

A Critical Analysis of Unsupervised Learning in the Context of Raven's Progressive Matrices

Rollin Omari¹, R I Mckay¹, Tom Gedeon¹, and Kerry Taylor¹

¹Affiliation not available

April 01, 2024

A Critical Analysis of Unsupervised Learning in the Context of Raven's Progressive Matrices

Rollin Omari, R. I. (Bob) McKay, Tom Gedeon and Kerry Taylor

Abstract—This paper undertakes a critical examination of unsupervised learning within the context of Raven's Progressive Matrices (RPMs). We trace the historical trajectory of computational models for RPMs, from early rule-based approaches to modern neural networks, and we focus on the innovative work of Zhuo et al. in introducing semi-supervised learning to RPMs. Our discussion highlights the nuances of unsupervised learning, emphasising the role of noisy labels as a form of guidance, albeit with a trade-off in precision compared to traditional supervised learning. In this paper, we recognise the challenge in formalising the distinction between supervised and unsupervised learning, but we underscore the importance of precision in communication and nomenclature, especially in regards to facilitating knowledge transfer and directing future research. We hope that this contribution enhances the discourse on unsupervised learning and offers valuable insights towards the challenges and opportunities in attaining human-level reasoning capabilities in machine learning and artificial intelligence.

Index Terms—Artificial intelligence, benchmark, critical analysis, demonstration, relational reasoning, Raven's progressive matrices, semi-supervised learning, supervised learning, unsupervised learning.

I. INTRODUCTION

RAVEN'S Progressive Matrices (RPMs) have long been a staple for measuring abstract reasoning and fluid intelligence in humans [1]. As such, it is natural that they would be subsequently adapted as test domains for artificial intelligence systems such as neural networks [2].

RPMs require subjects to solve problems in the absence of physical objects or concrete phenomena, and independent of their language, reading and writing skills, and arguably their cultural background [3], [4]. As shown in Fig. 1(a), an RPM consists of several visual geometric designs with a missing piece. One has to determine the underlying logical rules in the problem matrix and select, from an answer set of eight candidate choices, the most suitable choice that satisfies these hidden rules [5].

Early computational models for solving RPMs were highly reliant on pre-determined rules. The first example was introduced by Carpenter et al. [3]. This influential model operated through a production system that utilised shared memory to detect matrix patterns, construct rules, and derive answers,




Fig. 1. Two representations of an RPM problem taken from the I-RAVEN dataset. Figure (a) illustrates the standard representation (context set above, answer set below), while figure (b) illustrates the pseudo-label contrasting representation utilised by Zhuo and colleagues [6]. To complete the problem matrix, one has to select the best choice amongst the eight panels in the answer set that follows structural and analogical relations. In this problem, the *positions* and *colours* can vary freely as long as the *number*, *shapes* and *sizes* of the objects follow their respective underlying rules.

shedding light on the role of working memory in RPM performance. In 2003, Bringsjord and Schimanski built a theorem-prover and demonstrated its ability to solve selected RPM problems encoded in first-order logic [7]. Although Bringsjord and Schimanski did not report any specific results or provide technical details on their model, their work introduced one of the earliest robotic systems to be used in solving RPMs.

In mid-2010, Lovett et al. presented a model that focussed on visual processing of RPMs that operated on spatial relationships and predefined geometric transformation rules [8]. Meanwhile McGregor et al. introduced the first fractal technique for directly operating on visual inputs of RPMs, without any need to extract propositional representations as in earlier models [9]. Later in 2010, Cirillo and Ström introduced a system inspired by Lovett et al. [8] and Carpenter et al. [3]. Similarly to [8], Cirillo and Ström's system utilised hand-drawn vector graphics to generate hierarchical propositional representations of test problems, while similarly to [3], the system selected the best-fit pattern from pre-defined patterns, derived through *a priori* inspection of the Standard Progressive Matrices [10].

In 2011, neural networks were first applied to RPMs by Rasmussen and Eliasmith [11]. Their work introduced the first spiking model designed for solving RPMs. They manually encoded input images into propositional attribute-value pairs, and the spiking neuron model identified transformations among these vectors, generalising them to induce a rule for

R. Omari, R. McKay and K. Taylor are with the School of Cybernetics and Computer Science, CECCS, The Australian National University, Canberra, ACT, 2601 Australia.

