

Stratified Rule-Aware Network for Abstract Visual Reasoning

Sheng Hu,^{1,*} Yuqing Ma,^{1,*} Xianglong Liu,^{1,2,†} Yanlu Wei,¹ Shihao Bai¹

¹ State Key Laboratory of Software Development Environment, Beihang University, Beijing, China

² Beijing Advanced Innovation Center for Big Data-Based Precision Medicine, Beihang University, Beijing, China
husheng_7@163.com, {mayuqing,xlliu}@nlsde.buaa.edu.cn, {weiyalu,16061167}@buaa.edu.cn

Abstract

Abstract reasoning refers to the ability to analyze information, discover rules at an intangible level, and solve problems in innovative ways. Raven’s Progressive Matrices (RPM) test is typically used to examine the capability of abstract reasoning. The subject is asked to identify the correct choice from the answer set to fill the missing panel at the bottom right of RPM (e.g., a 3×3 matrix), following the underlying rules inside the matrix. Recent studies, taking advantage of Convolutional Neural Networks (CNNs), have achieved encouraging progress to accomplish the RPM test. However, they partly ignore necessary inductive biases of RPM solver, such as order sensitivity within each row/column and incremental rule induction. To address this problem, in this paper we propose a Stratified Rule-Aware Network (SRAN) to generate the rule embeddings for two input sequences. Our SRAN learns multiple granularity rule embeddings at different levels, and incrementally integrates the stratified embedding flows through a gated fusion module. With the help of embeddings, a rule similarity metric is applied to guarantee that SRAN can not only be trained using a triplet loss but also infer the best answer efficiently. We further point out the severe defects existing in the popular RAVEN dataset for RPM test, which prevent from the fair evaluation of the abstract reasoning ability. To fix the defects, we propose an answer set generation algorithm called Attribute Bisection Tree (ABT), forming an improved dataset named Impartial-RAVEN (I-RAVEN for short). Extensive experiments are conducted on both PGM and I-RAVEN datasets, showing that our SRAN outperforms the state-of-the-art models by a considerable margin.

Introduction

Abstract reasoning, also known as inductive reasoning, refers to the ability to analyze information, discover rules at an intangible level, and solve problems in innovative ways. This type of reasoning, as the foundation for human intelligence, helps human understand the world. It has been generally regarded and pursued as a critical component to the development of artificial intelligence during the past decades, and has attracted increasing attention in recent years. Raven’s Progressive Matrices (RPM) test (Raven

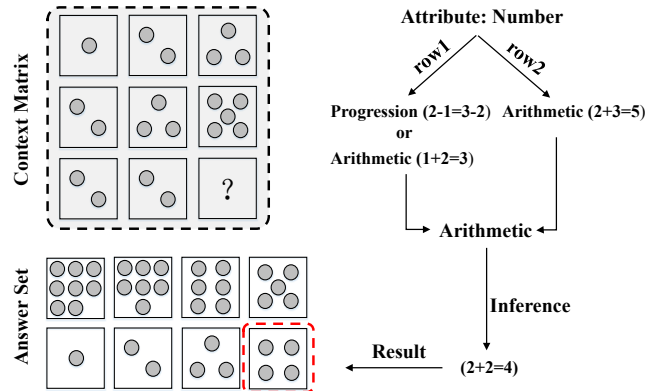


Figure 1: An example of RPM question and the human strategy to solve it. The underlying rule on the number of circles could be *Progression (2-1=3-2)* or *Arithmetic (1+2=3)* along row 1, and *Arithmetic (2+3=5)* along row 2. Therefore the dominant rule is *Arithmetic*. Apply it to the third row to figure out the answer ($2+2=4$). Besides, no viable rule can be found along the columns

1938; Carpenter, Just, and Shell 1990; Raven 2000; Kunda, McGregor, and Goel 2013; Strannegård, Cirillo, and Ström 2013) is one of the highly accepted and well-studied tools to examine the ability of abstract reasoning, which is believed as a good estimate of the real intelligence (Carpenter, Just, and Shell 1990). An illustration of RPM is shown in Figure 1, where usually the test-taker is presented with a 3×3 matrix with the bottom right panel left blank. The goal is to choose one image from an answer set of eight candidates to complete the matrix correctly, namely satisfying the underlying rules in the matrix. Subjects accomplish this by looking into the first two rows/columns and inducing the dominant rules which govern the attributes in those panels. The obtained rules can then be applied to the last row/column to figure out which answer belongs to the blank panel.

Computational models for RPM in the cognitive science community access symbolic representations of the images (Carpenter, Just, and Shell 1990; Lovett and Forbus 2017; Lovett, Forbus, and Usher 2010; Lovett et al. 2010). Recently there has been some success with end-to-end learning methods trying to accomplish abstract reasoning on

