

# Scale-Localized Abstract Reasoning

Yaniv Benny<sup>1</sup>Niv Pekar<sup>1</sup>Lior Wolf<sup>1,2</sup><sup>1</sup>The School of Computer Science, Tel Aviv University<sup>2</sup>Facebook AI Research (FAIR)

## Abstract

We consider the abstract relational reasoning task, which is commonly used as an intelligence test. Since some patterns have spatial rationales, while others are only semantic, we propose a multi-scale architecture that processes each query in multiple resolutions. We show that indeed different rules are solved by different resolutions and a combined multi-scale approach outperforms the existing state of the art in this task on all benchmarks by 5-54%. The success of our method is shown to arise from multiple novelties. First, it searches for relational patterns in multiple resolutions, which allows it to readily detect visual relations, such as location, in higher resolution, while allowing the lower resolution module to focus on semantic relations, such as shape type. Second, we optimize the reasoning network of each resolution proportionally to its performance, hereby we motivate each resolution to specialize on the rules for which it performs better than the others and ignore cases that are already solved by the other resolutions. Third, we propose a new way to pool information along the rows and the columns of the illustration-grid of the query. Our work also analyses the existing benchmarks, demonstrating that the RAVEN dataset selects the negative examples in a way that is easily exploited. We, therefore, propose a modified version of the RAVEN dataset, named RAVEN-FAIR. Our code and pretrained models are available at <https://github.com/yanivbenny/MRNet>.

## 1. Introduction

Raven’s Progressive Matrices (RPM) is a widely-used intelligence test [13, 3], which does not require prior knowledge in language, reading, or arithmetics. While IQ measurements are often criticized [21, 4], RPM is highly correlated with other intelligence-based properties [17] and has a high statistical reliability [12, 11]. Its wide acceptance by the psychological community led to an interest in the AI community. Unfortunately, as pointed out by [20, 25, 8],

applying machine learning models to it can sometimes result in shallow heuristics that have little to do with actual intelligence. Therefore, it is necessary to study the pitfalls of RPMs and the protocols to eliminate these exploits.

In Sec. 2.1, we present an analysis of the most popular machine learning RPM benchmarks, PGM [15] and RAVEN [26], from the perspective of biases and exploits. It is shown that RAVEN is built in a way that enables the selection of the correct answer with a high success rate without observing the question itself. To mitigate this bias, we construct, in Sec. 3, a fair variant of the RAVEN benchmark. To further mitigate the identified issues, we propose a new evaluation protocol, in which every choice out of the multiple-choice answers is evaluated independently. This new evaluation protocol leads to a marked deterioration in the performance of the existing methods and calls for the development of more accurate ones that also allows for a better understating of abstract pattern recognition. In Sec. 4, we propose a novel neural architecture, which, as shown in Sec. 5, outperforms the existing methods by a sizable margin. The performance gap is also high when applying the network to rules that were unseen during training. Furthermore, the structure of the new method allows us to separate the rules into families that are based on the scale in which reasoning occurs.

The success of the new method stems mostly from two components: (i) a multi-scale representation, which is shown to lead to a specialization in different aspects of the RPM challenge across the levels, and (ii) a new form of information pooling along the rows and columns of the challenge’s  $3 \times 3$  matrix.

To summarize, the contributions of this work are as follows: (i) An abstract multi-scale design for relational reasoning. (ii) A novel reasoning network that is applied on each scale to detect relational patterns between the rows and between the columns of the query grid. (iii) An improved loss function that balances between the single positive example and the numerous negative examples. (iv) A multi-head attentive loss function that prioritizes the different resolutions to specialize in solving different rules. (v) A new

balanced version of the existing RAVEN dataset, which we call RAVEN-FAIR.

## 2. Related Work and Dataset Analysis

The first attempt in an RPM-like challenge with neural networks involved simplified challenges [6]. The Wild Relation Network (WReN) [15], which relies on the Relation Module [16], was the first to address the full task and introduced the PGM benchmark. WReN considers the possible choices one-by-one and selects the most likely option. Two concurrently proposed methods, CoPINet [27] and LEN [28] have proposed to apply row-wise and column-wise relation operations. Also, both methods benefit from processing all eight possible choices at once, which can improve the reasoning capability of the model, but as we show in this work, has the potential to exploit biases in the data that can arise during the creation of the negative choices. CoPINet applies a contrast module, which subtracts a common factor from all choices, thereby highlighting the difference between the options. LEN uses an additional “global encoder”, which encodes both the context and choice images into a single global vector that is concatenated to the row-wise and column-wise representations. The latest methods, MXGNet [24] and Rel-AIR [18], have proposed different complex architectures to solve this task and also consider all choices at once.

Zhang et al. [28] also introduced a teacher-student training method. It selects samples in a specialized category- and difficulty-based training trajectory and improves performance. In a different line of work, variational autoencoders [10] were shown to disentangle the representations and improve generalization on held-out rules [19]. Our method shows excellent performance and better generalization to held-out rules without relying on either techniques.

When merging different outputs of intermediate paths, such as in residual [5] or shortcut [1, 14, 23] connections, most methods concatenate or sum the vectors into one. This type of pooling is used by the existing neural network RPM methods when pooling information across the grid of the challenge. In Siamese Networks [2], one compares the outputs of two replicas of the same network, by applying a distance measure to their outputs. Since the pooling of the rows and the columns is more akin to the task in siamese networks, our method generalizes this to perform a triple pair-wise pooling of the three rows and the three columns.

### 2.1. Datasets

The increased interest in the abstract reasoning challenges was enabled by the introduction of the PGM [15] and RAVEN [26] datasets. These datasets share the same overall structure. The participant is presented with the first eight images of a 3x3 grid of images, called the context images, and another eight images, called the choice images.

