
Probabilistic Abduction for Visual Abstract Reasoning via Learning Rules in Vector-symbolic Architectures

Michael Hersche^{1,2}, Francesco di Stefano^{1,2*}, Thomas Hofmann²,
Abu Sebastian¹, Abbas Rahimi^{1†}
¹IBM Research – Zurich, ²ETH Zurich

Abstract

Abstract reasoning is a cornerstone of human intelligence, and replicating it with artificial intelligence (AI) presents an ongoing challenge. This study focuses on efficiently solving Raven’s progressive matrices (RPM), a visual test for assessing abstract reasoning abilities, by using distributed computation and operators provided by vector-symbolic architectures (VSA). Instead of hard-coding the rule formulations associated with RPMs, our approach can learn the VSA rule formulations (hence the name Learn-VRF) with just one pass through the training data. Yet, our approach, with compact parameters, remains transparent and interpretable. Learn-VRF yields accurate predictions on I-RAVEN’s in-distribution data, and exhibits strong out-of-distribution capabilities concerning unseen attribute-rule pairs, significantly outperforming pure connectionist baselines including large language models. Our code is available at <https://github.com/IBM/learn-vector-symbolic-architectures-rule-formulations>.

1 Introduction

The Raven’s progressive matrices (RPM) [1] test assesses human’s fluid intelligence and abstract reasoning [2, 3]. It is a non-verbal test that involves perceiving pattern continuation, elemental abstraction, and finding relations between abstract elements based on underlying rules. Each RPM test is an analogy problem presented as a 3×3 pictorial matrix of context panels. Every panel in the matrix is filled with several geometric objects based on a certain rule, except the last panel, which is left blank. The task is to complete the missing panel by picking the correct answer from a set of candidate answer panels that matches the implicit rule (see Fig. 1).

Recently, RPM has become a widely used task for effectively testing AI capabilities in abstract reasoning, making analogies, and dealing with out-of-distribution (OOD) data [4–8]. Large language models (LLMs) have been recently shown zero-shot learning capability in solving RPMs when perceptual visual information is parsed externally and provided as symbolic inputs to LLMs [9, 10]. On the other hand, neuro-symbolic AI approaches combine subsymbolic perception with various forms of symbolic reasoning leading to state-of-the-art (SOTA) performance in visual [11–14], natural language [15], causal [16], mathematical [17], and RPM analogical [18–20] reasoning tasks. Specifically, for solving the RPM tests, as an upgrade of pure logical reasoning, a neuro-symbolic approach [18] implemented a *probabilistic* abduction reasoning. Abduction refers to the process of selectively inferring facts that give the best explanation to the perceptual observations based on prior background knowledge which is represented in a symbolic form [21]. The probabilistic abduction reasoning allowed for perceptual uncertainties in the symbol grounding which is particularly useful when connecting a trainable perception module to the reasoning module for the end-to-end training.

However, the probabilistic abduction reasoning incurs exhaustive symbolic searches. For instance, it requires to search over the symbols describing all possible rule realizations that could govern the

*Research conducted at IBM Research – Zurich.

†Corresponding author: abr@zurich.ibm.com

RPM tests to be able to abduce the probability distribution of the rules [18]. A viable option to reduce such search complexity is to use distributed representations provided by vector-symbolic architectures (VSA) [22–25], which is built on a rich algebra supporting well-defined operations, including multiplicative binding, unbinding, additive superposition, permutations, and similarity search that allow performing symbolic computations on top of high-dimensional distributed representations. Consequently, a neuro-vector symbolic architecture (NVSA) [20] was proposed to perform *tractable* probabilistic abductive reasoning with distributed VSA representations and operators which offer computation-in-superposition leading to two orders of magnitude faster inference. Nevertheless, like other neuro-symbolic approaches, NVSA had to carefully choose suitable operators and operands to formulate the individual RPM rules. While NVSA allows the incorporation of domain knowledge into the model, it requires complete knowledge of the problem.

We propose Learn-VRF, a novel probabilistic abduction reasoning approach that *learns* VSA rule formulations from examples by solving a convex optimization problem. Besides its higher flexibility, better generality, and one-pass learning, Learn-VRF maintains the advantages of neuro-symbolic approaches, including transparency and interpretability. Learn-VRF transparently operates in the rule space by abstracting away from the surface statistical patterns. On the I-RAVEN dataset [5], with extracted visual attribute labels, Learn-VRF achieves competitive accuracy on in-distribution data, and notably outperforms LLMs [10] and neural MLP approaches on OOD tests while having orders of magnitude lower trainable parameters.

2 Background

Vector-symbolic Architectures (VSA) VSA manipulates high-dimensional distributed representations by dimensionality-preserving operations, as opposed to tensor product representations which yield higher-order tensors [15, 17, 26]. Many VSA variants exist with different representations and operators (see [27] for a review). In this work, we use binary sparse block codes (SBC) [28], which induce a local blockwise structure that exhibits ideal variable binding properties [29]. In binary SBC, the binding is defined as the blockwise circular convolution (\circledast), and the approximate unbinding is the blockwise circular correlation (\circledcirc). The VSA binding has properties analogous to multiplication in the real number domain, including commutativity, associativity, and the existence of a neutral element (\mathbf{e}). The cosine similarity ($\text{cossim}(\cdot, \cdot)$) measures the similarity between two vectors.

