

# Graph-based Heuristic Search for Module Selection Procedure in Neural Module Network

Yuxuan Wu and Hideki Nakayama

The University of Tokyo  
`{wuyuxuan,nakayama}@nlab.ci.i.u-tokyo.ac.jp`

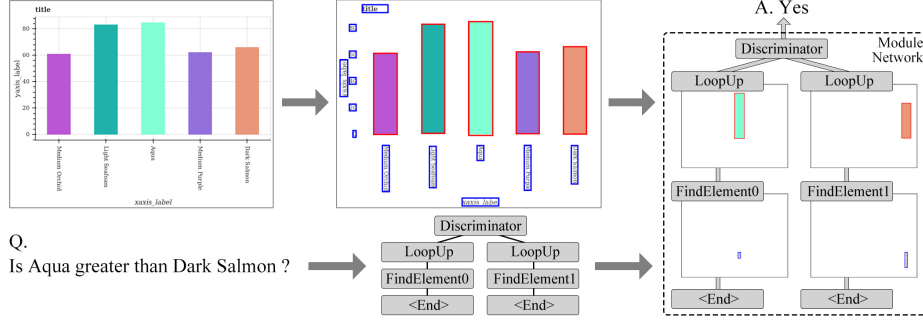
**Abstract.** Neural Module Network (NMN) is a machine learning model for solving the visual question answering tasks. NMN uses programs to encode modules' structures, and its modularized architecture enables it to solve logical problems more reasonably. However, because of the non-differentiable procedure of module selection, NMN is hard to be trained end-to-end. To overcome this problem, existing work either included ground-truth program into training data or applied reinforcement learning to explore the program. However, both of these methods still have weaknesses. In consideration of this, we proposed a new learning framework for NMN. Graph-based Heuristic Search is the algorithm we proposed to discover the optimal program through a heuristic search on the data structure named Program Graph. Our experiments on FigureQA and CLEVR dataset show that our methods can realize the training of NMN without ground-truth programs and achieve superior program exploring efficiency compared to existing reinforcement learning methods.<sup>1</sup>

## 1 Introduction

With the development of machine learning in recent years, more and more tasks have been accomplished such as image classification, object detection, and machine translation. However, there are still many tasks that human beings perform much better than machine learning systems, especially those in need of logical reasoning ability. Neural Module Network (NMN) is a model proposed recently targeted to solve these reasoning tasks [1, 2]. It first predicts a program indicating the required modules and their layout, and then constructs a complete network with these modules to accomplish the reasoning. With the ability to break down complicated tasks into basic logical units and to reuse previous knowledge, NMN achieved super-human level performance on challenging visual reasoning tasks like CLEVR [3]. However, because the module selection is a discrete and non-differentiable process, it is not easy to train NMN end-to-end.

To deal with this problem, a general solution is to separate the training into two parts: the program predictor and the modules. In this case, the program

<sup>1</sup> The code of this work is available at <https://github.com/evan-ak/gbhs>



**Fig. 1.** Our learning framework enables the NMN to solve the visual reasoning problem without ground-truth program annotation.

becomes a necessary intermediate label. The two common solutions to provide this program label are either to include the ground-truth programs into training data or to apply reinforcement learning to explore the optimal candidate program. However, these two solutions still have the following limitations. The dependency on ground-truth program annotation makes NMN’s application hard to be extended to datasets without this kind of annotation. This annotation is also highly expensive while being hand-made by humans. Therefore, program annotation cannot always be expected to be available for tasks in real-world environments. In view of this, methods relying on ground-truth program annotation cannot be considered as complete solutions for training NMN. On the other hand, the main problem in the approaches based on reinforcement learning is that with the growth of the length of programs and number of modules, the size of the search space of possible programs becomes so huge that a reasonable program may not be found in an acceptable time.

In consideration of this, we still regard the training of NMN as an open problem. With the motivation to take advantage of NMN on broader tasks and overcome the difficulty in its training in the meanwhile, in this work, we proposed a new learning framework to solve the non-differentiable module selection problem in NMN.

In this learning framework, we put forward the Graph-based Heuristic Search algorithm to enable the model to find the most appropriate program by itself. Basically, this algorithm is inspired by Monte Carlo Tree Search (MCTS). Similar to MCTS, our algorithm conducts a heuristic search to discover the most appropriate program in the space of possible programs. Besides, inspired by the intrinsic connection between programs, we proposed the data structure named Program Graph to represent the space of possible programs in a way more reasonable than the tree structure used by MCTS. Further, to deal with the cases that the search space is extremely huge, we proposed the Candidate Selection Mechanism to narrow down the search space.

With these proposed methods, our learning framework implemented the training of NMN regardless of the existence of the non-differentiable module

selection procedure. Compared to existing work, our proposed learning framework has the following notable characteristics:

- It can implement the training of NMN with only the triplets of {question, image, answer} and without the ground-truth program annotation.
- It can explore larger search spaces more reasonably and efficiently.
- It can work on both trainable modules with neural architectures and non-trainable modules with discrete processing.

