

Self-Assembling Modular Networks for Interpretable Multi-Hop Reasoning

Yichen Jiang and Mohit Bansal
UNC Chapel Hill
{yichenj, mbansal}@cs.unc.edu

Abstract

Multi-hop QA requires a model to connect multiple pieces of evidence scattered in a long context to answer the question. The recently proposed HotpotQA (Yang et al., 2018) dataset is comprised of questions embodying four different multi-hop reasoning paradigms (two bridge entity setups, checking multiple properties, and comparing two entities), making it challenging for a single neural network to handle all four. In this work, we present an interpretable, controller-based Self-Assembling Neural Modular Network (Hu et al., 2017, 2018) for multi-hop reasoning, where we design four novel modules (Find, Relocate, Compare, NoOp) to perform unique types of language reasoning. Based on a question, our layout controller RNN dynamically infers a series of reasoning modules to construct the entire network. Empirically, we show that our dynamic, multi-hop modular network achieves significant improvements over the static, single-hop baseline (on both regular and adversarial evaluation). We further demonstrate the interpretability of our model via three analyses. First, the controller can softly decompose the multi-hop question into multiple single-hop sub-questions to promote compositional reasoning behavior of the main network. Second, the controller can predict layouts that conform to the layouts designed by human experts. Finally, the intermediate module can infer the entity that connects two distantly-located supporting facts by addressing the sub-question from the controller.¹

1 Introduction

The task of multi-hop question answering (QA) requires the model to answer a natural language question by finding multiple relevant information

¹Our code is publicly available at: <https://github.com/jiangycTarheel/NMN-MultiHopQA>

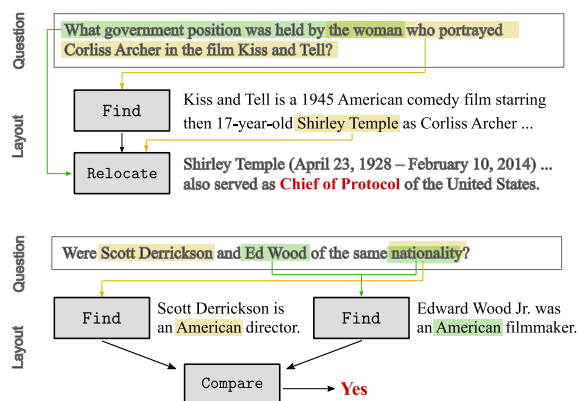


Figure 1: Two HotpotQA examples and the modular network layout predicted by the controller.

pieces scattered in a given natural language context. It has attracted more attention recently and multiple datasets have been proposed, including the recent HotpotQA (Yang et al., 2018) that is comprised of questions embodying four different multi-hop reasoning paradigms: inferring the bridge entity to complete the 2nd-hop question (first question in Fig. 1), inferring the answer through a bridge entity, checking multiple properties to select the answer, and comparing two entities (second question in Fig. 1). For the first question, it is necessary to first find the person “who portrayed Corliss Archer in the film Kiss and Tell”, and then find out the “government position” she held. For the second question, one may arrive at the answer by finding the country where Scott and Ed are from, and then comparing the two nationalities to conclude whether they are the same.

Multi-hop QA is more challenging than single-hop QA for two main reasons. First, the techniques used for single-hop QA are not sufficient to answer a multi-hop question. In single-hop QA tasks like SQuAD (Rajpurkar et al., 2016), the evidence necessary to answer the question is concentrated in a short context (Q: “What is the color of the grass”, Context: “The grass is green.”, An-

swer: “Green”). Such datasets emphasize the role of matching the information between the question and the short context surrounding the answer (“sky→sky, color→blue”), and can be solved by building a question-aware context representation (Seo et al., 2017; Xiong et al., 2017). In contrast, for multi-hop QA, directly matching the semantics of the question and context leads to the entity that bridges the two supporting facts (e.g., “Shirley Temple”), or the entities that need to be compared against each other (e.g., nationalities of Scott and Ed). In both cases, further action is required on top of direct semantic matching in order to get to the final answer. Second, HotpotQA is comprised of questions of four different types of multi-hop reasoning paradigms. In Fig. 1, the first question embodies a sequential reasoning path where the model has to solve a sub-question to get an entity “Shirley Temple”, which then leads to the answer to the main question about her “government position”. The second question, on the other hand, requires a tree-structured reasoning path where the model first locates two entities on the leaves and then compares them to get the answer. The difference in the required reasoning skills makes it hard for a single static model to handle all types of questions.

To automatically discover the multiple elaborated reasoning steps as required in HotpotQA, a model needs to dynamically adopt a sequence of different reasoning behaviors based on specific questions, which is still unexplored in the field of large-scale text-based QA. In our work, we present a highly-interpretable self-assembling Neural Modular Network (Andreas et al., 2016a; Hu et al., 2017) with three novel modules designed for multi-hop NLP tasks: `Find`, `Relocate`, `Compare`, where each module embodies a unique type of reasoning behavior. The `Find` module is similar to the previously-introduced 1-hop bi-attention model (Seo et al., 2017; Xiong et al., 2017), which produces an attention map over the context words given the context and the question representation. For the first example in Fig. 1, a `Find` module is used to find the answer (“Shirley Temple”) to the sub-question (“who portrayed Corliss Archer ...”). The `Relocate` module takes the intermediate answer to the previous sub-question (“Shirley Temple”), combines it with the current sub-question (“What government position was held by ...”), and then outputs a new

attention map conditioned on the previous one. The `Compare` module intuitively “compares” the outputs from two previous modules based on the question. For the second example in Fig. 1, after previous modules find the nationality of “Scott” and “Ed” (both American), the `Compare` module outputs the answer “Yes” based on the word “same” in the question. We also use the `NoOp` module that keeps the current state of the network unchanged when the model decides no more action is needed. After all the reasoning steps are performed, the final attention map is used to predict the answer span, and the vector output from the `Compare` module is used to predict whether the answer should be a context span or “Yes/No”.