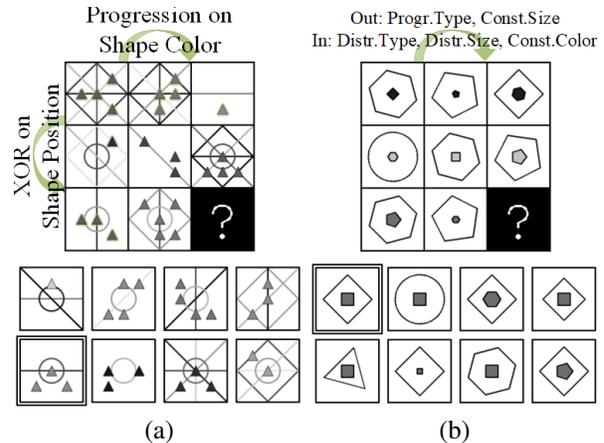


Figure 1. Dataset examples. (a) PGM. (b) RAVEN. The rules are annotated and the correct answer is highlighted.

The objective is to choose the missing ninth image of the grid out of the eight presented choices, by identifying one or more recurring patterns along the rows and/or columns of the grid. The correct answer is the one that fits the most patterns. See Fig. 1 for illustration.

PGM is a large-scale dataset consisting of 1.2M train, 20K validation, and 200K test questions. All images in the dataset are conceptually similar, each showing various amounts of lines and shapes of different types, colors, and sizes. Each question has between 1-4 rules along the rows or columns. Each applies to either the lines or the shapes in the image. Fig. 1(a) shows an example of the dataset.

RAVEN is a smaller dataset, consisting of 42K train, 14K validation, and 14K test questions divided into 7 distinct categories. Each question has between 4-8 rules along the rows only. Fig. 1(b) shows an example of the dataset.

**Evaluation protocols** Both datasets are constructed as a closed-ended test. It can be performed as either a multiple-choice test (MC) with eight choices, where the model can compare the choices, or as a single choice test (SC), where the model scores each choice independently and the choice with the highest score is taken. While this distinction was not made before, the previous works are divided in their approach. WReN follows the SC protocol and CoPINet, LEN, MXGNet, and Rel-AIR, all followed the MC protocol.

In the SC scenario, instead of presenting the agent with all the choices at once, it is presented with a single choice and has to predict if it is the right answer or not. The constraint of solving each image separately increases the difficulty since the model cannot make a decision by comparing the choices. While this eliminates the inter-choice biases, the agent may label multiple or zero images as the correct answer and opens the door to multiple success metrics. In order to directly compare models trained in the MC and SC protocols, we evaluate both types of models in a uniform

manner: the score for all models is the accuracy for the multiple-choice test, where for the SC-trained model, we consider the answer with the highest confidence, regardless of the number of positive answers.

**Dataset Analysis** When constructing the dataset for either the MC or the SC challenge, one needs to be very careful in how the negative examples are selected. If the negative examples are too obvious, the model can eliminate them and increase the probability of selecting the correct answer, without having to fully identify the underlying pattern. Negative examples are therefore constructed based on the question. However, when the negative examples are all conditioned on the correct answer, in the MC scenario, the agent might be able to retrieve the correct answer by looking at the choices without considering the context at all. For example, one can select the answer that has the most common properties with the other answers. See the supplementary for an illustration of such biases in a simple language-based multiple-choice test.

A context-blind test, which is conducted by training a model that does not observe the context images, can check whether the negative answers reveal the correct answer. Ideally, the blind test should return a uniform random accuracy, e.g. 12.5% for eight options. However, since the negative choices should not be completely random so that they will still form a challenge when the context is added, slightly higher accuracy is acceptable. When introduced, the PGM dataset achieved a blind test score of 22.4%. In the RAVEN benchmark, the negative examples are generated by changing a single attribute from the correct image for each example, making it susceptible to a majority-based decision. RAVEN was released without such a context-blind test, which we show in the supplementary that it fails at.

Concurrently to our work, Hu et al. [7] have also discovered the context-blind flaw of RAVEN. They propose an adjustment to the dataset generation scheme that eliminates this problem, which they call ‘Impartial-RAVEN’. Instead of generating the negative choices by changing a random attribute from the correct answer, they propose to sample in advance three independent attribute changes and generate seven images from all the possible combinations of them.

### 3. RAVEN-FAIR

In the supplementary, we analyze both PGM and RAVEN under the blind, SC, and MC settings. As we show, due to the biased selection of the negative examples, the RAVEN dataset fails the context-blind test, as it is solved with 80.17% accuracy by only looking at the choices, making it unsuitable for the MC protocol.

We, therefore, propose a modified dataset we term RAVEN-FAIR, generated by Algorithm 1. This algorithm starts with a set of choices that contains the correct answer

---

#### Algorithm 1: RAVEN-FAIR

---

```

input :  $C$  - 8 context components
          $a$  - the correct answer for  $C$ 
output:  $A$  - 8 choice components ( $a \in A$ )
 $A \leftarrow \{a\}$ ;
while  $length(A) < 8$  do
     $a' \leftarrow Choice(A)$ ;
     $\hat{a} \leftarrow Modify(a')$ ;
    if  $Solve(C, \hat{a}) = False$  then
         $A \leftarrow A \cup \{\hat{a}\}$ ;
    end
end

```

---

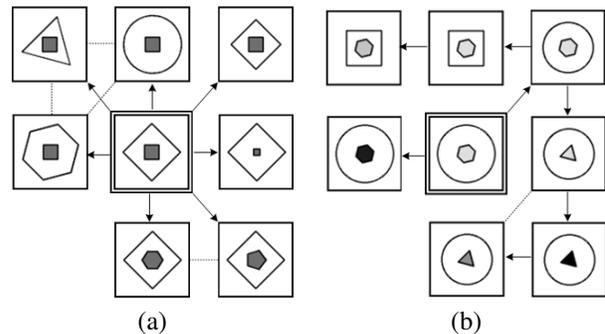


Figure 2. Illustration of negative examples’ generation processes. (a) RAVEN. (b) RAVEN-FAIR.

and iteratively enlarges this set, by generating one negative example at a time. At each round, one existing choice (either the correct one or an already generated negative one) is selected and a new choice is created by changing one of its attributes. The result is a connected graph of choices, where the edges represent an attribute change.