I-RAVEN Dataset The I-RAVEN dataset [5] provides RPM tests with unbiased candidate sets. Fig. 1 shows an I-RAVEN example test. Each RPM test consists of nine context panels, arranged in a 3x3 matrix, and eight candidate panels. The panels contain objects that are arranged according to one of seven different constellations. The object’s attributes (color, size, shape, number, position) are governed by different underlying rules: constant, arithmetic, progression, or distribute three.

3 Methods

This section presents Learn-VRF, which learns predictive VSA rule formulations based on examples, depicted in Fig. 1. First, the attribute values (provided by I-RAVEN metadata) are mapped to one-hot probability mass functions (PMFs), and translated to a VSA vector. Learn-VRF learns R -many rules, which are shared across the attributes. Each rule generates a prediction with a confidence value. The predicted representation of the empty panel ($\hat{\mathbf{a}}^{(3,3)}$) is either based on selecting a prediction of a rule with high confidence (sampling), or a weighted combination of all the rules. Finally, the candidate panel with the largest sum of cosine similarities to the predicted attribute vectors is the answer.

Translating PMFs to VSA This paper focuses exclusively on the RPM’s reasoning aspect; thus, we directly extract the attribute labels provided by I-RAVEN metadata, as practiced for LLMs [9, 10]. The attributes of each object in a panel are mapped to attribute one-hot PMFs, which are further combined, yielding panel-level PMFs describing the attribute’s distribution of a panel [18]. Each attribute yields eight PMFs (one per context panel), represented by $\mathbf{P}^c := (\mathbf{p}^{(1,1)}, \mathbf{p}^{(1,2)}, \dots, \mathbf{p}^{(3,2)})$. Next, we translate the panel-level PMFs to high-dimensional VSA representations:

$$\mathbf{a}^{(i,j)} = \sum_{k=1}^N \mathbf{p}^{(i,j)}[k] \cdot \mathbf{b}_k, \quad (1)$$

where N is the dimensionality of the panel-level PMF space. Similarly to NVSA [20], we use a binary SBC dictionary $\{\mathbf{b}_k\}_{k=1}^N$, where each basis vector (\mathbf{b}_i) has a dimension of 1024 and four

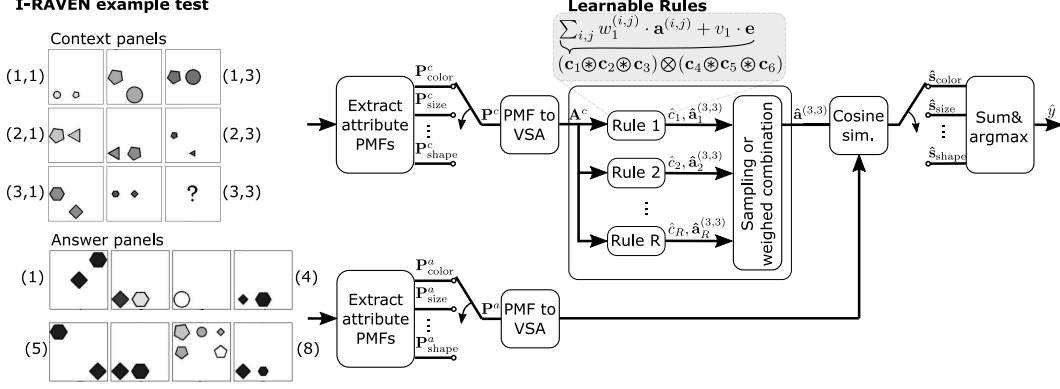


Figure 1: Learn-VRF for solving Raven’s progressive matrices.

Table 1: Translation of RPM rules into VSA rule formulations.

Rule	VSA formulation
Constant	$\hat{\mathbf{a}}^{(3,3)} = \mathbf{a}^{(3,2)} = \mathbf{a}^{(3,1)}$
Arithmetic plus	$\hat{\mathbf{a}}^{(3,3)} = \mathbf{a}^{(3,1)} \otimes \mathbf{a}^{(3,2)}$
Progression positive	$\hat{\mathbf{a}}^{(3,3)} = \mathbf{a}^{(3,2)} \otimes (\mathbf{a}^{(3,2)} \otimes \mathbf{a}^{(3,1)})$
Distribute three	$\hat{\mathbf{a}}^{(3,3)} = (\mathbf{a}^{(1,1)} \otimes \mathbf{a}^{(1,2)} \otimes \mathbf{a}^{(1,3)}) \otimes (\mathbf{a}^{(3,1)} \otimes \mathbf{a}^{(3,2)})$

blocks. When combined with fractional power encoding [24], this basis permits the representation of continuous attributes (e.g., color or size). The VSA representation of different context panels are represented by $\mathbf{A}^c := (\mathbf{a}^{(1,1)}, \mathbf{a}^{(1,2)}, \dots, \mathbf{a}^{(3,2)})$.

Learning VSA Rule Formulations (Learn-VRF) VSA allows us to predict the representation of the empty panel with a rule-specific combination of operands (i.e., a selection of context panels) and operations (binding and unbinding). Table 1 summarizes possible VSA formulations for each RPM rule. All the rules are represented as a sequence of binding and unbinding operations, which can be generalized to:

$$\hat{\mathbf{a}}^{(3,3)} = (\mathbf{c}_1 \otimes \mathbf{c}_2 \otimes \mathbf{c}_3) \otimes (\mathbf{c}_4 \otimes \mathbf{c}_5 \otimes \mathbf{c}_6), \quad (2)$$

where \mathbf{c}_i represents a context panel ($\mathbf{a}^{(i,j)}$) or the identity (\mathbf{e}). Eq. (2) can be interpreted as a fractional form, akin to a sequence of multiplicative and divisive operations among real numbers.

