Algebraic Machine Reasoning: How inductive reasoning reduces to algebraic computations

Kai Fong Ernest Chong*

Singapore University of Technology and Design, Singapore
ernest_chong@sutd.edu.sg

Abstract

In automated reasoning, nothing captures our excitement more than the prospect that machines would surpass humans in various intelligence tests and general reasoning tasks. For the well-known visual-based IQ test known as Raven's progressive matrices (RPMs), machines have just very recently surpassed human-level performance. This feat is based on "*algebraic machine reasoning*", a completely new approach to inductive reasoning, first introduced by the team led by Chong (the speaker of this talk). The crux is that algebraic machine reasoning (algMR) reduces the difficult process of novel problem-solving to *routine* algebraic computations. Crucially, in contrast to reasoning methods such as inductive logic programming, algMR is intrinsically not search-based. Using algMR, the task of solving RPMs is reduced to solving computational problems in algebra involving the ideals of polynomial rings. Specifically, the discovery of abstract patterns in RPMs can be realized very concretely as algebraic computations known as primary decompositions of ideals. In this talk, we shall introduce algMR, explain its underlying intuition and how it is used to solve the RPM task, as well as discuss the possibilities for using algMR to solve other reasoning tasks. Through this talk, we hope to initiate a discussion on how algMR could be useful for ATP, and more generally, on how algMR fits into the larger context of AGI.

Algebraic machine reasoning (algMR) is a fundamentally new approach to abstract reasoning. It was first introduced in [19] in the context of solving Raven's progressive matrices (RPMs), which are well-known inductive reasoning tests for measuring what cognitive scientists call fluid intelligence [1, 12]; see Fig. 1 for a typical example of an RPM. On the I-RAVEN dataset [6] (with 14,000 RPMs in the test set), algMR achieves an accuracy of 93.2% directly from raw image data of the RPMs, which significantly outperforms the previous state-of-the-art (SOTA) accuracy of 77.0%, and surpasses human-level performance at 84.4%. Previous SOTA methods on this dataset all rely on deep learning for reasoning and hence require a lot of training on answers to RPMs. In stark contrast, algMR does not require any training on such task-specific data, analogous to how a gifted child would be able to solve RPMs without needing any practice on RPMs. Perhaps most interestingly, algMR is also able to directly generate the correct answers to RPMs without needing any multiple-choice answer options to choose from.

Using algebraic methods in AI is not new. Indeed, it is well-known in automated theorem proving (ATP) that Gröbner basis computations can be used for proof-checking [14]. In statistical learning theory, methods from algebraic geometry [16] and algebraic statistics [3] are used to study statistical models via various suitably defined algebraic varieties [4, 8, 9, 10, 15, 17, 18, 20]. In deep learning, algebraic methods are used to study the expressivity of neural nets [2, 7, 11, 24]. In contrast, algMR establishes the first ever connection between commutative algebra and machine reasoning. As we have pointed out in [19], there is "over a century's worth of very deep results in commutative algebra that have not been tapped". More importantly, in contrast to

*This work is supported by the National Research Foundation, Singapore under its AI Singapore Program (AISG Award No: AISG-RP-2019-015) and under its NRFF Program (NRFFAI1-2019-0005), and by Ministry of Education, Singapore, under its Tier 2 Research Fund (MOE-T2EP20221-0016).

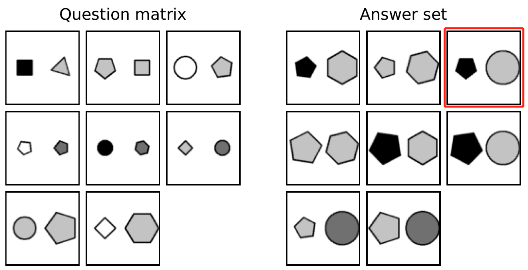

	Method	Conference	Accuracy
1	ResNet+DRT [21]	CVPR 2019	40.4%
2	LEN [25]	NeurIPS 2019	41.4%
3	Wild ResNet [13]	ICML 2018	44.3%
4	CoPINet [22]	NeurIPS 2019	46.1%
5	DCNet [26]	ICLR 2021	47.0%
6	SRAN [6]	AAAI 2021	60.8%
7	PrAE [23]	CVPR 2021	77.0%
8	**algMR (ours)** [19]	**CVPR 2023**	**93.2**%
	Human [21]	-	84.4%