As an example, we show the process of generating the negative examples for RAVEN in Fig. 2(a). For comparison, we also show the process of how our proposed algorithm generates the negative examples for RAVEN-FAIR in Fig. 2(b). Note that the two figures show the produced choices for different context questions. In each figure, eight choices for an initial question are shown. The center image, which is also highlighted, is the correct answer. Each arrow represents a newly generated negative example based on an already existing choice, by changing one arbitrary attribute. The dotted lines connect between choices that also differ by one attribute but were not generated with condition to each other. As can be seen, the RAVEN dataset generates the questions in a way that the correct answer is always the one with the most shared attributes with the rest of the examples. The negative option with the most neighbors in this example is the top-left one, which only has three neighbors, while the correct answer always has eight. Even without highlighting the correct answer, it would be easy to point out which one

it is, by selecting the one with the most neighbors, without looking at the question.

In contrast, our fair algorithm generates the negative examples in a more balanced way. Since each newly generated negative choice can now be conditioned on both the correct image or an already existing negative one, the correct choice cannot be tracked back by looking at the graph alone. In this example, the correct answer only has two neighbors, and both the right and top-right negative images have three neighbors. Because the algorithm is random, the number of neighbors that the correct image has is arbitrary across the dataset, and is between one and eight.

In the context-blind test (supplementary), RAVEN-FAIR returned a 17.24% accuracy, therefore passing the test and making it suitable for both SC and MC. The PGM dataset passed our context-blind test as well, with 18.64% accuracy.

## 4. Method

Our Multi-scale Relation Network (MRNet), depicted in Fig. 3(b), consists of five sub-modules: (i) a three-stage encoder  $E_t$ , where  $t \in \{h, m, l\}$ , which codes the input context and a single choice image into representations in three different resolutions:  $20 \times 20$  (high),  $5 \times 5$  (med) and  $1 \times 1$  (low), (ii) three relation modules  $RM_t$ , one for each resolution, which perform row-wise and column-wise operations on the encodings to detect relational patterns, (iii) three pattern modules  $PM_t$ , one for each resolution, which detects if similar patterns occur in all rows or columns, (iv) three bottleneck networks  $B_t$ , which merge the final features of each resolution, and (v) a predictor module  $MLP$ , which estimates the correctness of a given choice image to the context in question, based on the bottlenecks' outputs. The model is presented with a question in the form of 16 images: a set of context images  $I_C = \{I^n | n \in [1, 8]\}$  and a set of choice images  $I_A = \{I_a^n | n \in [1, 8]\}$ . Since the model operates in the SC protocol, it evaluates the choices separately. Therefore, the notations act as if there is only a single choice image  $I^a$  to answer. To solve all eight choices, the model is repeated eight times with a different choice image  $I_a \in I_A$ .

**Multi-scale encoder** The encoder is a three-stage Convolutional Neural Network that encodes all images (context or choice) into multi-resolution representations (high, middle, low). Every image  $I^n \in [-1, 1]^{1,80,80}$ , for  $n \in \{1, 2, \dots, 8, a\}$ , is passed sequentially through the three stages of the encoder, i.e., the middle resolution encoding is obtained by applying further processing to the output of the high resolution encoder and similarly the low resolution encoding is obtained by further processing the middle resolution one. Specifically, the encoding process results in three tensors.  $e_h^n \in \mathbb{R}^{64,20,20}$ ,  $e_m^n \in \mathbb{R}^{128,5,5}$  and  $e_l^n \in \mathbb{R}^{256,1,1}$ .

$$e_h^n = E_h(I^n), \quad e_m^n = E_m(e_h^n), \quad e_l^n = E_l(e_m^n). \quad (1)$$

**Relation Module** Based on the encoding of the images, we apply a relational module (RM) to detect patterns on each row and column of the query. There are three such modules ( $RM_h, RM_m, RM_l$ ). For each resolution  $t \in \{h, m, l\}$ , the 9 encodings  $e_t^n$  for  $n = 1, 2, \dots, 8, a$  are positioned on a  $3 \times 3$  grid, according to the underlying image  $I^n$ , see Fig. 3(c). The rows and columns are combined by concatenating three encodings on the channel dimension. The rows consist of the three triplets  $(e_t^1, e_t^2, e_t^3)$ ,  $(e_t^4, e_t^5, e_t^6)$ ,  $(e_t^7, e_t^8, e_t^a)$ . Similarly, the columns consist of the three triplets  $(e_t^1, e_t^4, e_t^7)$ ,  $(e_t^2, e_t^5, e_t^8)$ ,  $(e_t^3, e_t^6, e_t^a)$ . Each row and column is passed through the relation network (RN).

$$\begin{aligned} r_t^1 &= RN_t(e_t^1, e_t^2, e_t^3), & c_t^1 &= RN_t(e_t^1, e_t^4, e_t^7) \\ r_t^2 &= RN_t(e_t^4, e_t^5, e_t^6), & c_t^2 &= RN_t(e_t^2, e_t^5, e_t^8) \\ r_t^3 &= RN_t(e_t^7, e_t^8, e_t^a), & c_t^3 &= RN_t(e_t^3, e_t^6, e_t^a) \end{aligned} \quad (2)$$

Each RN consists of two residual blocks with two convolution layers inside each one. For the high and middle resolutions, the convolution has a kernel of size 3 with 'same' padding, while for the low resolution the kernel size is 1 without padding. The output of each relation block is of the same shape as its corresponding encoding  $e_t$ , i.e.  $RN_h : \mathbb{R}^{3 \times 64, 20, 20} \rightarrow \mathbb{R}^{64, 20, 20}$ ,  $RN_m : \mathbb{R}^{3 \times 128, 5, 5} \rightarrow \mathbb{R}^{128, 5, 5}$ , and  $RN_l : \mathbb{R}^{3 \times 256, 1, 1} \rightarrow \mathbb{R}^{256, 1, 1}$ . Note that we apply the same relation networks to all rows and columns. This approach allows comparison between the rows and between the columns to detect recurring patterns. In addition, it maintains the permutation and transpose invariance property we assume on the rows and columns.