T. Gedeon is with the School of Electrical Engineering, Computer and Mathematical Science EECMS, Curtin University, Perth, WA, 6102 Australia. (e-mail: {rollin.omari, robert.mckay, kerry.taylor}@anu.edu.au, and tom.gedeon@curtin.edu.au).

the specific problem. Although the authors claimed success in solving RPM problems with their model, they did not specify the particular tests or problems addressed in their results. In 2012, Kunda et al. introduced a model inspired by McGreggor et al. [9], which also operated purely with visual information. However, unlike fractals, Kunda et al.’s model utilised affine and set transformations to map image data between cells in a given RPM problem [12].

In 2013, Rasmussen and Eliasmith extended their spiking model and applied it to studying age-related cognitive decline in humans [13]. They reported a positive correlation between manipulations in the model and observed human behavioural data. Much like earlier work in the field, Rasmussen and Eliasmith emphasised the use of RPMs for demonstration purposes, especially in their use of computational frameworks in studying and elucidating factors in human cognition.

In 2018, Barret et al. introduced the first large-scale deployment of RPMs as a machine learning dataset known as the Procedurally Generated Matrices (PGM) dataset [2]. In an attempt to simplify the PGM dataset and ease training, the RAVEN dataset was introduced in 2019 by Zhang et al. [14]. However the generation mechanism for RAVEN was later shown by Hu et al. to be compromised [15]. They demonstrated that the generation mechanism for the answer panels resulted in substantially easier RPMs, such that it was unnecessary for an algorithm to consider the context panels in order to solve the RPM. To rectify this defect, Hu et al. introduced the I-RAVEN dataset in 2021 [15]. It should be noted that RAVEN still remains a valid machine learning dataset, especially for benchmarking, albeit far simpler than intended and without the intended correspondence with the human intelligence tests.

Most work in this area focusses on fully supervised learning, i.e., where the target classification of all training instances is made available to the learner. To our knowledge, only two papers on I-RAVEN, [16] and [6], have considered anything other than fully supervised learning. The first paper, [16], describes a semi-supervised learning method for RPMs, which efficiently trains neural networks with minimal labelled data, while the latter, [6], describes a highly innovative method as *unsupervised*, though there is some acknowledgement that it is not unsupervised in the strictest sense. The aim of this letter is to clarify the exact position of [6] in the research literature and thus alleviate any confusion that may arise in subsequent research.

To begin, we first discuss the characteristics and categories of benchmark and demonstrator problems with respect to RPMs in section II, considering both the presence and absence of supervision. Next, we carefully delineate the relationship between strict and broad uses of the term “unsupervised” in section III. In section IV, we discuss whether the approach proposed by Zhuo et al. [6] can be generalised across different classes of conceptual problems, of which RPMs are proxies or representatives. Finally we conclude the letter also in section IV, by distilling any insights into the effectiveness and adaptability of unsupervised methods in the RPM context.

II. DEMONSTRATION AND BENCHMARK PROBLEMS

Problem sets utilised in machine learning research generally fall into two large classes, i.e., demonstration problems and benchmark problems. Each class contains unique characteristics and serves a distinct role in the advancement and evaluation of machine learning algorithms.

Demonstration problems are primarily employed to elucidate and illustrate concepts, methods, or techniques. In machine learning research, these problems are primarily used to determine whether a particular problem class can be solved and if so, how well (e.g., protein folding [17]). Generally, the particular problem is not important, but is chosen to exemplify a general class of important problems that are likely to require similar techniques. Often, the problems are additionally chosen to have known solutions, to simplify the construction of training and testing sets.

Conversely, benchmark problems are primarily used in evaluation and comparing performance of different systems, algorithms or models [18], [19]. These problems are designed to establish a standardised and objective basis for assessing the capabilities of various solutions within a specific field (e.g., machine vision [20]). In machine learning, benchmark problems are commonly employed in testing and gauging the state-of-the-art within a particular domain (e.g., object recognition [21] or machine translation [22]). They may vary in complexity, ranging from relatively straightforward to highly intricate tasks. Importantly, benchmark problems often lack a single “correct” solution, shifting the focus towards evaluating the quality and effectiveness of different approaches. Their primary purpose is to measure progress, spur innovation, and provide a basis for comparison within a specific domain.

