.

References

1. Abbasnejad, E., Teney, D., Parvaneh, A., Shi, J., Hengel, A.V.D.: Counterfactual vision and language learning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 10044–10054 (2020)
2. Anderson, P., et al.: Bottom-up and top-down attention for image captioning and visual question answering. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 6077–6086 (2018)
3. Andreas, J., Rohrbach, M., Darrell, T., Klein, D.: Neural module networks. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 39–48 (2016)
4. Antol, S., et al.: VQA: visual question answering. In: Proceedings of the IEEE International Conference on Computer Vision, pp. 2425–2433 (2015)
5. Cadene, R., Ben-Younes, H., Cord, M., Thome, N.: MUREL: multimodal relational reasoning for visual question answering. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 1989–1998 (2019)
6. Chen, L., Zhang, H., Xiao, J., He, X., Pu, S., Chang, S.F.: Counterfactual critic multi-agent training for scene graph generation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision, pp. 4613–4623 (2019)
7. Chen, W., Gan, Z., Li, L., Cheng, Y., Wang, W., Liu, J.: Meta module network for compositional visual reasoning. arXiv preprint [arXiv:1910.03230](https://arxiv.org/abs/1910.03230) (2019)
8. Das, R., et al.: Go for a walk and arrive at the answer: reasoning over paths in knowledge bases using reinforcement learning. In: ICLR (2018)
9. Eyzaguirre, C.: NSM. <https://github.com/charlespwd/project-title> (2019)
10. Gan, Z., Chen, Y.C., Li, L., Zhu, C., Cheng, Y., Liu, J.: Large-scale adversarial training for vision-and-language representation learning. arXiv preprint [arXiv:2006.06195](https://arxiv.org/abs/2006.06195) (2020)
11. Glorot, X., Bengio, Y.: Understanding the difficulty of training deep feedforward neural networks. In: Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics, pp. 249–256 (2010)
12. Hildebrandt, M., Serna, J.A.Q., Ma, Y., Ringsquandl, M., Joblin, M., Tresp, V.: Reasoning on knowledge graphs with debate dynamics. arXiv preprint [arXiv:2001.00461](https://arxiv.org/abs/2001.00461) (2020)
13. Hochreiter, S., Schmidhuber, J.: Long short-term memory. *Neural Comput.* **9**(8), 1735–1780 (1997)
14. Hu, R., Andreas, J., Rohrbach, M., Darrell, T., Saenko, K.: Learning to reason: end-to-end module networks for visual question answering. In: Proceedings of the IEEE International Conference on Computer Vision, pp. 804–813 (2017)