**Pattern Module** At this point, at each level  $t \in \{h, m, l\}$  the representation of the panel is structured as three row features and three column features, which we want to merge into a single representation. The module dedicated to this is the pattern module (PM), which applies some operator  $F(\cdot)$  on the rows and columns. It is depicted in Fig. 3(d).

To promote order-invariance between the three rows or columns, a permutation-invariant operator  $F(\cdot)$  is recommended. One can perform sum pooling (SUM3). This kind of approach was employed by [15, 27, 28]. Due to the linearity of the sum pooling, a ReLU is employed.

$$\text{SUM3}(x_1, x_2, x_3) := \text{ReLU}(x_1 + x_2 + x_3) \quad (3)$$

The reduction is then applied on the rows and columns.

$$r_t = \text{SUM3}(r_t^1, r_t^2, r_t^3), \quad c_t = \text{SUM3}(c_t^1, c_t^2, c_t^3) \quad (4)$$

The SUM3 operator has a drawback in that it does not detect similarity between the rows and between the columns. Instead, we propose a novel method that is inspired by Siamese Networks, which compares multiple rows and

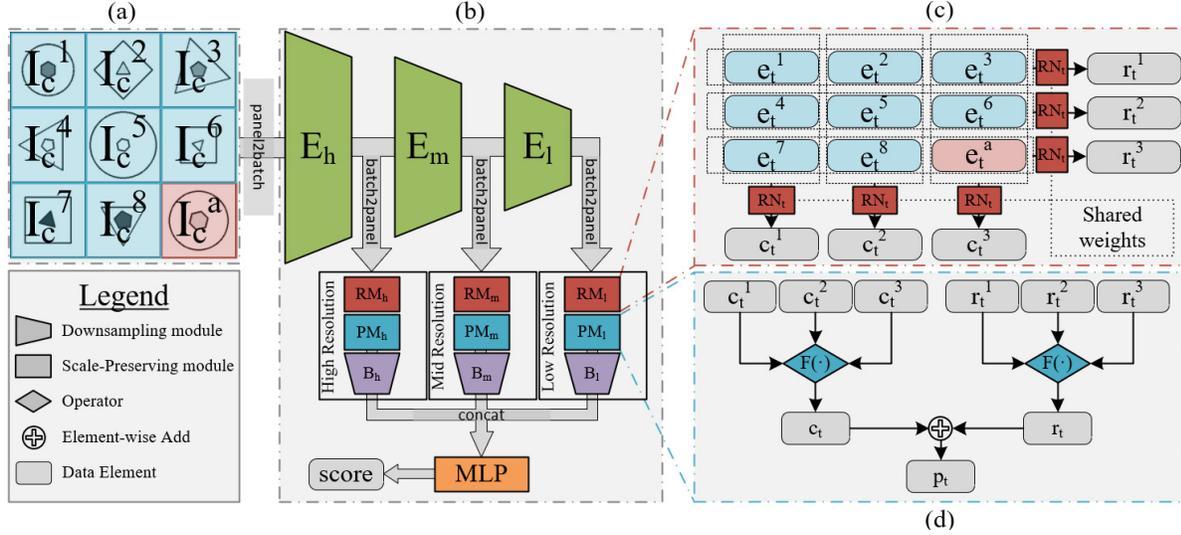


Figure 3. MRNet. (a) Depiction of the input query. (b) High-level diagram of the architecture. (c) The relation module ( $RM_t$ ). (d) The pattern module ( $PM_t$ ).

columns. This vector operator DIST3 is defined per vector index  $i$  and is applied to the rows and columns:

$$\text{DIST3}_i(x_1, x_2, x_3) := (x_{1,i} - x_{2,i})^2 + (x_{2,i} - x_{3,i})^2 + (x_{3,i} - x_{1,i})^2 \quad (5)$$

$$r_t = \text{DIST3}(r_t^1, r_t^2, r_t^3), \quad c_t = \text{DIST3}(c_t^1, c_t^2, c_t^3) \quad (6)$$

DIST3 does not contain an activation function, since the method is already non-linear. Its advantage is also visible in its backward path. The gradient of each operator with respect to element  $i$  of  $x_1$  is:

$$\frac{\partial \text{SUM3}_i(x_1, x_2, x_3)}{\partial x_{1,i}} = 1$$

$$\frac{\partial \text{DIST3}_i(x_1, x_2, x_3)}{\partial x_{1,i}} = 2 \cdot (2 \cdot x_{1,i} - x_{2,i} - x_{3,i}) \quad (7)$$

The gradient of SUM3 does not depend on the values of  $x_2$  and  $x_3$ , while the gradient of DIST3 does. Therefore, DIST3 has the potential to encourage increased coordination between the paths.

Finally, the row- and column-features are summed into a merged representation of the entire panel:

$$p_t = r_t + c_t \quad (8)$$

**Bottleneck module** The three relation modules and their pattern detectors return a latent representation of the panel in three resolutions ( $p_h, p_m, p_l$ ). This module collects the three representations into a single representation. For this

purpose, each representation is downsampled by a bottleneck network ( $B_t$ ) that encodes the panel representation  $p_t$  into feature vectors with 128 features.

$$v_t = B_t(p_t) \quad (9)$$

The three final features are then concatenated into a single vector with 384 features ( $3 \times 128$ ).

$$v = \text{Concat}(v_h, v_m, v_l) \quad (10)$$

**The predictor** The resulting merged feature vector from the bottleneck module is used to predict the correctness of a choice  $I^a \in I_A$  to the context  $I_C$ . A Multi-Layer Perceptron (MLP) predicts the score of the choice images. A sigmoid translates this score to the probability of the choice  $I^a$  to be a correct answer:

$$p(y = 1 | I^a, I_C) = \text{Sigmoid}(MLP(v)) \quad (11)$$

For each choice image, the loss is the Binary Cross Entropy loss. Because each grid has seven negative choices and only one correct answer, we weight the loss of each choice respectively, meaning that the correct answer has a weight of 7 and each negative answer has a weight of 1.