Next, to dynamically assemble these modules into a network based on the specific question, we use a controller RNN that, at each step, infers the required reasoning behavior from the question, outputs the sub-question, and predicts a soft combination of all available modules. As shown in Fig. 1, a modular network `Relocate(Find())` is constructed for the bridge-type question and another network with the layout `Compare(Find(), Find())` is built for the comparison-type question. In order to make the entire model end-to-end differentiable in gradient descent, we follow Hu et al. (2018) to use continuous and soft layout prediction and maintain a differentiable stack-based data structure to store the predicted modules’ output. This approach to optimize the modular network is shown to be superior to using a Reinforcement Learning approach which makes hard module decisions. Fig. 2 visualizes this controller that predicts the modular layout and the network assembled with the selected modules. We further apply intermediate supervision to the first produced attention map to encourage finding the bridge entity (“Shirley Temple” in Fig. 1). The entire model is trained end-to-end using cross entropy loss for all supervision.

Overall, our self-assembling controller-based Neural Modular Network achieves statistically significant improvements over both the single-hop bi-attention baseline and the original modular network (Hu et al., 2018) designed for visual-domain QA. We also present adversarial evaluation (Jiang and Bansal, 2019), where single-hop reasoning shortcuts are eliminated and compositional reasoning is enforced; and here our NMN again outperforms the BiDAF baseline significantly in the

EM score (as well as after adversarial training). We further demonstrate the interpretability of our modular network with three analyses. First, the controller understands the multi-hop semantics of the question and can provide accurate step-by-step sub-questions to the module to lead the main network to follow the reasoning path. Second, the controller can successfully predict layouts that conform to the layouts designed by human experts. Finally, before arriving at the final answer, the intermediate module can infer the bridge entity that connects two distantly-located supporting facts by leveraging the step-specific sub-question from the controller. All of these analyses show that our Modular Network is not operating as a black box, but demonstrates highly interpretable compositional reasoning behavior, which is beneficial in transparent model development and safe, trustworthy real-world applications. In summary, the contribution of this work is three-fold: 1) This is the first work to apply self-assembling modular networks to text-based QA; 2) We design three novel modules to handle multi-hop questions in HotpotQA; 3) The resulting network is interpretable in terms of the model’s intermediate outputs and the assembled layout.

2 Related Works

Multi-hop Reading Comprehension The last few years have witnessed significant progress on large-scale QA datasets including cloze-style tasks (Hermann et al., 2015), open-domain QA (Yang et al., 2015), and more (Rajpurkar et al., 2016, 2018). However, all of the above datasets are confined to a single-document context per question domain. Joshi et al. (2017) introduced a multi-document QA dataset with some questions requiring cross-sentence inferences to answer. The bAbI dataset (Weston et al., 2016) requires the model to combine multiple pieces of evidence in the synthetic text. Welbl et al. (2017) uses Wikipedia articles as the context and a subject-relation pair as the query, and constructs the multi-hop QAngaroo dataset by traversing a directed bipartite graph so that the evidence required to answer a query could be spread across multiple documents. HotpotQA (Yang et al., 2018) is a more recent multi-hop QA dataset that has crowd-sourced questions with more diverse syntactic and semantic features compared to QAngaroo. It includes four types of questions, each re-

quiring a different reasoning paradigm. Some examples require inferring the bridge entity from the question (Type I in Yang et al. (2018)), while others demand fact-checking or comparing subjects’ properties from two different documents (Type II and comparison question). Concurrently to our work, Min et al. (2019b) also tackle HotpotQA by decomposing its multi-hop questions into single-hop sub-questions to achieve better performance and interpretability. However, their system approaches the question decomposition by having a decomposer model trained via human labels, while our controller accomplishes this task automatically with the soft attention over the question-words’ representation and is only distantly supervised by the answer and bridge-entity supervision, with no extra human labels. Moreover, they propose a pipeline system with the decomposers, an answer-prediction model, and a decomposition scorer trained separately on the previous stage’s output. Our modular network, on the other hand, is an end-to-end system that is optimized jointly.

Neural Modular Network Neural Modular Network (NMN) is a class of models that are composed of a number of sub-modules, where each sub-module is capable of performing a specific subtask. In NMN (Andreas et al., 2016b), N2NMN (Hu et al., 2017), PG+EE (Johnson et al., 2017) and TbD (Mascharka et al., 2018), the entire reasoning procedure starts by analyzing the question and decomposing the reasoning procedure into a sequence of sub-tasks, each with a corresponding module. This is done by either a parser (Andreas et al., 2016b) or a layout policy (Hu et al., 2017; Johnson et al., 2017; Mascharka et al., 2018) that turns the question into a module layout. Then the module layout is executed with a neural module network. Overall, given an input question, the layout policy learns what sub-tasks to perform, and the neural modules learn how to perform each individual sub-task.