While NVSA hard-coded the rules as in Table 1, we *learn* the rule formulations by finding, for each element in the fractional form, a convex combination over the VSA vectors of the context panels’ attributes, augmented with the neutral element:

$$\mathbf{c}_k = \sum_{\text{panels } (i,j)} w_k^{(i,j)} \cdot \mathbf{a}^{(i,j)} + v_k \cdot \mathbf{e}, \quad (3)$$

where the following constraints apply to the weights:

$$\sum_{\text{panels } (i,j)} w_k^{(i,j)} + v_k = 1, \quad 0 \leq w_k^{(i,j)} \leq 1 \forall i, j, \quad 0 \leq v_k \leq 1 \forall k. \quad (4)$$

Here, $w_k^{(i,j)}$ and v_k are the trainable weights of the convex combination. A rule in the VSA space can manifest in multiple functionally equivalent representations within the proposed form. As a result, the optimization problem will have multiple global minima corresponding to different algebraic representations of the same rule.

Learn-VRF learns R -many rule formulations. Allowing more rules than actually embedded in I-RAVEN (i.e., $R > R^*$) yields Learn-VRF to come up with functionally equivalent duplications. Thus, by providing an overcomplete rule space, Learn-VRF can still operate and does not need to have the exact knowledge about the number of rules (R^*). In contrast, a related learnable neuro-symbolic approach [30] performed a grid search for the optimal number of attributes and rules in SATNet. For simplicity, we set the number of rules equal to the rules in I-RAVEN ($R = R^*$).

Confidence value computation Each rule (r) generates a VSA vector ($\hat{\mathbf{a}}_r^{(3,3)}$) accompanied by a confidence value (\hat{c}_r). The confidence is computed by deploying three formulations per rule, which predict all panels of the third column. At inference, the confidence value is calculated as the sum of the cosine similarities between the predicted VSA vectors of the first and second row with their respective ground-truth vector:

$$c_{r,\text{test}} = \text{cossim} \left(\mathbf{a}^{(1,3)}, \hat{\mathbf{a}}_r^{(1,3)} \right) + \text{cossim} \left(\mathbf{a}^{(2,3)}, \hat{\mathbf{a}}_r^{(2,3)} \right). \quad (5)$$

During training, the ground-truth of the last row is accessible, enabling the inclusion of the cosine similarity between the predicted panel and the ground-truth answer (\mathbf{a}_y) into the confidence score:

$$c_{r,\text{train}} = \text{cossim} \left(\mathbf{a}^{(1,3)}, \hat{\mathbf{a}}_r^{(1,3)} \right) + \text{cossim} \left(\mathbf{a}^{(2,3)}, \hat{\mathbf{a}}_r^{(2,3)} \right) + \text{cossim} \left(\mathbf{a}_y, \hat{\mathbf{a}}_r^{(3,3)} \right). \quad (6)$$

A softmax with a temperature parameter is applied to the confidence scores, producing a rule distribution.

Final estimation using sampling or weighted combination We distinguish between two variants to determine the final estimation of the empty panel. One variant is to *sample* a rule (\hat{r}) that has a high confidence value. During training, we sample from the distribution of the confidence values, whereas at inference, we select the rule with the maximizing confidence value. The estimated panel is then defined as:

$$\hat{\mathbf{a}}_{\text{sample}}^{(3,3)} = \hat{\mathbf{a}}_{\hat{r}}^{(3,3)}. \quad (7)$$

Alternatively, we propose to create a convex *weighted combination* of all the rules using the confidence values as weights:

$$\hat{\mathbf{a}}_{\text{weighted comb.}}^{(3,3)} = \sum_{r=1}^R c_r \cdot \hat{\mathbf{a}}_r^{(3,3)}. \quad (8)$$

Training loss The model weights are updated through stochastic gradient descent (SGD), reducing the inverse cosine similarity of the selected rule as the loss:

$$\mathcal{L} = 1 - \text{cossim} \left(\mathbf{a}^{(1,3)}, \hat{\mathbf{a}}^{(1,3)} \right) - \text{cossim} \left(\mathbf{a}^{(2,3)}, \hat{\mathbf{a}}^{(2,3)} \right) - \text{cossim} \left(\mathbf{a}_y, \hat{\mathbf{a}}^{(3,3)} \right), \quad (9)$$

where $\hat{\mathbf{a}}^{(i,3)}$ results either from the sampling or the weighted combination. The terms derived from the first and second rows refine the rule selection mechanism, whereas the term associated with the ground-truth answer adjusts the weights responsible for predicting the empty panel’s VSA vector.

4 Experimental Results

I-RAVEN: in-distribution Table 2 compares Learn-VRF’s accuracy with the SOTA deep neural nets SCL [31], probabilistic abductive neuro-symbolic approaches such as PrAE [18] and NVSA [20], and an LLM based on GPT-3 [10] on in-distribution (ID) I-RAVEN tasks. As an additional baseline, we trained a model that predicts the PMF for each attribute using a separate MLP. App. A describes the MLP baseline and App. B our experimental setup.