In terms of exemplifying abstract and relational reasoning, RPM problems clearly fall into the former class. They were undoubtedly originally introduced into machine learning precisely because they were well-studied in neuroscience and psychology as difficult conceptual problems. For machine learning, RPMs have become increasingly attractive since large datasets can be readily constructed and potentially because of their wide use as human tests make them more compelling for the general public.

III. SUPERVISED AND UNSUPERVISED LEARNING

The concepts of supervised and unsupervised learning have long existed in psychology [23], [24]. The earliest computational distinction between the two was first introduced by Spragins in 1966, albeit without modern terminology, i.e., Spragins distinguished the two as *learning with or without a teacher* [25]. The modern machine learning distinction was first introduced by Darling and Raudseps in 1970, with an emphasis on classification [26].

Since its inception, unsupervised learning has been pursued for two connected reasons, to provide impetus toward exploratory learning (and thus more completely mimic human learning), and to reduce the requirement of labelled training data, which is often prohibitively expensive to obtain [27]. Initially, the small toy training sets then used meant that the

labelling cost was low, so that the emphasis was on exploratory learning [28], and the formal definition correspondingly strict.

The expanding capabilities of neural nets has led to a demand for vast training sets and thus for mechanisms that reduce, but don't necessarily eliminate, the need for costly human-generated labels. This has led to a more relaxed definition of unsupervised learning, where some ground-truth labels and/or some noisy labels are permitted [29]. In machine learning literature, this loose form of unsupervised learning is sometimes referred to as few-shot learning or more often as *semi-supervised learning* [30].

Confusion can arise when *contrastive learning* is considered, as this method can be applied to both supervised and unsupervised settings [31]. When applied in unsupervised settings, no labels are utilised, noisy or otherwise. In such settings, a learner is tasked with discovering invariant representations of an input (usually an image) from some given distribution [32]. Generally, the level of supervision can be questioned in situations where contrastive examples are hand-crafted or human generated.

In its strictest sense, *unsupervised learning* refers to a category of machine learning algorithms designed to extract meaningful patterns, structures, or representations from unlabelled data. It is a process of discovering hidden patterns without any predefined labels or specific guidance from external sources [33]. The fundamental premise of unsupervised learning is to let the algorithm explore and identify inherent structures within the data itself [27].

The work presented by Zhuo et al. [6], as further clarified in the next section, eliminates this premise of exploration and discovery by communicating most of the hidden patterns to the algorithm and guiding its process of learning. Hence their work falls into the semi-supervised category, or at least in a category excluding strictly unsupervised learning. The importance of this exclusion comes from the observation that many (perhaps most) important conceptual learning problems, to which we wish to apply machine learning, may not have pre-determined class labels and so exploratory and strictly unsupervised learning will continue to be important. This is particularly relevant to RPMs as neural network demonstrators.

IV. DISCUSSION AND CONCLUSION

As illustrated in Fig. 1(b), the Zhuo et al., method applies an ingenious transformation to the original RPM problem, converting the original three-row problem into a ten-row problem [6]. They replicate the incomplete last row of the original problem with eight complete rows, one for each instance from the answer panels. They add pseudo labels to each row, marking the first two rows validly as correct, and the rest (not always validly), as contrastingly incorrect. In doing this, Zhuo et al. encourage a learning algorithm to learn and identify equal or similar properties between the positive rows and distinguish them from the negative rows. In practice, the learning algorithm is encouraged to assign higher estimated probabilities to the first two rows, while assigning lower values to dissimilar rows in the remaining eight rows.

Undoubtedly, Zhuo et al.'s transformation of an RPM problem results in an equivalent problem to the original, in the




Fig. 2. Illustrative example of an RPM highlighting the interplay between a first-order and a second-order relation. The second-order relation, ‘less strict than’, restricts how the first-order relation, ‘identical’, manifests in each successive row. In the first row, the first-order relation fully manifests. In the second row, the first-order relation manifests as “all objects are identical, except some objects are smaller”. In the third row, it manifests as “all objects are identical, except there are fewer objects and some might be smaller”. Solving the RPM requires one to consider the successive application of the second-order relation between and across the rows.

sense that there is an automated transformation between the two. However, is this transformation enough? For a benchmark problem, the answer is clear: a clever transformation such as this will be directly incorporated into methods and be applied across the board. However, for a demonstration problem, the matter is more nuanced.