However, since the the modular layout is sampled from the controller, the controller itself is not end-to-end differentiable and has to be optimized using Reinforcement Learning Algorithms like Reinforce (Williams, 1992). Hu et al. (2018) used soft program execution where the output of each step is the average of outputs from all modules weighted by the module distribution, and showed its superiority over hard-layout NMNs. All previous works in NMN, including Hu et al. (2018)

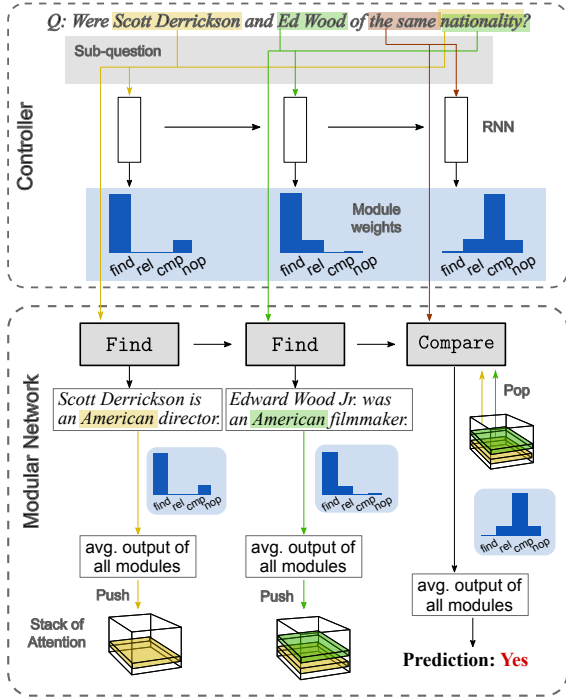


Figure 2: Modular network with a controller (top) and the dynamically-assembled modular network (bottom). At every step, the controller produces a sub-question vector and predicts a distribution to weigh the averages of the modules’ outputs.

targeted visual question answering or structured knowledge-based GeoQA, and hence the modules are designed to process image or KB inputs. We are the first to apply modular networks to unstructured, text-based QA, where we redesigned the modules for language-based reasoning by using bi-attention (Seo et al., 2017; Xiong et al., 2017) to replace convolution and multiplication of the question vector with the image feature. Our model also has access to the full-sized bi-attention vector before it is projected down to the 1-d distribution.

Architecture Learning Our work also shares the spirit of recent research on Neural Architecture Search (NAS) (Zoph and Le, 2016; Pham et al., 2018; Liu et al., 2018), since the architecture of the model is learned dynamically based on controllers instead of being manually-designed. However, Neural Architecture Search aims to learn the structure of the individual CNN/RNN cell with fixed inter-connections between the cells, while Modular Networks have preset individual modules but learns the way to assemble them into a larger network. Moreover, Modular Networks’ architectures are predicted dynamically on each data point, while previous NAS methods learn a single cell structure independent of the example.

3 Model

In this section, we describe how we apply the Neural Modular Network to the multi-hop HotpotQA task (Yang et al., 2018). In Sec. 3.2 and Sec. 3.3, we describe the controller, which sketches the model layout for a specific example, and introduce how it assembles the network with the predicted modules. In Sec. 3.4, we go into details of every novel module that we design for this task.

3.1 Encoding

We use a Highway Network (Srivastava et al., 2015) to merge the character embeddings and GloVe word embeddings (Pennington et al., 2014), building word representations for the context and the question as $\mathbf{x} \in \mathbb{R}^{S \times v}$ and $\mathbf{q} \in \mathbb{R}^{J \times v}$ respectively, where S and J are the lengths of the context and the question. We then apply a bi-directional LSTM-RNN (Hochreiter and Schmidhuber, 1997) of d hidden units to get the contextualized word representations for the context and question: $\mathbf{h} = \text{BiLSTM}(\mathbf{x})$; $\mathbf{u} = \text{BiLSTM}(\mathbf{q})$ so that $\mathbf{h} \in \mathbb{R}^{S \times 2d}$ and $\mathbf{u} \in \mathbb{R}^{J \times 2d}$. We also use a self attention layer (Zhong et al., 2019) to get qv , a fixed-sized vector representation of the question.

3.2 Model Layout Controller

In a modular network, the controller reads the question and predicts a series of our novel modules that could be executed in order to answer the given question. For multi-hop QA, each module represents a specific 1-hop reasoning behavior and the controller needs to deduce a chain of reasoning behavior that uncovers the evidence necessary to infer the answer. Similar to the model in Hu et al. (2018), at the step t , the probability of selecting module i is calculated as:

$$\begin{aligned}
 q_t &= W_{1,t} \cdot qv + b_{1,t} \\
 cq_t &= W_2 \cdot [c_{t-1}; q_t] + b_2 \\
 s_t &= W_3 \cdot cq_t \\
 p_{t,i} &= \text{Softmax}(s_t)
 \end{aligned} \tag{1}$$