Figure 1: An RPM (depicted on the left) is composed of a question matrix and an answer set. A question matrix is a 3×3 grid of panels, where the first 8 panels are filled with geometric entities, and the 9-th panel is "missing". These 9 panels obey certain relations, and the goal is to infer the correct answer for this last panel from an answer set consisting of 8 possible answers. In the depicted RPM, the correct answer is indicated in red. RPMs can be solved via algMR in a fully automated manner, directly from raw images, at an accuracy of 93.2%, which surpasses human-level performance at 84.4%. Prior to algMR, the SOTA accuracy was 77.0%. The table on the right gives a comparison of algMR with previous methods, evaluated on the I-RAVEN dataset.

previous applications of algebraic methods in AI, we are *not* solving a system of polynomial equations. In particular, there is no notion of satisfiability. We do *not* assign numerical values (or truth values) to variables/literals. Instead, ideals are treated as "actual objects of study", whereby reasoning is realized as non-numerical computational problems involving ideals.

The starting point of algMR is to define concepts algebraically. Let $R = \mathbb{R}[x_1, \dots, x_n]$ be the ring of polynomials in variables $x_1, \dots, x_n$, with real coefficients. A subset $I \subseteq R$ is called an *ideal* if there are polynomials $g_1, \dots, g_k$ in R such that $I = \{f_1 g_1 + \cdots + f_k g_k | f_1, \dots, f_k \in R\}$ contains all polynomial combinations of $g_1, \dots, g_k$. We say that $\mathcal{G} = \{g_1, \dots, g_k\}$ is a *generating set* for I. If I has a generating set consisting only of monomials, then we say that I is a *monomial ideal*. (A monomial is a polynomial with a single term.) In algMR, a *concept* is defined to be a monomial ideal of R. Thus, we can apply ideal-theoretic operations on concepts, from basic operations such as sums, products, and intersections, to "advanced" operations such as primary decompositions. All these operations can be computed via well-studied algebraic subroutines, which are already implemented in most existing computer algebra systems (e.g. Macaulay2 [5]).

Primary decompositions of ideals are of fundamental importance in commutative algebra; they are a vast generalization of the idea of prime factorization for integers. Informally, every ideal $J \subseteq R$ has a decomposition $J = J_1 \cap \cdots \cap J_k$ as an intersection of finitely many *primary* ideals. This intersection is called a *primary decomposition* of J, and each J_j is called a *primary component* of the decomposition. We proved in [19, Thm. 3.1] that if $J \neq R$ is a concept, then J has a *unique* minimal primary decomposition $J = J_1 \cap \cdots \cap J_k$ such that $J_1, \dots, J_k$ are concepts satisfying a nice structural property. In essence, computing the primary decompositions of concepts serves as a rather general and effective algebraic subroutine for extracting abstract patterns from multiple concepts. A significant portion of the talk will focus on the algebraic intuition on how we use primary decompositions to solve RPMs, and more generally, how we use primary decompositions to solve inductive reasoning tasks.

Although we initialize R as a polynomial ring, there is nothing inherently special about this initialization. We could instead initialize R as a quotient ring. Concepts would still be defined as monomial ideals of this quotient ring R, and we could still compute the primary decompositions of concepts in R. In the last part of our talk, we shall explain what quotient rings are, and how they relate to reasoning on concepts. We shall also discuss how further extensions and variations of algMR (where R is initialized as a quotient ring) could be used to solve other kinds of reasoning tasks.

References

[1] Patricia A Carpenter, Marcel A Just, and Peter Shell. What one intelligence test measures: a theoretical account of the processing in the Raven progressive matrices test. *Psychological review*, 97(3):404, 1990.

[2] Kai Fong Ernest Chong. A closer look at the approximation capabilities of neural networks. In *International Conference on Learning Representations*, 2020.

[3] Mathias Drton, Bernd Sturmfels, and Seth Sullivant. *Lectures on algebraic statistics*, volume 39. Springer Science & Business Media, 2008.

[4] Eliana Duarte, Orlando Marigliano, and Bernd Sturmfels. Discrete statistical models with rational maximum likelihood estimator. *Bernoulli*, 27(1):135–154, 2021.