*Equal contribution

†Corresponding author

RPM test (Hoshen and Werman 2017; Barrett et al. 2018; Steenbrugge et al. 2018; Zhang et al. 2019a,b; Zheng, Zha, and Wei 2019; van Steenkiste et al. 2019; Wang, Jamnik, and Lio 2020), inspired by the progress of computer vision tasks (Krizhevsky, Sutskever, and Hinton 2012; Simonyan and Zisserman 2015; Szegedy et al. 2015; He et al. 2016) and boosted by the large-scale PGM (Barrett et al. 2018) and RAVEN (Zhang et al. 2019a) datasets. Typical works including CoPINet (Zhang et al. 2019a), LEN (Zheng, Zha, and Wei 2019), and MXGNet (Wang, Jamnik, and Lio 2020) followed the paradigm that predicts a classification score for each multiple-choice panel based on the relations inside each row/column, showing great potential to solve RPM test. However, these models partly ignore the important characteristics for RPM, such as the permutation invariance (Zhang et al. 2019b), the order sensitivity of panels inside a row/column, etc. Previous work (Wang, Jamnik, and Lio 2020) specially mentions that they do not choose a permutation-invariant structure because it leads to severe ‘overfitting’ on the RAVEN dataset. We will discuss this phenomenon in later sections. What is even worse, directly extracting the relations, without considering the incremental rule induction mechanism widely adopted in human cognitive systems (Carpenter, Just, and Shell 1990), inevitably leads to inferior performance.

To achieve reliable and efficient abstract reasoning, in this paper we develop a powerful architecture called Stratified Rule-Aware Network (SRAN) that naturally integrates the indispensable inductive biases, including order sensitivity, permutation invariance, and incremental rule induction. SRAN takes two rows/columns as input and learns stratified rule embeddings at different levels, namely cell-wise, individual-wise, and ecological hierarchy. These multiple granularity embeddings are incrementally integrated via a gate fusion module, which naturally preserves the order sensitivity of panels and maps the inputs to a rule embedding space. With the help of the embeddings, we further introduce a rule similarity metric, based on which SRAN can not only be well trained using a triplet loss but also infer the best answer efficiently. This framework resembles the human strategy for RPM shown in Figure 1.

To fairly evaluate the abstract reasoning ability, we also design a general algorithm named Attribute Bisection Tree (ABT) to generate an impartial answer set for any attribute-based RPM question. We point out and further fix the underlying defects of the commonly-used RAVEN (Zhang et al. 2019a) dataset, where the correct answer could be inferred even without the presence of the context matrix. Therefore, we introduce an improved dataset named Impartial-RAVEN (I-RAVEN) to fairly evaluate the abstract reasoning capability of RPM solvers.

To the best of our knowledge, the proposed SRAN is the first RPM solver to induce rule embeddings which are discriminative and measurable. We are also the first to point out and fix the defects of the misleading benchmark RAVEN, and generate an impartial dataset I-RAVEN based on our ABT algorithm. Extensive experiments conducted on widely used dataset PGM and our improved I-RAVEN show that SRAN outperforms state-of-the-art methods by a consider-

able margin, e.g. 60.8% accuracy compared to the second best 46.1% on I-RAVEN.

Our Approach

In this section, we first give a formal definition of the abstract reasoning task on the RPM test. Then we introduce the inductive-biased framework, and present our Stratified Rule-Aware Network (SRAN). Finally, we demonstrate the learning and inference process of the proposed model.

Preliminary

For a common RPM question, usually a 3×3 matrix M^- is given, with bottom right context panel left blank. Ω denotes the answer set with N multiple-choice panels, where typically $N=8$. The dominant rules governing the features inside the matrix could be induced from the first two intact rows/columns. The goal is to select a multiple-choice panel $\omega \in \Omega$ to complete the context matrix M^- , maintaining the dominant rule inside of the context matrix.

We define the completed matrix with a multiple-choice panel ω infilled as M , where M_i is denoted as the i -th row, and $m_{i,j}$ indicates the panel in i -th row and j -th column. Intuitively, M is almost the same as M^- , except for $m_{33} = \omega$ while the corresponding element missing in M^- . In fact, whether rules exist in rows or columns is uncertain. Therefore, our framework induces both the row-wise rule representation and the column-wise representation in the same way. In order to simplify the notation, we only take the induction of the row-wise rule representation as example.

The Reasoning Framework

Based on the necessary inductive biases for RPM, we develop a novel abstract reasoning architecture named Stratified Rule-Aware Network (SRAN). Given two input rows M_i, M_j , the proposed framework embeds the input into multiple granularity embeddings using a stratified rule embedding module \mathbb{E} . Named after biological organizations (Parent 1996), \mathbb{E} consists of three hierarchies including cell-wise network \mathbb{E}_{cell} , individual-wise network \mathbb{E}_{ind} , and ecological network \mathbb{E}_{eco} . With the multiple granularity rule embeddings, the gated embedding fusion module \mathbb{G} will incrementally integrate these stratified embedding flows and map the two input sequences M_i and M_j to a discriminative rule embedding $\mathbf{r}_{ij}^{(3)}$, while maintaining the order sensitivity and permutation invariance. We further introduce a rule similarity metric \mathcal{D} to estimate the similarity between the rule representations. The correct answer can be predicted by choosing the multiple-choice panel within the shortest distance to the dominant rule generated by the first two rows in the matrix.

Stratified Rule Embedding

As we all know, organization of behaviour into a nested hierarchy of tasks is characteristic of purposive cognition in humans. The prevalent Convolution Neural Network inspired by the human visual system, is a stratified model itself, with the projection from each layer showing the hierarchical nature of features. The bottom layers extract low-level features,

such as texture, edge, etc, while the top layers abstract high-level semantic information from the low-level information transmitted from the bottom layers.