**Multi-head attentive predictor** Optimizing the three heads with the main loss works relatively well, but we have found that adding an objective that optimizes each head separately improves the final results and reduces the overall training time. To do so, three additional predictors were created, one for each bottleneck output  $v_h, v_m, v_l$ . The architecture of the heads is the same as that of the main head and the loss function is identical. To encourage each resolution

to specialize in the rules that the other resolutions are having difficulty with, we propose an attentive weight distribution on the loss functions of the three heads. Consider the prediction of each head  $p_t(y = y^*|I^a, I_C)$  for  $t \in \{h, m, l\}$ , where  $y^* \in \{0, 1\}$  represent whether the choice  $I^a$  is correct or not. The weight applied to each head is:

$$w_t = \frac{\exp\{p_t(y = y^*|I^a, I_C)\}}{\sum_{\tau} \exp\{p_{\tau}(y = y^*|I^a, I_C)\}} \quad (12)$$

Let  $\mathcal{L}_t$  be the Binary Cross Entropy loss for head  $t \in \{h, m, l\}$ . The final multi-head loss is:

$$\mathcal{L}_3 = \sum_t w_t \mathcal{L}_t \quad (13)$$

**Training method** We trained with a batch size of 32 queries with an early stopping scheme of 20 epochs, measuring the accuracy on the validation set. The reported accuracy was measured on the test set, for the checkpoint with the best validation accuracy. See supplementary for a detailed description of the architecture. We did not tune any hyper-parameter of the model. Adam [9] was used with  $\text{lr}=1\text{e-}3$ ,  $\beta=(0.9, 0.999)$ , and weight decay of  $1\text{e-}6$ .

## 5. Experiments

We experiment on PGM, RAVEN, and RAVEN-FAIR. Each dataset has a preset train, validation, and test splits of the data. We used the train set for training and the test set for the evaluation. The validation set was used to pick the best checkpoint during training for evaluation. PGM has additional metadata of the rules logic, which can be used as additional supervision during training. The use of the metadata has been shown to improve performance on all previous works. We refer to this more informative setting as PGM\_meta. Following previous works, the utilization of the metadata was done by extending the predictor module with an additional head that predicts the 12 bits of the metadata, and training the model with an additional (auxiliary) Cross Entropy loss. Similar to the previous works, the weight for the auxiliary loss is  $\beta = 10$ .

**Baselines** We compare our method to the state of the art methods: WReN [15], LEN [28], CoPINet [27], MXGNet [24], and Rel-AIR [18]. We also employed the ResNet [5] models described in the supplementary: ResNet-SC and ResNet-MC. The former evaluates each choice separately and the latter all choices at once.

As noted in Sec. 2, LEN and CoPINet follow the MC protocol. To evaluate these baselines without the possibility of exploiting the weaknesses of RAVEN, we have created additional versions of them for the SC protocol. In LEN, the ‘global-CNN’ originally accepted all 16 images, including the choices. We have changed it to accept only the context images (‘global-8’). In CoPINet, we removed

		Test Accuracy (%)				
		Model	PGM		RAVEN	
	Name	Version	SC base	meta	FAIR	orig
SC Baselines	<i>ResNet-SC</i>		✓ †48.9	-	†58.3	†40.4
	<i>WReN’18</i>	normal	✓ 62.6	76.9	†30.3	†16.8
	<i>V-WReN’18</i>	Normal	✓ 64.8	-	-	-
	<i>CoPINet’19</i>	no-contrast	✓ †42.7	†45.2	†36.5	†18.4
	<i>LEN’19</i>	global-8	✓ †65.6	†79.6	†50.9	†29.9
<i>MRNet</i>			✓ <b>93.4</b>	<b>92.6</b>	<b>86.8</b>	<b>84.0</b>
Ablation	no $\mathcal{L}_3$		✓ 84.4	88.4	84.0	80.1
	no $\mathcal{L}_3$ , no wb		✓ 75.2	84.9	80.6	78.6
	no $\mathcal{L}_3$ , no wb, with SUM3		✓ 74.3	79.0	77.0	69.6
	no wb		✓ 87.6	88.9	83.4	80.2
	with SUM3		✓ 83.2	85.3	79.5	78.2
MC Baselines	<i>ResNet-MC</i>		✗ †41.1	-	†24.5	†72.5
	<i>CoPINet’19</i>	normal	✗ 56.4	†51.1	†50.6	91.4
	<i>LEN’19</i>	normal	✗ 68.1	82.3	†51.0	72.9
	<i>LEN’19</i>	teacher-model	✗ 79.8	85.8	-	78.3
	<i>T-LEN’19</i>	normal	✗ 70.3	84.1	✗	✗
	<i>T-LEN’19</i>	teacher-model	✗ 79.8	88.9	✗	✗
	<i>MXGNet’20</i>		✗ 66.7	89.6	-	83.9
<i>Rel-AIR’20</i>		✗ 85.5	-	-	94.1	
<i>MRNet</i> with-contrast			✗ <b>94.5</b>	<b>92.8</b>	<b>88.4</b>	<b>96.6</b>
Ablation	no $\mathcal{L}_3$		✗ 85.7	89.0	85.2	95.5
	no $\mathcal{L}_3$ , no wb		✗ 76.4	85.4	81.3	94.3
	no wb		✗ 87.4	89.8	86.1	95.0

Table 1. Evaluation on all datasets. †Baseline was run by us, due to missing results.