Learn-VRF achieves an average accuracy of 81.3% when using a sampling-based rule selection and 84.1% with the weighted rule combination, while only requiring 5 k trainable parameters. The superior performance of the weighted combination over the sampling approach may stem from its uniform update mechanism, which adjusts all the rules (including the “correct” rule) instead of only the sampled one. Learn-VRF remains accurate when training it with a single pass over the training data, achieving 80.0% and 81.0% using the sampling and weighted combination, respectively. Learn-VRF generates accurate predictions in constellations with no position attribute (i.e., center, left-right, up-down, out-in center). In these constellations, it outperforms GPT-3 and PrAE, and is on par with NVSA and SCL. Instead, the 2x2 grid, 3x3 grid, and out-in grid require the position attribute that employs logical rules on its bit-code vector, which the VSA rule formulations cannot support (see App. D). While NVSA circumvented this limitation by treating some rules of the position attribute in the original PMF domain, Learn-VRF uniformly implements all rules in the VSA space, at some degree of accuracy degradation. The MLP baseline can improve the accuracy in the position-related constellations, but it requires $60\times$ more parameters (300 k vs. 5 k) while also sacrificing transparency and interpretability.

Table 2: ID accuracy (%) on I-RAVEN. Selective approaches require the answer panels as model inputs, while predictive approaches compare their prediction against all possible answers.

Method	Predictive/ Selective	#Parms	#Epochs	I-RAVEN accuracy (%)							
				Avg.	C	2x2	3x3	L-R	U-D	O-IC	O-IG
SCL [31]	Selective	961 k	200	84.3	99.9	68.9	43.0	98.5	99.1	97.7	82.6
GPT-3 [10]	Selective	175 B	n/a	86.5	86.4	83.2	81.8	83.4	84.6	92.8	93.0
PrAE ¹ [18]	Predictive	-	-	71.1	83.8	82.9	47.4	94.8	94.8	56.6	37.4
NVSA ¹ [20]	Predictive	-	-	88.1	99.8	96.2	54.3	100	99.9	99.6	67.1
MLP baseline	Predictive	300 k	50	87.1	97.6	87.1	61.8	99.4	99.4	98.7	65.6
Learn-VRF (sampling)	Predictive	5 k	50	81.3	98.7	62.9	48.4	98.7	99.9	99.9	60.8
Learn-VRF (weighted comb.)	Predictive	5 k	50	84.1	99.1	73.5	50.9	99.9	99.9	100	65.2
Learn-VRF (sampling)	Predictive	5 k	1	80.0	99.1	52.3	51.1	98.7	99.9	99.9	59.5
Learn-VRF (weighted comb.)	Predictive	5 k	1	81.0	99.1	53.1	50.9	99.9	100	100	63.7

¹ No trainable parameters in the hard-coded reasoning; hence, learnable parameters and epochs are zero for reasoning module.

I-RAVEN: out-of-distribution We test GPT-3, the MLP baseline, and the Learn-VRF on unseen rule-attribute pairs in the center constellation, validating their OOD generalization capabilities (see App. B). In this OOD task, Learn-VRF has been trained for 50 epochs. As shown in Table 3, the MLP baseline faces challenges in OOD even though it was accurate on ID. Moreover, the GPT-3 shows a low accuracy, particularly on the arithmetic rule, despite the potential for data contamination [32], and our additional efforts to perform in-context learning (see App. E). In contrast, Learn-VRF generalizes to OOD, achieving almost perfect accuracy on most rule-attribute pairs thanks to Learn-VRF’s rule sharing, which is enabled by the uniform, expressive PMF representation in the VSA space using fractional power encoding.

Table 3: OOD accuracy (%) on unseen rule-attribute pairs on I-RAVEN in the center constellation.

	Type			Size			Color				
	Const.	Progr.	Dist.3	Const.	Progr.	Dist.3	Arith.	Const.	Progr.	Dist.3	Arith.
GPT-3 zero-shot	88.5	86.0	88.6	93.6	93.2	92.6	71.6	94.2	94.7	94.3	65.8
MLP baseline	14.8	14.9	30.2	22.8	75.0	47.0	46.6	56.3	60.5	44.4	48.9
Learn-VRF (weighted comb.)	100	100	99.7	100	100	99.8	99.8	100	98.8	100	100

Interpretation of learned rule formulations Learn-VRF provides transparency and interpretability into the learned rules. We can derive the most dominant formula by selecting the VSA vector that exhibited the highest weight for each term within the Learn-VRF model’s internal representations. Table 4 shows the learned rule formulations for the center constellation. Analyzing these rules and performing straightforward algebraic simplifications, we discern the equivalences shown in the second column. App. C provides more details on the interpretation of the learned rule formulations.

Table 4: Rule formulations learned by Learn-VRF with the respective rule interpretation.

Extracted Learned Rules	Corresponding RPM Rule
$\hat{\mathbf{a}}^{(3,3)} = (\mathbf{a}^{(3,2)} \otimes \mathbf{a}^{(1,1)} \otimes \mathbf{a}^{(3,1)}) \otimes (\mathbf{e} \otimes \mathbf{a}^{(1,2)} \otimes \mathbf{a}^{(3,2)})$	dist. three right & constant
$\hat{\mathbf{a}}^{(3,3)} = (\mathbf{a}^{(3,2)} \otimes \mathbf{a}^{(1,2)} \otimes \mathbf{a}^{(2,3)}) \otimes (\mathbf{a}^{(2,3)} \otimes \mathbf{e} \otimes \mathbf{a}^{(1,1)})$	dist. three left & constant
$\hat{\mathbf{a}}^{(3,3)} = (\mathbf{a}^{(3,2)} \otimes \mathbf{a}^{(3,1)} \otimes \mathbf{a}^{(3,1)}) \otimes (\mathbf{e} \otimes \mathbf{a}^{(3,1)} \otimes \mathbf{e})$	arithmetic plus
$\hat{\mathbf{a}}^{(3,3)} = (\mathbf{a}^{(3,2)} \otimes \mathbf{a}^{(3,2)} \otimes \mathbf{a}^{(3,1)}) \otimes (\mathbf{a}^{(3,2)} \otimes \mathbf{a}^{(3,2)} \otimes \mathbf{a}^{(3,2)})$	arithmetic minus
$\hat{\mathbf{a}}^{(3,3)} = (\mathbf{a}^{(1,3)} \otimes \mathbf{a}^{(3,1)} \otimes \mathbf{a}^{(3,2)}) \otimes (\mathbf{e} \otimes \mathbf{a}^{(3,2)} \otimes \mathbf{a}^{(1,1)})$	progression

