

Learning to be Simple

Yang-Hui He^{a,b*}, Vishnu Jejjala^{c†}, Challenger Mishra^{d‡}, Max Sharnoff^{e,f§}

^a *London Institute for Mathematical Sciences, Royal Institution, London W1S 4BS, UK*

^b *Merton College, University of Oxford, OX1 4JD, UK*

^c *Mandelstam Institute for Theoretical Physics, School of Physics, NITheCS, and CoE-MaSS, University of the Witwatersrand, Johannesburg, South Africa*

^d *Department of Computer Science & Technology, University of Cambridge, CB3 0FD, UK*

^e *Department of Computer Science, University of Oxford, OX1 3QG, UK*

^f *Christ Church, University of Oxford, OX1 1DP, UK*

Abstract

In this work we employ machine learning to understand structured mathematical data involving finite groups and derive a theorem about necessary properties of generators of finite simple groups. We create a database of all 2-generated subgroups of the symmetric group on n -objects and conduct a classification of finite simple groups among them using shallow feed-forward neural networks. We show that this neural network classifier can decipher the property of simplicity with varying accuracies depending on the features. Our neural network model leads to a natural conjecture concerning the generators of a finite simple group. We subsequently prove this conjecture. This new toy theorem comments on the necessary properties of generators of finite simple groups. We show this explicitly for a class of sporadic groups for which the result holds. Our work further makes the case for a machine motivated study of algebraic structures in pure mathematics and highlights the possibility of generating new conjectures and theorems in mathematics with the aid of machine learning.

*hey@maths.ox.ac.uk

†v.jejjala@wits.ac.za

‡cm2099@cam.ac.uk

§github@max.sharnoff.org

Contents

1	Introduction: Machine learning and symmetries	1
2	Machine learning simplicity	3
3	Learning outcomes and a conjecture	4
3.1	Machine learning experiments	4
3.2	A machine guided mathematical conjecture	8
4	Discussion	13
A	Machine learning experiments	16
A.1	2-generated subgroups of S_n	16
A.2	Datasets	16
A.3	Neural network architectures	20

1 Introduction: Machine learning and symmetries

Machine learning is an increasingly ubiquitous tool for studying a wide range of problems from self-driving cars and drug design to many electron systems in quantum chemistry and protein folding *in vivo*. However, thus far machine learning has played a smaller rôle in developing pure mathematics. Since the injection of machine learning into investigations of algebraic geometry in the context of theoretical physics [1–8], there has been an explosion of activity to machine learn various aspects of the topology and geometry of Calabi–Yau manifolds [9–20], algebra [21–24], knot theory [25–30], combinatorics [31, 32], and number theory [33–35], etc. The present work has a two-fold purpose. We investigate the following questions: (i) can one learn the structure of mathematics and let artificial intelligence help the intuition of a mathematician along [36, 37]? (ii) how does one develop machine learning architectures that can identify structures in mathematical datasets which are difficult to observe with the human eye? The aim is to develop new machine driven methodologies and architectures that can study such *synthetic* data, as opposed to *real world* data.

One fundamental algebraic structure to examine through the lens of machine learning is a group. Groups provide a mathematical description of the symmetries of a system and are a guiding principle in our descriptions of Nature. Noether’s theorems [38] establish that conserved quantities arise from the symmetries of a theory. For example, the conservation of energy is the consequence of the translational invariance of a system in time, the conservation of momentum is the consequence of the translational invariance of a system in space, and the conservation of angular momentum is the consequence of the invariance of a system under spatial rotations.

Similarly, conserved currents in electromagnetism originate from the $U(1)$ gauge symmetry of quantum electrodynamics. The particles in the Standard Model are organized according to how they transform, namely in representations of certain Lie groups associated to the gauge symmetries. Identifying the underlying symmetries of a system and assessing their meaning is of paramount importance in understanding the physics [39, 40].

The analogous argument can be made for many machine learning endeavours. Knowledge of the symmetries of a dataset, real world or otherwise, is central to identifying correlations within the data. Indeed, making machine learning inferences using datasets is made more tractable by incorporating known invariances of the dataset into the models *ab initio*. This could, for instance, be accomplished by embedding the invariances into the architecture of a neural network [41, 42], or in the kernel of a Gaussian process. This way, various seemingly disconnected parts of the parameter space at play, albeit connected by these invariances, inform each other. This has the practical advantage of reducing computational costs, improving generalisation, and has led to many real world applications. The abiding principle is that a cat is a cat regardless of how it is viewed and this equivalence should be in built where possible. These applications exploit the relationship between model invariances and the dynamics of optimisation, leading to improved generalisation.

Preliminary studies of machine learning the algebraic structure of groups and rings were initiated for finite groups [21, 36] and for Lie groups [43]. For specific algebraic structures, the reader is also referred to [13, 22–24, 31, 44, 45]. One key motivation of our present work is to advance these investigations to a deeper level. With a view towards building a machine driven detector of algebraic structures (and groups in particular), we ask if *interpretable* neural networks can study different properties of a group. Such properties can range from the order of a group (or group element), to more involved computations such as the invariant ring of a group. In this work, we concern ourselves with finite simple groups.

In Section 2, we review some general results on finite simple groups, and describe the subclass of groups we study in this paper, *i.e.*, two generated subgroups of the symmetric group S_n with examples. We also explain motivations behind our various representations for these groups, and their limitations. In Section 3, we present our machine learning outcomes as well as a proposition that we were able to extract from our the machine learning investigations. Finally, we conclude with the discussion in Section 4. Appendices A.1 and A.2 describe details of the datasets of groups we employ for our machine learning endeavours, while Appendix A.3 presents the neural network architectures.

2 Machine learning simplicity

Recall that a group is *simple* if it does not admit any non-trivial normal subgroups. Following decades of effort, the finite simple groups are completely classified: they are cyclic groups of prime order \mathbb{Z}_p , or alternating groups A_n with $n > 4$, or belong to one of 16 infinite families of groups of Lie type (plus the related Tits group), or are one of 26 exceptional cases called sporadic groups, the largest of which is the Fischer–Griess Monster, the source of Moonshine [46]. Indeed, the classification of the finite simple groups initially relied on computational technology to establish the existence and uniqueness of certain sporadic groups.