[5] Daniel R. Grayson and Michael E. Stillman. Macaulay2, a software system for research in algebraic geometry. Available at `http://www.math.uiuc.edu/Macaulay2/`.

[6] Sheng Hu, Yuqing Ma, Xianglong Liu, Yanlu Wei, and Shihao Bai. Stratified rule-aware network for abstract visual reasoning. In *Proceedings of AAAI Conference on Artificial Intelligence (AAAI)*, 2021.

[7] Joa Kileel, Matthew Trager, and Joan Bruna. On the expressive power of deep polynomial neural networks. *Advances in neural information processing systems*, 2019.

[8] Franz J Király, Paul Von Bünau, Frank C Meinecke, Duncan AJ Blythe, Klaus-Robert Müller, and Kenji Fukumizu. Algebraic geometric comparison of probability distributions. *Journal of Machine Learning Research*, 13(3), 2012.

[9] Paul Larsen and Franz Király. Fano schemes of generic intersections and machine learning. *International Journal of Algebra and Computation*, 24(07):923–933, 2014.

[10] Shaowei Lin. *Algebraic methods for evaluating integrals in Bayesian statistics*. PhD thesis, UC Berkeley, 2011.

[11] Petros Maragos, Vasileios Charisopoulos, and Emmanouil Theodosis. Tropical geometry and machine learning. *Proceedings of the IEEE*, 2021.

[12] Patrick Perret. Children's inductive reasoning: Developmental and educational perspectives. *Journal of Cognitive Education and Psychology*, 14(3):389–408, 2015.

[13] Adam Santoro, Felix Hill, David Barrett, Ari Morcos, and Timothy Lillicrap. Measuring abstract reasoning in neural networks. In *International Conference on Machine Learning*, pages 4477–4486, 2018.

[14] Jürgen Stuber. Superposition theorem proving for commutative rings. In *Automated Deduction—A Basis for Applications*, pages 31–55. Springer, 1998.

[15] Seth Sullivant. Algebraic geometry of gaussian bayesian networks. *Advances in Applied Mathematics*, 40(4):482–513, 2008.

[16] Sumio Watanabe. *Algebraic geometry and statistical learning theory*. Number 25. Cambridge university press, 2009.

[17] Sumio Watanabe. A widely applicable bayesian information criterion. *Journal of Machine Learning Research*, 14(3):867–897, 2013.

[18] Sumio Watanabe and Manfred Opper. Asymptotic equivalence of bayes cross validation and widely applicable information criterion in singular learning theory. *Journal of machine learning research*, 11(12), 2010.

[19] Jingyi Xu, Tushar Vaidya, Yufei Wu, Saket Chandra, Zhangsheng Lai, and Kai Fong Ernest Chong. Abstract visual reasoning: An algebraic approach for solving raven's progressive matrices. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023.

[20] Keisuke Yamazaki and Sumio Watanabe. Singularities in mixture models and upper bounds of stochastic complexity. *Neural Networks*, 16(7):1029–1038, 2003.

[21] Chi Zhang, Feng Gao, Baoxiong Jia, Yixin Zhu, and Song-Chun Zhu. RAVEN: A dataset for relational and analogical visual reasoning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5317–5327, 2019.

[22] Chi Zhang, Baoxiong Jia, Feng Gao, Yixin Zhu, Hongjing Lu, and Song-Chun Zhu. Learning perceptual inference by contrasting. In *Advances in Neural Information Processing Systems*, pages 1075–1087, 2019.

[23] Chi Zhang, Baoxiong Jia, Song-Chun Zhu, and Yixin Zhu. Abstract spatial-temporal reasoning via probabilistic abduction and execution. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021.

[24] Liwen Zhang, Gregory Naitzat, and Lek-Heng Lim. Tropical geometry of deep neural networks. In *International Conference on Machine Learning*, pages 5824–5832. PMLR, 2018.

[25] Kecheng Zheng, Zheng-Jun Zha, and Wei Wei. Abstract reasoning with distracting features. In *Advances in Neural Information Processing Systems*, volume 32, 2019.

[26] Tao Zhuo and Mohan Kankanhalli. Effective abstract reasoning with dual-contrast network. In *International Conference on Learning Representations*, 2021.