## 2 Related Work

### 2.1 Visual Reasoning

Generally, Visual Reasoning can be considered as a kind of Visual Question Answering (VQA) [4]. Besides the requirement of understanding information from both images and questions in common VQA problems, Visual Reasoning further asks for the capacity to recognize abstract concepts such as spatial, mathematical, and logical relationships. CLEVR [5] is one of the most famous and widely used datasets for Visual Reasoning. It provides not only the triplets of {question, image, answer} but also the functional programs paired with each question. FigureQA [6] is another Visual Reasoning dataset we focus on in this work. It provides questions in fifteen different templates asked on five different types of figures.

To solve Visual Reasoning problems, a naive approach would be the combination of Convolutional Neural Network (CNN) and Recurrent Neural Network (RNN). Here, CNN and RNN are responsible for extracting information from images and questions, respectively. Then, the extracted information is combined and fed to a decoder to obtain the final answer. However, this methodology of treating Visual Reasoning simply as a classification problem sometimes cannot achieve desirable performance due to the difficulty of learning abstract concepts and relations between objects [4, 6, 3]. Instead, more recent work applied models based on NMN to solve Visual Reasoning problems [3, 7–12].

### 2.2 Neural Module Network

Neural Module Network (NMN) is a machine learning model proposed in 2016 [1, 2]. Generally, the overall architecture of NMN can be considered as a controller and a set of modules. Given the question and the image, firstly, the controller of NMN takes the question as input and outputs a program indicating the required modules and their layout. Then, the specified modules are concatenated with each other to construct a complete network. Finally, the image is fed to the assembled network and the answer is acquired from the root module. As far as we are concerned, the advantage of NMN can be attributed to the ability to break down complicated questions into basic logical units and the ability to reuse previous knowledge efficiently.

By the architecture of modules, NMN can further be categorized into three subclasses: the feature-based, attention-based, and object-based NMN.

For feature-based NMNs, the modules apply CNNs and their calculations are directly conducted on the feature maps. Feature-based NMNs are the most concise implementation of NMN and were utilized most in early work [3].

For attention-based NMNs, the modules also apply neural networks but their calculations are conducted on the attention maps. Compared to feature-based NMNs, attention-based NMNs retain the original information within images better so they achieved higher reasoning precision and accuracy [1, 2, 7, 9].

For object-based NMNs, they regard the information in an image as a set of discrete representations on objects instead of a continuous feature map. Correspondingly, their modules conduct pre-defined discrete calculations. Compared to feature-based and attention-based NMNs, object-based NMNs achieved the highest precision on reasoning [10, 11]. However, their discrete design usually requires more prior knowledge and pre-defined attributes on objects.

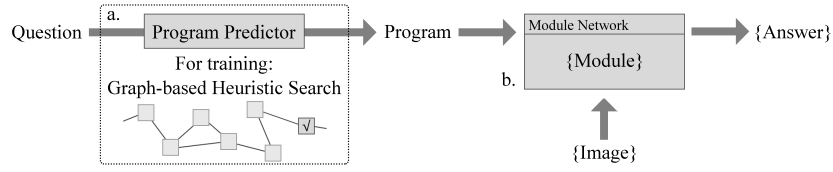
### 2.3 Monte Carlo Methods

Monte Carlo Method is the general name of a group of algorithms that make use of random sampling to get an approximate estimation for a numerical computing [13]. These methods are broadly applied to the tasks that are impossible or too time-consuming to get exact results through deterministic algorithms. Monte Carlo Tree Search (MCTS) is an algorithm that applied the Monte Carlo Method to the decision making in game playing like computer Go [14, 15]. Generally, this algorithm arranges the possible state space of games into tree structures, and then applies Monte Carlo estimation to determine the action to take at each round of games. In recent years, there also appeared approaches to establish collaborations between Deep Learning and MCTS. These work, represented by AlphaGo, have beaten top-level human players on Go, which is considered to be one of the most challenging games for computer programs [16, 17].

## 3 Proposed Method

### 3.1 Overall Architecture

The general architecture of our learning framework is shown as Fig.2. As stated above, the training of the whole model can be divided into two parts: a. Program Predictor and b. modules. The main difficulty of training comes from the side of Program Predictor because of the lack of expected programs as training labels. To overcome this difficulty, we proposed the algorithm named Graph-based Heuristic Search to enable the model to find the optimal program by itself through a heuristic search on the data structure Program Graph. After this searching process, the most appropriate program that was found is utilized as the program label so that the Program Predictor can be trained in a supervised manner. In other words, this searching process can be considered as a procedure targeted to provide training labels for the Program Predictor.



**Fig. 2.** Our Graph-based Heuristic Search algorithm assists the learning of the Program Predictor.

The abstract of the total training workflow is presented as Algorithm 1. Note that here  $q$  denotes the question,  $p$  denotes the program,  $\{module\}$  denotes the set of modules available in the current task,  $\{img\}$  denotes the set of images that the question is asking on,  $\{ans\}$  denotes the set of answers paired with images. Details about the *Sample* function are provided in Appendix A.