In [21], relatively shallow neural networks as well as support vector machines were used to distinguish simple groups from non-simple one to 99% accuracy without the AI knowing anything about the usual techniques; this beckoned the question as to whether there is underlying new mathematics and constituted a motivation for the present study. In particular, the Cayley multiplication tables of all finite groups up to size 70 (there are 602) were taken. Next, random permutations were performed on each (since Cayley tables are only defined up to permutations) such that more permutations are included for the simple groups (since there are many more finite groups that are non-simple). This created a *balanced* database of 60 000 examples of 70×70 matrices (a group of size n would have all entries in $\{1, 2, \dots, n\}$ and all tables are padded with 0 where necessary) labelled as “simple” or “non-simple”, with 50% each. Remarkably, when flattened and represented as points in \mathbb{R}^{70^2} , a support vector machine with Gaussian kernel was able to *separate* them to 99% accuracy. This led to a proto-conjecture that in the space of finite groups, the simple and non-simple groups are thus separated and can be so classified.

One shortcoming of using the Cayley table is that these grow as the square of the group order and working in \mathbb{R}^{n^2} limits the computational power. In the paper, we focus on building a multi-layer perceptron model that can classify whether finite groups are simple using a much more succinct representation. This gives a two-fold advantage: (i) it will allow the exploration of more groups; and (ii) it will cross-check whether the simple/non-simple separation is truly underlying some deep mathematics and not just an artifact of Cayley tables.

Now, in the literature there are some age old results regarding the simplicity of groups. Burnside’s theorem, Sylow’s test, and detecting zeroes from the character table are notable examples. Such classical results are helpful in assessing a small sample of possible finite simple groups. One naïve algorithm that would determine simplicity of a group would list out all possible non-trivial subgroups of a given group and sequentially check if any of those are normal. State of the art deterministic algorithms that test the simplicity of a finite group are computationally non-trivial even though they might incorporate known classical theorems about finite simple groups.¹

¹One method is to compute the character table and spotting the positions of 0s. We grant that these are

Non-deterministic algorithms are faster at establishing simplicity. In this paper, we provide an example of a non-deterministic multi-layer perceptron classifier that works in polynomial time in the size of the inputs.

Conjecture (Dixon 1969) [48]: Two randomly chosen elements of a finite simple group G generate G with probability $\rightarrow 1$ as $|G| \rightarrow \infty$.

Noting that every finite group is a subgroup of the symmetric group S_n for some n and that every simple group is 2-generated [49], Dixon’s conjecture motivates our study of 2-generated subgroups of S_n , and consider criterion about their simplicity [48,50]. In the experiments, groups were randomly generated by two permutations drawn at random from S_n . We use these as generators of a finite group and query a neural network classifier with the question of simplicity of the resulting group. We conduct a number of experiments with different representations. In one set of experiments, we use the full permutation representation of both the generators (which are required to be unequal to each other and identity). In the second set of experiments, we only use traces and determinants of generators (as they are representation invariant quantities). In a final experiment we only use the orders of the group elements and the order of the group as features. This is motivated by Theorem 1, which says that finite simple groups can be characterised by these integers. Further details about the experiments are in Appendix A.

3 Learning outcomes and a conjecture

3.1 Machine learning experiments

Experiment 1

Our first experiment takes, as input, a pair of matrices as elements of the symmetric group S_n . We mark whether the group generated by the pair as simple or non-simple accordingly. All results described here were generated via the cross-validation process described in Appendix A.2. Figure 1 shows the individual average validation accuracies at the end of training for each portion of dataset for n . Only a single data point is available for $n = 8$ (at 4.6%), due to computational restrictions from such a large dataset.

In general, there seem to be two different patterns in the validation accuracies above: for $n = 5$ and 6, the validation accuracy appears to increase in a roughly linear fashion with the percent of the dataset given, ending between 85% to 95% accuracy. For $n = 7$, however — and this effect is still visible with $n = 8$ to some extent — the validation accuracy jumps up

 polynomial complexity [47], but our main motivation is to uncover new structures in simple groups using AI rather than to find faster algorithms in their detection.




Figure 1: Validation accuracy for varied amount of datasets for each n . Percentages are of the total dataset, after balancing but before splitting to allow k -fold cross validation.

to 82% at only 5% of the full dataset. There is a slight decrease as the percent of dataset increases for $n = 7$ from 60% to 80%; this is likely due to the variability in the final accuracy for cross-validation runs at $n = 7$. We will discuss this further shortly.

Putting these patterns aside, the models appear to become highly accurate when given greater portions of the dataset — finishing at 89%, 96%, and 96% validation accuracy for $n = 5, 6$, and 7 respectively. Excluding an outlier in the final result for $n = 7$ gives an average of 99% validation accuracy on the full dataset. Figure 2 displays the models’ loss and accuracy on the training set during training, for each cross-validation run combined. The curves for $n = 5$ fit with the typically expected images; the other two sets merit further comments.

A particular effect discovered in the training runs for $n = 6$ is that the model appears to take some time to “get off the ground” — *i.e.*, some amount of training with slow progress is required before learning can accelerate. This effect was equally visible in the validation loss and accuracy. Multiple initializers were tried as replacements in an attempt to mitigate this effect, but it did not improve. Once training accelerated, the models appeared to behave normally.

Because epochs count the number of times the entire dataset has been used for training, there were also different requirements for each n in the number of epochs run for. For example, $n = 7$ only required six epochs, which partially explains why the per-epoch variability is more visible. Also present in $n = 7$ were occasional regressions during training, where one of the cross-validation runs would suddenly increase in loss and decrease in both training and validation accuracy. This consistently affected one or two cross-validation runs, entirely at random; even though the individual datasets used for each cross-validation run were kept consistent across repeated experimentation, the particular “failing” runs were not consistent. In the particular case shown in Figure 2, the “failing” run finished with a heavy bias for negative results (*i.e.*,




Figure 2: Loss and accuracy on training data while training on full datasets, for $n \in \{5, 6, 7\}$. Bold curves give the average value at that epoch across all cross-validation runs.

indicating that a group is not simple). The class accuracies for that trial were 98% and 4%, respectively. We were unable to identify a cause, though it happens to be simply to manually identify when training a model, given knowledge of the previous behavior. There were insufficient data to determine whether this effect would occur at $n = 8$.