5 Conclusion

We proposed Learn-VRF, a novel VSA-based approach for probabilistic abduction which can learn rule formulations. The model’s structure inherently provides transparency: by identifying the index with the highest probability mass for each term in the weight distribution, we could reconstruct the rule formula chosen by the Learn-VRF. Moreover, Learn-VRF features a low parameter count, single-pass learning, and OOD generalization with respect to unseen attribute-rule combinations.

Acknowledgement

This work is supported by the Swiss National Science foundation (SNF), grant 200800.

References

- [1] J. Raven, J. Court, and J. Raven, *Raven’s progressive matrices*. Oxford Psychologists Press, 1938.
- [2] P. A. Carpenter, M. A. Just, and P. Shell, “What one intelligence test measures: a theoretical account of the processing in the Raven progressive matrices test,” *Psychological review*, 1990.
- [3] W. B. Bilker, J. A. Hansen, C. M. Brensinger, J. Richard, R. E. Gur, and R. C. Gur, “Development of abbreviated nine-item forms of the Raven’s standard progressive matrices test,” *Assessment*, 2012.
- [4] C. Zhang, F. Gao, B. Jia, Y. Zhu, and S.-C. Zhu, “Raven: A dataset for relational and analogical visual reasoning,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019.
- [5] S. Hu, Y. Ma, X. Liu, Y. Wei, and S. Bai, “Stratified rule-aware network for abstract visual reasoning,” in *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, 2021.
- [6] Y. Benny, N. Pekar, and L. Wolf, “Scale-localized abstract reasoning,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2021.
- [7] M. Mitchell, “Abstraction and analogy-making in artificial intelligence,” *Annals of the New York Academy of Sciences*, vol. 1505, no. 1, pp. 79–101, 2021.
- [8] Małkiński, Mikołaj and Mańdziuk, Jacek, “Deep learning methods for abstract visual reasoning: A survey on Raven’s progressive matrices,” *arXiv preprint arXiv:2201.12382*, 2022.
- [9] T. Webb, K. J. Holyoak, and H. Lu, “Emergent analogical reasoning in large language models,” *Nature Human Behaviour*, vol. 7, no. 9, pp. 1526–1541, Sep 2023.
- [10] X. Hu, S. Storks, R. L. Lewis, and J. Chai, “In-context analogical reasoning with pre-trained language models,” in *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Long Paper)*, 2023.
- [11] K. Yi, J. Wu, C. Gan, A. Torralba, P. Kohli, and J. Tenenbaum, “Neural-symbolic VQA: Disentangling reasoning from vision and language understanding,” in *Advances in Neural Information Processing Systems (NeurIPS)*, 2018.
- [12] J. Mao, C. Gan, P. Kohli, J. B. Tenenbaum, and J. Wu, “The neuro-symbolic concept learner: Interpreting scenes, words, and sentences from natural supervision,” in *International Conference on Learning Representations (ICLR)*, 2019.
- [13] C. Han, J. Mao, C. Gan, J. Tenenbaum, and J. Wu, “Visual concept-metaconcept learning,” in *Advances in Neural Information Processing Systems (NeurIPS)*, 2019.
- [14] L. Mei, J. Mao, Z. Wang, C. Gan, and J. B. Tenenbaum, “FALCON: Fast visual concept learning by integrating images, linguistic descriptions, and conceptual relations,” in *International Conference on Learning Representations (ICLR)*, 2022.
- [15] I. Schlag and J. Schmidhuber, “Learning to reason with third-order tensor products,” in *Advances in Neural Information Processing Systems (NeurIPS)*, 2018.
- [16] K. Yi, C. Gan, Y. Li, P. Kohli, J. Wu, A. Torralba, and J. B. Tenenbaum, “CLEVRER: Collision events for video representation and reasoning,” in *International Conference on Learning Representations (ICLR)*, 2020.
- [17] I. Schlag, P. Smolensky, R. Fernandez, N. Jovic, J. Schmidhuber, and J. Gao, “Enhancing the transformer with explicit relational encoding for math problem solving,” *arXiv preprint arXiv:1910.06611*, 2019.
- [18] C. Zhang, B. Jia, S.-C. Zhu, and Y. Zhu, “Abstract spatial-temporal reasoning via probabilistic abduction and execution,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021.
- [19] V. Shah, A. Sharma, G. Shroff, L. Vig, T. Dash, and A. Srinivasan, “Knowledge-based analogical reasoning in neuro-symbolic latent spaces,” in *Proceedings of the 16th International Workshop on Neural-Symbolic Learning and Reasoning (NeSy)*, 2022.