where c_{t-1} is the controller’s hidden state and q_t is the vector representation of the question coming out of the encoding LSTM. W_3 projects the output to the dimension of $1 \times I$, where I is the number of available modules. When predicting the next module, the controller also supplies the current module, which we design for the language domain, with a fixed-sized vector as the sub-question at the

current reasoning step (the gray-shaded area of the controller in Fig. 2). Consider the first example in Fig. 1. At the first reasoning step, the `Find` module’s sub-question is “the woman who portrayed Corliss Archer in the film Kiss and Tell”. To generate this sub-question vector, the controller first calculates a distribution over all question words and then computes the weighted average of all the word representations:

$$\begin{aligned} e_{t,j} &= W_4 \cdot (c_{q_t} \cdot u_j) + b_4 \\ c_{v_t} &= \text{Softmax}(e_t) \\ c_t &= \sum_j c_{v_t,j} \cdot u_j \end{aligned} \quad (2)$$

The sub-question vector c_t also serves as the controller’s hidden state and is used to calculate the module probability in the next reasoning step. Similarly, in the second reasoning step, the sub-question of the `Relocate` module is “What government position was held by” and the answer from the `Find` module at the first step.

3.3 Stack NMN with Soft Program Execution

In our NMN, some modules may interact with the attention output from previous modules to collect information that could benefit the current reasoning step. For tree-structured layouts, such as `Compare(Find(), Find())` as shown in Fig. 1, the latter `Compare` module requires access to previous outputs. We follow Hu et al. (2018) to use a stack to store attention outputs generated by our modules. In the tree-structured layout, the two `Find` modules push two attention outputs onto the stack, and the `Compare` module pops these two attention outputs to compare the content they each focus on. To make the entire model differentiable, we also adopt soft program execution from Hu et al. (2018), where the output of each step is the sum of outputs from all modules weighted by the module distribution (blue-shaded area in Fig. 2) computed by the controller.

3.4 NMN Modules

We next describe all the modules we designed for our Neural Modular Network. All modules take the question representation \mathbf{u} , context representation \mathbf{h} , and sub-question vector \mathbf{c}_t as input. The core modules that produce the attention map over the context are based on bi-attention (Seo et al., 2017) between the question and context, instead of relying on convolution as in previous NMNs (Hu

et al., 2017, 2018). Every module outputs a fixed-size vector/zero vector as the memory state of the NMN, and is also able to push/pop attention maps onto/from the stack. These modules, each solving a single-hop sub-question, are chained together according to the layout predicted by the controller to infer the final answer to the multi-hop question.

Find This module performs the bidirectional attention (Seo et al., 2017) between question \mathbf{u} and context $\mathbf{h}' = \mathbf{h} \cdot \mathbf{c}_t$. By multiplying the original context representation by \mathbf{c}_t , we inject the sub-question into the following computation. The model first computes a similarity matrix $M^{J \times S}$ between every question and context word and use it to derive context-to-query attention:

$$\begin{aligned} M_{j,s} &= W_1 \mathbf{u}_j + W_2 \mathbf{h}'_s + W_3 (\mathbf{u}_j \odot \mathbf{h}'_s) \\ p_{j,s} &= \frac{\exp(M_{j,s})}{\sum_{j=1}^J \exp(M_{j,s})} \\ \mathbf{c}_{q_s} &= \sum_{j=1}^J p_{j,s} \mathbf{u}_j \end{aligned} \quad (3)$$

where W_1, W_2 and W_3 are trainable parameters, and \odot is element-wise multiplication. Then the query-to-context attention vector is derived as:

$$\begin{aligned} m_s &= \max_{1 \leq s \leq S} M_{j,s} \\ p_s &= \frac{\exp(m_s)}{\sum_{s=1}^S \exp(m_s)} \\ \mathbf{q}_c &= \sum_{s=1}^S p_s \mathbf{h}_s \end{aligned} \quad (4)$$

We then obtain the question-aware context representation as: $\tilde{\mathbf{h}} = [\mathbf{h}'; \mathbf{c}_q; \mathbf{h}' \odot \mathbf{c}_q; \mathbf{c}_q \odot \mathbf{q}_c]$, and push this bi-attention result onto the stack. This process is visualized in the first two steps in Fig. 2.

Relocate For the first example in Fig. 1, the second reasoning step requires finding the “government position” held by the woman, who is identified in the first step. We propose the `Relocate` module to model this reasoning behavior of finding the answer based on the information from the question as well as the previous reasoning step. The `Relocate` module first pops an attention map att_1 from the stack and computes the bridge entity’s representation \mathbf{b} as the weighted average over \mathbf{h} using the popped attention. It then computes a bridge-entity-aware representation of context \mathbf{h}_b , and then applies a `Find`

module between \mathbf{h}_b and question \mathbf{u} :

$$\begin{aligned} \mathbf{b} &= \sum_s att_{1,s} \cdot \mathbf{h}_s \\ \mathbf{h}_b &= \mathbf{h} \odot \mathbf{b} \\ \tilde{\mathbf{h}} &= \text{Find}(\mathbf{u}, \mathbf{h}_b, \mathbf{c}_t) \end{aligned} \quad (5)$$

The output $\tilde{\mathbf{h}}$ is then pushed onto the stack.