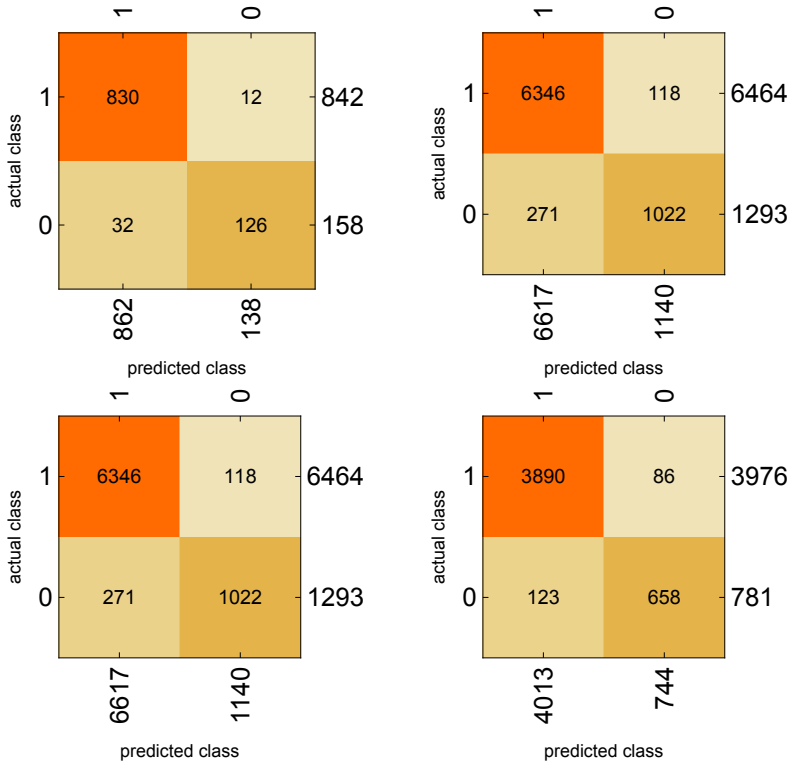


Figure 3: Error in predictions of accuracy from simplicity (dashed) versus parity (solid) during training for $n = 5$ (left) and $n = 6$ (right).

In order to gain further insights into understanding simplicity of finite groups, we now describe further machine learning experiments we conducted using alternate features.



(L) Training performance; (R) Test performance

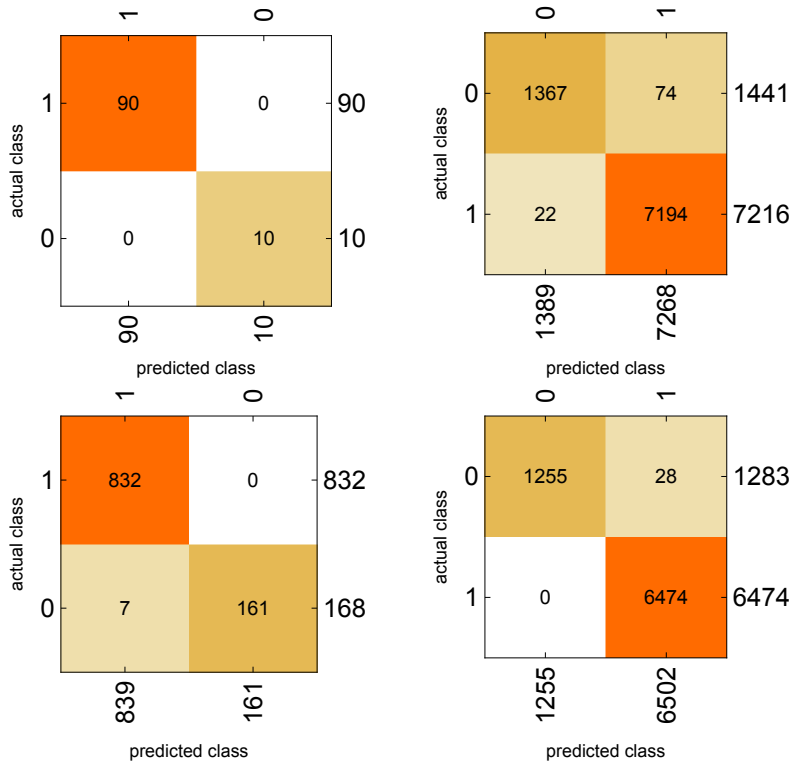
Figure 4: A classifier predicts simplicity of a group based on orders and traces of group generators and the property of being Abelian. The plot shows a confusion plot where the class 0 stands for simple groups and 1 for nonsimple groups. Using a neural network model and 1000 training points, class accuracies were found to be $\approx 97\%$ and 85% , respectively, on the training set (left) and $\approx 97\%$ and 86% on the test set (right). We use 1000 (top row) and 4000 (bottom row) points for training.

Experiment 2

In this experiment we choose the features to be the traces and determinants of the two generators and the binary property of the group being Abelian. We used a multi-layer perceptron model with three hidden layers with 1000, 500, and 200 nodes respectively, and logistic sigmoid activation. We used an ADAM optimiser. The dataset of size ≈ 8500 was split into a training set of sizes size 1000, 4000. The remaining data was used for validation in both cases. The results are described in Figure 4.

Experiment 3

In this experiment we chose the features to be the orders of the group elements and the order of the group. We used a multi-layer perceptron model with three hidden layers with 1000, 500, and 200 nodes respectively, and logistic sigmoid activation. We used an ADAM optimiser. The



(L) Training performance; (R) Test performance

Figure 5: A classifier predicts simplicity of a group based on orders of group elements and order of the group. The plot shows a confusion plot where the class 0 stands for simple groups and 1 for nonsimple groups. Using a neural network model when using only 100 training points (top) and 1000 training points (bottom), class accuracies on the test set were found to be $\approx 99\%$.

dataset of size ~ 8500 was split into a training set of size 1000, and a test set of ~ 7500 . We repeated the experiment using a smaller architecture with 10 times fewer nodes and a training set of size 100. The results are described in Figure 5.

The second and third experiments were particularly illuminating in terms of choice of features. The new features in these experiments were the traces and determinants of the generators as opposed to the full generators; and some combination of two other features, the group order and the property of being Abelian.

With these outcomes, it is interesting to consider a mathematical statement that might capture these experiments. We therefore consider necessary conditions on traces and determinants of generators of finite simple groups; hinted by our machine learning outcomes above. This manner of investigation has yielded different conjectures using an alternative approach [51].

3.2 A machine guided mathematical conjecture

In Experiments 1–3 of Section 3.1 we have built simple feed-forward multi-layer perceptron models to predict the simplicity of a 2-generated group given merely the generators, or their

properties. Our learning outcomes were modest with predictive accuracies $\sim 80\%$. In a traditional approach, one can conceive of a straightforward methodology for resolving the question of the simplicity of the resulting group. This could be seen to involve two key steps, the first of which is to freely generate a group from the two generators. Following this, one could list all non-trivial subgroups of this group, and check for normal subgroups. Both these key steps are computationally intensive. Therefore, it is impressive that a simple feed-forward multi-layer perceptron model should address this question to any reasonable degree of accuracy. We now ask the following question: can similar learning outcomes be attained using alternate representations for the generators? These could involve high level properties of the generators given a representation — traces and determinants are examples of such properties. If a similar or better learning outcome is obtained using such the above properties as alternate features, this is perhaps indicative of an underlying mathematical relationship between such generator properties and the property of simplicity. This indeed turns out to be the case: [Figure 4](#) and [Figure 5](#) show learning outcomes using different features. In this section, we present such a mathematical relationship as a conjecture motivated by the above learning outcome, for which we then provide a proof. This AI-guided mathematical conjecture formulation is very much in the spirit of [\[6, 37, 51\]](#).

