

One Self-Configurable Model to Solve Many Abstract Visual Reasoning Problems

Mikołaj Małkiński¹, Jacek Mańdziuk^{1,2}

¹Faculty of Mathematics and Information Science, Warsaw University of Technology, Warsaw, Poland

²Faculty of Computer Science, AGH University of Krakow, Krakow, Poland
mikolaj.malkinski.dokt@pw.edu.pl, mandziuk@mini.pw.edu.pl

Abstract

Abstract Visual Reasoning (AVR) comprises a wide selection of various problems similar to those used in human IQ tests. Recent years have brought dynamic progress in solving particular AVR tasks, however, in the contemporary literature AVR problems are largely dealt with in isolation, leading to highly specialized task-specific methods. With the aim of developing universal learning systems in the AVR domain, we propose the *unified model for solving Single-Choice Abstract visual Reasoning tasks* (SCAR), capable of solving various single-choice AVR tasks, without making any a priori assumptions about the task structure, in particular the number and configuration of panels. The proposed model relies on a novel *Structure-Aware dynamic Layer* (SAL), which adapts its weights to the structure of the considered AVR problem. Experiments conducted on Raven’s Progressive Matrices, Visual Analogy Problems, and Odd One Out problems show that SCAR (SAL-based models, in general) effectively solves diverse AVR tasks, and its performance is on par with the state-of-the-art task-specific baselines. What is more, SCAR demonstrates effective knowledge reuse in multi-task and transfer learning settings. To our knowledge, this work is the first successful attempt to construct a general single-choice AVR solver relying on self-configurable architecture and unified solving method. With this work we aim to stimulate and foster progress on task-independent research paths in the AVR domain, with the long-term goal of development of a general AVR solver.

Introduction

For many years, Abstract Visual Reasoning (AVR) has been a highly challenging area of artificial intelligence (AI) research (Hernández-Orallo et al. 2016). In a typical AVR task, the solver has to identify a set of rules (visual patterns) that govern the distribution of a set of 2D shapes with certain attributes. The shapes are scattered across several *image panels*, which are arranged in a specific 2D structure. Popularised by human IQ tests, the prevailing form of task representation is a single-choice matrix, where a test-taker has to select one of the panels as a task answer. A common example are Raven’s Progressive Matrices (RPMs) (Raven 1936; Raven and Court 1998) that consist of a 3×3 grid of image panels with the bottom-right panel being missing. The task

is to complete the grid with one of the panels picked from a separate set of *answer panels* that best fits the arrangement of the image panels. Top-left part of Fig. 1 presents an RPM with 9 image panels and 8 answer panels (A – H). Additional examples are presented in Appendix B in the supplementary material (Małkiński and Mańdziuk 2023a).

Recently, the AVR field has witnessed an increasing interest from the deep learning (DL) community (Małkiński and Mańdziuk 2023b). While early works reported successful use of DL models to solve easier instances of AVR tasks, e.g. certain RPMs (Hoshen and Werman 2017; Mańdziuk and Żychowski 2019), soon after the effectiveness of DL models in solving more demanding matrices has been questioned (Barrett et al. 2018; Zhang et al. 2019a). Concurrently, a wide suite of AVR benchmarks, extending beyond the concept of RPMs, has been proposed. For instance, Visual Analogy Problems (VAPs) (Hill et al. 2019) present a conceptual abstraction challenge similar to RPMs, but with different numbers of context and answer panels. Another example are Odd One Out (O3) tests (Mańdziuk and Żychowski 2019) that require the subject to identify the panel that stands out from the remaining ones. Though all the above AVR problems have abstract reasoning at their core, they differ in the number of panels, task structure, and the rules (visual patterns) that conceptually underlie their matrix representations.

Due to the variety of AVR problems, numerous approaches to tackle them have been proposed (Małkiński and Mańdziuk 2022). In effect, for the most popular AVR tasks a steady research progress has been observed, and the performance of the methods reached, or even surpassed, the human level (Wu et al. 2020). Despite impressive performance, contemporary methods largely focus on solving single tasks, often by embedding task-specific design choices into the model’s architecture, and cannot be easily evaluated on other, even similar problems. Consequently, the generality and wider applicability of these solution methods remains unclear. What is more, the mainstream research commonly utilizes well-established benchmarks with large amounts of available training instances, as the performance of DL AVR methods, when evaluated in low-data regimes, rapidly decreases (Zhuo and Kankanhalli 2021). On the contrary, human IQ tests which incorporate AVR tasks often assume that the solver has limited prior experience with the

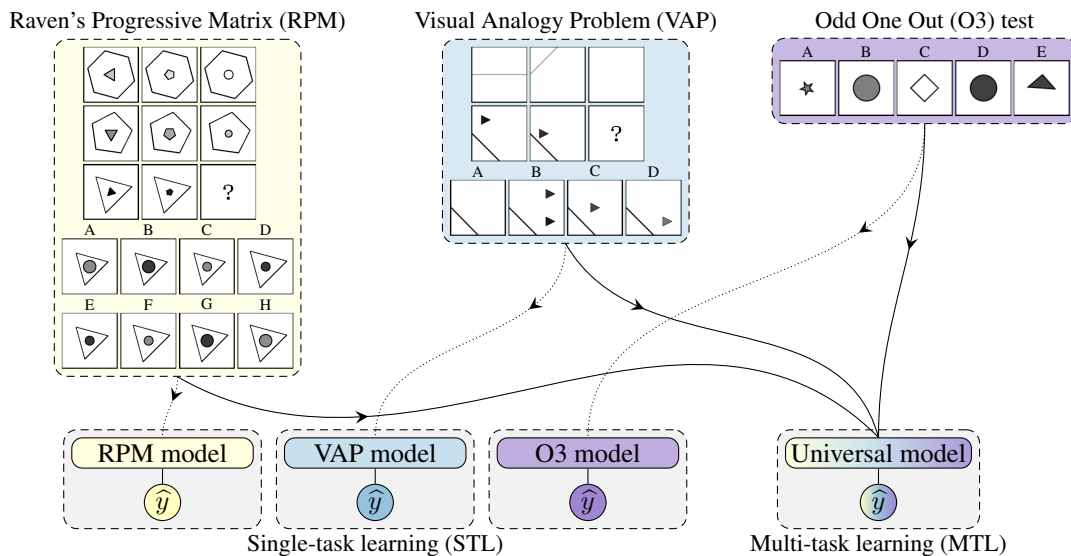


Figure 1: STL vs MTL. In this paper, we consider a diverse set of AVR tasks including RPMs, VAPs, and O3 tests (top). Existing literature deals with AVR tasks in isolation, which leads to the development of task-specific methods with limited applicability to other, even related problems (bottom left). Instead, we propose a single self-configurable model capable of dealing with diverse AVR tasks, thus intrinsically facilitating MTL settings (bottom right).

tests, enabling less biased evaluation of *fluid intelligence* (as opposed to *crystallised intelligence*, which depends on existing knowledge that can be acquired) (Hofstadter 1995).