---

**Algorithm 1** Total Training Workflow

---

```

1: function TRAIN()
2:   Program_Predictor,  $\{module\} \leftarrow \text{Initialize}()$ 
3:   for loop in range( $Max\_loop$ ) do
4:      $q, \{img\}, \{ans\} \leftarrow \text{Sample}(\text{Dataset})$ 
5:      $p \leftarrow \text{Graph-based\_Heuristic\_Search}(q, \{img\}, \{ans\}, \{module\})$ 
6:     Program_Predictor.train( $q, p$ )
7:   end for
8: end function

```

---

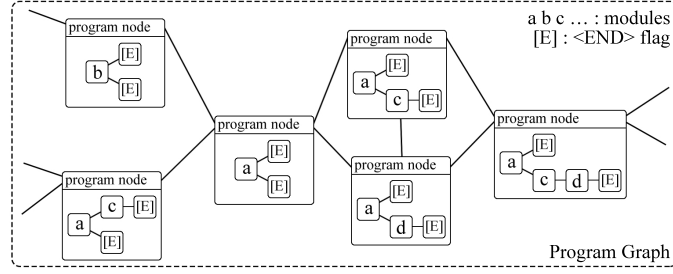
### 3.2 Program Graph

To start with, we first give a precise definition of the program we use. Note that each of the available modules in the model has a unique name, fixed numbers of inputs, and one output. Therefore, a program can be defined as a tree meeting the following rules :

- i) Each of the non-leaf nodes stands for a possible module, each of the leaf nodes holds a  $\langle \text{END} \rangle$  flag.
- ii) The number of children that a node has equal to the number of inputs of the module that the node represents.

For the convenience of representation in prediction, a program can also be transformed into a sequence of modules together with  $\langle \text{END} \rangle$  flags via pre-order tree traversal. Considering that the number of inputs of each module is fixed, the tree form can be rebuilt from such sequence uniquely.

Then, as for the Program Graph, Program Graph is the data structure we use to represent the relation between all programs that have been reached throughout the searching process, and it is also the data structure that our algorithm



**Fig. 3.** Illustration of part of a Program Graph

Graph-based Heuristic Search works on. A Program Graph can be built meeting the following rules :

- i) Each graph node represents a unique program that has been reached.
- ii) There is an edge between two nodes if and only if the edit distance of their programs is one. Here, insertion, deletion, and substitution are the three basic edit operations whose edit distance is defined as one. Note that the edit distance between programs is judged on their tree form.
- iii) Each node in the graph maintains a score. This score is initialized as the output probability of the program of a node according to the Program Predictor when the node is created, and can be updated when the program of a node is executed.

Fig.3 is an illustration of a Program Graph consisting of several program nodes together with their program trees as examples. To distinguish the node in the tree of a program and the node in the Program Graph, the former will be referred to as  $m\_n$  for “module node” and the latter will be referred to as  $p\_n$  for “program node” in the following discussion. Details about the initialization of the Program Graph are provided in Appendix B.

### 3.3 Graph-based Heuristic Search

Graph-based Heuristic Search is the core algorithm in our proposed learning framework. Its basic workflow is presented as the *Main* function in line 1 of Algorithm 2. After Program Graph  $g$  gets initialized, the basic workflow can be described as a recurrent exploration on the Program Graph consisting of the following four steps :

- i) Collecting all the program nodes in Program Graph  $g$  that have not been fully explored yet as the set of candidate nodes  $\{p\_n\}_c$ .
- ii) Calculating the Expectation for all the candidate nodes.
- iii) Selecting the node with the highest Expectation value among all the candidate nodes.
- iv) Expanding on the selected node to generate new program nodes and update the Program Graph.

The details about the calculation of Expectation and expanding strategy are as follows.