Compare As shown in the final step in Fig. 2, the `Compare` module pops two attention maps att_1 and att_2 from the stack and computes two weighted averages over \mathbf{h} using the attention maps. It then merges these two vectors with \mathbf{c}_t to update the NMN’s memory state m .

$$\begin{aligned} hs_1 &= \sum_s att_{1,s} \cdot \mathbf{h}_s; \quad hs_2 = \sum_s att_{2,s} \cdot \mathbf{h}_s \\ o_{in} &= [c_t; hs_1; hs_2; c_t \cdot (hs_1 - hs_2)] \\ m &= W_1 \cdot (\text{ReLU}(W_2 \cdot o_{in})) \end{aligned} \quad (6)$$

The intuition is that the `Compare` module aggregates the information from two attention maps and compares them according to the sub-question \mathbf{c}_t .

NoOp This module updates neither the stack nor the memory state. It can be seen as a skip command when the controller decides no more computation is required for the current example.

3.5 Bridge Entity Supervision

Even with the multi-hop architecture to capture the hop-specific distribution over the question, there is no supervision on the controller’s output distribution c about which part of the question is related to the current reasoning step, thus preventing the controller from learning the composite reasoning skill. To address this problem, we supervise the first `Find` module to predict the bridge entity (“Shirley Temple” in Fig. 1), which indirectly encourages the controller to look for question information related to this entity (“the woman who portrayed Corliss Archer...”) at the first step. Since the original dataset does not label the bridge entity, we apply a simple heuristic to detect the bridge entity that is the title of one supporting document while also appearing in the other supporting document.²

3.6 Prediction Layer

After all the reasoning steps have completed, to predict a span from the context as the answer, we

²There are two supporting documents per example.

pop the top attention map from the stack, apply self-attention (Yang et al., 2018), and project it down to get the span’s start index and end index. To predict yes/no, we take the final memory (m in Eqn. 6) and project it down to a 2-way classifier. We concatenate the question vector qv and memory m and then project down to a 2-way classifier to decide whether to output a span or yes/no.

3.7 Optimization

Previous works on modular network (Hu et al., 2017, 2018) optimize the controller parameters θ jointly with the modules parameters ω on the training set. However, we found that our controller converges³ in less than 20 iterations under this training routine. It is also likely that this setup causes the model to overfit to the single-hop reasoning shortcuts in the training set. Hence, to prevent such premature convergence and reasoning-shortcut overfitting, we adopt the 2-phase training strategy that is widely used in Neural Architecture Search (Zoph and Le, 2016; Pham et al., 2018). The first phase updates ω for an entire epoch, followed by the second phase that updates θ over the entire dev set. We alternate between the two phases until the entire system converges.

4 Experimental Setup

Dataset and Metric We conduct our training and evaluation on the HotpotQA (Yang et al., 2018) dataset’s distractor setting, which has 90,447 training examples (72991 bridge-type questions and 17456 comparison-type questions) and 7,405 dev examples (5918 bridge-type questions and 1487 comparison-type questions). Since there is no public test set available, we split the original dev set into a dev set of 3,702 examples on which we tune our models, and a test set of 3,703 examples that is only used for the final test results. Each question is paired with 10 documents, among which 2 are supporting documents that contain the evidence necessary to infer the answer, and the other 8 are distracting documents. We evaluate our models based on the exact-match and F1 score between the prediction and ground-truth answer.

Data Augmentation We augment the comparison-type questions in the training set to promote the model’s robustness on this reasoning paradigm. We start with getting the

³Generating modular layout with probability near 1.

ORIGINAL question	answer	MUTATED question	answer
Were Pavel Urysohn and Levin known for the same type of work?	No	Were Henry Cavill and Li Na known for the same type of work?	Yes
Were Pavel Urysohn and Levin known for the same type of work?	No	Were Pavel Urysohn and Levin known for the different type of work?	Yes
Is Rohan Bopanna older than Sherwood Stewart?	Yes	Is Rohan Bopanna younger than Sherwood Stewart?	No
Was Howard Hawks a screenwriter of more productions than Arthur Berthelet?	Yes	Was Arthur Berthelet a screenwriter of more productions than Howard Hawks?	No

Table 1: Examples of mutated comparison-type questions and answers from HotpotQA training set.

	Dev		Test	
	EM	F1	EM	F1
BiDAF Baseline	44.68	57.19	42.7	55.81
NMN	31.04	40.28	30.87	39.90
Our NMN	50.67	63.35	49.58	62.71
+ Data aug.	50.63	63.29	49.46	62.59
- Bridge sup.	46.56	58.60	45.91	57.22
- Relocate	47.81	60.22	46.75	59.23
- Compare	50.29	63.30	48.87	62.52
- NoOp	49.11	61.79	48.56	62.10

Table 2: EM and F1 scores on HotpotQA dev set and test set. All models are tuned on dev set.

part-of-speech tags and NER tag of the question using CoreNlp (Manning et al., 2014). For each question whose answer is either “yes” or “no”, we generate a new question by randomly sampling two titles of the same type (based on POS and NER) from the training set to substitute the original entities in the question and corresponding supporting documents (1st row in Table 1). We then employ three strategies to generate 5,342 extra questions by mutating original questions. First, if the question contains the word “same” and the answer is yes or no, we substitute “same” with “different” and vice versa (2nd row in Table 1). Second, we detect the comparative and superlative adjectives/adverbs in the original question, transform them into their antonyms using wordnet, and then transform the antonyms back to their comparative form (3rd row in Table 1). Finally, if the question has a comparative adj./adv., we flip the order of the two entities compared (4th row in Table 1). In all three cases, the answer to the mutated question is also flipped.