To address the above limitations and facilitate development of universal AVR solvers, we:

1. Propose a *unified model for solving Single-Choice Abstract visual Reasoning tasks* (SCAR), with no assumptions regarding the number of panels or the task structure and layout. The key concept of the model is a linear *Structure-Aware dynamic Layer* (SAL), which adapts its structure to the current problem instance, thus enabling to solve diverse AVR tasks with a common model architecture.
2. Explore multi-task learning (MTL) (Caruana 1997) and transfer learning (TL) formulations in the AVR domain, with knowledge reuse between diverse tasks.
3. Experimentally evaluate the proposed SCAR and SCL+SAL models on five AVR benchmarks, including RPMs, VAPs, and O3 tests, in STL (single-task learning), MTL and TL setups. In all three setups, the results prove high effectiveness of SAL-based models, which visibly outperform two universal learning benchmark models, and are on par with the state-of-the-art task-specific SCL model in the case of solving RPMs – the most popular AVR task.

One of the aims of this work is to shift the focus of future AVR research from task-specific methods towards general models applicable to diverse AVR tasks.

Related Work

Tasks. Early AVR works concentrate on RPMs, a well-established benchmark used in measuring human intel-

ligence. PGM (Barrett et al. 2018), I-RAVEN (Zhang et al. 2019a; Hu et al. 2021), and G-set (Mańdziuk and Żychowski 2019; Tomaszewska, Żychowski, and Mańdziuk 2022) datasets present RPMs composed of a 3×3 grid of context panels, with up to 8 answer panels to choose from. The VAP benchmark (Hill et al. 2019), structurally similar to RPMs, comprises matrices with a 2×3 context grid and up to 4 answer panels. Differently, O3 tests (Mańdziuk and Żychowski 2019) align up to 6 images in a single row and require the subject to detect the odd panel, which stands out from the remaining ones. While in this work we concentrate on these five datasets which present sufficient diversity to evaluate knowledge reuse in the AVR domain, there are also other relevant tasks that could be of interest for future work. In the SVRT dataset (Fleuret et al. 2011) context panels are split into two groups, each of them containing shapes with certain attributes that conform to a distinct rule (e.g. one shape inside the other vs two separated shapes). Given two test panels, the role of the subject is to assign them to a matching group. A similar setting is considered in Bongard Problems (Bongard 1968) and the Bongard-LOGO (Nie et al. 2020) dataset.

AVR solvers. A frequent design choice made in models for solving AVR tasks is to employ a convolutional backbone which processes each panel separately, producing a vector embedding. Next, various ways have been explored to aggregate these panel embeddings into the model’s prediction. SRAN (Hu et al. 2021) involves three panel representation pathways covering single images, rows/cols, and pairs of rows/cols, each processed with a dedicated convolutional module. These representations are then aggregated using a sequence of multi-layer perceptrons (MLPs). SCL (Wu et al. 2020) proposes the scattering transforma-

tion which merges panel representations with the help of an MLP. However, together with other related methods (Zheng, Zha, and Wei 2019; Spratley, Ehinger, and Miller 2020; Wang, Jamnik, and Lio 2020; Benny, Pekar, and Wolf 2021; Zhuo and Kankanhalli 2021), these approaches incorporate strong task-specific biases into their design (e.g. by utilizing a dense layer with the number of input neurons being determined by the number of panels in the considered problem instance), limiting their applicability to other tasks with a different structure or number of panels. Some methods take a more universal approach by incorporating neural modules that operate on a set of panel embeddings, without making explicit assumption on the number of input objects. WRn (Barrett et al. 2018) employs the Relation Network (RN) (Santoro et al. 2017), which groups input objects into pairs and computes their joint representation. The Recurrent Neural Network (RNN) (Hill et al. 2019) aligns inputs in a sequence and processes them successively, updating their internal representation along the way. Nevertheless, a limitation of these more universal methods is their limited awareness of relative panel positions. This issue is partly mitigated in (Barrett et al. 2018) by concatenating position encodings with panel embeddings. Contrary to the existing approaches, we propose a model capable of processing AVR tasks with diverse structure and variable number of panels through exploiting a neural layout-aware layer, which dynamically adapts its structure to a particular input instance.

Dynamic neural modules. The sample-wise structural layer adaptability of the proposed model relates to the branch of DL research that focuses on designing dynamic neural modules (Han et al. 2021), which adapt their computation in a sample-dependent manner, as opposed to static modules, which perform the same operation for each data point. In particular, conditional computation methods adapt the network’s architecture by gating access to specific neurons based on the input, leading to reduced inference cost and greater expressivity (Bengio, Léonard, and Courville 2013; Shazeer et al. 2017). In contrast, the layer proposed in this work in each forward pass utilizes all available parameters in a differentiable manner, alleviating certain optimization challenges related to sparse models (Bengio et al. 2015). Other dynamic modules adapt kernel sampling locations in convolutional networks (Dai et al. 2017; Zhu et al. 2019), or directly adjust kernel weights (Gao et al. 2020). Alternatively, sample-wise direct model parameter prediction has been explored (Schmidhuber 1992) and has found its applications to linear (Bello 2021) and convolutional (Jia et al. 2016) layers. In contrast, we rely on a 2D structure of the input to form a linear projection layer with dynamic weights, without taking into account input values, resulting in a simpler and parameter-efficient module.