		Accuracy (%)				
		Baselines			<i>MRNet</i>	
Regime		<i>WReN</i>	<i>V-WReN</i>	<i>MXGNet</i>	(a)	(b)
Neutral		62.6	64.2	66.7	<b>75.2</b>	<b>93.4</b>
Interpolation		64.4	-	65.4	<b>67.1</b>	<b>68.1</b>
Extrapolation		17.2	-	18.9	<b>19.0</b>	<b>19.2</b>
HO Pairs		27.2	36.8	33.6	<b>37.8</b>	<b>38.4</b>
HO Triple Pairs		41.9	43.6	43.3	<b>53.4</b>	<b>55.3</b>
HO Triples		19.0	24.6	19.9	<b>25.3</b>	<b>25.9</b>
HO line-type		14.4	-	16.7	<b>27.0</b>	<b>30.1</b>
HO shape-color		12.5	-	16.6	<b>16.9</b>	<b>16.9</b>

Table 2. Generalization evaluation on the held-out regimes of PGM. (a) Without weight balancing and  $\mathcal{L}_3$ . (b) With them.

the contrasting module (‘no-contrast’), which allows information to pass between the choices. We also compared to LEN with their proposed teacher-model training schedule, and to T-LEN, which has dedicated prediction paths for line

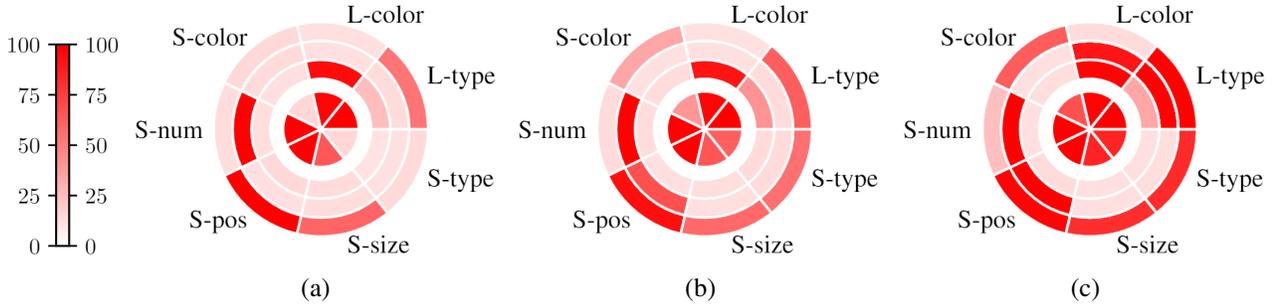


Figure 4. Ablation on the role of each scale in PGM, using a fully trained MRNet on the SC protocol. (a) Without auxiliary loss (PGM), no  $\mathcal{L}_3$ . (b) With auxiliary loss (PGM.meta), no  $\mathcal{L}_3$ . (c) Without auxiliary loss (PGM), with  $\mathcal{L}_3$ . We measured the accuracy for each type, using only one of the scales and the other ones are masked. The inner radius shows the accuracy of the full model. The three rings show the accuracy of each resolution. The inner ring is the low resolution, then the middle, and the high resolution is the outer one.

Scales			Configurations. Accuracy(%)							
h	m	l	All	Center	2×2	3×3	L-R	U-D	O-IC	O-IG
✓	✗	✗	72.0	84.1	47.2	48.3	<b>90.0</b>	<b>91.5</b>	<b>87.6</b>	55.5
✗	✓	✗	<b>73.3</b>	90.5	56.1	53.4	83.5	83.0	82.5	<b>64.1</b>
✗	✗	✓	58.9	<b>97.7</b>	<b>62.2</b>	<b>62.8</b>	42.3	42.5	56.8	48.1
✓	✓	✗	77.7	82.8	61.7	59.5	<b>91.7</b>	<b>92.2</b>	<b>89.9</b>	<b>66.3</b>
✓	✗	✓	<b>78.3</b>	91.0	63.2	64.7	88.8	92.2	87.7	60.3
✗	✓	✓	76.6	<b>98.3</b>	<b>67.5</b>	<b>65.4</b>	81.6	83.4	81.5	58.3
✓	✓	✓	80.6	87.7	64.7	64.3	94.2	94.3	90.9	68.0
with $\mathcal{L}_3$			<b>86.8</b>	<b>97.0</b>	<b>72.7</b>	<b>69.5</b>	<b>98.7</b>	<b>98.9</b>	<b>97.6</b>	<b>73.3</b>

Table 3. Ablation on the role of each scale in RAVEN-FAIR

rules and shape rules and is only applicable to PGM. The results can be seen in Tab. 1 separately for each protocol.

**Results** Before evaluating our model, two important things can be noticed by observing at the baselines. First, the performance of LEN and CoPINet declines in SC, where unlike MC, the models cannot compare the choices within the model. This is especially noticeable in RAVEN, where CoPINet practically solved it in MC and failed it in SC. On both RAVEN and RAVEN-FAIR, all SC baselines performed worse than the simple ResNet-SC, which does not have any dedicated modules for relational reasoning. Second, by comparing each models’ performance between RAVEN and RAVEN-FAIR, we can see that MC models perform significantly lower on RAVEN-FAIR than on RAVEN. This confirms that RAVEN-FAIR indeed fixes the biases in the data and is much more suitable as a benchmark than RAVEN, especially for the MC protocol. These observations align with the dataset analysis presented in the supplementary, which conclusions were discussed in Sec. 2.1.

One can observe that our model outperforms the baselines in both protocols and across all datasets. In SC, our method outperforms previous baselines by 27.8% in PGM, 13.0% in PGM.meta, 35.9% in RAVEN-FAIR, and 54.1% in RAVEN. In MC, we outperform the baselines by 9.0%,

3.2%, 37.4% and 4.9% respectively. Except for RAVEN, our SC protocol model even performed better than both SC and MC protocols of the baselines, despite the SC protocol having less information. We noticed that our model’s performance has reached a point where the auxiliary task of the metadata hurts performance since the added task creates a burden on the optimization of the main task (Vapnik et al. (1998) [22]) and does no longer benefit the model. Therefore the accuracy on PGM.meta is lower than that of PGM, which is the first time a model shows such behavior.