First, we recall a theorem for characterising finite simple groups [\[52\]](#). Let us define the set of orders of group elements as

$$\pi(G) := \{\text{Ord}(g) : g \in G\} . \tag{1}$$

Theorem 1. *A finite simple group G is completely characterised by the order of group elements $\pi(G)$ and the order (size) of the group G .*

Theorem 1 implies that for G and H simple finite groups, $G \simeq H$ iff $|G| = |H|$ and $\pi(G) = \pi(H)$ (as sets). This is indicative that perhaps determinants of the generator could play a crucial role in discovering a new mathematical relationship between the generators and the property of simplicity. In fact, this also provides an alternate representation for our classification problem. Computing $\pi(G)$ for a group 2-generated G requires further computation. When our ML models are trained using these alternate features, it is not surprising that the learning outcomes are nearly perfect. [Figure 5](#) shows the learning outcome. However, this is at additional computational cost, since generating such representations involves computing $\pi(G)$ and computing $|G|$ requires.

Furthermore, we make an observation from studying our dataset. Before doing this, we introduce a preliminary notion [\[53\]](#). Let G be a subgroup of a symmetric group acting as permutation of a set Ω , as is with the case with all groups we consider here. We recall the standard definitions that for an element $\alpha \in \Omega$, the stabilizer G_α of α is the set of all group elements that fix α . On the contrary, the set of all fixed points for a group element $x \in G$ is

denoted $C_\Omega(x)$. In short,

$$\begin{aligned} G_\alpha &:= \{x \in G : x(\alpha) = \alpha\} , \\ C_\Omega(x) &:= \{\alpha \in \Omega : x(\alpha) = \alpha\} . \end{aligned} \tag{2}$$

Then,

Definition 1. *The fixed point ratio of $x \in G$, denoted by $fpr(x)$ (which is of course implicitly dependent upon Ω), is the proportion of points in Ω fixed by x , i.e.,*

$$fpr(x) = \frac{|C_\Omega(x)|}{|\Omega|} .$$

For our dataset of all the 2-generated subgroups of S_n up to $n \leq 10$, we notice that whenever a resulting group is simple, the number of fixed points of either of its two generators is never $n - 2$ or $n - 4$ for $n \geq 5$. Now, all our group elements are $n \times n$ permutation matrices (in particular matrices with only 0 and 1) acting on $\{1, 2, \dots, n\}$. Hence, the number of fixed points is counted by the number of 1s on the diagonal. Hence, this observation hints towards the existence of a conjecture regarding traces (or equivalently, number of fixed points) and signs of permutations.

Proposition 1. *(conjectured from machine learning) Consider 2-generated subgroups $\mathbb{H} \subset S_n$ ($n \geq 5$) with distinct non-trivial generators in the permutation representation. That is, consider $\sigma_{i=1,2} (\neq e) \in S_n$ as $n \times n$ permutation matrices. If \mathbb{H} is a simple group, then*

1. $\det(\sigma_i) = 1$, and,
2. $\text{tr}(\sigma_i) \in \{1, 2, \dots, n\} \setminus \{n - 4, n - 2, n - 1, n\}$.

In particular, if $\text{tr}(\sigma_i) = n - 4$, then $\mathbb{H} = \mathbb{D}_{2n} \subseteq \mathbb{A}_n \triangleleft \mathbb{S}_n, n \geq 5$. That is \mathbb{H} is the dihedral group, and thus not simple.

We can prove this conjecture which is inspired from our machine learning experiments.

Proof. First, since \mathbb{H} is a subgroup of S_n as represented by permutation matrices, the determinant of all group elements is equal to ± 1 . Suppose, without loss of generality, $\det(\sigma_1) = -1$. It is straightforward to check that the following defines a homomorphism from \mathbb{H} to $(\mathbb{Z}/(2\mathbb{Z}), + \text{mod } 2)$.

$$\begin{aligned} \psi : \mathbb{H} &\longrightarrow \{0, 1\} \\ \psi(\sigma) &= \begin{cases} 0, & \text{if } \det(\sigma) = 1 , \\ 1, & \text{if } \det(\sigma) = -1 . \end{cases} \end{aligned}$$

By the first isomorphism theorem, $\mathbb{H}/\ker(\psi) \simeq \text{im}(\psi)$, with $\ker(\psi) \triangleleft \mathbb{H}$. Clearly, $\psi(e) = 0$, and $\psi(\sigma_1) = 1$ (by assumption). Therefore, $\text{im}(\psi) = \{0, 1\} = \mathbb{Z}/(2\mathbb{Z})$. As such, $\ker(\psi)$ is an index two subgroup of \mathbb{H} , and therefore normal. Hence, \mathbb{H} is not simple. Thus the determinant of both σ_1 should be 1. This proves the first part of the proposition.

To prove the second part, let's assume \mathbb{H} is simple. We begin by observing that $\text{tr}(\sigma_i) \neq n$ since e is the only element that has trace n and $\sigma_i \neq e$ by assumption. Progressing further, $\text{tr}(\sigma_i) \neq n - 1$ as that would require a single 0 and 1s everywhere else on a diagonal, which is a singular matrix. Finally, we in fact have that $\text{tr}(\sigma_i) \neq n - 2$ since then $\det(\sigma_i) = -1$, coming from the $\begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$ block.

Now consider the case of $\text{tr}(\sigma_i) = n - 4$ and $\det(\sigma_i) = 1$. We will show that $\sigma_i^2 = e$. If this were to be the case, \mathbb{H} would be a group generated by involutions. This would imply that \mathbb{H} is the dihedral group of order $2m$, \mathbb{D}_{2m} (where for $m = 1, 2$ they are, $\mathbb{Z}/(2\mathbb{Z})$ and $\mathbb{Z}/(2\mathbb{Z}) \times \mathbb{Z}/(2\mathbb{Z})$). In any case, \mathbb{H} would not simple (except for the trivial case of $m = 1$, which we excluded by having $n \geq 5$).