Method

Considering each AVR dataset as a separate task $t \in \mathcal{T}$, we define it as $t = (\{M_i\}_{i=1}^{N_t}, S_t)$. Here, $\{M_i\}_{i=1}^{N_t}$ is a set of N_t problem instances, $M_i = (X_i, y_i)$ is a specific problem instance, $X_i = \{x_{i,j}\}_{j=1}^{P_t} = \{x_{i,j}^C\}_{j=1}^{C_t} \cup \{x_{i,j}^A\}_{j=1}^{A_t}$ is a set of P_t matrix panels which can be partitioned into C_t

context panels and A_t answer panels, $x_{i,j} \in [0, 1]^{h \times w}$ is a greyscale image with height h and width w , $y_i \in [1, A_t]$ is an index of the correct answer, and S_t is the task structure which specifies the layout of the panels (e.g. that context panels are arranged in an $r \times c$ grid, where r/c is the number of rows/columns, resp.). In general, a model for solving matrices from t has the following form:

$$\mathcal{F}_t(\{x_{i,j}\}_{j=1}^{P_t}) = \hat{y}_i \quad (1)$$

A common approach to solve single-choice AVR tasks, is to arrange panels X_i into A_t groups $\{\chi_{i,k}\}_{k=1}^{A_t}$, such that the group $\chi_{i,k} = \{x_{i,j}\}_{j=1}^{I_t} \subset X_i$ corresponds to the answer k , where I_t is the number of panels in the group. Taking RPMs as an example, for a given RPM instance M_i , the groups can be formed by filling in the context grid with each answer panel, giving $\chi_{i,k} = \{x_{i,j}^C\}_{j=1}^{C_t} \cup \{x_{i,k}^A\}$ and $I_t = C_t + 1 = 9$. The group that conforms to the highest number of RPM rules defines the completed form of the (solved) matrix M_i .

To reflect this approach in a DL framework, contemporary approaches proposed to: 1) embed each input panel $x_{i,j}$ with a convolutional encoder \mathcal{E}_t to a latent representation $h_{i,j} \in \mathbb{R}^{d_h}$, 2) arrange the embeddings into A_t groups $\{H_{i,k}\}_{k=1}^{A_t}$, where group $H_{i,k} \subset \{h_{i,j}\}_{j=1}^{P_t}$ corresponds to answer k , 3) fuse the embeddings inside each group with a reasoning module \mathcal{G}_t to generate a joint embedding $g_{i,k} \in \mathbb{R}^{d_g}$, 4) employ a decoder \mathcal{D} to convert the joint embedding to an alignment score $s_{i,k} \in \mathbb{R}$, 5) transform the resultant scores into a probability distribution \hat{p}_i over the set of available answers, 6) return the answer corresponding to the highest probability as the model’s prediction \hat{y}_i :

$$\begin{aligned} 1) \mathcal{E}_t(x_{i,j}) &= h_{i,j} & 2) \text{ARRANGE}(\{h_{i,j}\}_{j=1}^{P_t}) &= \{H_{i,k}\}_{k=1}^{A_t} \\ 3) \mathcal{G}_t(H_{i,k} | S_t) &= g_{i,k} & 4) \mathcal{D}(g_{i,k}) &= s_{i,k} \\ 5) \text{SOFTMAX}(\{s_{i,k}\}_{k=1}^{A_t}) &= \hat{p}_i & 6) \text{ARG MAX } \hat{p}_i &= \hat{y}_i \end{aligned} \quad (2)$$

Insofar, models were developed with the aim of tackling a specific task, without much consideration for generalizing to other problems. To a large extent, this led to task-specific models which assume I_t and S_t in advance, and incorporate them directly into their architecture. To lift these assumptions and bolster universality of AVR models, we i) propose a neural module which dynamically adapts to matrices with various I_t and S_t , ii) embody it in an end-to-end architecture for solving AVR tasks, and iii) demonstrate how the method can be applied to solving tasks with diverse structures.

Dynamic Weight Adaptation

To solve various AVR tasks with a single unified model architecture, we start by designing a differentiable computation layer $\psi(H_{i,k} | S_t) = v_{i,k}$ which produces a latent representation $v_{i,k} \in \mathbb{R}^{d_v}$ for a set of I_t panel embeddings $H_{i,k}$, with a layout defined by the task structure S_t . As we discuss in the next section, ψ will serve as the first layer in \mathcal{G} , and $v_{i,k}$ will be transformed to $g_{i,k}$ by the subsequent layers of \mathcal{G} . To instantiate ψ , we propose the *Structure-Aware dynamic Layer* (SAL) – see Fig. 2:

$$\begin{aligned} \mathcal{W}(W^* | S_t) &= W & \text{CONCAT}(H_{i,k}) &= G_{i,k} \\ W^T G_{i,k} + B &= v_{i,k} \end{aligned} \quad (3)$$

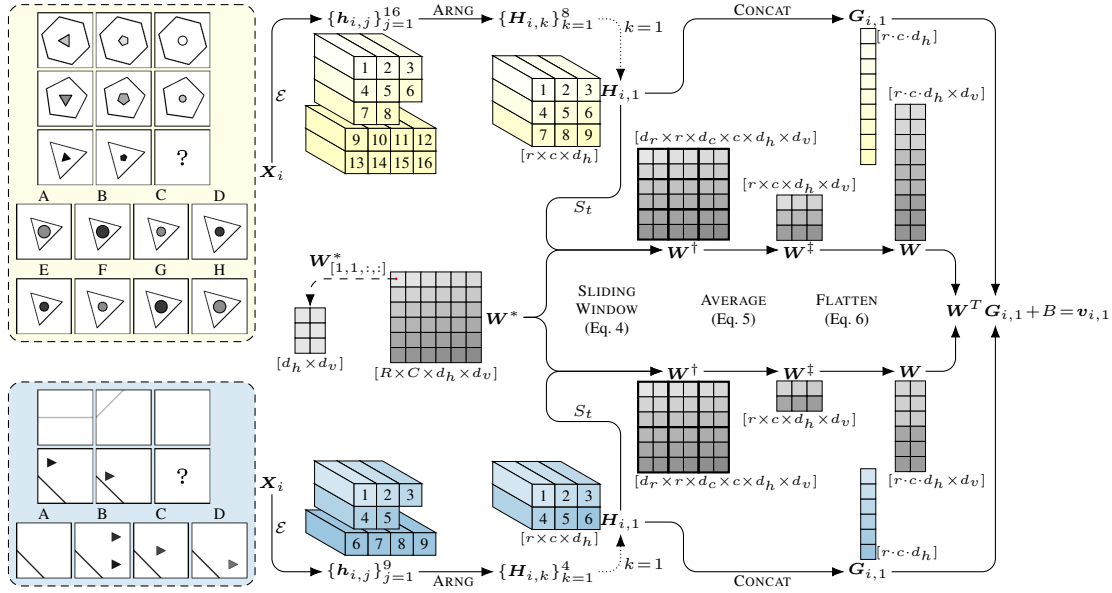


Figure 2: Structure-Aware dynamic Layer. SAL enables processing of AVR tasks with diverse structures by adapting its weights to the problem instance. The figure illustrates two separate forward passes through the layer for an RPM (top) and a VAP matrix (bottom). From left, the panels X_i belonging to a single problem instance are processed with an encoder \mathcal{E} . The resultant embeddings $\{h_{i,j}\}$ are arranged into A_t groups $\{H_{i,k}\}$. Next, the figure illustrates further processing of a single group $H_{i,1}$ for both matrices. Thanks to SAL’s adaptability, the resultant vector $v_{i,1} \in \mathbb{R}^{d_v}$ has the same dimensions irrespectively of the considered task, thus enabling uniform processing in the subsequent model layers.