For **ablation**, we trained our model with the SUM3 operator instead of DIST3, without the weight balancing of the positive choices versus the negative ones (no wb), and without the multi-head attentive loss ( $\mathcal{L}_3$ ). The results are shown in Tab. 1 below those of the full model. For SC, we show an aggregated ablation on these modifications, i.e., each following instance applies an additional modification. Each added modification results in a performance decline compared to the previous version. The experiments which remove one component at a time also support that both multi-head attentive loss and the weight balancing both greatly improve the training of the main head. It is also evident that DIST3 is a superior pooling operator over SUM3. For MC, we performed an ablation with either or both the weight balancing and the  $\mathcal{L}_3$  removed. The simplified models without these contributions are not competitive.

Aside from test accuracy, model capacity and training resource consumption is also an important factor. While our method outperforms the other baselines, it is not bigger or slower to train. Our method trains in about 40 minutes per epoch on PGM on a single 1080TI GPU, which is about the same speed as WReN and CoPINet, and about 3x faster than LEN. We did not measure the runtime of MXGNet, Rel-AIR, or LEN with the teacher-model framework, but we expect them to require substantial resources and time.

**Generalization** An important property of a model that is good at abstract reasoning is to be able to generalize over rules that were not seen during training. PGM has specially

built versions of the dataset for this purpose. Aside from the ‘Neutral’ regime, there are seven more regimes, where different rule types were omitted from the train and validation sets, and the test set consists solely of the held-out rules. For more details, please refer to Santoro et al. [15].

Tab. 2 shows that our method generalizes better in all regimes. Most notably are the ‘HO line-type’ and ‘HO Triple Pair’ regimes, where our method has a gap of 10% from the baselines. We do note, however, that other regimes, such as ‘HO shape-color’ and ‘Extrapolation’, only made a negligible improvement on the baselines and are still only slightly better than random. This reassures that there is still much to be done for future work in this regard. In the generalization experiments, training with  $\mathcal{L}_3$  only showed minor improvement. We, therefore, conclude that the performance gain for out-of-domain rules has more to do with the architecture, i.e. the multi-scale design and the relation and pattern modules.

**Understanding the Role of Each Scale** The premise of the importance of multi-scale design is incomplete without gaining additional knowledge on the contribution of each resolution. We, therefore, trained the model multiple times on RAVEN-FAIR with different combinations of scales  $(h, m, l)$  removed. We did not use weight balancing or the multi-head attentive loss for this evaluation. In Tab. 3, we show the average accuracy per rule type. The results clearly show that each resolution has an advantage in different rules. The lower resolution solves the ‘Center’, ‘2×2’, and ‘3×3’ configurations better, and the upper resolution is more adept for ‘L-R’, ‘U-D’ and ‘O-IC’. The ‘O-IG’ configuration is best solved by the middle resolution, which suggests it requires a combination of high-resolution and low-resolution features. When using two scales, the combination of the lower and upper resolutions was better than the ones with the middle resolution, even though the middle resolution is better when only one scale is permitted. This shows that the combination of upper and lower resolutions is more informative. Finally, the full model performs best on average, and the added  $\mathcal{L}_3$  improves on all configurations by optimizing each scale separately.

A similar analysis was done on PGM. Since PGM is a very large dataset that takes longer to train, instead of re-training, we measured the accuracy of prediction using the output of a single bottleneck by masking the outputs of two different bottlenecks each time. Questions with multiple rules were ignored in this analysis since it wouldn’t be clear which rules assisted in answering the question correctly and would contaminate the analysis. Fig. 4 visualizes the accuracy per resolution for each type of rule. As with RAVEN-FAIR, there is a clear role for each resolution. The lower resolution is responsible for the ‘line’ rules, which are more semantic, while the upper resolution is for the ‘shape’ rules, which have a strong spatial dependency. The middle resolu-

tion is specialized in shape-number, likely because the  $3\times 3$  grid alignment of the shapes in each individual image  $I_C^n$  of PGM (not to be confused with the  $3\times 3$  question alignment) is close to the  $5\times 5$  shaped matrix of the encodings  $e_m$ . The analysis of the model trained on PGM\_meta shows specifically where the added auxiliary loss contributes. All resolutions improved their area of expertise from the setting without the metadata and learned to solve additional tasks. Some rules were already solved without the auxiliary loss, such as ‘shape-number’ and ‘line-type’, but others, such as ‘shape-color’ and ‘shape-type’, received a substantial increase in accuracy.  $\mathcal{L}_3$  showed a large improvement on all rules and better utilization of the middle resolution, without having to use the metadata. The exact results on each rule can be found in the supplementary.

Both analyses come to the same conclusion: Each scale is naturally advantageous for a different set of tasks. Rules that require full embedding of the image, to recognize the shape of a large object (RAVEN ‘Center’) or detecting a pattern in arbitrary located lines (PGM ‘line-color’), require low-resolution encoding. Rules on small objects and specific positions (RAVEN ‘L-R’, PGM ‘shape-position’) are better solved in high resolution, before the spatial coherency is lost. We noticed during training that the model converges in steps, i.e. that several times the improvement on the validation set stagnates and then starts to rise again. We hypothesized that these steps occur when learning new rules and have conducted a per-rule analysis during training. The results, presented in the supplementary, indicate that this is indeed the case. A different set of rules is learned in each step and the learned rules usually have a common property.