**Algorithm 2** Graph-based Heuristic Search

---

```

1: function MAIN( $q, \{img\}, \{ans\}, \{module\}$ )
2:    $g \leftarrow \text{InitializeGraph}(q)$ 
3:   for step in range( $Max\_step$ ) do
4:      $\{p\_n\}_c \leftarrow p\_n$  for  $p\_n$  in  $g$  and  $p\_n.fully\_explored == \text{False}$ 
5:      $p\_n.Exp \leftarrow \text{FindExpectation}(p\_n, g)$  for  $p\_n$  in  $\{p\_n\}_c$ 
6:      $p\_n_e \leftarrow p\_n$  s.t.  $p\_n.Exp = \max\{p\_n.Exp \text{ for } p\_n \text{ in } \{p\_n\}_c\}$ 
7:     Expand( $p\_n_e, g, \{img\}, \{ans\}, \{module\}$ )
8:   end for
9:    $p\_n_{best} \leftarrow p\_n$  s.t.  $p\_n.score = \max\{p\_n.score \text{ for } p\_n \text{ in } \{p\_n\}\}$ 
10:  return  $p\_n_{best}.program$ 
11: end function

12: function EXPAND( $p\_n_e, g, \{img\}, \{ans\}, \{module\}$ )
13:   $p\_n_e.visit\_count \leftarrow p\_n_e.visit\_count + 1$ 
14:  if  $p\_n_e.visited == \text{False}$  then
15:     $p\_n_e.score \leftarrow \text{accuracy}(p\_n_e.program, \{img\}, \{ans\}, \{module\})$ 
16:     $p\_n_e.visited \leftarrow \text{True}$ 
17:  end if
18:   $\{m\_n\}_c \leftarrow m\_n$  for  $m\_n$  in  $p\_n_e.program$  and  $m\_n.expanded == \text{False}$ 
19:   $m\_n_m \leftarrow \text{Sample}(\{m\_n\}_c)$ 
20:   $\{program\}_{new} \leftarrow \text{Mutate}(p\_n_e.program, m\_n_m, \{module\})$ 
21:  for  $program_i$  in  $\{program\}_{new}$  do
22:    if  $\text{LegalityCheck}(program_i) == \text{True}$  then
23:       $g.update(program_i)$ 
24:    end if
25:  end for
26:   $m\_n_m.expanded \leftarrow \text{True}$ 
27:   $p\_n_e.fully\_explored \leftarrow \text{True}$  if  $\{m\_n\}_c.remove(m\_n_m) == \emptyset$ 
28: end function

```

---

**Expectation** Expectation is a measurement defined on each program node to determine which node should be selected for the following expanding step. This Expectation is calculated through the following Equation 1.

$$\begin{aligned}
p\_n_i.Exp = & \sum_{d=0}^D w_d * \max\{p\_n_j.score \mid p\_n_j \text{ in } g, \text{distance}(p\_n_i, p\_n_j) \leq d\} \\
& + \frac{\alpha}{p\_n_i.visit\_count + 1}
\end{aligned} \tag{1}$$

Intuitively, this equation estimates how desirable a program is to guide the modules to answer a given question reasonably. Here,  $D$ ,  $w_d$ , and  $\alpha$  are hyper-parameters indicating the max distance in consideration, a sequence of weight coefficients while summing best scores within different distances  $d$ , and the scale coefficient to encourage visiting less-explored nodes, respectively.

In this equation, the first term observes the nodes nearby and find the highest score within each different distance  $d$  from 0 to  $D$ . Then, these scores are weighted by  $w_d$  and summed up. Note that the distance here is measured on the Program Graph, which also equals to the edit distance between two programs. The second term in this equation is a balance term negatively correlated to the number of times that a node has been visited and expanded on. This term balances the Expectation levels that unexplored or less-explored nodes get.

**Expansion Strategy** Expansion is another essential procedure of our proposed algorithm as shown in line 12 of Algorithm 2. The main purpose of this procedure is to generate new program nodes and update the Program Graph. To realize this, the five main steps are as follows:

- i) If the program node  $p_{n_e}$  is visited for the first time, try its program by building the model with specified modules to answer the question. The question answering accuracy is used to update its score. If there are modules with neural architecture, these modules should also be trained here, but the updated parameters are retained only if the new accuracy exceeds the previous one.
- ii) Collect the module nodes that have not been expanded on yet within the program, then sample one from them as the module node  $m_{n_m}$  to expand on.
- iii) Mutate the program at module  $m_{n_m}$  to generate a set of new programs  $\{\text{program}\}_{new}$  with three edit operations: insertion, deletion, and substitution.
- iv) For each new program judged to be legal, if there is not yet a node representing the same program in the Program Graph  $g$ , then create a new program node representing this program and add it to  $g$ . The related edge should also be added to  $g$  if it does not exist yet.
- v) If all of the module nodes have been expanded on, then mark this program node  $p_{n_e}$  as fully explored.

For the Mutation in step iii), the three edit operations are illustrated by Fig.4. Here, insertion adds a new module node between the node  $m_{n_m}$  and its parent. The new module can be any of the available modules in the model. If the new module has more than one inputs,  $m_{n_m}$  should be set as one of its children, and the rest of the children are set to leaf nodes with  $\langle \text{END} \rangle$  flag.

Deletion removes the node  $m_{n_m}$  and set its child as the new child of  $m_{n_m}$ 's parent. If  $m_{n_m}$  has more than one child, only one of them should be retained and the others are abandoned.

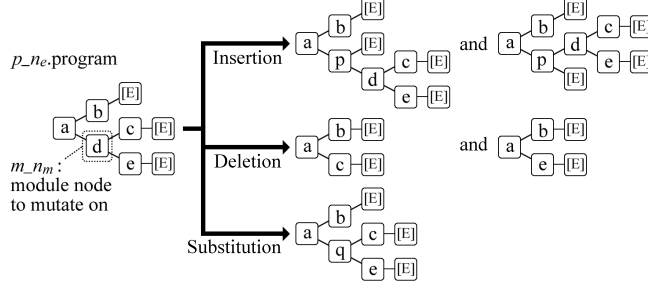
Substitution replaces the module of  $m_{n_m}$  with another module. The new module can be any of the modules that have the same number of inputs as  $m_{n_m}$ .