From a high-level perspective, SAL can be viewed as a linear layer with weights $W \in \mathbb{R}^{r \cdot c \cdot d_h \times d_v}$, where r/c is the number of rows/columns in the context grid, and $B \in \mathbb{R}^{d_v}$ are optional biases. In contrast to the widely-adopted linear layer with static weights, the matrix W in SAL is computed dynamically by the transformation \mathcal{W} , based on the underlying weight matrix $W^* \in \mathbb{R}^{R \times C \times d_h \times d_v}$, where R and C are its row and column dimensions, resp. Specifically, \mathcal{W} utilizes a sliding window, which is adapted according to the task structure S_t , to unfold the first two dimensions of W from $R \times C$ to $d_r \times r \times d_c \times c$ (Eq. 4), averages along the first and third dimensions (Eq. 5), and flattens the matrix into 2D form (Eq. 6):

$$\text{SLIDINGWINDOW}(W^* | S_t) = W^\dagger \in \mathbb{R}^{d_r \times r \times d_c \times c \times d_h \times d_v} \quad (4)$$

$$\frac{1}{d_r} \frac{1}{d_c} \sum_{i=1}^{d_r} \sum_{k=1}^{d_c} W_{ijklmn}^\dagger = W_{jlmn}^\ddagger \in \mathbb{R}^{r \times c \times d_h \times d_v} \quad (5)$$

$$\text{FLATTEN}(W^\ddagger, 1 : 3) = W \in \mathbb{R}^{r \cdot c \cdot d_h \times d_v} \quad (6)$$

where $d_r = \frac{R}{r}$ and $d_c = \frac{C}{c}$. The SAL input is constructed by concatenating the embedding vectors $H_{i,k}$, to obtain $G_{i,k} \in \mathbb{R}^{r \cdot c \cdot d_h}$.

Multi-head SAL. On top of the default formulation, we propose a multi-head version of SAL. In this setting, each embedding vector $h_{i,j}$ from $\{h_{i,j}\}_{j=1}^{P_i}$ is partitioned into a set of L vectors $\{h_{i,j,l}\}_{l=1}^L$, where $h_{i,j,l} = h_{i,j}[(l-1)d_l : l d_l] \in \mathbb{R}^{d_l}$, and $d_l = \frac{d_h}{L}$. The resultant set of panel embeddings $\{\{h_{i,j,l}\}_{l=1}^L\}_{j=1}^{P_i}$ is processed in parallel for each l with the arrangement operator (ARRANGE in Eq. 2) and ψ , leading to $\{v_{i,k,l}\}_{l=1}^L$. The vectors are then concatenated producing $v_{i,k}$, which can be processed by the subsequent layers of \mathcal{G} . The above multi-head SAL variant, with $L = 20$, was used in all conducted experiments.

Model Architecture

We propose the *unified model for solving Single-Choice Abstract visual Reasoning tasks* (SCAR), which is an end-to-end architecture based on SAL, illustrated in Fig. 3. A detailed description of all SCAR hyperparameters is provided in Appendix C (Małkiński and Mańdziuk 2023a).

SCAR starts with a panel encoder \mathcal{E} which processes each $x_{i,j}$ with a 4-layer convolutional module to extract low-level image features. Each layer comprises a 2D 3×3 convolution, 2D Batch Normalization (Ioffe and Szegedy 2015), and the ReLU (Nair and Hinton 2010) activation. The output of the module is flattened along the spatial dimensions to a 2D matrix and processed with a token mixer module, which fuses feature representations from different spatial regions of the considered panel. Token mixer is composed of Layer Normalization (Ba, Kiros, and Hinton 2016), two linear layers with the GELU (Hendrycks and Gimpel 2016) nonlinearity in-between, and a residual connection. Next, a channel mixer module consisting of two 1D 1×1 convolutions interleaved with GELU is employed to mix features from a given spatial region and flatten the matrix to a 1D vector. Another

token mixer transforms the flattened vector into the panel embedding $h_{i,j}$.

Next, the resultant panel embeddings $\{h_{i,j}\}_{j=1}^{P_t}$ are arranged into A_t groups $\{H_{i,k}\}_{k=1}^{A_t}$. Each group is then processed separately by the reasoning module \mathcal{G} , which consists of SAL, GELU, channel mixer, token mixer, and a linear layer that produces $g_{i,k}$. Lastly, two linear layers are interleaved with GELU to map $g_{i,k}$ onto the predicted score $s_{i,k}$.

Unified Reasoning

The proposed model is applied to three challenging AVR problems of different structure – for RPMs $r = c = 3$, for VAPs $r = 2$ and $c = 3$, and for O3 tests $r = 1$ and $c \in \{4, 5, 6\}$. To enable solving these tree diverse tasks with a common SCAR model, the following design choices have been made.