- [20] M. Hersche, M. Zeqiri, L. Benini, A. Sebastian, and A. Rahimi, “A neuro-vector-symbolic architecture for solving Raven’s progressive matrices,” *Nature Machine Intelligence*, vol. 5, no. 4, pp. 363–375, 2023.
- [21] L. Magnani, *Abductive Cognition: The Epistemological and Eco-Cognitive Dimensions of Hypothetical Reasoning*. Springer Berlin, Heidelberg, 2009.
- [22] R. W. Gayler, “Vector symbolic architectures answer Jackendoff’s challenges for cognitive neuroscience,” in *Joint International Conference on Cognitive Science (ICCS/ASCS)*, 2003.
- [23] T. A. Plate, “Holographic reduced representations,” *IEEE Transactions on Neural Networks*, vol. 6, no. 3, pp. 623–641, 1995.
- [24] —, *Holographic Reduced Representations: Distributed Representation for Cognitive Structures*. Center for the Study of Language and Information, Stanford, 2003.
- [25] P. Kanerva, “Hyperdimensional computing: An introduction to computing in distributed representation with high-dimensional random vectors,” *Cognitive Computation*, vol. 1, no. 2, pp. 139–159, 2009.
- [26] P. Smolensky, R. T. McCoy, R. Fernandez, M. Goldrick, and J. Gao, “Neurocompositional computing: From the central paradox of cognition to a new generation of ai systems,” *AI Magazine*, vol. 43, no. 3, pp. 308–322, 2022.
- [27] D. Kleyko, D. A. Rachkovskij, E. Osipov, and A. Rahimi, “A survey on hyperdimensional computing aka vector symbolic architectures, part I: models and data transformations,” *ACM Comput. Surv.*, 2022.
- [28] M. Laiho, J. H. Poikonen, P. Kanerva, and E. Lehtonen, “High-dimensional computing with sparse vectors,” in *2015 IEEE Biomedical Circuits and Systems Conference (BioCAS)*, 2015.
- [29] E. P. Frady, D. Kleyko, and F. T. Sommer, “Variable binding for sparse distributed representations: Theory and applications,” *IEEE Transactions on Neural Networks and Learning Systems*, 2021.
- [30] S. Topan, D. Rolnick, and X. Si, “Techniques for symbol grounding with SATNet,” *Advances in Neural Information Processing Systems (NeurIPS)*, vol. 34, pp. 20 733–20 744, 2021.
- [31] Y. Wu, H. Dong, R. Grosse, and J. Ba, “The scattering compositional learner: Discovering objects, attributes, relationships in analogical reasoning,” *arXiv preprint arXiv:2007.04212*, 2020.
- [32] A. Saparov and H. He, “Language models are greedy reasoners: A systematic formal analysis of chain-of-thought,” in *The Eleventh International Conference on Learning Representations (ICLR)*, 2023.
- [33] T. Brown, B. Mann, N. Ryder, M. Subbiah, J. D. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell *et al.*, “Language models are few-shot learners,” *Advances in Neural Information Processing Systems (NeurIPS)*, vol. 33, pp. 1877–1901, 2020.
- [34] L. Qiu, P. Shaw, P. Pasupat, T. Shi, J. Herzig, E. Pitler, F. Sha, and K. Toutanova, “Evaluating the impact of model scale for compositional generalization in semantic parsing,” in *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, 2022, pp. 9157–9179.

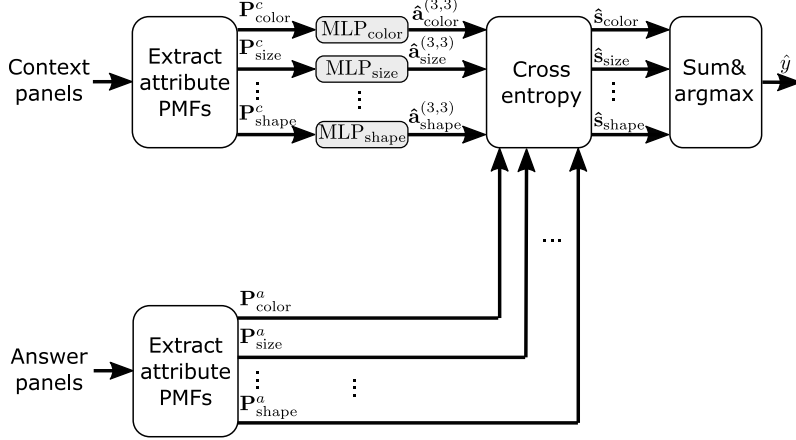


Figure A1: MLP baseline which uses one MLP per each attribute.

A MLP Baseline

This baseline, illustrated in Fig. A1, is built upon the use of A instances of MLPs, where A denotes the total number of attributes. The MLP directly models the transformations of the attribute PMFs. For every attribute, there is a dedicated MLP, ensuring that each attribute is modeled independently. Since the PMFs of the attributes have different dimensionality (i.e., different number of possible values), the MLPs cannot be shared across the attributes, while it is possible with our Learn-VRF thanks to the use of VSA-based fractional power encoding that allows to map any PMF vectors with arbitrary sizes to the fixed-width, yet expressive, distributed representation. Thanks to this *uniform* interface, Learn-VRF can be plugged-and-played from one attribute to another.

The first step of MLP baseline involves processing the derived attribute values. These values are converted into one-hot PMFs. Once translated to the PMF format, these representations are processed by a 3-layer MLP with 1024 hidden dimensions in each layer. The output from this MLP is another PMF spanning the attribute values. For each attribute, its associated PMF vector is compared against the PMF vectors of all candidate panels using the cross entropy. For every attribute, we obtain eight scores corresponding to the eight candidate panels. Lastly, we compute the average for each attribute, resulting in eight overall candidate scores. The candidate panel with the highest score is selected.

This baseline does not make any distinction between rule selection and execution. Unlike Learn-VRF, MLPs can potentially capture non-linear relationships, such as the logical rules present in the position attribute in the 2x2 grid, 3x3 grid, and the out-in grid.