15. Hudson, D., Manning, C.D.: Learning by abstraction: the neural state machine. In: *Advances in Neural Information Processing Systems*, pp. 5901–5914 (2019)
16. Hudson, D.A., Manning, C.D.: Compositional attention networks for machine reasoning. arXiv preprint [arXiv:1803.03067](https://arxiv.org/abs/1803.03067) (2018)
17. Hudson, D.A., Manning, C.D.: GQA: a new dataset for real-world visual reasoning and compositional question answering. arXiv preprint [arXiv:1902.09506](https://arxiv.org/abs/1902.09506) (2019)
18. Johnson, J., Hariharan, B., van der Maaten, L., Fei-Fei, L., Lawrence Zitnick, C., Girshick, R.: CLEVR: a diagnostic dataset for compositional language and elementary visual reasoning. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 2901–2910 (2017)
19. Johnson, J., et al.: Image retrieval using scene graphs. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 3668–3678 (2015)
20. Kim, J.H., Jun, J., Zhang, B.T.: Bilinear attention networks. In: *Advances in Neural Information Processing Systems*, pp. 1564–1574 (2018)
21. Kingma, D.P., Ba, J.: Adam: a method for stochastic optimization. arXiv preprint [arXiv:1412.6980](https://arxiv.org/abs/1412.6980) (2014)
22. Kipf, T.N., Welling, M.: Semi-supervised classification with graph convolutional networks. arXiv preprint [arXiv:1609.02907](https://arxiv.org/abs/1609.02907) (2016)
23. Koner, R., Sinhamahapatra, P., Tresp, V.: Relation transformer network. arXiv preprint [arXiv:2004.06193](https://arxiv.org/abs/2004.06193) (2020)
24. Krishna, R., et al.: Visual genome: connecting language and vision using crowd-sourced dense image annotations. *Int. J. Comput. Vis.* **123**(1), 32–73 (2017)
25. Lin, T.Y., et al.: Microsoft COCO: common objects in context. In: Fleet, D., Pajdla, T., Schiele, B., Tuytelaars, T. (eds.) *ECCV 2014*. LNCS, vol. 8693, pp. 740–755. Springer, Cham (2014). https://doi.org/10.1007/978-3-319-10602-1_48
26. Mao, J., Gan, C., Kohli, P., Tenenbaum, J.B., Wu, J.: The neuro-symbolic concept learner: interpreting scenes, words, and sentences from natural supervision. arXiv preprint [arXiv:1904.12584](https://arxiv.org/abs/1904.12584) (2019)
27. Nickel, M., Murphy, K., Tresp, V., Gabrilovich, E.: A review of relational machine learning for knowledge graphs. *Proc. IEEE* **104**(1), 11–33 (2015)
28. Paszke, A., et al.: Automatic differentiation in PyTorch (2017)
29. Pennington, J., Socher, R., Manning, C.D.: GloVe: global vectors for word representation. In: *Empirical Methods in Natural Language Processing (EMNLP)*, pp. 1532–1543 (2014). <http://www.aclweb.org/anthology/D14-1162>
30. Qiao, S., Chen, L.C., Yuille, A.: DetectorRS: detecting objects with recursive feature pyramid and switchable Atrous convolution. arXiv preprint [arXiv:2006.02334](https://arxiv.org/abs/2006.02334) (2020)
31. Shi, J., Zhang, H., Li, J.: Explainable and explicit visual reasoning over scene graphs. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 8376–8384 (2019)
32. Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., Salakhutdinov, R.: Dropout: a simple way to prevent neural networks from overfitting. *J. Mach. Learn. Res.* **15**(1), 1929–1958 (2014)
33. Tang, K., Niu, Y., Huang, J., Shi, J., Zhang, H.: Unbiased scene graph generation from biased training. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 3716–3725 (2020)
34. Teney, D., Liu, L., van Den Hengel, A.: Graph-structured representations for visual question answering. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 1–9 (2017)
35. Vaswani, A, et al.: Attention is all you need. In: *Advances in Neural Information Processing Systems*, pp. 5998–6008 (2017)

36. Veličković, P., Cucurull, G., Casanova, A., Romero, A., Lio, P., Bengio, Y.: Graph attention networks. arXiv preprint [arXiv:1710.10903](https://arxiv.org/abs/1710.10903) (2017)
37. Williams, R.J.: Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Mach. Learn.* **8**(3–4), 229–256 (1992)
38. Xiong, W., Hoang, T., Wang, W.Y.: DeepPath: a reinforcement learning method for knowledge graph reasoning. In: Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing (EMNLP 2017), Copenhagen, Denmark. ACL (2017)
39. Yang, Z., He, X., Gao, J., Deng, L., Smola, A.: Stacked attention networks for image question answering. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 21–29 (2016)
40. Yu, Z., Yu, J., Fan, J., Tao, D.: Multi-modal factorized bilinear pooling with co-attention learning for visual question answering. In: Proceedings of the IEEE International Conference on Computer Vision, pp. 1821–1830 (2017)
41. Zhu, C., Zhao, Y., Huang, S., Tu, K., Ma, Y.: Structured attentions for visual question answering. In: Proceedings of the IEEE International Conference on Computer Vision, pp. 1291–1300 (2017)

Open Access This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter’s Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter’s Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