First, a critical question is whether this transformation can be applied across the class of problems for which the original is a demonstrator. This is clearly not the case here. The method in [6] depends heavily on the specific structure of the RPM problem, i.e., a multi-relation problem with multiple-choice answers. It is not at all clear that the transformation of Zhuo et al. can be applied outside this restricted class, and equally unclear how many interesting conceptual learning problems actually share this structure.

Second, there is the question of the sense in which the original and transformed are equivalent. With sufficient knowledge, the two can be inter-transformed. In the case of RPMs, we the supervisors or instance constructors have this knowledge. But the learner does not: it has to learn it. Even if the problems are equivalent given certain knowledge, it does not imply that the problems are equivalent from the perspective of a learner.

Formally, the classic representation of an RPM, as in Fig. 1(a) and 2, is most naturally represented as a *second-order* logic problem. The content of each row can be represented as a relation between the three images (i.e., a first order problem). Thus the content of the whole three rows becomes a second order assertion, that the relations in the three rows satisfy a second order relationship ‘same-relation’. The overall problem is then to discover such a relation and choose which of the

eight answer panels makes that overall second-order relation correctly describe the first-order relations of the three rows. Thus it is in essence a second-order learning problem.

From the perspective of an RPM solver, human or computer, the second order relation to be discovered could be any second-order relation, it does not have to be ‘same-relation’. For example, it could alternatively be the relation ‘less strict than’, where the relation in the first row is stricter than in the second and the second is stricter than in the third. That is, an RPM can be validly constructed such that the first-order relation varies with each successive row, as opposed to those in traditional RPMs where they remain constant. Fig. 2 offers a concrete example of this, where the second-order relation ‘less-strict-than’ acts on the underlying first-order relation ‘identical’ and varies successively with each row. In this illustrative example, a learner is tasked with identifying the properties of the individual elements, i.e., colour, position, shape and size; the first-order relationships between the elements, i.e., ‘identical’, and the second-order relationships between these first-order relationships, i.e., ‘less strict than’.

Comparing Fig. 1(b) and 2, it is clear that the transformed problem naturally loses the second-order property and becomes best represented as a *first-order* logic problem. Namely, as in the original RPM, each row has a first order relation between the three images, however in the transformed problem, some are labelled as correct or contrastingly incorrect. The result of this transformation is that the second-order aspect of the RPM becomes embedded in the underlying assumption of machine learning, that the rows have something in common. In the transformed case, the ‘same relation’ predicate becomes privileged in a way that it was not in the original.

At a formal level, Zhuo et al.’s ingenious method in [6] results in a different problem. That problem seems less representative of the class of difficult conceptual reasoning problems RPMs are supposed to represent. Namely, when a learner is tasked with solving an RPM without guidance, they are tasked with self-discovering higher-order relations.

For humans, such a task, at its highest level, is often awarded the special name of “research”. It is also often associated with abstract reasoning and fluid intelligence. These two cognitive traits are related to how quickly a human subject is able to reason with information to solve new, unfamiliar problems, independent of any prior knowledge [34]. Furthermore, these traits also include a human subject’s ability to think laterally and flexibly, to reason logically, and to extrapolate rules or relationships beyond the most obvious and to other possible scenarios.

For machine learning algorithms, simulating these traits offers advantages beyond simply solving RPMs. Fluid intelligence and abstract reasoning are also implicated in action perception and production, as well as in physical and observational learning [35]. That is, these traits are implicated in cognitive processes strongly emulated by reinforcement learning and autonomous vehicles as reported in machine learning literature [36]. Furthermore, fluid intelligence, and by extension observational learning, overlap with key action representation systems in the human brain, such as the mirror neuron system [37], as well as domain-general control pro-

cesses that have been associated with the multiple demand system [38].

It is exactly in relation to these cognitive traits that Raven’s Progressive Matrices serve as proxies for general conceptual reasoning problems. It is clear that RPM problems do not capture all aspects of general conceptual reasoning, but offer an objective measure of an entity’s ability to reason abstractly, recognise patterns and draw logical conclusions. The skills assessed by RPM tasks in a strictly unsupervised setting align closely with the cognitive abilities required for general conceptual reasoning, making them valuable tools for assessing and understanding this broader cognitive domain.