However, without specifying information from different levels, it is hard for CNN to figure out different hierarchies, and thus fail to obtain robust and discriminative representations. Therefore, it would be better to feed the input of different hierarchies explicitly and extract rule representations from different granularity with artificial guidance. Motivated by that, we deploy a stratified rule embedding module, consisting of cell-wise hierarchy, individual-wise hierarchy, and ecological hierarchy.

Cell-wise Hierarchy The network of the cell-wise hierarchy \mathbb{E}_{cell} takes each panel as input and recognize the attributes of inside graphical elements. It handles each panel independently without considering the difference or correlations among panels inside the matrix. Therefore, it observes the information from the most detailed perspective. We obtain the cell-wise rule representation for each input panel:

$$\mathbf{x}_{ij} = \mathbb{E}_{\text{cell}}(\mathbf{m}_{ij}). \quad (1)$$

Individual-wise Hierarchy Moreover, the network of individual hierarchy takes each row as input. It begins to take the correlations among panels of the same row into consideration, and encode the entire row with a compact embedding, rather than simply combining each panel. In this way, the rule embedding process for each panel is coupled and interacts with each other. Intuitively, each row may contain multiple plausible rules. In this hierarchy, the framework extracts intermediate rule embedding for each row individually, which still ignores the comprehensive information from the matrix perspective, especially the correlations across rows. The individual-wise rule embedding \mathbf{y}_i is denoted as:

$$\mathbf{y}_i = \mathbb{E}_{\text{ind}}(\mathbf{M}_i). \quad (2)$$

Ecological Hierarchy Furthermore, the network of the ecological hierarchy takes the two rows together as input and jointly learns the rule patterns underlying the two rows. As we mentioned before, in the individual hierarchy, the framework extracts intermediate rule embedding for each row, without considering the interaction between two rows. The rule that exists in one row may not lie in another. Therefore, to obtain the shared rule patterns between the two rows, it is essential to put these two rows together and jointly learn the features from an ecological level. Thus the shared rule embedding is obtained as follows:

$$\mathbf{z}_{ij} = \mathbb{E}_{\text{eco}}([\mathbf{M}_i, \mathbf{M}_j]), \quad (3)$$

where $[\cdot, \cdot]$ denotes the concatenating operation.

Gated Embedding Fusion

Since the rule embeddings at different levels focus on different attributes or patterns, to generate one discriminative representation for the rule, we should aggregate the multiple granularity embeddings. Due to the requirement that the aggregation should preserve the order of cell-wise rule embeddings and be permutation-invariant to the individual-wise ones, we propose a stratified rule embedding learning

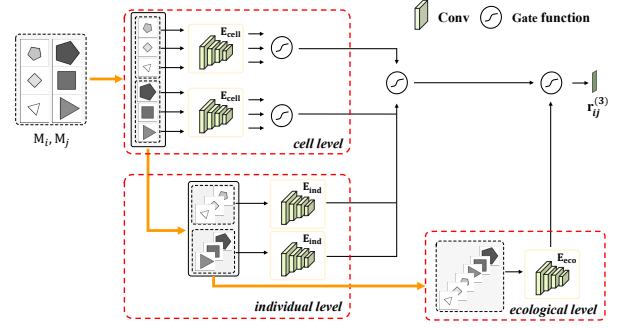


Figure 2: The architecture of SRAN, consisting of a hierarchical rule embedding module and a gated embedding fusion module. Given two row sequences as input, it outputs the rule embedding

method named gated embedding fusion module, which is responsible for gradually aggregating the multiple granularity embeddings.

Specifically, we define a gate function φ to fuse the rule embeddings from different hierarchies. It concatenates all the inputs and encodes into a single embedding using fully connected layers. The gate function is similar to the attention mechanism, which detects and concentrates on the useful features according to the task. Even for the same attribute, they may focus on different facets. Based on the gate function, our gated embedding fusion module could regulate the flow of rule embeddings into the framework and make the utmost of their complementary information.

At the cell level, after obtaining cell-wise rule embeddings for panels in i -th row \mathbf{M}_i , the module aggregates them to infer a row-wise rule embedding $\mathbf{r}_i^{(1)}$:

$$\mathbf{r}_i^{(1)} = \varphi_1(\mathbf{x}_{i1}, \mathbf{x}_{i2}, \mathbf{x}_{i3}), \quad (4)$$

Similarly, we obtain $\mathbf{r}_j^{(1)}$ for the j -th row \mathbf{M}_j . The fused embedding integrates different types of attributes in the panels.

At the individual level, intuitively both $\mathbf{r}_i^{(1)}$ and \mathbf{y}_i are the row-wise embeddings corresponding to the i -th row, but convey the different granularity rule information. We further fuse them, and jointly mine the shared rules contained in the i -th and j -th row :

$$\mathbf{r}_{ij}^{(2)} = \varphi_2(\mathbf{r}_i^{(1)}, \mathbf{y}_i, \mathbf{r}_j^{(1)}, \mathbf{y}_j). \quad (5)$$

At the ecological level, similarly we can further combine fused embedding $\mathbf{r}_{ij}^{(2)}$ and \mathbf{z}_{ij} using the gate fusion function, abstracting the final rule embedding:

$$\mathbf{r}_{ij}^{(3)} = \varphi_3(\mathbf{r}_{ij}^{(2)}, \mathbf{z}_{ij}). \quad (6)$$

To make sure the framework is permutation-invariant to the input rows, we exchange the concatenation order of the two input rows and average the output rule embeddings. This invariance ensures that, the rule embedding respects the characteristic of RPM and thus distills the representative information of the relations existing in the inputs.