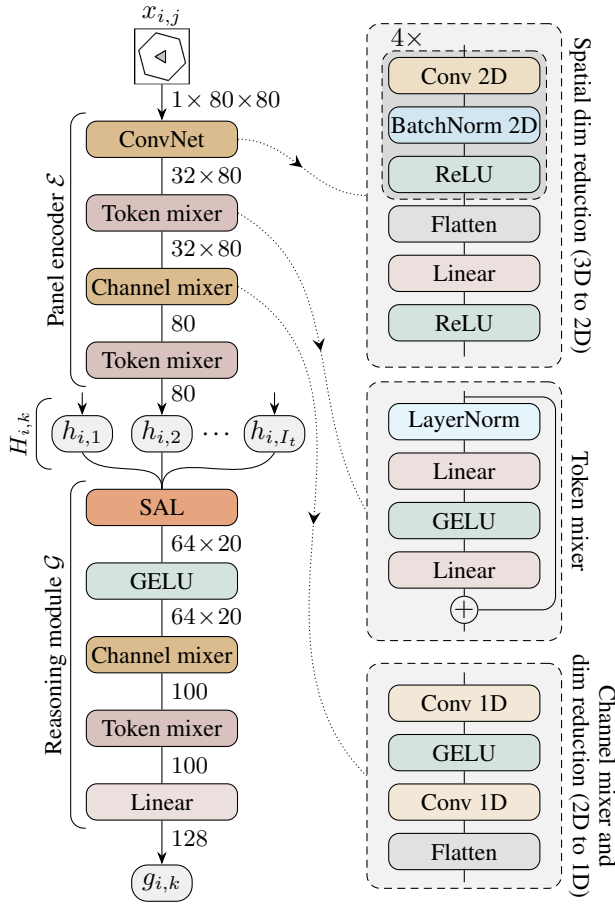


Figure 3: SCAR architecture. Each input panel $x_{i,j}$ is embedded separately with the panel encoder \mathcal{E} to a latent representation $h_{i,j}$. The embeddings are arranged into A_t groups $\{H_{i,k}\}_{k=1}^{A_t}$. The figure depicts processing of a single group $H_{i,k}$ corresponding to answer k . Vectors belonging to this group are fused and processed with the reasoning module \mathcal{G} . The resultant representation $g_{i,k}$ is used to predict the index of the correct answer and optionally the encoded rules.

SCAR adaptation to diverse structures. The proposed formulation of dynamic weight adaptation in SAL utilizes an underlying weight matrix W^* , for which R and C (its first two dimensions) have to be evenly divisible by r and c , see Eq. 4. To instantiate SCAR, we calculate the least common multiple of all possible values for $r \in \{1, 2, 3\}$ and $c \in \{3, 4, 5, 6\}$, which gives $R = 6$ and $C = 60$.

Solving O3 tests. O3 tests have a fundamentally different goal than RPMs and VAPs. Instead of selecting an answer that correctly completes the context grid, in O3 tests it is required to select a panel that differentiates most from the remaining ones. To handle these opposing settings in one unified model, we cast O3 to the setting where a subset of $P_t - 1$ panels with the most common features has to be chosen. Putting this into the introduced framework, we arrange the panel embeddings $\{h_{i,j}\}_{j=1}^{P_t}$ into groups $\{H_{i,k}\}_{k=1}^{A_t}$ (for O3 tests $P_t = A_t$) such that $H_{i,k} = \{h_{i,j} \mid j \neq k\}_{j=1}^{P_t}$. This allows to interpret $g_{i,k}$ as a latent representation of semantic congruency of the considered subset of panels.

Training objectives. To optimize parameters of the model, we compute cross-entropy \mathcal{L}^{CE} between the predicted distribution \hat{p} and ground-truth p obtained by one-hot encoding of the target label k . In addition, some considered datasets provide auxiliary metadata referring to the problem instances (Barrett et al. 2018). In particular, matrices from PGM, I-RAVEN, and VAP provide ground-truth annotations of the underlying matrix rules: $M_i = (X_i, y_i, r_i)$. To utilize this additional training signal, for each task we employ a separate rule prediction head $\mathcal{R}_t(g_{i,k}) = \hat{r}_{i,k}$ composed of two linear layers with GELU in-between. \mathcal{R} predicts the rule representation r_i encoded with sparse encoding (Małkiński and Mańdziuk 2024). Binary cross-entropy \mathcal{L}^{AUX} is used to assess the fitness between \hat{r}_i and r_i . In effect, the model is trained to minimise $\mathcal{L}^{\text{CE}} + \beta \mathcal{L}^{\text{AUX}}$, where $\beta = 10$ is a fixed balancing coefficient.

Mini-batch training. In MTL experiments, each training batch is composed of matrices belonging to a single randomly-chosen task. This ensures that the operation $W^T G_{i,k}$ (Eq. 3) can be efficiently computed for the whole batch with a single matrix multiplication.

Experiments

The experimental evaluation of SCAR focuses on the following three settings: (1) single-task learning (STL), where the model is trained and evaluated on the same task, (2) multi-task learning (MTL), where the model is simultaneously trained on several tasks and then fine-tuned, separately for each task, and (3) transfer learning (TL), where the model is pre-trained on a base set of tasks and fine-tuned on a novel target task. In each scenario, the performance of SCAR is compared to leading AVR literature models, both task-specific and universal. The code required to reproduce the experiments is available online¹. Extended results are provided in Appendix A (Małkiński and Mańdziuk 2023a).

¹www.github.com/mikomel/sal

| MODEL | TEST ACCURACY (%) | | | | |
|---------|-------------------|-------------|-------------|-------------|-------------|
| | G-SET | PGM-S | I-RAVEN | VAP-S | O3 |
| WREN | 34.7 | 25.1 | 21.9 | 73.6 | <u>65.6</u> |
| RNN | 49.6 | 29.4 | 38.9 | 88.1 | 51.2 |
| SCL+SAL | 89.9 | 51.5 | 95.5 | <u>92.6</u> | 63.1 |
| SCAR | <u>82.1</u> | <u>46.1</u> | <u>94.7</u> | 92.9 | 86.7 |
| COPINET | 63.5 | 37.3 | 47.1 | N/A | N/A |
| SRAN | <u>82.3</u> | 39.4 | 66.9 | N/A | N/A |
| RELBASE | 81.3 | <u>50.3</u> | <u>93.7</u> | N/A | N/A |
| SCL | 83.4 | 51.1 | 95.5 | N/A | N/A |

Table 1: Single-task learning. Test accuracy of universal (WReN, RNN, SCL+SAL, SCAR) and task-specific (CoPINet, SRAN, RelBase, SCL) models in solving a single task. Due to architectural constraints, the considered task-specific models created with RPMs in mind can’t be directly applied to solve VAPs and O3 tests. Best results in each group are marked in bold, and second best are underlined.