For insertion and deletion, if there are multiple possible mutations because the related node has more than one child as shown in Fig.4, all of them are retained.

These rules ensure that newly generated programs consequentially have legal structures, but there are still cases that these programs are not legal in the sense of semantics, e.g., the output data type of a module does not match the input data type of its parent. Legality check is conducted to determine whether



a program is legal and should be added to the Program Graph, more details about this function are provided in Appendix C.



**Fig. 4.** Example of the mutations generated by the three operations insertion, deletion, and substitution.

### 3.4 Candidate Selection Mechanism for Modules

The learning framework presented above is already a complete framework to realize the training of the NMN. However, in practice we found that with the growth of the length of programs and the number of modules, the size of search space explodes exponentially. This brings trouble to the search. To overcome this problem, we further proposed the Candidate Selection Mechanism (CSM), which is an optional component within our learning framework. Generally speaking, if CSM is activated, it selects only a subset of modules from the whole of available modules. Then, only these selected modules are used in the following Graph-based Heuristic Search. The abstract of the training workflow with CSM is presented as Algorithm 3.

Here, we included another model named Necessity Predictor into the learning framework. This model takes the question as input, and predicts a  $N_m$ -dimensions vector as shown in Fig.5. Here,  $N_m$  indicates the total number of

---

#### Algorithm 3 Training Workflow with Candidate Selection Mechanism

---

```

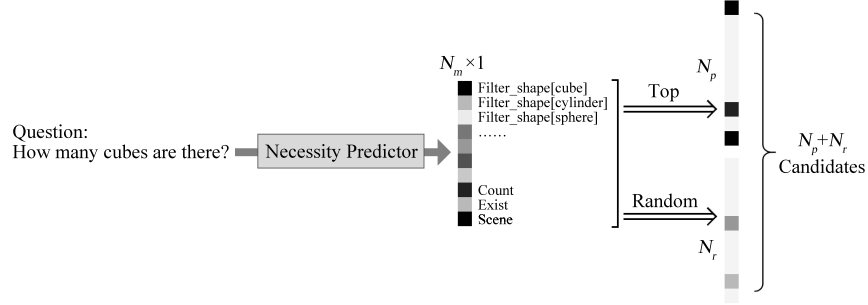
1: function TRAIN()
2:   Program_Predictor, Necessity_Predictor, {module}  $\leftarrow$  Initialize()
3:   for loop in range(Max_Loop) do
4:      $q, \{img\}, \{ans\} \leftarrow$  Sample(Dataset)
5:     {module}candidate  $\leftarrow$  Necessity_Predictor( $q, \{module\}$ )
6:      $p \leftarrow$  Graph-based_Heuristic_Search( $q, \{img\}, \{ans\}, \{module\}_{candidate}$ )
7:     Necessity_Predictor.train( $q, p$ )
8:     Program_Predictor.train( $q, p$ )
9:   end for
10: end function

```

---

modules. Each component of this output is a real number between zero and one indicating the possibility that each module is necessary for solving the given question.  $N_p$  and  $N_r$  are the two hyperparameters for the candidate modules selection procedure.  $N_p$  indicates the number of modules to select according to the predicted possibility value. The  $N_p$  modules with the top  $N_p$  values of predictions are selected.  $N_r$  indicates the number of modules that are selected randomly besides the  $N_p$  selected ones. The union of these two selections with  $N_p + N_r$  modules becomes the candidate modules for the following search.

For the training of this Necessity Predictor, the optimal program found in the search is transformed into a  $N_m$ -dimensions boolean vector indicating whether each module appeared in the program. Then, this boolean vector is set as the training label so that the Necessity Predictor can also be trained in a supervised manner as Program Predictor does.



**Fig. 5.** The process to selecte the  $N_p + N_r$  candidate modules

## 4 Experiments and Results

Our experiments are conducted on the FigureQA and the CLEVR dataset. Their settings and results are presented in the following subsections respectively.

### 4.1 FigureQA Dataset

The main purpose of the experiment on FigureQA is to certify that our learning framework can realize the training of NMN on a dataset without ground-truth program annotations and outperform the existing methods with models other than NMN.

An overview of how our methods work on this dataset is shown in Fig.6. Considering that the size of the search space of the programs used in FigureQA is relatively small, the CSM introduced in Section 3.4 is not activated.