On the whole, the SRAN can be formulated in its simplest form as follows:

$$\begin{aligned} \mathbf{r}_{ij}^{(3)} &= \text{SRAN}(\mathbf{M}_i, \mathbf{M}_j) \\ &= \mathbb{G}(\mathbf{x}_i, \mathbf{x}_j, \mathbf{y}_i, \mathbf{y}_j, \mathbf{z}_{ij}), \end{aligned} \quad (7)$$

where $\mathbf{r}_{ij}^{(3)}$ is the shared rule embedding of the \mathbf{M}_i and \mathbf{M}_j . An illustration of SRAN is shown in Figure 2.

Learning and Inference

With SRAN framework, the question turns to how we train the network, and apply it to infer the correct answer to RPM test. The key to address the question lies in the similarity measure between two rule embeddings, based on which we can define the loss function for SRAN training, and meanwhile determine the best choice during inference.

Similarity function We introduce similarity function \mathcal{D} to measure the closeness between two rules in the embedding space. In this paper, we adopt inner product similarity for good experimental results:

$$\mathcal{D}(\mathbf{r}, \mathbf{r}') = \mathbf{r}^T \mathbf{r}'. \quad (8)$$

Training For a given RPM question, the first two rows $\mathbf{M}_1, \mathbf{M}_2$ are fed into our proposed SRAN and produce the shared rule embedding \mathbf{g} :

$$\mathbf{g} = \mathbf{r}_{12}^{(3)} = \text{SRAN}(\mathbf{M}_1, \mathbf{M}_2), \quad (9)$$

which represents the dominant pattern of the matrix.

Intuitively, the rule extracted from the first two rows can be treated as the reference rule, and we name it the dominant rule in the matrix. Subsequently, the correct answer can be found by checking whether its corresponding rule embedding is similar to the dominant rule. Specifically, given a multiple-choice panel $\omega_k \in \Omega$, where $k \in \{1, \dots, N\}$, we denote $\bar{\mathbf{r}}_k$ as the new rule embedding inside \mathbf{M} caused by the k -th multiple-choice panel:

$$\bar{\mathbf{r}}_k = \frac{1}{2} \left(\mathbf{r}_{13}^{(3)} + \mathbf{r}_{23}^{(3)} \right). \quad (10)$$

This procedure is illustrated in Figure 3. In practice, we generate the column-wise rule representation just as the row-wise one, and concatenate the two representations together as the final representation.

For the rule embedding $\bar{\mathbf{r}}^*$ generated by rows/columns infilled with correct answer, the desirable SRAN should enforce it to be more similar to the dominant rule \mathbf{g} , compared to the other rules $\bar{\mathbf{r}}_k$ corresponding to the wrong answers, where $\bar{\mathbf{r}}_k \neq \bar{\mathbf{r}}^*$. Subsequently, the generated rules of N candidates, alongside with the dominant rule, form a tuple containing $N+1$ elements. Based on the similarity function, the $(N+1)$ -tuple loss (Sohn 2016) can be defined for SRAN training:

$$\mathcal{L} = \log \left(1 + \sum_{k=1, \bar{\mathbf{r}}_k \neq \bar{\mathbf{r}}^*}^N \exp(\mathcal{D}(\mathbf{g}, \bar{\mathbf{r}}_k) - \mathcal{D}(\mathbf{g}, \bar{\mathbf{r}}^*)) \right), \quad (11)$$

which means the SRAN can be trained in a fully end-to-end manner. The architecture of the SRAN (Figure 2) is well matched to the problem of abstract reasoning, because it leverages human strategies and explicitly generates the rules governing the matrix.




Figure 3: The similarity score for a candidate answer. A multiple-choice panel from the answer set is infilled in the blank panel (row 3), generating a rule embedding $\bar{\mathbf{r}}_k$ through SRAN. The similarity score for the candidate answer can be estimated based on $\bar{\mathbf{r}}_k$ and the dominant rule embedding \mathbf{g} extracted from row 1 and 2

Inference Once the training of SRAN is finished, we could make the inference of the newly given RPM question. Initially, the intact rows/columns of the RPM are fed into the framework to get the dominant rule \mathbf{g} . After that, each multiple-choice panel is filled to the blank position to complete the matrix, and the framework will generate the rule embeddings $\bar{\mathbf{r}}_k$ for all candidate answers, given the current completed matrix. We can accomplish the abstract reasoning by choosing the correct multiple-choice as follows:

$$k^* = \arg \max_k \mathcal{D}(\mathbf{g}, \bar{\mathbf{r}}_k). \quad (12)$$

Note that since we investigate each panel independently, the above inference process promises that our model's output stays invariant if the answer set is shuffled.