Training Details We use 300-d GloVe pre-trained embeddings (Pennington et al., 2014). The model is supervised to predict either the start and end index of a span or “Yes/No” for specific questions. The entire model (controller + modular net-

work) is trained end-to-end using the Adam optimizer (Kingma and Ba, 2014) with an initial learning rate of 0.001.

5 Results

5.1 Baseline

Our baseline is the bi-attention (Seo et al., 2017) + self-attention model as in (Clark and Gardner, 2018; Yang et al., 2018), which was shown to be able to achieve strong performance on single-hop QA tasks like SQuAD (Rajpurkar et al., 2016) and TriviaQA (Joshi et al., 2017). Our baseline shares the preprocessing and encoding layer with the modular network.

5.2 Primary NMN Results

We first present our model’s performance on the HotpotQA (Yang et al., 2018) dev and test set of our split. As shown in the first three rows of Table 2, our modular network achieves significant improvements over both the baseline and the convolution-based NMN (Hu et al., 2018) on our test set. In Table 3, we further break down the dev-set performance on different question types,⁴ and our modular network obtains higher scores in both question types compared to the BiDAF baseline.

5.3 Ablation Studies

Data Augmentation: We also conduct ablation studies on our modeling choices. Augmenting the comparison-type questions in the training set boosts the performance on the comparison-type questions in the dev set (comparing row 3 and 4 in Table 3) without harming the scores on the bridge-type questions too much.

Bridge-entity Supervision: Supervising the bridge entity (described in 3.5) proves to be

⁴The 3 types of question other than “comparison” are all classified as “bridge-type” questions by Yang et al. (2018)

	Bridge		Compare	
	EM	F1	EM	F1
BiDAF Baseline	43.17	57.74	45.26	51.73
NMN	30.15	41.49	34.27	40.07
Our NMN	49.85	64.49	51.24	57.20
+ Data aug.	48.97	63.85	54.40	60.03
- Bridge sup.	45.59	59.45	48.82	51.73
- Relocate	46.11	60.13	51.92	58.10
- Compare	49.37	64.46	49.76	56.00
- NoOp	48.46	63.46	49.63	55.26

Table 3: EM and F1 scores on bridge-type and comparison-type questions from HotpotQA dev set.

Train	Reg		Adv	
	Reg	Adv	Reg	Adv
BiDAF Baseline	43.12	34.00	45.12	44.65
Our NMN	50.13	44.70	49.33	49.25

Table 4: EM scores after training on the regular data or on the adversarial data from Jiang and Bansal (2019), and evaluation on the regular dev set or the adv-dev set.

beneficial for the modular network to achieve good performance (row 5 in Table 2 and Table 3).

Modules: As shown in the 6th to 8th row, removing either the `Compare`, `Relocate` or `NoOp` module also causes drops in the metric scores. Specifically, removing the `Relocate` module causes significant degrade in bridge-type questions, which solidifies our claim that relocating the attention based on the inferred bridge entity is important for compositional reasoning. Similarly, removing the `Compare` module harms the model’s performance on comparison-type questions, suggesting the effectiveness of the module in addressing questions that require comparing two entities’ properties. These results demonstrate the contribution of each module toward achieving a self-assembling modular network with the strong overall performance.

5.4 Comparison with Original NMN Modules

One primary contribution of this work is that we adapt neural modular networks (NMN) (Andreas et al., 2016a; Hu et al., 2017, 2018), which were designed for visual-domain QA, to text-domain QA by rebuilding every reasoning module. We substitute convolution and multiplication between question vectors and context features with bi-attention as the basic reasoning component in the `Find` and `Relocate`. Moreover, our model maintains a stack of attention outputs before it is projected down to 1-d, thus enabling skip connections when predicting the answer span. As shown

in Table 2 and Table 3, our adapted modular network outperforms the original NMN significantly.

5.5 Adversarial Evaluation

Multiple previous works (Chen and Durrett, 2019; Min et al., 2019a) have shown that models performing strongly on HotpotQA are not necessarily capable of compositional reasoning. Jiang and Bansal (2019) proposed to construct adversarial distractor documents to eliminate the reasoning shortcut and necessitate compositional reasoning on HotpotQA dataset. To test whether our modular network can perform robust multi-hop reasoning against such adversaries, we evaluate our models on the adversarial dev set. The second column of Table 4 shows that our NMN outperforms the baseline significantly (+10 points in EM score) on the adversarial evaluation, suggesting that our NMN is indeed learning stronger compositional reasoning skills compared to the BiDAF baseline. We further train both models on the adversarial training set, and the results are shown in the last two columns of Table 4. We observe that after adversarial training, both the baseline and our NMN obtain significant improvement on the adversarial evaluation, while our NMN maintains a significant advantage over the BiDAF baseline.

6 Analysis

In this section, we present three analysis methods to show that our multi-hop NMN is highly interpretable and makes decisions similar to human actions. Combining the analyses below, we can understand how our model reasons to break the multi-step task into multiple single-step sub-tasks that are solvable by specific modules, and harnesses intermediate outputs to infer the answer.