B Experimental Setup

Training setup

In our experimental setup with the MLP baselines, we used three hidden layers, each containing 1024 neurons. Layer normalization was applied between consecutive layers. We set the learning rate to $1e-4$, used a batchsize of 32, and trained the model for 50 epochs. The model with the lowest validation loss from these epochs was chosen for further evaluation on the test dataset.

In Learn-VRF, we used VSA vectors of 1024 dimensions with 4 blocks. Moreover, we used a learning rate of 0.1 and batchsize of 4. We set the number of learnable rules (R) equal to the number of rules in I-RAVEN. Finally, we used a softmax temperature of 0.01, favoring exploitation over exploration, allowing us to converge quickly.

OOD test setup

In the out-of-distribution (OOD) experiments, we evaluate whether the model can solve an unseen target attribute-rule pair (e.g., the constant rule on the type attribute) when it has been trained on the examples containing all of the attribute-rule pairs except the specific target one (e.g., the constant rule on size and color, the progression rule on all attributes, and the distribute rule on all attributes).

We generate a new training and validation set containing all examples except those with the target attribute-rule pair and a test set containing examples exclusively with the target attribute-rule pair.

In these experiments, we focus on the center single constellation because it encompasses a single panel. In contrast, configurations such as up-down, left-right, and in-out entail two center single constellations, introducing ambiguity in applying attribute-rule filtering.

The datasets are generated by filtering the existing splits I-RAVEN. The training sets encompass between 2622 and 3437 samples, the validation sets range from 841 to 1160 samples, while the test sets span from 803 to 1117 samples.

C Interpretation of Learned Rule Formulation in Learn-VRF

Within Learn-VRF, a singular rule in the VSA space can be depicted through various synonymous representations. As an example, let us examine the arithmetic plus rule for the last row:

$$\hat{\mathbf{a}}^{(3,3)} = \mathbf{a}^{(3,1)} \otimes \mathbf{a}^{(3,2)}. \quad (10)$$

This rule can be reformulated in the VSA fractional form in multiple manners, an example subset of all the possible algebraic forms of this rule is the following one:

$$\hat{\mathbf{a}}^{(3,3)} = (\mathbf{a}^{(3,1)} \otimes \mathbf{a}^{(3,2)} \otimes \mathbf{e}) \otimes (\mathbf{e} \otimes \mathbf{e} \otimes \mathbf{e}) \quad (11)$$

$$\hat{\mathbf{a}}^{(3,3)} = (\mathbf{a}^{(3,2)} \otimes \mathbf{a}^{(3,1)} \otimes \mathbf{e}) \otimes (\mathbf{e} \otimes \mathbf{e} \otimes \mathbf{e}) \quad (12)$$

$$\hat{\mathbf{a}}^{(3,3)} = (\mathbf{a}^{(3,1)} \otimes \mathbf{a}^{(3,2)} \otimes \mathbf{a}^{(3,2)}) \otimes (\mathbf{a}^{(3,2)} \otimes \mathbf{e} \otimes \mathbf{e}) \quad (13)$$

When employing a regression loss function to the resultant VSA vector from the Learn-VRF and to its related target VSA vector, the loss curve’s global minima align with diverse algebraic interpretations of the identical rule.

D Model Limitations

The primary constraint of the Learn-VRF model lies in its expressiveness. To understand this, consider that in the VSA space, the sum of two attribute values is represented by the binding of their respective VSA vectors, while the difference is depicted by the unbinding operation. Given the VSA fractional form underpinning the Learn-VRF:

$$\hat{\mathbf{a}}^{(3,3)} = (\mathbf{a}_1 \otimes \mathbf{a}_2 \otimes \mathbf{a}_3) \otimes (\mathbf{a}_4 \otimes \mathbf{a}_5 \otimes \mathbf{a}_6) \quad (14)$$

In the numeric attribute space, this translates to:

$$\hat{a}^{(3,3)} = (a_1 + a_2 + a_3) - (a_4 + a_5 + a_6) \quad (15)$$

This implies that the model is limited to representing sums and differences of attributes, precluding the possibility of modeling linear combinations. This limitation becomes evident in complex constellations like the 2x2 grid, the 3x3 grid, or the out-in grid, where the position attribute employs logical rules on its bit-code vector.

Table A1 and Table A2 show the number of rule occurrences of the number/position and type/size/color attributes, respectively. On the position attribute, the logical rules (progression and arithmetic) make up 23.9%, 26.7%, and 26.6% on the 2x2 grid, 3x3 grid, and out-in grid constellation, respectively. Since Learn-VRF cannot represent these logical rules, an upper bound for the accuracy can be determined by the share of non-logical rules (i.e., 76.1%, 73.3%, and 73.4%). Indeed, Learn-VRF (weighted combination, 50 epochs) reaches almost the upper bound on the 2x2 grid (73.5% vs. 76.1%), whereas the performances on the 3x3 grid (50.9% vs. 73.3%) and out-in grid (65.2% vs. 73.4%) are slightly behind.

E LLM Experiments

This appendix describes our additional experiments testing the OOD generalization of GPT-3 on I-RAVEN. Our implementation is based on [33], which uses a `text-davinci-002` model.

Table A1: Number of rule occurrences on the position and number attribute in the I-RAVEN test set (2000 samples). The progression and arithmetic rule on position attribute are logical rules, e.g., I-RAVEN implements the arithmetic plus rule as logical OR.