Attribute Bisection Tree for Impartial Dataset

RAVEN (Zhang et al. 2019a) is a popular RPM-style dataset adopted by all recent studies (Zhang et al. 2019b; Zheng, Zha, and Wei 2019; Wang, Jamnik, and Lio 2020). However, we find severe defects in its answer sets, making RAVEN incompetence as a measurement of abstract reasoning. In this section, we first give a brief review of RAVEN, and then explain the defects with analysis and experiments. Finally we introduce a general algorithm to generate an impartial answer set for any attribute-based RPM question. Thus we fix the defects of RAVEN and propose an improved dataset.

Defects of RAVEN

RAVEN dataset consists of 70,000 RPM questions distributed in 7 different figure configurations. Panels are constructed with 5 attributes (Number, Position, Type, Size, Color). Each attribute is governed by one of 4 rules and takes a value from a predefined set. Rules are applied only row-wise in RAVEN.

After carefully examining the data in RAVEN, we find unexpected pattern among the eight multiple-choice panels. Each distractor in the answer set is generated by randomly modifying one attribute of the correct answer (see Figure 4(a)). As a consequence, the panel with the most common values for each attribute will be the correct answer. This means the correct answer can be found by simply scanning

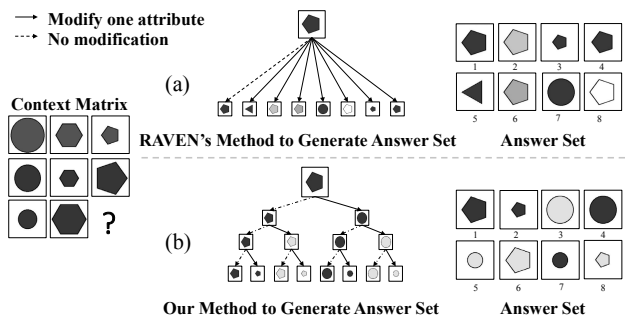


Figure 4: Comparison between RAVEN and I-RAVEN

Model	RAVEN	I-RAVEN
ResNet (Zhang et al. 2019a)	53.4	40.3
CoPINet (Zhang et al. 2019b)	91.4	46.1
Context-blind ResNet	71.9	12.2
Context-blind CoPINet	94.2	14.2

Table 1: Test on RAVEN and I-RAVEN

the answer set without considering the context images. An example is also shown on the right of Figure 4(a). Among the answer set, the most common `Color` and `Type` are black (1, 3, 4, 5, and 7) and pentagon (1, 2, 3, 4, 6, and 8). Besides, multiple-choice panel 1, 2, 5, 6, 7, and 8 are in the same `Size`. Therefore, multiple-choice panel 1, which is the panel with the most common attribute values, is inferred as (and indeed is) the correct answer, even without considering the context matrix.

Note that understanding of the context matrix is the cornerstone of RPM test. The RAVEN dataset, where the correct answer can be found without the context, is obviously against the essence of abstract reasoning, and thus is incapable of evaluating abstract reasoning ability.

More severely, such underlying patterns can also be captured by neural networks, especially for models which combine features of eight multiple-choice panels. We train two models with context-blind (Barrett et al. 2018) setting, including a simple ResNet-based classifier (Zhang et al. 2019a) and the competitive CoPINet (Zhang et al. 2019b). These context-blind models are trained with only eight multiple-choice panels as input, and should have predicted the answer randomly, if the dataset is logical. However, as shown in Table 1, the context-blind models can achieve even better performance to the normal models, which proves that RAVEN contains illogical patterns where the correct answer can be found when only presented with the answer set. This type of back-door solutions is quite hidden and has also been discussed in other relational reasoning benchmarks, such as Visual Question Answering (VQA) (Johnson et al. 2017) and symbolic analogy (Hill et al. 2019). We can conclude that the 'overfitting' phenomenon on RAVEN reported by (Wang, Jamnik, and Lio 2020) is not caused by a permutation-invariant structure but by the dataset itself.

Algorithm 1 Attribute Bisection Tree

Input: the correct answer ω^*

- 1: Initialize the answer set $\Omega = \{\omega^*\}$
- 2: Sample 3 attributes a_1, a_2, a_3 according to ω^*
- 3: Sample new value v_i for each a_i
- 4: **for** $i = 1$ to 3 **do**
- 5: Initialize $\Gamma = \{\}$
- 6: **for** each w_k in the current answer set Ω **do**
- 7: $\gamma \leftarrow$ modifying attribute a_i of w_k with v_i
- 8: $\Gamma \leftarrow \Gamma \cup \{\gamma\}$
- 9: **end for**
- 10: $\Omega \leftarrow \Omega \cup \Gamma$
- 11: **end for**

Output: the answer set Ω ($|\Omega| = 2^3 = 8$)

Attribute Bisection Tree

We design a general algorithm named Attribute Bisection Tree (ABT) to generate an impartial answer set for any attribute-based RPM question. The ABT ensures attribute modifications among the answer set are well balanced. Thus, no clue can be found to guess the correct answer only depending on the answer set, and no distractor can be eliminated without reasoning from the context matrix as well.

Figure 4(b) demonstrates the generation process using a tree structure. Each node indicates a multiple-choice panel, and the root of the tree structure is the correct answer. Different levels indicate different iterations, where nodes of this level are the candidate answers of current answer set. The generation process flows in a top-down manner. For each iteration, only one attribute will be modified. At each level, a node has two children nodes, where one node remains the same with the father node, the other changes the value of the attribute sampled for this iteration of the father node. Finally, at the bottom level, we could obtain the whole answer set. Algorithm 1 summarizes the key steps of the answer generation process.