To finish the proof, it remains show that

$$P(n) := ((\det(\sigma) = 1) \wedge (\text{tr}(\sigma) = n - 4)) \implies \sigma_i = \sigma^T$$

holds. We do this by induction. We can enumerate to see that $P(4)$ holds. Now assume $P(n)$ holds. Let $\hat{\sigma} \in \mathbb{A}_{n+1}$, so that $\det(\hat{\sigma}) = 1$ (by definition of the alternating group), and $\text{tr}(\hat{\sigma}) = (n + 1) - 4$. We must show that $P(n + 1)$ holds by showing $\hat{\sigma} = \hat{\sigma}^T$. Note that each $\hat{\sigma}$ can be obtained from a σ with the introduction of a 1 in the diagonal at one of $(n + 1)$ possible positions (reflecting the fact that $|\mathbb{A}_{n+1}| = (n + 1)|\mathbb{A}_n|$). Note that this doesn't alter the determinant but increases the trace by 1. It therefore follows that $\hat{\sigma} = \hat{\sigma}^T$ since $\sigma = \sigma^T$, establishing that $P(n + 1)$ holds, completing the proof. \square

The following corollary is a restatement of the above written in terms of fixed point ratios.

Corollary 1. *Let \mathbb{H} be a finite simple group with the generating set containing only two distinct group elements in a permutation representation of degree n . Then, the fixed point ratio (fpr) of any such generator cannot equal $2^i/n$, $\forall i \in \{0, 1, 2\}$.*

Sporadic group	generators	trace	$\{n - 2^k\}_{k=0}^2$
M ₉	(1,4,9,8)(2,5,3,6), (1,6,5,2)(3,7,9,8)	1, 1	5, 7, 8
M ₁₀	(1, 9, 6, 7, 5)(2,10, 3, 8, 4), (1,10, 7, 8)(2, 9, 4, 6)	0, 2	6,8,9
M ₁₁	(1, 2, 3, 4, 5, 6, 7, 8, 9,10,11), (3,7,11,8)(4,10,5,6)	0, 3	7,9,10
M ₁₂	(1,2,3,4,5,6,7,8,9,10,11), (3,7,11,8)(4,10,5,6), (1,12)(2,11)(3,6)(4,8)(5,9)(7,10)	1, 4, 0	8,10,11
M ₂₁	(1, 4, 5, 9, 3)(2, 8,10, 7, 6) (12,15,16,20,14)(13,19,21,18,17), (1,21, 5,12,20)(2,16, 3, 4,17) (6,18, 7,19,15)(8,13, 9,14,11)	0, 3	17,19,20
M ₂₂	(1, 2, 3, 4, 5, 6, 7, 8, 9,10,11) (12,13,14,15,16,17,18,19,20,21,22), (1, 4, 5, 9, 3)(2, 8,10, 7, 6) (12,15,16,20,14)(13,19,21,18,17), (1,21)(2,10, 8, 6)(3,13, 4,17) (5,19, 9,18)(11,22)(12,14,16,20)	1, 1, 1	18,20,21
M ₂₃	(1, 2, 3, 4, 5, 6, 7, 8, 9,10,11,12, 13,14,15,16,17,18,19,20,21,22,23), (3,17,10, 7, 9)(4,13,14,19, 5) (8,18,11,12,23)(15,20,22,21,16)	1, 1	19, 21,22
M ₂₄	(1, 2, 3, 4, 5, 6, 7, 8, 9,10,11,12,13, 14,15,16,17,18,19,20,21,22,23), (3,17,10, 7, 9)(4,13,14,19, 5) (8,18,11,12,23)(15,20,22,21,16), (1,24)(2,23)(3,12)(4,16) (5,18)(6,10)(7,20)(8,14) (9,21)(11,17)(13,22)(15,19)	1, 4, 0	20, 22, 23

Table 1: Properties of Mathieu Group generators. All the generators are of positive signature and the number of fixed points is never $n - 2^k$, $k \in \{0, 1, 2\}$, ala proposition 1.

4 Discussion

In this work, we demonstrate that standard off the shelf machine learning tools such as neural networks can help determine the simplicity of a group with relatively high accuracies and produce novel insights. We conduct a number of machine learning experiments with different mathematical features to determine simplicity. When using permutation representations of group generators, we find modest learning outcomes. In another experiment, using element orders as features alongside the order of the group, we demonstrate remarkable class accuracies of close to 99% in Figure 5, reflective of a known result for finite simple groups, *viz.*, Theorem 1. The success our experiments gives further confidence that the ability of machine-learning (in particular support vector machines) to distinguish simple/non-simple groups when trained on Cayley tables [21] was not a mere artifact of data representation, but truly underlies interesting mathematics.

From these experiments, we distill Proposition 1, which we then prove using properties of finite simple groups. The result restricts choices of 2-generating sets for finite simple groups. We recast this result in terms of fixed point ratios of group generators in Corollary 1. Bounds on fixed point ratios are an important topic of consideration. The above observation feeds directly into studies of fixed point ratios of groups, which have been studied extensively over many decades [53]. The result places restrictions on the traces of generators of any finite simple group. The computational advantage this observation yields is $\approx 10^{-4}\%$ when looking for a generator of the Monster group in the smallest faithful irreducible representation which occurs at $n = 196\,883$. The more interesting point is that the determinant and trace of the generators conspire in some cases to form a dihedral group. Corollary 1 holds for all finite simple groups in the permutation representation. As such it holds for sporadic groups in their permutation representations. In Table 1, we check consistency of our results for the Mathieu group. For this, we note that the traces and determinants of known generators of the Mathieu group are consistent with the allowed values in Corollary 1. We can as well study other presentations of finite groups, and we leave this to future work.

The prospect of finding novel mathematical results is tantalising in the age of AI. There are a number of emerging pathways for mathematical research in light of machine learning. In this work, we exploit one such pathway which involves learning a mathematical property (in this case simplicity) using standard machine learning architectures, and exploiting insights from the learning process to distill a theorem. Often, tools from machine interpretability such as relevance scores can aid in this process resulting in new mathematical insights or theorems [6, 28, 36, 37]. Complementary approaches have been proposed in the recent literature, relying on an organisational principle for mathematical statements, resulting in novel conjectures which could often be proved using domain expertise [51].