	Number/Position		Number				Position			Share of non-logical rules	Acc. Learn-VRF (weighted comb., 50 epochs)
	Const.	Progr.	Arith.	Dist.3	Progr. (logical)	Arith. (logical)	Dist.3				
2x2 grid	320	178	335	326	163	316	362	76.1%	73.5%		
3x3 grid	267	304	312	269	247	287	314	73.3%	50.9%		
O-IG (inner grid)	325	158	321	342	175	357	322	73.4%	65.2%		

Table A2: Number of rule occurrences on the type, size, and color attribute in the I-RAVEN test set (2000 samples). The sum of occurrences per attribute is 2000 for single constellations (center, 2x2 grid, 3x3 grid), and 4000 for dual constellations (L-R, U-D, O-IC, O-IG).

	Type				Size				Color			
	Const.	Progr.	Arith.	Dist.3	Const.	Progr.	Arith.	Dist.3	Const.	Progr.	Arith.	Dist.3
Center	660	665	0	675	514	485	500	501	520	494	474	512
2x2 grid	625	676	0	669	515	503	492	490	507	495	519	479
3x3 grid	686	657	0	657	466	546	464	524	494	476	506	524
L-R	1294	1306	0	1400	1031	983	972	1014	1006	1019	943	1032
U-D	1332	1364	0	1304	989	1030	994	987	983	1015	983	1019
O-IC	1278	1359	0	1363	1286	934	503	1277	2490	493	501	516
O-IG	1334	1646	0	1320	1383	709	551	1357	2493	534	516	457

The attribute labels are transformed into prompts using two different approaches. In the entangled setting (aka entity attribute naming), the different attributes of an object are represented in one entity (e.g., packed in brackets), yielding one prompt. Conversely, the attributes are fully decomposed into one prompt per attribute in the disentangled setting (aka attribute decomposition). Both approaches perform selective classification, i.e., prompting the LLM with every answer panel in the context matrix and selecting the candidate panel as the answer that yielded the highest sum of log-probabilities in the answer part. On ID data, explicit disentanglement improves LLM’s reasoning: in the center constellation, the GPT-3 achieves an accuracy of 86.4% with disentanglement and 80.8% without disentanglement [33].

We tested the GPT-3 on our OOD test set using the disentangled setting. Note that we do not have complete knowledge about the data used to train GPT-3. Hence, data contamination is possible, i.e., the model could have been trained on RPMs. As a result, there is no certainty that this test is truly OOD; creating new synthetic datasets would be needed to have full control over the OOD tests as it has been done in a synthetic question-answering dataset [32]. Table 3 shows that GPT-3 mainly faces challenges with the arithmetic rule, achieving 71.6% on the size attribute and 65.8% on color.

As an additional experiment, we tested the effect of in-context learning on the arithmetic rules. We randomly selected in-context examples that contained the arithmetic rule on an attribute different from the one under test, e.g., for the OOD test with the arithmetic on the attribute size, we showed additional examples where arithmetic was applied on color (but not size). As no in-context examples could be shown in the disentangled configuration, we switched to entangled setting. Fig. A2 shows an example. To reduce the number of prompts, we only tested on a subset of 50 samples, which still gives an indication of the trends. The results shown in Table A3 indicate that in-context learning is rather detrimental in this application. Further efforts such as prompt-tuning (e.g., the in-context selection and arrangement) may improve the LLM’s in-context learning performance [34].

Table A3: Classification accuracy (%) on a subset (50 samples) of I-RAVEN OOD test set using GPT-3 (entangled setting) using a different number of in-context examples.

#In-context samples	Size	Color
0	70.0	70.0
3	70.0	52.0
5	60.0	48.0
10	58.0	46.0

<p>let's think step by step.</p> <p>row 1: (5,0.5,40), (5,0.3,0), (5,0.1,40);</p> <p>row 2: (6,0.6,20), (6,0.4,20), (6,0.2,0);</p> <p>row 3: (3,0.6,80), (3,0.4,20), (3,0.2,60);</p> <p>row 1: (5,0.1,60), (6,0.5,30), (7,0.6,30);</p> <p>row 2: (4,0.5,60), (5,0.6,20), (6,0.1,40);</p> <p>row 3: (4,0.6,50), (5,0.1,10), (6,0.5,40);</p> <p>row 1: (6,0.5,70), (4,0.3,10), (3,0.1,80);</p> <p>row 2: (4,0.6,60), (3,0.4,10), (6,0.2,70);</p> <p>row 3: (3,0.6,60), (6,0.4,10), (4,0.2,70);</p> <p>row 1: (6,0.2,60), (5,0.1,0), (4,0.1,60);</p> <p>row 2: (6,0.3,60), (5,0.1,40), (4,0.2,20);</p> <p>row 3: (7,0.6,80), (6,0.3,80), (4,0.3,0);</p>	<p>In-context example #1: type: constant size: progression (-0.2) color: arithmetic (minus)</p> <p>In-context example #2: type: progression (+1) size: distribute three color: arithmetic (minus)</p> <p>In-context example #3: type: distribute three size: progression (-0.2) color: arithmetic (plus)</p> <p>Test example (with on answer panel): type: progression (-1) size: arithmetic (minus) color: arithmetic (minus)</p>
---	--

Figure A2: Prompt for GPT-3 experiment with three in-context examples in entangled setting. The single object in the center constellation is encoded into (type, size color). The OOD test focuses on the arithmetic rule on the attribute size. In-context examples show examples of arithmetic rule on attribute color. In this prompt, the empty panel in the text example is filled with the first candidate panel, which is not the correct one.