Generally, the workflow consists of three main parts. Firstly, the technique of object detection [18] together with optical character recognition [19] are applied

**Table 1.** Setting of hyperparameters in our experiment

| $Max\_loop$ | $Max\_step$ | $D$ | $w_d$                  | $\alpha$ |
|-------------|-------------|-----|------------------------|----------|
| 100         | 1000        | 4   | (0.5, 0.25, 0.15, 0.1) | 0.05     |

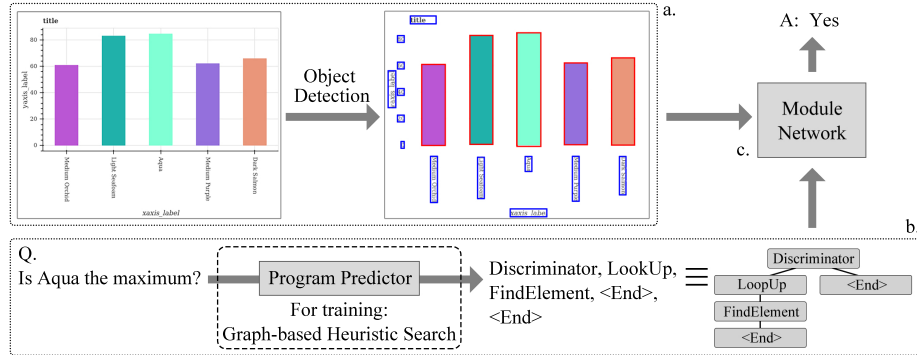
to transform the raw image into discrete element representations as shown in Fig.6.a. For this part, we applied Faster R-CNN [20, 21] with ResNet 101 as the backbone for object detection and Tesseract open source OCR engine [22, 23] for text recognition. All the images are resized to 256 by 256 pixels before following calculations.

Secondly, for the part of program prediction as shown in Fig.6.b., we applied our Graph-based Heuristic Search algorithm for the training. The setting of the hyperparameters for this part are shown in Table 1. The type of figure is treated as an additional token appended to the question.

Thirdly, for the part of modules as shown in Fig.6.c., we designed some pre-defined modules with discrete calculations on objects. Their functions are corresponded to the reasoning abilities required by FigureQA. These pre-defined modules are used associatively with other modules with neural architecture. Details of all these modules are provided in Appendix D.

Table 2 shows the results of our methods compared with baseline and existing methods. “Ours” is the primitive result from the experiment settings presented above. Besides, we also provide the result named “Ours + GE” where “GE” stands for ground-truth elements. In this case, element annotations are obtained directly from ground-truth plotting annotations provided by FigureQA instead of the object detection results. We applied this experiment setting to measure the influence of the noise in object detection results.

Through the result, firstly it can be noticed that both our method and our method with GE outperform all the existing methods. In our consideration, the superiority of our method mainly comes from the successful application of NMN. As stated in Section 2.2, NMN has shown outstanding capacity in solving logical

**Fig. 6.** An example of the inference process on FigureQA

**Table 2.** Comparison of accuracy with previous methods on the FigureQA dataset.

| Method                   | Accuracy        |               |               |               |
|--------------------------|-----------------|---------------|---------------|---------------|
|                          | Validation Sets |               | Test Sets     |               |
|                          | Set 1           | Set 2         | Set 1         | Set 2         |
| Text only [6]            | 50.01%          |               | 50.01%        |               |
| CNN+LSTM [6]             | 56.16%          |               | 56.00%        |               |
| Relation Network [6, 24] | 72.54%          |               | 72.40%        |               |
| Human [6]                |                 |               | 91.21%        |               |
| FigureNet [25]           |                 |               | 84.29%        |               |
| PTGRN [26]               | 86.25%          |               | 86.23%        |               |
| PReFIL [27]              | 94.84%          | 93.26%        | 94.88%        | 93.16%        |
| Ours                     | 95.74%          | 95.55%        | <b>95.61%</b> | <b>95.28%</b> |
| Ours + GE                | <b>96.61%</b>   | <b>96.52%</b> |               |               |

problems. However, limited by the non-differentiable module selection procedure, the application of NMN can hardly be extended to those tasks without ground-truth program annotations like FigureQA. In our work, the learning framework we proposed can realize the training of NMN without ground-truth programs so that we succeeded to apply NMN on this FigureQA. This observation can also be certified through the comparison between our results and PReFIL.

Compared to PReFIL, considering that we applied the nearly same 40-layer DenseNet to process the image, the main difference we made in our model is the application of modules. The modules besides the final Discriminator ensure that the inputs fed to the Discriminator are related to what the question is asking on more closely.

Here, another interesting fact shown by the result is the difference between accuracies reached on set 1 and set 2 of both validation sets and test sets. Note that in FigureQA, validation set 1 and test set 1 adopted the same color scheme as the training set, while validation set 2 and test set 2 adopted an alternated color scheme. This difference leads to the difficulty of the generalization from the training set to the two set 2. As a result, for PReFIL the accuracy on each set 2 drops more than 1.5% from the corresponding set 1. However, for our method with NMN, this decrease is only less than 0.4%, which shows a better generalization capacity brought by the successful application of NMN.

Also, Appendix E reports the accuracies achieved on test set 2 by different question types and figure types. It is worth mentioning that our work is the first one to exceed human performance on every question type and figure type.

## 4.2 CLEVR Dataset

The main purpose of the experiment on CLEVR is to certify that our learning framework can achieve superior program exploring efficiency compared to the classic reinforcement learning method.

For this experiment, we created a subset of CLEVR only containing the training data whose questions appeared at least two times in the whole training set.