Tasks. The experiments are conducted on three challenging AVR problems. Firstly, we consider RPMs belonging to three datasets: G-set (Mańdziuk and Żychowski 2019) with 1 100 matrices; PGM-S, a subset of the Neutral regime of PGM (Barrett et al. 2018) where the training split was obtained by uniformly sampling instances from the parent dataset preserving the original ratio of object–attribute–rule triplets; and I-RAVEN (Zhang et al. 2019a; Hu et al. 2021) with 70K samples. Secondly, we consider the VAP dataset (Hill et al. 2019) to construct the VAP-S dataset with training matrices sampled analogously to PGM-S. Thirdly, the O3 tests (Mańdziuk and Żychowski 2019) with 1 000 instances are utilized. In G-set, I-RAVEN, and O3 datasets we uniformly allocate 60% / 20% / 20% samples to train / val / test splits, resp. PGM-S contains 42K / 20K / 200K matrices and VAP-S has 42K / 10K / 100K instances in train / val / test splits, resp. Altogether, the selection of tasks comprises a diverse set of AVR matrices ranging from visually simple instances from G-set and O3, through more challenging ones from PGM and VAP, up to matrices from I-RAVEN with hierarchical structure. In addition, the variability of the component dataset sizes allows to evaluate the models’ performance in data-constrained scenarios. At the same time, considering subsets of PGM and VAP allows the construction of a joint dataset for MTL where the size of these initially large datasets doesn’t dominate the smaller ones.

Baselines. SCAR is compared with WReN (Barrett et al. 2018) and RNN (Hill et al. 2019), the two alternative universal models, capable of processing AVR tasks with variable structure, which is a prerequisite for their application to multi-task scenarios. We also consider a variant of SCL (Wu et al. 2020) where the first layer in its relationship network \mathcal{N}^r is replaced with SAL (denoted as SCL+SAL). In addition, we compare SCAR with selected state-of-the-art task-specific models: SRAN (Hu et al. 2021), CoPINet (Zhang et al. 2019b), RelBase (Spratley, Ehinger, and Miller 2020) and SCL. These additional baselines are included only in the RPM settings, as the architectural constraints prevent their

direct application to tasks with different structure.

Experimental setting. We use batches of size 32 for G-set and O3, and 128 for the remaining datasets. Early stopping is performed after the model’s performance stops improving on a validation set for 17 successive epochs. We use the Adam (Kingma and Ba 2014) optimizer with $\beta_1 = 0.9$, $\beta_2 = 0.999$ and $\epsilon = 10^{-8}$, the learning rate is initialized to 0.001 and reduced 10-fold if the validation loss doesn’t improve for 5 subsequent epochs. Data augmentation is employed with 50% probability of being applied to a particular sample. When applied, a pipeline of transformations (vertical flip, horizontal flip, rotation by 90 degrees, rotation, transposition) is constructed, each with 25% probability of being adopted. The resultant pipeline is applied to each image in the matrix in the same way. Each training job was run on a node with a single NVIDIA DGX A100 GPU.

Single-Task Learning

Table 1 presents STL results, where each model is trained and evaluated on the same dataset. Among universal methods, either SCL+SAL or SCAR achieved the best results on each considered dataset, demonstrating SAL’s applicability to solving AVR tasks with diverse structures, and high efficacy even in tasks with limited number of training samples (G-set and O3). When compared to task-specific models, SCAR performed better than two selected baselines (CoPINet and SRAN) and was slightly inferior to SCL. In contrast to the universal methods, the dedicated RPM models cannot be evaluated on VAPs and O3 tests due to their architectural design.

Multi-Task Learning

Next, we evaluate knowledge reuse capabilities of the adopted methods. Toward this end, in the MTL setting, the models are first pre-trained on a base set of tasks and then separately fine-tuned on each task. Two settings for the selection of tasks are considered: (1) RPMs only, which allows measuring knowledge reuse within a single problem; and (2) RPMs combined with VAPs, to evaluate knowledge reuse between problems. In the pre-training phase batches of size 128 are used.

The results are presented in Table 2. Among universal models, either SCAR or SCL+SAL achieved the highest scores. In both settings MTL helped to increase their performance on the most challenging tasks, i.e. G-set and PGM-S, showing SAL’s ability to effectively reuse the gained knowledge. On I-RAVEN and VAP-S their results are on-par with these of STL, however, the models achieved high accuracy already in the STL setup, leaving small room for potential improvement. For the two other universal methods, MTL proved beneficial in some settings (G-set and VAP-S), while it hindered the performance in other cases (I-RAVEN). This, contrary to SCAR and SCL+SAL, suggests insufficient capacity of WReN and RNN to leverage previously acquired knowledge. Regarding the task-specific methods, MTL improved their performance in most experiments, especially on the data-constrained G-set. Due to architectural constraints, VAP could not be included in the experiments concerning

| MODEL | TEST ACCURACY (%) | | | | | | |
|---------|---------------------|---------------------|--------------------|------------------------------|---------------------|--------------------|--------------------|
| | PRE-TRAINED ON RPMs | | | PRE-TRAINED ON RPMs AND VAPs | | | |
| | G-SET | PGM-S | I-RAVEN | G-SET | PGM-S | I-RAVEN | VAP-S |
| WREN | 41.6 (+6.9) | 24.5 (-0.6) | 19.4 (-2.5) | 48.8 (+14.1) | 26.1 (+1.0) | 20.2 (-1.7) | 80.6 (+7.0) |
| RNN | 66.4 (+16.8) | 27.8 (-1.6) | 34.0 (-4.9) | 61.9 (+12.3) | 26.5 (-2.9) | 27.8 (-11.1) | 88.6 (+0.5) |
| SCL+SAL | 94.9 (+5.0) | 58.1 (+6.6) | 95.4 (-0.1) | 94.8 (+4.9) | 68.6 (+17.1) | 95.3 (-0.2) | 92.8 (+0.2) |
| SCAR | <u>91.0</u> (+8.9) | <u>53.7</u> (+7.6) | <u>94.3</u> (-0.4) | <u>93.2</u> (+11.1) | <u>54.8</u> (+8.7) | <u>94.7</u> (+0.0) | 93.0 (+0.1) |
| COPINET | 74.9 (+11.4) | 35.2 (-2.1) | 50.3 (+3.2) | N/A | N/A | N/A | N/A |
| SRAN | 87.2 (+4.9) | 40.7 (+1.3) | 68.3 (+1.4) | N/A | N/A | N/A | N/A |
| RELBASE | <u>95.7</u> (+14.4) | <u>64.0</u> (+13.7) | <u>89.3</u> (-4.4) | N/A | N/A | N/A | N/A |
| SCL | 96.2 (+12.8) | 69.6 (+18.5) | 95.3 (-0.2) | N/A | N/A | N/A | N/A |