6.1 Controller’s Attention on Questions

Different from single-hop questions, a multi-hop question usually encodes a sequence of sub-questions, among which the final sub-question requires an explicit answer. Therefore, the first step of solving a multi-hop question is to identify the hidden sub-questions and to sort them according to the correct reasoning order. Our controller handles this task by computing an attention distribution over all question words (see details in Sec. 3.2) at every reasoning step, which signifies which part of the question forms a sub-question that should be answered at the current step. Fig. 3 visualizes four

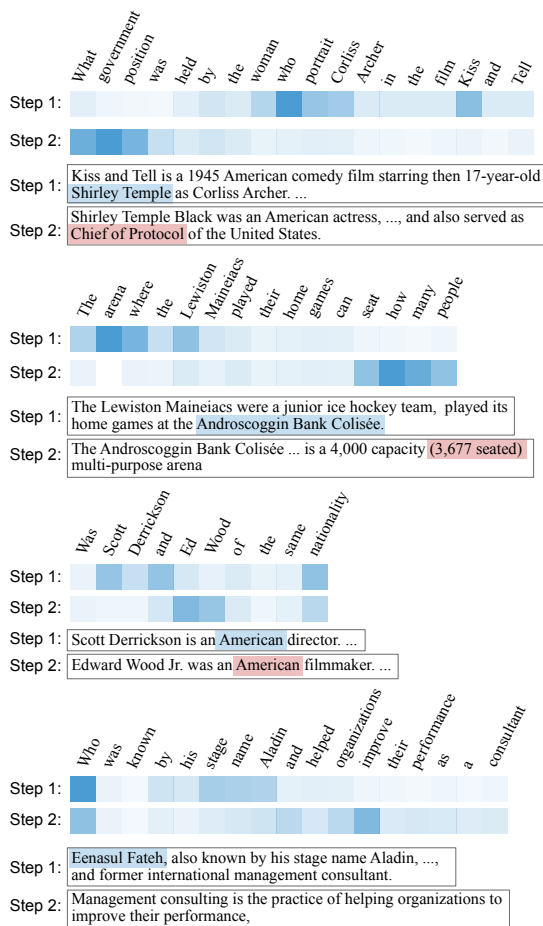


Figure 3: The controller’s multi-hop attention map on the question (randomly chosen from the first 10 dev examples). This attention is used to compute the sub-question representation that is passed to the modules.

attention maps from the controller when dealing with questions of different reasoning paradigms (randomly chosen from the first 10 dev examples), suggesting that our controller is able to accurately identify the sub-questions in the correct order. In the first example, the controller is able to attend to the first sub-question “who portrayed Corliss Archer in ...”, and then relocate the attention to the next sub-question “What government position was held by ...” to complete the 2-step reasoning.

6.2 Predicting Network Layouts

After the controller decides the sub-question at the current reasoning step, it then predicts a soft layout of the the current step by outputting a distribution cv over all available modules (Find, Relocate, Compare, NoOp). As elaborated in Sec. 3.3, the output vector and attention map of all modules are weighted by cv . To quantitatively evaluate the layout predicted by the controller, we label the expert layout for each ques-

tion in the dev set (Find-Relocate for bridge-type questions and Find-Find/Relocate-Compare for comparison-type questions). Next, we convert the soft module predictions into hard modular layouts by picking the module with the largest prediction probability at each step. Intuitively, we want the controller to assign higher weights to the module that can perform the desired reasoning. We then compute the percentage of dev set examples where the predicted layout matches the expert layout exactly. The empirical study shows that the layouts predicted by our NMN controller match the expert layout in 99.9% of bridge-type questions and 68.6% of comparison-type questions with Yes/No answer (examples shown in Fig. 1 and Fig. 2).

6.3 Finding Intermediate Bridge-Entities

The final action of an intermediate reasoning step is the predicted module reasoning through the question and context to answer the sub-question. For bridge-type questions (e.g., the first example in Fig. 3), the required action is to find the name (“Shirley Temple”) of “the woman who portrait Corliss Archer ...” to complete the next sub-question “What government position was held by Shirley Temple”. For comparison questions (e.g., the third example in Fig. 3), the model needs to infer the nationalities of “Scott” and “Ed” that will be compared in the following step.

7 Conclusion

In this work, we proposed a self-assembling neural modular network for multi-hop QA. We designed three modules that reason between the question and text-based context. The resulting model outperforms both the single-hop baseline and the original NMN on HotpotQA (Yang et al., 2018). Because of the interpretable nature of our model, we presented analyses to show that our model does in fact learn to perform compositional reasoning and can dynamically assemble the modular network based on the question.

Acknowledgments

We thank the reviewers for their helpful comments. This work was supported by ARO-YIP Award #W911NF-18-1-0336, DARPA #YFA17-D17AP00022 and awards from Google, Facebook, Salesforce, and Adobe (plus Amazon and Google GPU cloud credits). The views are those of the authors and not of the funding agency.