This subset contains 31252 different questions together with their corresponding annotated programs. The reason for applying such a subset is that the size of the whole space of possible programs in CLEVR is approximately up to  $10^{40}$ , which is so huge that no existing method can realize the search in it without any prior knowledge or simplification on programs. Considering that the training of modules is highly time-consuming, we only activated the part of program prediction in our learning framework, which is shown as Fig.6.b. With this setting, the modules specified by the program would not be trained actually. Instead, a boolean value indicating whether the program is correct or not is returned to the model as a substitute for the question answering accuracy. Here, only the programs that are exactly the same as the ground-truth programs paired with given questions are considered as correct.

In this experiment, comparative experiments were made on the cases of both activating and not activating the CSM. The structures of the models used as the Program Predictor and the Necessity Predictor are as follows. For Program Predictor, we applied a 2-layer Bidirectional LSTM with hidden state size of 256 as the encoder, and a 2-layer LSTM with hidden state size of 512 as the decoder. Both the input embedding size of encoder and decoder are 300. The setting of hyperparameters are the same as FigureQA as shown in Table.1 except that *Max.loop* is not limited. For Necessity Predictor, we applied a 4-layer MLP. The input of the MLP is a boolean vector indicating whether each word in the dictionary appears in the question, the output of the MLP is a 39-dimensional vector for there are 39 modules in CLEVR, the size of all hidden layers is 256. The hyperparameters  $N_p$  and  $N_r$  are set to 15 and 5 respectively. For the sentence embedding model utilized in the initialization of the Program Graph, we applied the GenSen model with pre-trained weights [28, 29].

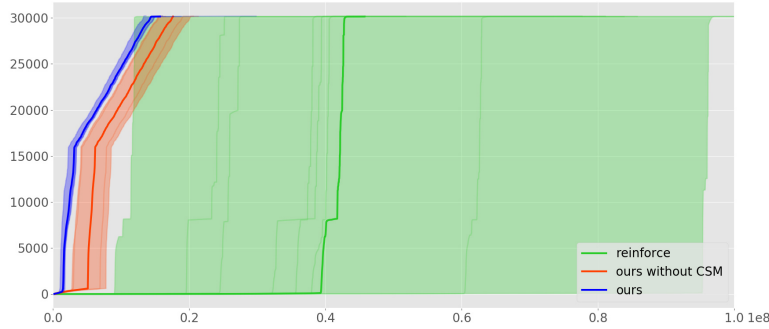
For the baseline, we applied REINFORCE [30] as most of the existing work [3, 12] did to train the same Program Predictor model.

The searching processes of our method, our method without CSM, and REINFORCE are shown by Fig.7. Note that in this figure, the horizontal axis indicates the times of search, the vertical axis indicates the number of correct programs found. The experiments on our method and our method without CSM are repeated four times each, and the experiment on REINFORCE is repeated eight times. Also, we show the average results as the thick solid lines in this figure. They indicate the average times of search consumed to find specific numbers of correct programs. Although in this subset of CLEVR, the numbers of correct programs that can be finally found are quite similar for the three methods, their searching processes show great differences. From this result, three main conclusions can be drawn.

Firstly, in terms of the average case, our method shows a significantly higher efficiency in exploring appropriate programs.

Secondly, the searching process of our method is much more stable while the best case and worst case of REINFORCE differ greatly.

Thirdly, the comparison between the result of our method and our method without CSM certified the effectiveness of the CSM.



**Fig. 7.** Relation between the times of search and the number of correct programs found within the searching processes of three methods.

## 5 Conclusion

In this work, to overcome the difficulty of training the NMN because of its non-differentiable module selection procedure, we proposed a new learning framework for the training of NMN. Our main contribution in this framework can be summarized as follows.

Firstly, we proposed the data structure named Program Graph to represent the search space of programs more reasonably.

Secondly and most importantly, we proposed the Graph-based Heuristic Search algorithm to enable the model to find the most appropriate program by itself to get rid of the dependency on the ground-truth programs in training.

Thirdly, we proposed the Candidate Selection Mechanism to improve the performance of the learning framework when the search space is huge.

Through the experiment, the experiment on FigureQA certified that our learning framework can realize the training of NMN on a dataset without ground-truth program annotations and outperform the existing methods with models other than NMN. The experiment on CLEVR certified that our learning framework can achieve superior efficiency in searching programs compared to the classic reinforcement learning method. In view of this evidence, we conclude that our proposed learning framework is a valid and advanced approach to realize the training of NMN.

Nevertheless, our learning framework still shows weakness in dealing with the extremely huge search spaces, e.g., the whole space of possible programs in CLEVR. We leave further study on methods that can realize the search in such enormous search spaces for future work.

## Acknowledgment

This work was supported by JSPS KAKENHI Grant Number JP19K22861.