Table 2: Multi-task learning. Test accuracy of universal and task-specific models pre-trained simultaneously on multiple tasks and then separately fine-tuned on each task. The results shown in the left part of the table were achieved by pre-training only on RPMs (G-set, PGM-S, and I-RAVEN), while those in the right part additionally included VAPs in the pre-training set. Values in parentheses show the difference w.r.t. the STL setup (cf. Table 1).

| MODEL | TEST ACCURACY (%) | | |
|---------|--------------------|--------------------|---------------------|
| | R → V | R → O3 | R+V → O3 |
| WREN | 83.3 (+9.7) | <u>65.4</u> (-0.2) | 65.5 (-0.1) |
| RNN | 88.1 (+0.0) | 55.9 (+4.7) | 55.3 (+4.1) |
| SCL+SAL | 93.6 (+1.0) | <u>65.4</u> (+2.3) | <u>75.5</u> (+12.4) |
| SCAR | <u>93.0</u> (+0.1) | 88.6 (+1.9) | 89.0 (+2.3) |

Table 3: Transfer learning. Test accuracy of universal models after transferring knowledge from a set of pre-training tasks to a novel target task. R and V denote RPMs and VAPs, resp. In parentheses there are p.p. changes w.r.t. STL (cf. Table 1).

| | TEST ACCURACY (%) | | | | |
|---|-------------------|--------------|--------------|-------------|--------------|
| | G-SET | PGM-S | I-RAVEN | VAP-S | O3 |
| R | 36.5 (-45.6) | 33.2 (-12.9) | 56.1 (-38.6) | 90.5 (-2.4) | 69.5 (-17.2) |
| L | 43.5 (-38.6) | 26.0 (-20.1) | 41.7 (-53.0) | 90.7 (-2.2) | 52.4 (-34.3) |

Table 4: Ablation study. Test accuracy of SCAR variants where SAL is replaced with RN (row R) and LSTM (row L), resp. Both models were trained in the STL setup. Values in parentheses denote p.p. change w.r.t. the STL training of the default variant of SCAR with SAL – see Table 1.

the RPM-specialised methods. Generally, adopting MTL approach to solving AVR tasks, proposed in this paper, has proven effective, especially for the data-limited tasks. At the same time, the results show that MTL need to be designed with care, as in certain settings its application may deteriorate the results, a problem known in the literature as the *negative transfer* effect (Pan and Yang 2010; Cao et al. 2018).

Transfer Learning

To further study knowledge reuse capabilities across different AVR problems, we conduct TL experiments where the models are pre-trained on a set of tasks, and then fine-tuned on a novel target task. Three such settings are considered: taking all three RPM datasets as pre-training tasks and fine-tuning on VAP-S and O3, respectively, and pre-training on

RPMs and VAP-S and fine-tuning on O3. Table 3 shows the results. TL allowed to significantly increase the performance of WREN on VAP-S and improved the results of RNN, SCL+SAL and SCAR on O3. The results amplify the importance of developing universal methods in the AVR domain capable of processing tasks with diverse structures.

Ablation Study

Lastly, we perform an ablation study to validate the usefulness of the proposed SAL in the SCAR architecture. To this end, we develop two variants of SCAR in which SAL is replaced with RN (Santoro et al. 2017) and LSTM (Hochreiter and Schmidhuber 1997), resp., and evaluate them in the STL setting. The results are presented in Table 4. First, we observe that both alternatives are competitive to other methods trained with STL (see Table 1). For instance, the RN variant achieved better results than universal WREN and RNN on G-set, PGM-S, VAP-S, and O3, and outperformed task-specific CoPINet on I-RAVEN. Still, both variants perform significantly worse than the default version of SCAR, which confirms the critical importance of SAL in the proposed model.

Conclusion

In this work we consider the problem of solving diverse AVR tasks with a unified model architecture. To this end we propose SCAR, a DL model capable of solving a variety of single-choice AVR problems. SCAR design principles differ from the vast majority of the contemporary AVR models, which are developed with a particular task in mind. The core idea of the proposed model is SAL, a novel structure-aware layer which adapts its weights to the structure of the considered AVR instance. In the experiments conducted on RPMs, VAPs, and O3 tests, SAL-based models significantly surpassed the performance of other universal methods and even of some task-specific models. At the same time, the models demonstrated effective knowledge reuse capabilities in MTL and TL settings. On a more general note, with this work we aim to encourage AVR researchers to shift the research focus from task-specific methods towards universal AVR solvers applicable to a wide selection of AVR tasks.

Acknowledgements

This research was carried out with the support of the Laboratory of Bioinformatics and Computational Genomics and the High Performance Computing Center of the Faculty of Mathematics and Information Science Warsaw University of Technology. Mikołaj Małkiński was funded by the Warsaw University of Technology within the Excellence Initiative: Research University (IDUB) programme.