Acknowledgments

YHH would like to thank STFC for grant ST/J00037X/2, the Leverhulme Trust for a project grant, as well as Joseph Chuang and Radha Kessar for many helpful discussions. VJ is supported by the South African Research Chairs Initiative of the Department of Science and Innovation and the National Research Foundation. CM is supported by the Accelerate Programme for Scientific Discovery, at the Computer Laboratory, University of Cambridge. CM would like to thank Damián Kaloni Mayorga Penã, Aditya Ravuri, Subhayan Roy Moulik for helpful discussions. The authors would like to thank the Isaac Newton Institute for Mathematical Sciences for support and hospitality during the program “Black holes: bridges between number theory and holographic quantum information” when work on this paper was undertaken; this work was supported by EPSRC grant number EP/R014604/1.

Group Id	Name	Simple?	Generators	
[1, 1]	1	No	$\begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}$	$\begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}$
[2, 1]	\mathbb{Z}_2	Yes	$\begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}$	$\begin{bmatrix} 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}$
[3, 1]	\mathbb{Z}_3	Yes	$\begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}$	$\begin{bmatrix} 0 & 0 & 1 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}$
[4, 1]	\mathbb{Z}_4	No	$\begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}$	$\begin{bmatrix} 0 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{bmatrix}$
[6, 1]	S_3	No	$\begin{bmatrix} 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}$	$\begin{bmatrix} 0 & 0 & 1 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}$
[24, 12]	S_4	No	$\begin{bmatrix} 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}$	$\begin{bmatrix} 0 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{bmatrix}$
[8, 3]	\mathbb{D}_8	No	$\begin{bmatrix} 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}$	$\begin{bmatrix} 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \end{bmatrix}$
[4, 2]	$\mathbb{Z}_2 \times \mathbb{Z}_2$	No	$\begin{bmatrix} 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}$	$\begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{bmatrix}$
[12, 3]	A_4	No	$\begin{bmatrix} 0 & 0 & 1 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}$	$\begin{bmatrix} 0 & 0 & 0 & 1 \\ 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{bmatrix}$

A Machine learning experiments

There were therefore many entries in the dataset for groups equivalent under isomorphism, and the number of entries per-group was not changed to be less imbalanced. This was intentional.²

A.1 2-generated subgroups of S_n

The input data to our models were the permutation matrices corresponding to the generators of each group. We found this to be the most effective representation of the data, as it allowed for fixed-size inputs and produced better performing models than encoding the generators as permutation vectors.

Example: 2-generated subgroups of S_3

Group Id	Name	Simple?	Generators
[1, 1]	1	No	$\begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}$ $\begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}$
[2, 1]	\mathbb{Z}_2	Yes	$\begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}$ $\begin{bmatrix} 0 & 1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 1 \end{bmatrix}$
[3, 1]	\mathbb{Z}_3	Yes	$\begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}$ $\begin{bmatrix} 0 & 0 & 1 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \end{bmatrix}$
[6, 1]	S_3	No	$\begin{bmatrix} 0 & 1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 1 \end{bmatrix}$ $\begin{bmatrix} 0 & 0 & 1 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \end{bmatrix}$

Table 2: Subgroups of S_3 .

Model training and validation were performed on subsets of the *balanced* datasets for 5-fold cross-validation. The mean final validation accuracy for models trained on different portions of the balanced dataset for each n is displayed in [Figure 1](#), from “Results”.

A.2 Datasets

The datasets are given by the generating permutation matrices for each subgroup of some S_n . Only the two-generated subgroups were selected, but all permutations were included, giving a

²To give a sense of scale, the number of subgroups of S_n for $n = [4, 7]$ are 11, 19, 56, 96. The number of entries generated by our methods for each S_n would be 560, 14375, 518364, 25401551.

Group Id	Name	Simple?	Count (filtered)	Count (unfiltered)
[2, 1]	C2	Yes	18	27
[3, 1]	C3	Yes	24	32
[4, 1]	C4	No	30	36
[4, 2]	C2 × C2	No	24	24
[6, 1]	S3	No	72	72
[8, 3]	D8	No	72	72
[12, 3]	A4	No	96	96
[24, 12]	S4	No	216	216

Table 3: Number of each subgroup of S_4 that would be present in the dataset.

total $(n!)^2 - n^2$ individual entries in each dataset. The orders of the groups in these datasets were therefore not evenly distributed; the number of each group included in the dataset for S_4 and S_5 is in [Table 3](#) and [Table 4](#). It was computationally infeasible to include the same information for S_6 and larger.

For all datasets, we used the **SageMath** interface to **GAP** to analyze the groups generated by each pair of permutation matrices. The categorization of whether a group is simple was paired with the generating matrices to produce the entry for that group.

With most of the symmetric groups examined here, it was feasible to produce the categorization from **GAP** for all permutation pairs, which was done by simple iteration. Experiments on S_8 used a random subset of all possible permutation pairs by taking a random sample of integers 1 to $(8!)^2$, mapping these values to permutation matrices, and filtering out pairs with the same permutation. The component of this mapping that produces the permutation³ is specified in [algorithm 1](#).

Group Id	Name	Simple?	Count (filtered)	Count (unfiltered)
[2, 1]	C2	Yes	50	75
[3, 1]	C3	Yes	60	80
[4, 1]	C4	No	150	180
[4, 2]	C2 x C2	No	120	120
[5, 1]	C5	Yes	120	144
[6, 2]	C6	No	220	240
[6, 1]	S3	No	360	360
[8, 3]	D8	No	360	360
[10, 1]	D10	No	360	360
[12, 3]	A4	No	480	480
[12, 4]	D12	No	360	360
[20, 3]	C5 : C4	No	1440	1440
[24, 12]	S4	No	1080	1080
[60, 5]	A5	Yes	2280	2280
[120, 34]	S5	No	6840	6840

Table 4: Number of each subgroup of S_5 present in the dataset.

Algorithm 1: Conversion from $k \in [0, n!)$ to a permutation of n elements.

Input: An integer $k \in [0, n!)$
Output: A permutation of the elements $0, 1, \dots, n$
`vs := [] // The values $0, \dots, n$`
`ps := [] // The final permutation`
for $i \leftarrow 0$ **to** $n - 1$ **do**
 `vs.append(i)`
end
for $i \leftarrow n$ **to** 1 **do**
 `$r :=$ remainder of $\frac{k}{i}$`
 `$k \leftarrow k \bmod i$`
 `$v :=$ vs[r]`
 Remove vs[r] and shift remaining elements to the left
 `ps.append(v)`
end
return `ps`

A single sample in one of the datasets could be constructed as follows. Suppose we are generating data for the subgroups of S_4 , with the permutations given by the integer pair $(6, 19)$. The corresponding permutations, as generated by the aforementioned algorithm, are:

$$(2, 1, 0, 3) \text{ and } (3, 1, 2, 0) \quad (3)$$

which generate the following permutation matrices, respectively:

$$\begin{pmatrix} 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} \text{ and } \begin{pmatrix} 0 & 0 & 0 & 1 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 1 & 0 & 0 & 0 \end{pmatrix} \quad (4)$$

The inputs for the training sample are just the concatenation of each row in the two matrices into a single vector, as shown below:

$$\langle 0, 0, 1, 0, 0, 1, 0, 0, 1, 0, 0, 0, 0, 0, 0, 1, 0, 0, 0, 1, 0, 1, 0, 0, 0, 0, 1, 0, 1, 0, 0, 0 \rangle \quad (5)$$

The particular group corresponding to these inputs S_3 , which is not simple.