While most research has focused on supervised learning, the application of unsupervised learning to RPMs has added complexity and innovation to the field. This discussion has highlighted some of these complexities. Firstly, it has noted the nuances between demonstration and benchmark problems, where we have emphasised the need to assess the transferability of methods across problem classes. Secondly, we have underscored the critical differences between supervised, semi-supervised and unsupervised learning, with contrastive learning taken into account. In doing so we have clarified the categorisation of Zhuo et al.’s work reported in [6].

Finally, while considering supervised, semi-supervised and unsupervised learning, this discussion has highlighted the subtle interaction between guidance and difficulty in solving RPMs. Zhuo et al.’s method introduces a novel perspective on unsupervised learning and addresses the challenges of obtaining labelled training data, but careful consideration is needed to assess its generalisability and its ability to capture the essence of difficult conceptual reasoning problems.

REFERENCES

- [1] C. James, “Raven. mental tests used in genetic studies: The performance of related individuals on tests mainly educative and mainly reproductive,” *Unpublished master’s thesis, University of London*, 1936.
- [2] D. Barrett, F. Hill, A. Santoro, A. Morcos, and T. Lillicrap, “Measuring abstract reasoning in neural networks,” in *International conference on machine learning*. PMLR, 2018, pp. 511–520.
- [3] 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*, vol. 97, no. 3, p. 404, 1990.
- [4] J. Raven, “The raven’s progressive matrices: change and stability over culture and time,” *Cognitive psychology*, vol. 41, no. 1, pp. 1–48, 2000.
- [5] G. Domino and M. L. Domino, *Psychological testing: An introduction*. Cambridge University Press, 2006.
- [6] T. Zhuo, Q. Huang, and M. Kankanhalli, “Unsupervised abstract reasoning for raven’s problem matrices,” *IEEE Transactions on Image Processing*, vol. 30, pp. 8332–8341, 2021.
- [7] S. Bringsjord and B. Schimanski, “What is artificial intelligence? psychometric ai as an answer,” in *Proceedings of the 18th international joint conference on Artificial intelligence*, 2003, pp. 887–893.
- [8] A. Lovett, K. Forbus, and J. Usher, “A structure-mapping model of raven’s progressive matrices,” in *Proceedings of the Annual Meeting of the Cognitive Science Society*, vol. 32, no. 32, 2010.
- [9] K. McGregor, M. Kunda, and A. K. Goel, “A fractal analogy approach to the raven’s test of intelligence,” in *Visual Representations and Reasoning*, 2010.
- [10] S. Cirillo and V. Ström, “An anthropomorphic solver for raven’s progressive matrices,” 2010.
- [11] D. Rasmussen and C. Eliasmith, “A neural model of rule generation in inductive reasoning,” *Topics in Cognitive Science*, vol. 3, no. 1, pp. 140–153, 2011.