I-RAVEN

With ABT, we generate an alternative answer set for each RPM question in the RAVEN dataset, forming an improved dataset named Impartial-RAVEN (I-RAVEN). Next, we will show that compared with RAVEN, I-RAVEN is more rigorous and fair for evaluating abstract reasoning capability.

Taking Figure 4(b) for example, each attribute has two different values which distribute evenly in the answer set. The attribute `Color` of half answer candidates (1, 2, 4, and 7) are black, while the other half (3, 5, 6, and 8) are light grey. Similarly, the attribute `Type` of half answer candidates (1, 2, 6, and 8) are pentagon, while the other half (3, 4, 5, and 7) are circle. Half of the answer set (1, 3, 4, and 6) are in the same size, which are different from the other same-sized half (2, 5, 7, and 8). As a result, there is no candidate with the most common values for each attribute. In other words, the back-door solution on RAVEN can no longer be applied to the new answer set.

To better explain the superiority of I-RAVEN over RAVEN, as shown in Figure 5, we use undirected graphs

Model	Acc	Center	2×2G	3×3G	O-IC	O-IG	L-R	U-D
LSTM (Zhang et al. 2019a)	18.9	26.2	16.7	15.1	21.9	21.1	14.6	16.5
WReN (Barrett et al. 2018)	23.8	29.4	26.8	23.5	22.5	21.5	21.9	21.4
ResNet (Zhang et al. 2019a)	40.3	44.7	29.3	27.9	46.2	35.8	51.2	47.4
ResNet+DRT (Zhang et al. 2019a)	40.4	46.5	28.8	27.3	46.0	34.2	50.1	49.8
LEN (Zheng, Zha, and Wei 2019)	41.4	56.4	31.7	29.7	52.1	31.7	44.2	44.2
Wild ResNet (Barrett et al. 2018)	44.3	50.9	33.1	30.8	50.9	38.7	53.1	52.6
CoPINet (Zhang et al. 2019b)	46.1	54.4	36.8	31.9	52.2	42.8	51.9	52.5
SRAN (Ours)	60.8	78.2	50.1	42.4	68.2	46.3	70.1	70.3

Table 2: Test accuracy of different models on I-RAVEN. Acc denotes the mean accuracy, while other columns show accuracy across seven figure configurations. 2×2G, 3×3G, O-IC, O-IG, L-R, and U-D denote 2×2Grid, 3×3Grid, Out-InCenter, Out-InGrid, Left-Right, and Up-Down, respectively

Model	LSTM	ResNet	Wild ResNet	CoPINet	WReN	MXGNet	LEN	SRAN
Acc	35.8	42.0	48.0	56.4	62.6	66.7	68.1	71.3

Table 3: Test accuracy of different models on PGM




Figure 5: Characterizing answer sets of RAVEN and I-RAVEN using graphs

to characterize typical answer sets of the two datasets respectively, where each candidate answer is represented by a node with its degree inflated. An edge between two nodes represents that the corresponding candidates differ in one attribute. In Figure 5(a), there is always a central node with a degree of 7 and the other nodes with less degrees. The back-door solution is to find the central node, which is indeed the correct answer. By contrast, in Figure 5(b), due to balanced modifications of attributes, each node always has the same degree of 3, which is indistinguishable from each other without the context matrix. Moreover, inspired by (Hill et al. 2019), we make the noise attribute `Uniformity` of each distractor stay consistent with the correct answer, so that each distractor is perceptually plausible and cannot be eliminated simply by attribute mismatching. This setting encourages models to reason from the context.

We also train ResNet and CoPINet on I-RAVEN alongside with their context-blind versions to verify the fairness of the proposed dataset. The right column in Table 1 lists the results which are in stark contrast with those on the original RAVEN dataset. The performance of context-blind models is almost at a random guess level (12.5%), while the normal models relying on the context can perform much better.

Experiments

Experimental Setup

With I-RAVEN, we first compare our method with several state-of-the-art models using public implementations, including LSTM (Hochreiter and Schmidhuber 1997), ResNet-based (He et al. 2016) image classifier (ResNet), ResNet with DRT (Zhang et al. 2019a), Wild ResNet (Barrett et al. 2018), WReN (Barrett et al. 2018), CoPINet (Zhang et al. 2019b), and LEN (Zheng, Zha, and Wei 2019). We adopt the public implementations of LSTM, ResNet, and DRT in (Zhang et al. 2019a). Other variants of LEN are not included in this section because they require additional training labels.

PGM (Barrett et al. 2018) is another RPM dataset consisting of 1.42M questions. Rules in a matrix are composed with 1 to 4 relation-object-attribute tuples and can be applied along the rows or columns. SRAN is compared with results on PGM reported in (Barrett et al. 2018; Zhang et al. 2019b; Zheng, Zha, and Wei 2019; Wang, Jamnik, and Lio 2020).