Training datasets

The generated datasets were unbalanced; only approximately 20% of the entries correspond to simple groups in each dataset.⁴ From here on, we refer only to the *balanced* datasets, which are

⁴Constructing the permutation matrix from the corresponding permutation of n elements is trivial; we leave this undiscussed.

chosen by a random selection of the non-simple entries in equal number to the set of simple ones.

All experiments were done with 5-fold cross-validation, with each train/validation split taking from the same balanced subset of the unbalanced dataset. For experimentation on differing amounts of data (*e.g.*, using 50% of the “full”, balanced dataset), all balanced datasets were subsets of the “full” balanced dataset.

A.3 Neural network architectures




Figure 6: Generic architecture of the models used. The generating permutation matrices are flattened and provided as input. Softmax is applied to the two outputs, which give probabilities for whether the group is simple or not.

The models used for experimentation were fully-connected neural networks with the two input permutation matrices flattened into a single vector of $2n^2$ boolean values for each S_n . Predictions are chosen by the greater of two output values, scaled by the softmax activation function. Hidden layers for all final models consist of 256 nodes, with all nodes using the ReLU activation function. See Figure 6 above for an illustration.

Optimization was done with stochastic gradient descent (SGD) with Nesterov-accelerated momentum and Mean Squared Error as the loss function. The hyperparameters for SGD (both learning-rate and momentum), the number of hidden layers for each n , and the size of the hidden layers were selected by manual hyperparameter search. The results with the best mean validation accuracy across all five validation datasets at the end of training. A hidden layer size of 256 was chosen for all n as a common best size for $n = 5$ and $n = 6$.

For $n = 7$ and $n = 8$, rigorous hyperparameter search was infeasible due to the size of the

⁴For each value of $n \in \{5, \dots, 8\}$, these fractions are 0.1758, 0.2027, 0.2315, and 0.2133, respectively.

n	Learning rate	Momentum	Num. hidden layers	γ
5	0.05	0	1	0.1
6	0.001	0.01	3	0.05
7*	0.01	0.1	9	0.05
8*	0.01	0.1	9	0.05

Table 5: Hyperparameters used for each dataset.

datasets and the resulting computational restrictions – less so with the former than the latter. For these datasets, the hyperparameters were initially chosen by extrapolation from those used with the lower two values for n , without attempting to produce optimal results. Adjustments were made only until the model was successfully better than random chance. This less rigorous iteration was required as, for some sets of hyperparameters, the models completely fail to perform any meaningful learning (often skewing entirely towards one output or another). This occurred most prominently at $n = 7$, when testing with five or fewer hidden layers. This effect was avoided by increasing the number of hidden layers, which did not prove to be additionally necessary for $n = 8$ (*i.e.*, there was no increase from $n = 7$ to 8).

In addition to the hyperparameters previously mentioned, learning rate decay was also provided, with the learning rate at each epoch calculated from the previous by:

$$\text{lr}_{i+1} = (1 - \gamma)\text{lr}_i \tag{6}$$

with γ as the hyperparameter controlling the rate of decay. The hyperparameters used for each n are given above, in [Table 5](#).

References

- [1] Y.-H. He, *Deep-Learning the Landscape*, [1706.02714](#).
- [2] J. Carifio, J. Halverson, D. Krioukov and B. D. Nelson, *Machine learning in the string landscape*, *JHEP* **09** (2017) 157, [[1707.00655](#)].
- [3] D. Krefl and R.-K. Seong, *Machine learning of Calabi-Yau volumes*, *Phys. Rev.* **D96** (2017) 066014, [[1706.03346](#)].
- [4] F. Ruehle, *Evolving neural networks with genetic algorithms to study the string landscape*, *JHEP* **08** (2017) 038, [[1706.07024](#)].
- [5] Y.-H. He, *Machine-learning the string landscape*, *Phys. Lett.* **B774** (2017) 564–568.
- [6] Y.-H. He, *The Calabi–Yau Landscape: From Geometry, to Physics, to Machine Learning*. Lecture Notes in Mathematics. 5, 2018, [10.1007/978-3-030-77562-9](#).
- [7] F. Ruehle, *Data science applications to string theory*, *Phys. Rept.* **839** (2020) 1–117.
- [8] A. Jain, C. Mishra and P. Liò, *A physics-informed search for metric solutions to ricci flow, their embeddings, and visualisation*, *arXiv preprint arXiv:2212.05892* (2022) .
- [9] K. Bull, Y.-H. He, V. Jejjala and C. Mishra, *Machine Learning CICY Threefolds*, *Phys. Lett. B* **785** (2018) 65–72, [[1806.03121](#)].
- [10] K. Bull, Y.-H. He, V. Jejjala and C. Mishra, *Getting CICY High*, *Phys. Lett. B* **795** (2019) 700–706, [[1903.03113](#)].
- [11] P. Berglund, B. Campbell and V. Jejjala, *Machine Learning Kreuzer-Skarke Calabi-Yau Threefolds*, [2112.09117](#).
- [12] A. Ashmore, Y.-H. He and B. A. Ovrut, *Machine Learning Calabi–Yau Metrics*, *Fortsch. Phys.* **68** (2020) 2000068, [[1910.08605](#)].
- [13] D. Peifer, M. Stillman and D. Halpern-Leistner, *Learning selection strategies in buchberger’s algorithm*, in *International Conference on Machine Learning*, pp. 7575–7585, PMLR, 2020.
- [14] L. B. Anderson, M. Gerdes, J. Gray, S. Krippendorf, N. Raghuram and F. Ruehle, *Moduli-dependent Calabi-Yau and $SU(3)$ -structure metrics from Machine Learning*, [2012.04656](#).
- [15] M. R. Douglas, S. Lakshminarasimhan and Y. Qi, *Numerical Calabi-Yau metrics from holomorphic networks*, [2012.04797](#).