References

- Jacob Andreas, Marcus Rohrbach, Trevor Darrell, and Dan Klein. 2016a. Learning to compose neural networks for question answering. *arXiv preprint arXiv:1601.01705*.
- Jacob Andreas, Marcus Rohrbach, Trevor Darrell, and Dan Klein. 2016b. Neural module networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 39–48.
- Jifan Chen and Greg Durrett. 2019. Understanding dataset design choices for multi-hop reasoning. In *NAACL*.
- Christopher Clark and Matt Gardner. 2018. Simple and effective multi-paragraph reading comprehension. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*.
- Karl Moritz Hermann, Tomas Kocisky, Edward Grefenstette, Lasse Espeholt, Will Kay, Mustafa Suleyman, and Phil Blunsom. 2015. Teaching machines to read and comprehend. In *Advances in Neural Information Processing Systems*, pages 1693–1701.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation*, 9(8):1735–1780.
- Ronghang Hu, Jacob Andreas, Trevor Darrell, and Kate Saenko. 2018. Explainable neural computation via stack neural module networks. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 53–69.
- Ronghang Hu, Jacob Andreas, Marcus Rohrbach, Trevor Darrell, and Kate Saenko. 2017. Learning to reason: End-to-end module networks for visual question answering. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 804–813.
- Yichen Jiang and Mohit Bansal. 2019. Avoiding reasoning shortcuts: Adversarial evaluation, training, and model development for multi-hop qa. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*.
- Justin Johnson, Bharath Hariharan, Laurens van der Maaten, Judy Hoffman, Li Fei-Fei, C Lawrence Zitnick, and Ross Girshick. 2017. Inferring and executing programs for visual reasoning. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2989–2998.
- Mandar Joshi, Eunsol Choi, Daniel S Weld, and Luke Zettlemoyer. 2017. Triviaqa: A large scale distantly supervised challenge dataset for reading comprehension. *arXiv preprint arXiv:1705.03551*.
- Diederik P. Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *CoRR*.
- Hanxiao Liu, Karen Simonyan, and Yiming Yang. 2018. Darts: Differentiable architecture search. *arXiv preprint arXiv:1806.09055*.
- Christopher Manning, Mihai Surdeanu, John Bauer, Jenny Finkel, Steven Bethard, and David McClosky. 2014. The stanford corenlp natural language processing toolkit. In *Proceedings of 52nd annual meeting of the association for computational linguistics: system demonstrations*, pages 55–60.
- David Mascharka, Philip Tran, Ryan Soklaski, and Arjun Majumdar. 2018. Transparency by design: Closing the gap between performance and interpretability in visual reasoning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4942–4950.
- Sewon Min, Eric Wallace, Sameer Singh, Matt Gardner, Hannaneh Hajishirzi, and Luke Zettlemoyer. 2019a. Compositional questions do not necessitate multi-hop reasoning. In *ACL*.
- Sewon Min, Victor Zhong, Luke Zettlemoyer, and Hannaneh Hajishirzi. 2019b. Multi-hop reading comprehension through question decomposition and rescoring. In *ACL*.
- Jeffrey Pennington, Richard Socher, and Christopher D Manning. 2014. Glove: Global vectors for word representation. In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- Hieu Pham, Melody Y Guan, Barret Zoph, Quoc V Le, and Jeff Dean. 2018. Efficient neural architecture search via parameter sharing. *arXiv preprint arXiv:1802.03268*.
- P. Rajpurkar, R. Jia, and P. Liang. 2018. Know what you don’t know: Unanswerable questions for SQuAD. In *Association for Computational Linguistics (ACL)*.
- Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. 2016. Squad: 100,000+ questions for machine comprehension of text. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- Minjoon Seo, Aniruddha Kembhavi, Ali Farhadi, and Hannaneh Hajishirzi. 2017. Bidirectional attention flow for machine comprehension. In *International Conference on Learning Representations (ICLR)*.
- Rupesh Kumar Srivastava, Klaus Greff, and Jürgen Schmidhuber. 2015. Highway networks. In *International Conference on Machine Learning (ICML)*.
- Johannes Welbl, Pontus Stenetorp, and Sebastian Riedel. 2017. Constructing datasets for multi-hop reading comprehension across documents. In *TACL*.
- Jason Weston, Antoine Bordes, Sumit Chopra, Alexander M Rush, Bart van Merriënboer, Armand Joulin, and Tomas Mikolov. 2016. Towards ai-complete question answering: A set of prerequisite toy tasks. In *ICLR*.

- Ronald J Williams. 1992. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning*, 8(3-4):229–256.
- Caiming Xiong, Victor Zhong, and Richard Socher. 2017. Dynamic coattention networks for question answering. In *ICLR*.
- Yi Yang, Wen-tau Yih, and Christopher Meek. 2015. Wikiqa: A challenge dataset for open-domain question answering. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 2013–2018.
- Zhilin Yang, Peng Qi, Saizheng Zhang, Yoshua Bengio, William W Cohen, Ruslan Salakhutdinov, and Christopher D Manning. 2018. Hotpotqa: A dataset for diverse, explainable multi-hop question answering. In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- Victor Zhong, Caiming Xiong, Nitish Shirish Keskar, and Richard Socher. 2019. Coarse-grain fine-grain coattention network for multi-evidence question answering. In *ICLR*.
- Barret Zoph and Quoc V Le. 2016. Neural architecture search with reinforcement learning. *arXiv preprint arXiv:1611.01578*.