For our SRAN, we adopt three ResNet-18 (He et al. 2016) as the embedding networks for the three hierarchies, by modifying the input channels. The gate fusion φ_1 and φ_2 are 2-layer fully connected networks, while φ_3 is a 4-layer fully connected network with dropout (Srivastava et al. 2014) of 0.5 applied on the last layer. We adopt stochastic gradient descent using ADAM (Kingma and Ba 2014) optimizer. The exponential decay rate parameters are $\beta_1 = 0.9$, $\beta_2 = 0.999$, $\epsilon = 10^{-8}$. Each reported accuracy is averaged over 5 runs.

Comparisons with State-of-the-art Methods

Table 2 and Table 3 list the test accuracy of different models trained on I-RAVEN and PGM, respectively. From the table, it is obvious that our proposed SRAN outperforms other methods by a considerable margin. Besides, we observe that models benefit from inductive biases, such as the competitive CoPINet, LEN, Wild ResNet, and our SRAN. Such in-

ductive biases (even partly) can encourage models to explore the underlying rules. For more detailed comparison, Table 2 also reports the accuracy on seven figure configurations of I-RAVEN. We can observe that accuracy on different configurations is not uniform, possibly due to the difficulty of configurations. But compared with other models, our SRAN consistently achieves the best performance on all the configurations, which proves that our model can work stably, even facing diverse conditions and complex rules.

We observe that the accuracy of SRAN on I-RAVEN is very close to that on the original RAVEN dataset (60.8% vs. 60.7%). This phenomenon is as expected since our method mainly focuses on rules in the context, and thus is robust to the answer set.

Ablation Study

As aforementioned, our method mainly gains from the inductive-biased architecture. To validate this point, we study the effects of different components in our SRAN. Table 4 lists the results.

We analyze the stratified strategy to incrementally induce rules, using the performance of different choices of hierarchies. Specifically, we set the rule embedding of certain hierarchy as a zero vector before gate function φ . Thus, the gate function regulates the flow of features into the gated embedding fusion module. We observe that combing more hierarchies always leads to better performance on I-RAVEN, which shows all hierarchies contribute to our framework. We make \mathbb{E}_{cell} orderless by summing all cell-wise embeddings and observe a major drop in performance. These observations prove the effectiveness of the proposed inductive biases of order sensitivity and incremental rule induction.

We further discover that the effectiveness of our model can be put down to its attention to different attributes. We conduct experiments only utilizing single-hierarchy rule embeddings from \mathbb{E}_{cell} , \mathbb{E}_{ind} , \mathbb{E}_{eco} , with respect to three attributes (Type, Size, and Color) of I-RAVEN. As shown in Figure 6, \mathbb{E}_{cell} has strong capacity to infer attributes Type and Size, but struggles to distinguish attribute Color. By contrast, \mathbb{E}_{ind} and \mathbb{E}_{eco} have modest ability to infer attributes Type and Size, and are efficient for attribute Color.

The Advantage of Rule Embeddings

In the real RPM test, it is not clear whether the rule exists in rows or columns. Therefore, it is important to check whether the proposed model can discover the knowledge without any guidance. Rule induction for columns is normally left out when trained on I-RAVEN, given the prior knowledge that rules are applied only row-wise. In order to test the ability of distinguishing whether the rules are applied along rows or columns, we train a SRAN model on I-RAVEN which the induction for column rules is reintegrated into. As a result, there is only a bit drop in accuracy (from 60.8% to 59.6%). This indicates that our model can neglect the distraction brought by columns on its own.

Conclusion

In this paper, we introduced necessary inductive biases for abstract visual reasoning task, such as order sensitivity and

Model	I-RAVEN
\mathbb{E}_{cell} (orderless)	23.5
\mathbb{E}_{cell}	36.7
\mathbb{E}_{ind}	48.7
\mathbb{E}_{eco}	51.6
$\mathbb{E}_{\text{cell}} + \mathbb{E}_{\text{eco}}$	52.9
$\mathbb{E}_{\text{ind}} + \mathbb{E}_{\text{eco}}$	57.0
$\mathbb{E}_{\text{cell}} + \mathbb{E}_{\text{ind}}$	57.8
$\mathbb{E}_{\text{cell}} + \mathbb{E}_{\text{ind}} + \mathbb{E}_{\text{eco}}$	60.8

Table 4: SRAN ($\mathbb{E}_{\text{cell}} + \mathbb{E}_{\text{ind}} + \mathbb{E}_{\text{eco}}$) and the results of eliminating different components




Figure 6: Accuracy of single hierarchy with respect to the different attributes

incremental rule induction. We further proposed a novel Stratified Rule-Aware Network, which could extract multiple granularity rule embeddings at different level and integrate them through a gated embedding fusion module. A rule similarity metric was further introduced based on the embeddings, so that SRAN can not only be trained using a tuple loss but also infer the best answer according to the similarity score. We also designed an algorithm named Attribute Bisection Tree to fix the defects of the popular dataset RAVEN, and generated a more rigorous dataset based on the algorithm. Extensive experiments conducted on PGM dataset and our improved dataset I-RAVEN proved that, our proposed framework could significantly outperform other state-of-the-art approaches. Moreover, we studied the effects of each component of our proposed model and evaluated the advantage of our induced rule embeddings.

Acknowledgments

This work was supported by National Natural Science Foundation of China (62022009, 61872021), Beijing Nova Program of Science and Technology (Z191100001119050), State Key Lab of Software Development Environment (SKLSDE-2020ZX-06), Fundamental Research Funds for Central Universities (YWF-20-BJ-J-646), and the Academic Excellence Foundation of BUAA for PhD Students.