References

- Ba, J. L.; Kiros, J. R.; and Hinton, G. E. 2016. Layer normalization. *arXiv preprint arXiv:1607.06450*.
- Barrett, D.; Hill, F.; Santoro, A.; Morcos, A.; and Lillicrap, T. 2018. Measuring abstract reasoning in neural networks. In *International Conference on Machine Learning*, 511–520. PMLR.
- Bello, I. 2021. LambdaNetworks: Modeling long-range Interactions without Attention. In *International Conference on Learning Representations*.
- Bengio, E.; Bacon, P.-L.; Pineau, J.; and Precup, D. 2015. Conditional computation in neural networks for faster models. *arXiv preprint arXiv:1511.06297*.
- Bengio, Y.; Léonard, N.; and Courville, A. 2013. Estimating or propagating gradients through stochastic neurons for conditional computation. *arXiv preprint arXiv:1308.3432*.
- Benny, Y.; Pekar, N.; and Wolf, L. 2021. Scale-localized abstract reasoning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 12557–12565.
- Bongard, M. M. 1968. The recognition problem. Technical report, Foreign Technology Div Wright-Patterson AFB Ohio.
- Cao, Z.; Long, M.; Wang, J.; and Jordan, M. I. 2018. Partial transfer learning with selective adversarial networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2724–2732.
- Caruana, R. 1997. Multitask learning. *Machine learning*, 28(1): 41–75.
- Dai, J.; Qi, H.; Xiong, Y.; Li, Y.; Zhang, G.; Hu, H.; and Wei, Y. 2017. Deformable convolutional networks. In *Proceedings of the IEEE international conference on computer vision*, 764–773.
- Fleuret, F.; Li, T.; Dubout, C.; Wampler, E. K.; Yantis, S.; and Geman, D. 2011. Comparing machines and humans on a visual categorization test. *Proceedings of the National Academy of Sciences*, 108(43): 17621–17625.
- Gao, H.; Zhu, X.; Lin, S.; and Dai, J. 2020. Deformable Kernels: Adapting Effective Receptive Fields for Object Deformation. In *International Conference on Learning Representations*.
- Han, Y.; Huang, G.; Song, S.; Yang, L.; Wang, H.; and Wang, Y. 2021. Dynamic neural networks: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(11): 7436–7456.
- Hendrycks, D.; and Gimpel, K. 2016. Gaussian error linear units (gelus). *arXiv preprint arXiv:1606.08415*.
- Hernández-Orallo, J.; Martínez-Plumed, F.; Schmid, U.; Siebers, M.; and Dowe, D. L. 2016. Computer models solving intelligence test problems: Progress and implications. *Artificial Intelligence*, 230: 74–107.
- 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*.
- Hochreiter, S.; and Schmidhuber, J. 1997. Long short-term memory. *Neural computation*, 9(8): 1735–1780.
- Hofstadter, D. R. 1995. *Fluid concepts and creative analogies: Computer models of the fundamental mechanisms of thought*. Basic books.
- Hoshen, D.; and Werman, M. 2017. IQ of neural networks. *arXiv preprint arXiv:1710.01692*.
- Hu, S.; Ma, Y.; Liu, X.; Wei, Y.; and Bai, S. 2021. Stratified Rule-Aware Network for Abstract Visual Reasoning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, 1567–1574.
- Ioffe, S.; and Szegedy, C. 2015. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *International Conference on Machine Learning*, 448–456. PMLR.
- Jia, X.; De Brabandere, B.; Tuytelaars, T.; and Gool, L. V. 2016. Dynamic filter networks. *Advances in neural information processing systems*, 29.
- Kingma, D. P.; and Ba, J. 2014. Adam: A method for stochastic optimization. *International Conference on Learning Representations*.
- Małkiński, M.; and Mańdziuk, J. 2022. Deep Learning Methods for Abstract Visual Reasoning: A Survey on Raven’s Progressive Matrices. *arXiv preprint arXiv:2201.12382*.
- Małkiński, M.; and Mańdziuk, J. 2023a. One Self-Configurable Model to Solve Many Abstract Visual Reasoning Problems. *arXiv preprint arXiv:2312.09997*.
- Małkiński, M.; and Mańdziuk, J. 2023b. A Review of Emerging Research Directions in Abstract Visual Reasoning. *Information Fusion*, 91: 713–736.
- Małkiński, M.; and Mańdziuk, J. 2024. Multi-Label Contrastive Learning for Abstract Visual Reasoning. *IEEE Transactions on Neural Networks and Learning Systems*, 35(2): 1941–1953.
- Mańdziuk, J.; and Żychowski, A. 2019. DeepIQ: A Human-Inspired AI System for Solving IQ Test Problems. In *2019 International Joint Conference on Neural Networks*, 1–8. IEEE.
- Nair, V.; and Hinton, G. E. 2010. Rectified linear units improve restricted boltzmann machines. In *Proceedings of the 27th International Conference on International Conference on Machine Learning*, 807–814. Omnipress.
- Nie, W.; Yu, Z.; Mao, L.; Patel, A. B.; Zhu, Y.; and Anandkumar, A. 2020. Bongard-logo: A new benchmark for human-level concept learning and reasoning. *Advances in Neural Information Processing Systems*, 33: 16468–16480.

- Pan, S. J.; and Yang, Q. 2010. A survey on transfer learning. *IEEE Transactions on knowledge and data engineering*, 22(10): 1345–1359.
- Raven, J. C. 1936. Mental tests used in genetic studies: The performance of related individuals on tests mainly educative and mainly reproductive. *Master's thesis, University of London*.
- Raven, J. C.; and Court, J. H. 1998. *Raven's progressive matrices and vocabulary scales*. Oxford psychologists Press Oxford, England.
- Santoro, A.; Raposo, D.; Barrett, D. G.; Malinowski, M.; Pascanu, R.; Battaglia, P.; and Lillicrap, T. 2017. A simple neural network module for relational reasoning. *Advances in neural information processing systems*, 30: 4967–4976.
- Schmidhuber, J. 1992. Learning to control fast-weight memories: An alternative to dynamic recurrent networks. *Neural Computation*, 4(1): 131–139.
- Shazeer, N.; Mirhoseini, A.; Maziarz, K.; Davis, A.; Le, Q.; Hinton, G.; and Dean, J. 2017. Outrageously Large Neural Networks: The Sparsely-Gated Mixture-of-Experts Layer. In *International Conference on Learning Representations*.
- Spratley, S.; Ehinger, K.; and Miller, T. 2020. A closer look at generalisation in raven. In *European Conference on Computer Vision*, 601–616. Springer.
- Tomaszewska, P.; Żychowski, A.; and Mańdziuk, J. 2022. Duel-based Deep Learning system for solving IQ tests. In *International Conference on Artificial Intelligence and Statistics*, 10483–10492. PMLR.
- Wang, D.; Jamnik, M.; and Lio, P. 2020. Abstract diagrammatic reasoning with multiplex graph networks. In *International Conference on Learning Representations*.
- Wu, Y.; Dong, H.; Grosse, R.; and Ba, J. 2020. The Scattering Compositional Learner: Discovering Objects, Attributes, Relationships in Analogical Reasoning. *arXiv preprint arXiv:2007.04212*.
- 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/CVF Conference on Computer Vision and Pattern Recognition*, 5317–5327.
- Zhang, C.; Jia, B.; Gao, F.; Zhu, Y.; Lu, H.; and Zhu, S.-C. 2019b. Learning perceptual inference by contrasting. *Advances in neural information processing systems*, 32: 1075–1087.
- Zheng, K.; Zha, Z.-J.; and Wei, W. 2019. Abstract reasoning with distracting features. In *Advances in Neural Information Processing Systems*, 5842–5853.
- Zhu, X.; Hu, H.; Lin, S.; and Dai, J. 2019. Deformable convnets v2: More deformable, better results. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 9308–9316.
- Zhuo, T.; and Kankanhalli, M. 2021. Effective Abstract Reasoning with Dual-Contrast Network. In *International Conference on Learning Representations*.