## References

1. Andreas, J., Rohrbach, M., Darrell, T., Klein, D.: Neural module networks. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. (2016) 39–48
2. Andreas, J., Rohrbach, M., Darrell, T., Klein, D.: Learning to compose neural networks for question answering. In: Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies. (2016) 1545–1554
3. Johnson, J., Hariharan, B., van der Maaten, L., Hoffman, J., Fei-Fei, L., Lawrence Zitnick, C., Girshick, R.: Inferring and executing programs for visual reasoning. In: Proceedings of the IEEE International Conference on Computer Vision. (2017) 2989–2998
4. Antol, S., Agrawal, A., Lu, J., Mitchell, M., Batra, D., Lawrence Zitnick, C., Parikh, D.: Vqa: Visual question answering. In: Proceedings of the IEEE international conference on computer vision. (2015) 2425–2433
5. 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. (2017) 2901–2910
6. Kahou, S.E., Michalski, V., Atkinson, A., Kádár, Á., Trischler, A., Bengio, Y.: Figureqa: An annotated figure dataset for visual reasoning. In: International Conference on Learning Representations. (2018)
7. 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. (2017) 804–813
8. Hu, R., Andreas, J., Darrell, T., Saenko, K.: Explainable neural computation via stack neural module networks. In: Proceedings of the European conference on computer vision (ECCV). (2018) 53–69
9. Mascharka, D., Tran, P., Soklaski, R., Majumdar, A.: Transparency by design: Closing the gap between performance and interpretability in visual reasoning. In: Proceedings of the IEEE conference on computer vision and pattern recognition. (2018) 4942–4950
10. 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. (2019) 8376–8384
11. Yi, K., Wu, J., Gan, C., Torralba, A., Kohli, P., Tenenbaum, J.: Neural-symbolic vqa: Disentangling reasoning from vision and language understanding. In: Advances in Neural Information Processing Systems. (2018) 1031–1042
12. Mao, J., Gan, C., Kohli, P., Tenenbaum, J.B., Wu, J.: The neuro-symbolic concept learner: Interpreting scenes, words, and sentences from natural supervision. In: International Conference on Learning Representations. (2019)
13. Metropolis, N., Ulam, S.: The monte carlo method. *Journal of the American statistical association* **44** (1949) 335–341
14. Kocsis, L., Szepesvári, C.: Bandit based monte-carlo planning. In: European conference on machine learning, Springer (2006) 282–293
15. Coulom, R.: Efficient selectivity and backup operators in monte-carlo tree search. In: International conference on computers and games, Springer (2006) 72–83

16. Silver, D., Huang, A., Maddison, C.J., Guez, A., Sifre, L., Van Den Driessche, G., Schrittwieser, J., Antonoglou, I., Panneershelvam, V., Lanctot, M., et al.: Mastering the game of go with deep neural networks and tree search. *nature* **529** (2016) 484
17. Silver, D., Schrittwieser, J., Simonyan, K., Antonoglou, I., Huang, A., Guez, A., Hubert, T., Baker, L., Lai, M., Bolton, A., et al.: Mastering the game of go without human knowledge. *Nature* **550** (2017) 354
18. Girshick, R., Donahue, J., Darrell, T., Malik, J.: Rich feature hierarchies for accurate object detection and semantic segmentation. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. (2014) 580–587
19. Singh, S.: Optical character recognition techniques: a survey. *Journal of emerging Trends in Computing and information Sciences* **4** (2013) 545–550
20. Ren, S., He, K., Girshick, R., Sun, J.: Faster R-CNN: Towards real-time object detection with region proposal networks. In: *Advances in Neural Information Processing Systems (NIPS)*. (2015)
21. Yang, J., Lu, J., Batra, D., Parikh, D.: A faster pytorch implementation of faster r-cnn. <https://github.com/jwyang/faster-rcnn.pytorch> (2017)
22. Smith, R.: An overview of the tesseract ocr engine. In: *Ninth International Conference on Document Analysis and Recognition (ICDAR 2007)*. Volume 2., IEEE (2007) 629–633
23. Smith, R.: Tesseract open source ocr engine. <https://github.com/tesseract-ocr/tesseract> (2019)
24. Santoro, A., Raposo, D., Barrett, D.G., Malinowski, M., Pascanu, R., Battaglia, P., Lillicrap, T.: A simple neural network module for relational reasoning. In: *Advances in neural information processing systems*. (2017) 4967–4976
25. Reddy, R., Ramesh, R., Deshpande, A., Khapra, M.M.: FigureNet: A deep learning model for question-answering on scientific plots. In: *2019 International Joint Conference on Neural Networks (IJCNN)*, IEEE (2019) 1–8
26. Cao, Q., Liang, X., Li, B., Lin, L.: Interpretable visual question answering by reasoning on dependency trees. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2019)
27. Kafle, K., Shrestha, R., Price, B., Cohen, S., Kanan, C.: Answering questions about data visualizations using efficient bimodal fusion. *arXiv preprint arXiv:1908.01801* (2019)
28. Subramanian, S., Trischler, A., Bengio, Y., Pal, C.J.: Learning general purpose distributed sentence representations via large scale multi-task learning. In: *International Conference on Learning Representations*. (2018)
29. Subramanian, S., Trischler, A., Bengio, Y., Pal, C.J.: Gensen. <https://github.com/Maluuba/gensen> (2018)
30. Williams, R.J.: Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning* **8** (1992) 229–256