References

- Barrett, D.; Hill, F.; Santoro, A.; Morcos, A.; and Lillicrap, T. 2018. Measuring abstract reasoning in neural networks. In *Proceedings of International Conference on Machine Learning (ICML)*, 511–520.
- Carpenter, P. A.; Just, M. A.; and Shell, P. 1990. What One Intelligence Test Measures: A Theoretical Account of the Processing in the Raven Progressive Matrices Test. *Psychological Review* 97(3): 404–431.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 770–778.
- Hill, F.; Santoro, A.; Barrett, D.; Morcos, A.; and Lillicrap, T. 2019. Learning to Make Analogies by Contrasting Abstract Relational Structure. In *International Conference on Learning Representations (ICLR)*.
- Hochreiter, S.; and Schmidhuber, J. 1997. Long short-term memory. *Neural computation* 9(8): 1735–1780.
- Hoshen, D.; and Werman, M. 2017. Iq of neural networks. *arXiv preprint arXiv:1710.01692*.
- Johnson, J.; Hariharan, B.; van der Maaten, L.; Fei-Fei, L.; Lawrence Zitnick, C.; and Girshick, R. 2017. Clevr: A diagnostic dataset for compositional language and elementary visual reasoning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2901–2910.
- Kingma, D. P.; and Ba, J. 2014. Adam: A method for stochastic optimization. *International Conference on Learning Representations (ICLR)*.
- Krizhevsky, A.; Sutskever, I.; and Hinton, G. E. 2012. ImageNet classification with deep convolutional neural networks. In *Advances in Neural Information Processing Systems (NeurIPS)*, 1097–1105.
- Kunda, M.; McGregor, K.; and Goel, A. K. 2013. A computational model for solving problems from the Raven’s Progressive Matrices intelligence test using iconic visual representations. *Cognitive Systems Research* 22: 47–66.
- Lovett, A.; and Forbus, K. 2017. Modeling visual problem solving as analogical reasoning. *Psychological Review* 124(1): 60–90.
- Lovett, A.; Forbus, K.; and Usher, J. 2010. A structure-mapping model of Raven’s Progressive Matrices. In *Proceedings of the Annual Meeting of the Cognitive Science Society*, volume 32.
- Lovett, A.; Tomai, E.; Forbus, K.; and Usher, J. 2010. Solving Geometric Analogy Problems Through Two-Stage Analogical Mapping. *Cognitive Science* 33(7): 1192–1231.
- Parent, E. 1996. The living systems theory of James Grier Miller. Retrieved November 29: 2007.
- Raven, J. 2000. The Raven’s progressive matrices: change and stability over culture and time. *Cognitive psychology* 41(1): 1–48.
- Raven, J. C. 1938. Raven’s progressive matrices. *Western Psychological Services*.
- Simonyan, K.; and Zisserman, A. 2015. Very deep convolutional networks for large-scale image recognition. In *International Conference on Learning Representations (ICLR)*.
- Sohn, K. 2016. Improved deep metric learning with multi-class n-pair loss objective. In *Advances in Neural Information Processing Systems (NeurIPS)*, 1857–1865.
- Srivastava, N.; Hinton, G.; Krizhevsky, A.; Sutskever, I.; and Salakhutdinov, R. 2014. Dropout: a simple way to prevent neural networks from overfitting. *The journal of machine learning research* 15(1): 1929–1958.
- Steenbrugge, X.; Leroux, S.; Verbelen, T.; and Dhoedt, B. 2018. Improving generalization for abstract reasoning tasks using disentangled feature representations. *arXiv preprint arXiv:1811.04784*.
- Strannegård, C.; Cirillo, S.; and Ström, V. 2013. An anthropomorphic method for progressive matrix problems. *Cognitive Systems Research* 22: 35–46.
- Szegedy, C.; Liu, W.; Jia, Y.; Sermanet, P.; Reed, S.; Anguelov, D.; Erhan, D.; Vanhoucke, V.; and Rabinovich, A. 2015. Going deeper with convolutions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 1–9.
- van Steenkiste, S.; Locatello, F.; Schmidhuber, J.; and Bachem, O. 2019. Are Disentangled Representations Helpful for Abstract Visual Reasoning? In *Advances in Neural Information Processing Systems (NeurIPS)*, 14222–14235.
- Wang, D.; Jamnik, M.; and Lio, P. 2020. Abstract diagrammatic reasoning with multiplex graph networks. In *International Conference on Learning Representations (ICLR)*.
- Zhang, C.; Gao, F.; Jia, B.; Zhu, Y.; and Zhu, S.-C. 2019a. RAVEN: A dataset for relational and analogical visual reasoning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 5317–5327.
- Zhang, C.; Jia, B.; Gao, F.; Zhu, Y.; Lu, H.; and Zhu, S.-C. 2019b. Learning perceptual inference by contrasting. In *Advances in Neural Information Processing Systems (NeurIPS)*, 1073–1085.
- Zheng, K.; Zha, Z.-J.; and Wei, W. 2019. Abstract Reasoning with Distracting Features. In *Advances in Neural Information Processing Systems (NeurIPS)*, 5834–5845.