## 6. Conclusions

The novel method we introduce outperforms the state of the art methods across tasks, protocols, benchmarks, and splits. It also generalizes better on unseen rules. Since the MC protocol can be readily exploited, we advocate for the SC protocol, which more directly tests the ability to infer the pattern from the question. For MC, evaluating on the RAVEN-FAIR variant mitigates the weaknesses of the original RAVEN benchmark. We expect both multi-scale encoders and pairwise distance-based pooling to be beneficial also for other multi-faceted tasks that involve reasoning such as causality analysis based on temporal sequences, visual question answering, and physics-based reasoning.

## Acknowledgements

The contribution of the first author is part of a Ph.D. thesis research conducted at Tel Aviv University.

This project has received funding from the European Research Council (ERC) under the European Unions Horizon 2020 research and innovation programme (grant ERC CoG 725974).

## References

- [1] Christopher M Bishop et al. *Neural networks for pattern recognition*. Oxford university press, 1995. [2](#)
- [2] Jane Bromley, Isabelle Guyon, Yann LeCun, Eduard Säckinger, and Roopak Shah. Signature verification using a “siamese” time delay neural network. In *Advances in neural information processing systems*, pages 737–744, 1994. [2](#)
- [3] Patricia A Carpenter, Marcel A Just, and Peter Shell. What one intelligence test measures: a theoretical account of the processing in the raven progressive matrices test. *Psychological review*, 97(3):404, 1990. [1](#)
- [4] James R Flynn. Massive iq gains in 14 nations: What iq tests really measure. *Psychological bulletin*, 101(2):171, 1987. [1](#)
- [5] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. [2](#), [6](#)
- [6] Dokhyam Hoshen and Michael Werman. IQ of neural networks. *arXiv preprint arXiv:1710.01692*, 2017. [2](#)
- [7] Sheng Hu, Yuqing Ma, Xianglong Liu, Yanlu Wei, and Shihao Bai. Stratified rule-aware network for abstract visual reasoning. In *AAAI Conference on Artificial Intelligence (AAAI)*, 2021. [3](#)
- [8] Jason Jo and Yoshua Bengio. Measuring the tendency of cnns to learn surface statistical regularities. *arXiv preprint arXiv:1711.11561*, 2017. [1](#)
- [9] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. [6](#)
- [10] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013. [2](#)
- [11] Nicholas Mackintosh and Nicholas John Mackintosh. *IQ and human intelligence*. Oxford University Press, 2011. [1](#)
- [12] Ulric Neisser, Gwyneth Boodoo, Thomas J Bouchard Jr, A Wade Boykin, Nathan Brody, Stephen J Ceci, Diane F Halpern, John C Loehlin, Robert Perloff, Robert J Sternberg, et al. Intelligence: knowns and unknowns. *American psychologist*, 51(2):77, 1996. [1](#)
- [13] Jean Raven et al. Raven progressive matrices. In *Handbook of nonverbal assessment*, pages 223–237. Springer, 2003. [1](#)
- [14] Brian D Ripley. *Pattern recognition and neural networks*. Cambridge university press, 2007. [2](#)
- [15] Adam Santoro, Felix Hill, David Barrett, Ari Morcos, and Timothy Lillicrap. Measuring abstract reasoning in neural networks. In *International Conference on Machine Learning*, pages 4477–4486, 2018. [1](#), [2](#), [4](#), [6](#), [8](#)
- [16] Adam Santoro, David Raposo, David G Barrett, Mateusz Malinowski, Razvan Pascanu, Peter Battaglia, and Timothy Lillicrap. A simple neural network module for relational reasoning. In *Advances in neural information processing systems*, pages 4967–4976, 2017. [2](#)
- [17] Richard E Snow, Patrick C Kyllonen, Brachia Marshalek, et al. The topography of ability and learning correlations. *Advances in the psychology of human intelligence*, 2(S 47):103, 1984. [1](#)
- [18] Steven Spratley, Krista Ehinger, and Tim Miller. A closer look at generalisation in raven. *ECCV*, 2020. [2](#), [6](#)
- [19] Xander Steenbrugge, Sam Leroux, Tim Verbelen, and Bart Dhoedt. Improving generalization for abstract reasoning tasks using disentangled feature representations. *arXiv preprint arXiv:1811.04784*, 2018. [2](#)
- [20] Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian Goodfellow, and Rob Fergus. Intriguing properties of neural networks. *arXiv preprint arXiv:1312.6199*, 2013. [1](#)
- [21] Jan Te Nijenhuis, Olga F Voskuijl, and Natasja B Schijve. Practice and coaching on iq tests: Quite a lot of g. *International Journal of Selection and Assessment*, 9(4):302–308, 2001. [1](#)
- [22] Vladimir N. Vapnik. *Statistical Learning Theory*. Wiley-Interscience, 1998. [7](#)
- [23] William N Venables and Brian D Ripley. *Modern applied statistics with S-PLUS*. Springer Science & Business Media, 2013. [2](#)
- [24] Duo Wang, Mateja Jamnik, and Pietro Lio. Abstract diagrammatic reasoning with multiplex graph networks. *ICLR*, 2020. [2](#), [6](#)
- [25] Chiyuan Zhang, Samy Bengio, Moritz Hardt, Benjamin Recht, and Oriol Vinyals. Understanding deep learning requires rethinking generalization. *arXiv preprint arXiv:1611.03530*, 2016. [1](#)
- [26] Chi Zhang, Feng Gao, Baoxiong Jia, Yixin Zhu, and Song-Chun Zhu. Raven: A dataset for relational and analogical visual reasoning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. [1](#), [2](#)
- [27] Chi Zhang, Baoxiong Jia, Feng Gao, Yixin Zhu, Hongjing Lu, and Song-Chun Zhu. Learning perceptual inference by contrasting. In *Advances in Neural Information Processing Systems*, pages 1073–1085, 2019. [2](#), [4](#), [6](#)
- [28] Kecheng Zheng, Zheng-Jun Zha, and Wei Wei. Abstract reasoning with distracting features. In *Advances in Neural Information Processing Systems*, pages 5834–5845, 2019. [2](#), [4](#), [6](#)