- [16] V. Jejjala, D. K. Mayorga Pena and C. Mishra, *Neural Network Approximations for Calabi-Yau Metrics*, [2012.15821](#).
- [17] M. Larfors, A. Lukas, F. Ruehle and R. Schneider, *Learning Size and Shape of Calabi-Yau Spaces*, [2111.01436](#).
- [18] A. Ashmore, L. Calmon, Y.-H. He and B. A. Ovrut, *Calabi-Yau Metrics, Energy Functionals and Machine-Learning*, [2112.10872](#).
- [19] M. Larfors, A. Lukas, F. Ruehle and R. Schneider, *Numerical metrics for complete intersection and Kreuzer–Skarke Calabi–Yau manifolds*, *Mach. Learn. Sci. Tech.* **3** (2022) 035014, [[2205.13408](#)].
- [20] P. Berglund, G. Butbaia, T. Hübsch, V. Jejjala, D. Mayorga Peña, C. Mishra et al., *Machine Learned Calabi–Yau Metrics and Curvature*, [2211.09801](#).
- [21] Y.-H. He and M. Kim, *Learning Algebraic Structures: Preliminary Investigations*, [1905.02263](#).
- [22] J. Bao, S. Franco, Y.-H. He, E. Hirst, G. Musiker and Y. Xiao, *Quiver Mutations, Seiberg Duality and Machine Learning*, *Phys. Rev. D* **102** (2020) 086013, [[2006.10783](#)].
- [23] L. Amorós, O. Gasanova and L. Jakobsson, *A machine learning approach to commutative algebra: Distinguishing table vs non-table ideals*, *arXiv preprint arXiv:2109.11417* (2021) .
- [24] P.-P. Dechant, Y.-H. He, E. Heyes and E. Hirst, *Cluster Algebras: Network Science and Machine Learning*, [2203.13847](#).
- [25] M. C. Hughes, *A neural network approach to predicting and computing knot invariants*, *Journal of Knot Theory and Its Ramifications* **29** (2020) 2050005.
- [26] V. Jejjala, A. Kar and O. Parrikar, *Deep Learning the Hyperbolic Volume of a Knot*, *Phys. Lett. B* **799** (2019) 135033, [[1902.05547](#)].
- [27] S. Gukov, J. Halverson, F. Ruehle and P. Sulkowski, *Learning to Unknot*, [2010.16263](#).
- [28] J. Craven, V. Jejjala and A. Kar, *Disentangling a Deep Learned Volume Formula*, [2012.03955](#).
- [29] J. Craven, M. Hughes, V. Jejjala and A. Kar, *Learning knot invariants across dimensions*, [2112.00016](#).
- [30] J. Craven, M. Hughes, V. Jejjala and A. Kar, *Illuminating new and known relations between knot invariants*, [2211.01404](#).

- [31] Y.-H. He and S.-T. Yau, *Graph Laplacians, Riemannian Manifolds and their Machine-Learning*, [2006.16619](#).
- [32] J. Bao, Y.-H. He, E. Hirst, J. Hofscheier, A. Kasprzyk and S. Majumder, *Polytopes and Machine Learning*, [2109.09602](#).
- [33] L. Alessandretti, A. Baronchelli and Y.-H. He, *Machine Learning meets Number Theory: The Data Science of Birch-Swinnerton-Dyer*, [1911.02008](#).
- [34] Y.-H. He, E. Hirst and T. Peterken, *Machine-learning dessins d'enfants: explorations via modular and Seiberg–Witten curves*, *J. Phys. A* **54** (2021) 075401, [[2004.05218](#)].
- [35] Y.-H. He, K.-H. Lee and T. Oliver, *Machine-Learning the Sato–Tate Conjecture*, [2010.01213](#).
- [36] Y.-H. He, *Machine-Learning Mathematical Structures*, [2101.06317](#).
- [37] A. Davies, P. Veličković, L. Buesing, S. Blackwell, D. Zheng, N. Tomašev et al., *Advancing mathematics by guiding human intuition with ai*, *Nature* **600** (2021) 70–74.
- [38] E. Noether, *Invariant Variation Problems*, *Gott. Nachr.* **1918** (1918) 235–257, [[physics/0503066](#)].
- [39] V. I. Arnol'd, *Mathematical methods of classical mechanics*, vol. 60. Springer Science & Business Media, 2013.
- [40] S. Weinberg, *The quantum theory of fields*, vol. 1,2. Cambridge university press, 1995.
- [41] T. S. Cohen and M. Welling, *Group equivariant convolutional networks*, *Proceedings of The 33rd International Conference on Machine Learning, Proceedings of Machine Learning Research* **48** (2016) 2990–2999.
- [42] R. Kondor and S. Trivedi, *On the generalization of equivariance and convolution in neural networks to the action of compact groups*, in *International Conference on Machine Learning*, pp. 2747–2755, PMLR, 2018.
- [43] H.-Y. Chen, Y.-H. He, S. Lal and S. Majumder, *Machine learning Lie structures & applications to physics*, *Phys. Lett. B* **817** (2021) 136297, [[2011.00871](#)].
- [44] M.-W. Cheung, P.-P. Dechant, Y.-H. He, E. Heyes, E. Hirst and J.-R. Li, *Clustering Cluster Algebras with Clusters*, [2212.09771](#).
- [45] K.-H. Lee, *Data-scientific study of kronecker coefficients*, 2023.
- [46] D. Gorenstein, *Finite groups*, vol. 301. American Mathematical Soc., 2007.

- [47] D. Bernstein, *The computational complexity of rules for the character table of sn* , *Journal of Symbolic Computation* **37** (2004) 727–748.
- [48] J. D. Dixon, *The probability of generating the symmetric group*, *Mathematische Zeitschrift* **110** (1969) 199–205.
- [49] M. Aschbacher and R. Guralnick, *Some applications of the first cohomology group*, *Journal of Algebra* **90** (1984) 446–460.
- [50] A. Shalev, *Asymptotic group theory*, *Notices of the AMS April* (2001) 383–389.
- [51] C. Mishra, S. R. Moulik and R. Sarkar, *Mathematical conjecture generation using machine intelligence*, *arXiv preprint arXiv:2306.07277* (2023) .
- [52] A. V. Vasil'ev, M. A. Grechkoseeva and V. D. Mazurov, *Characterization of the finite simple groups by spectrum and order*, *Algebra and Logic* **48** (2009) 385–409.
- [53] T. C. Burness, *Simple groups, fixed point ratios and applications*, *arXiv preprint arXiv:1707.03564* (2017) .