- [12] M. Kunda, K. McGregor, and A. Goel, "Reasoning on the raven's advanced progressive matrices test with iconic visual representations," in *Proceedings of the Annual Meeting of the Cognitive Science Society*, vol. 34, no. 34, 2012.
- [13] D. Rasmussen and C. Eliasmith, "A spiking neural model applied to the study of human performance and cognitive decline on raven's advanced progressive matrices," *Intelligence*, vol. 42, pp. 53–82, 2014.
- [14] 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*, 2019, pp. 5317–5327.
- [15] 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*, vol. 35, no. 2, 2021, pp. 1567–1574.
- [16] Y. Xu, L. Yang, H. You, Z. Zhen, D.-H. Wang, X. Wan, X. Xie, and R.-Y. Zhang, "Rulematch: matching abstract rules for semi-supervised learning of human standard intelligence tests," in *Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence*, 2023, pp. 1613–1621.
- [17] J. Jumper, R. Evans, A. Pritzel, T. Green, M. Figurnov, O. Ronneberger, K. Tunyasuvunakool, R. Bates, A. Žídek, A. Potapenko *et al.*, "Highly accurate protein structure prediction with alphafold," *Nature*, vol. 596, no. 7873, pp. 583–589, 2021.
- [18] T. Eftimov, G. Petelin, G. Cenikj, A. Kostovska, G. Ispirova, P. Korošec, and J. Bogatinovski, "Less is more: Selecting the right benchmarking set of data for time series classification," *Expert Systems with Applications*, vol. 198, p. 116871, 2022.
- [19] G. Papadakis, N. Kirielle, P. Christen, and T. Palpanas, "A critical re-evaluation of benchmark datasets for (deep) learning-based matching algorithms," *arXiv preprint arXiv:2307.01231*, 2023.
- [20] O. Firschein, M. Fischler, and T. Kanade, "Creating benchmarking problems in machine vision: Scientific challenge problems," in *DARPA Image Understanding Workshop*, 1993, pp. 177–182.
- [21] A. Athar, J. Luiten, P. Voigtlaender, T. Khurana, A. Dave, B. Leibe, and D. Ramanan, "Burst: A benchmark for unifying object recognition, segmentation and tracking in video," in *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 2023, pp. 1674–1683.
- [22] F. Stahlberg, "Neural machine translation: A review," *Journal of Artificial Intelligence Research*, vol. 69, pp. 343–418, 2020.
- [23] G. Grindley, "Experiments on the influence of the amount of reward on learning in young chickens," *British Journal of Psychology*, vol. 20, no. 2, p. 173, 1929.
- [24] L. Morrisett Jr and C. I. Hovland, "A comparison of three varieties of training in human problem solving," *Journal of Experimental Psychology*, vol. 58, no. 1, p. 52, 1959.
- [25] J. Spragins, "Learning without a teacher," *IEEE Transactions on Information Theory*, vol. 12, no. 2, pp. 223–230, 1966.
- [26] E. M. Darling Jr and J. G. Raudseps, "Non-parametric unsupervised learning with applications to image classification," *Pattern Recognition*, vol. 2, no. 4, pp. 313–335, 1970.
- [27] G. James, D. Witten, T. Hastie, R. Tibshirani, and J. Taylor, "Unsupervised learning," in *An Introduction to Statistical Learning: with Applications in Python*. Springer, 2023, pp. 503–556.
- [28] R. O. Duda, P. E. Hart *et al.*, *Pattern classification and scene analysis*. Wiley New York, 1973, vol. 3.
- [29] S. Li, X. Xia, S. Ge, and T. Liu, "Selective-supervised contrastive learning with noisy labels," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 316–325.
- [30] H. Xu, H. Xiao, H. Hao, L. Dong, X. Qiu, and C. Peng, "Semi-supervised learning with pseudo-negative labels for image classification," *Knowledge-Based Systems*, vol. 260, p. 110166, 2023.
- [31] P. H. Le-Khac, G. Healy, and A. F. Smeaton, "Contrastive representation learning: A framework and review," *Ieee Access*, vol. 8, pp. 193 907–193 934, 2020.
- [32] C.-H. Ho and N. Nvasconcelos, "Contrastive learning with adversarial examples," *Advances in Neural Information Processing Systems*, vol. 33, pp. 17 081–17 093, 2020.
- [33] Z. Ghahramani, "Unsupervised learning," in *Summer School on Machine Learning*. Springer, 2003, pp. 72–112.
- [34] J. R. Gray, C. F. Chabris, and T. S. Braver, "Neural mechanisms of general fluid intelligence," *Nature neuroscience*, vol. 6, no. 3, pp. 316–322, 2003.
- [35] K. Sakreida, S. Higuchi, C. Di Dio, M. Ziessler, M. Turgeon, N. Roberts, and S. Vogt, "Cognitive control structures in the imitation learning of spatial sequences and rhythms — an fMRI study," *Cerebral Cortex*, vol. 28, no. 3, pp. 907–923, 2018.
- [36] B. R. Kiran, I. Sobh, V. Talpaert, P. Mannion, A. A. Al Sallab, S. Yogamani, and P. Pérez, "Deep reinforcement learning for autonomous driving: A survey," *IEEE Transactions on Intelligent Transportation Systems*, 2021.
- [37] G. Rizzolatti and C. Sinigaglia, "The mirror mechanism: a basic principle of brain function," *Nature Reviews Neuroscience*, vol. 17, no. 12, pp. 757–765, 2016.
- [38] J. Duncan, "The multiple-demand (MD) system of the primate brain: mental programs for intelligent behaviour," *Trends in cognitive sciences*, vol. 14, no. 4, pp. 172–179, 2010.