Distinguishing elliptic fibrations with AI

Yang-Hui He [a,b,c], Seung-Joo Lee [d]

[a] *Department of Mathematics, City, University of London, EC1V0HB, UK*
[b] *Merton College, University of Oxford, OX14JD, UK*
[c] *School of Physics, NanKai University, Tianjin, 300071, China*
[d] *CERN, Theory Department, 1 Esplande des Particules, Geneva 23, CH-1211, Switzerland*

ABSTRACT

We use the latest techniques in machine-learning to study whether from the landscape of Calabi-Yau manifolds one can distinguish elliptically fibred ones. Using the dataset of complete intersections in products of projective spaces (CICY3 and CICY4, totalling about a million manifolds) as a concrete playground, we find that a relatively simple neural network with forward-feeding multi-layers can very efficiently distinguish the elliptic fibrations, much more so than using the traditional methods of manipulating the defining equations. We cross-check with control cases to ensure that the AI is not randomly guessing and is indeed identifying an inherent structure. Our result should prove useful in F-theory and string model building as well as in pure algebraic geometry.
© 2019 The Authors. Published by Elsevier B.V. This is an open access article under the CC BY license (http://creativecommons.org/licenses/by/4.0/). Funded by SCOAP³.

1. Introduction and summary

Ever since the birth of string theory its compactifications on compact Calabi-Yau manifolds have been a subject of constant interest to theoretical physicists and algebraic geometers alike. Such compactifications provide a well-controlled setup to study supersymmetric effective theories, many physical properties of which can be calculated via algebro-geometric tools without explicit Ricci flat metrics at hand. Tremendous efforts to find realistic models of string theory in this setup have been one of the most important aims of string phenomenology; see [1] for an earliest attempt. Furthermore, in line with the so-called swampland program [2], geometric origins of quantum gravity constraints have recently started to be addressed from the universal properties of general Calabi-Yau manifolds in various string theoretic setups [3–11].

More generally, non-Ricci-flat manifolds may also lead to supersymmetric compactifications, with an appropriately varying axio-dilaton profile turned on, in the framework of F-theory [12]. This provides arguably the most general geometric approach currently available to study non-perturbative string compactifications. Interestingly, elliptically fibered Calabi-Yau manifolds of one-dimension higher, which are fictitious at least from the Type IIB point of view, can completely specify such compactifications. It is therefore

of utmost interest to string theorists to better understand elliptic Calabi-Yau manifolds.

Elliptic fibration structures often bring in a computational power as well. For instance, the first exact MSSM particle content directly from string theory compactifications arose from an elliptically fibred Calabi-Yau threefold, a quotient of the so-called Schön manifold, whose elliptic-fibration structure was what allowed the explicit computation of equivariant cohomology in order to obtain the matter representations [13–16].

In the meantime various dualities amongst Calabi-Yau compactifications have been influential in exciting development in enumerative and algebraic geometry [17,18] as well as many other branches of pure mathematics. Interestingly, such string dualities oftentimes base on elliptic fibration structures within the internal Calabi-Yau manifolds; see e.g. [19].

While it still remains an open question if Calabi-Yau n-folds form a finite set for $n \geq 3$, finiteness of elliptic ones with dimension $n = 3$ has been established [20,21] (see also [22] for relevant results for $n = 4$ and 5). At the same time, recent investigations of vast Calabi-Yau datasets have observed ubiquity of elliptic fibrations [23–30]. Hence, also in the hope of providing a meaningful measure for the potential paucity of non-ellipticity, there arises a pressing question of distinguishing the elliptic Calabi-Yau manifolds from the non-elliptic.

In principle, there are methods to determine whether a Calabi-Yau manifold of dimension not bigger than 3 is an elliptic fibration [31,32], which were conjectured to work also for a higher-

dimensional manifold [33]. However, the computation, as is with any computation in algebraic geometry, could often become rather intense in practice because one needs to first determine the Kähler cone of the geometry in question to apply the methods. Another source of difficulty is that manipulating high degree polynomials in many variables is a non-linear and complicated matter [34].

In [35,39] a paradigm was proposed to attempt to use artificial intelligence (AI) to bypass expensive algorithms in computational geometry, in particular to study the string landscape and beyond. It was found that central problems such as computing cohomology of vector bundles appear to be machine-learnable to very high precision. Indeed, [35–39] brought machine learning to the landscape and there has subsequently been a host of activity successfully addressing various problems in the string landscape using AI techniques and machine-learning [36–38,40–57] (cf. [58] for a pedagogical introduction). It is therefore natural to ask whether machine-learning techniques can be employed into our present problem of recognizing elliptic fibrations.

The purpose of this letter is to report that this problem of recognizing elliptic fibrations within Calabi-Yau manifolds (and presumably more arbitrary dataset of algebraic varieties), using complete intersection 3-folds and 4-folds within products of projective spaces as a playground, appear to belong to the class of problems which can be addressed by machine-learning to high precision. Thus, like computing cohomology of bundles over varieties, distinguishing elliptic fibrations appears to be a pattern recognizable by the likes of a neural network - completely without any knowledge of algebraic geometry or expensive algorithms needed to deterministically address the problem.

The outline of the letter is as follows. In Section 2, we briefly set the scene by reminding the reader of the CICY (complete intersection Calabi-Yau manifolds in products of projective spaces) dataset on which we will focus for concreteness, and of the elliptic fibration structures within the set. Then, in Section 3, we establish a neural network to efficiently learn and identify which of the set are (not) elliptically fibred, for the 3-fold and 4-fold cases, as well as a contrived control case which gives a good sanity check. Finally, in Section 4, we end with conclusions and outlook.

Acknowledgments

2. CICYs and elliptic fibrations

As stated in the introduction, we will focus on the so-called CICY dataset for concreteness. These are smooth Calabi-Yau manifolds embedded as complete intersection in products of projective spaces; we start by briefly recalling their construction.

Let us first consider a complete intersection $\mathcal{M}$ of K polynomials, $p_{\alpha=1,\cdots,K}$, in the product ambient space, $\mathcal{A} = \mathbb{P}^{n_1} \times \cdots \times \mathbb{P}^{n_m}$. By construction the complex dimension n of $\mathcal{M}$ is given as $n = \sum_{r=1}^{m} n_r - K$. Denoting the degrees of the defining polynomials p_α in the r-th projective spaces by a_α^r, we can describe a family of such geometries by a *configuration* matrix M of the form,

$$M := [\, \mathbf{n} \,|\, \{\mathbf{a}_\alpha\} \,] = \begin{bmatrix} \mathbb{P}^{n_1} & a_1^1 & \cdots & a_K^1 \\ \vdots & \vdots & \ddots & \vdots \\ \mathbb{P}^{n_m} & a_1^m & \cdots & a_K^m \end{bmatrix}, \tag{1}$$

Table 1
Obvious elliptic fibrations (OEF) and elliptic fibrations (EF) in the CICY3 and CICY4 datasets with $SU(3)$ and $SU(4)$ holonomies, respectively.

	Total	not OEF	not EF	Remark
CICY3	7,868	53	53	$h^{1,1} > 4 \Rightarrow$ OEF
CICY4	905,684	462	not known	$h^{1,1} > 12 \Rightarrow$ OEF

with non-negative integer entries a_α^r (cf. see [59] for a generalization to include "negative degrees", dubbed as gCICYs, and also [60] for an example study of fibration structures thereof). Such a complete intersection is a Calabi-Yau n-fold when $\sum_{\alpha=1}^{K} a_\alpha^r = n_r + 1$ for every $r = 1, \ldots, m$.

In this letter, we will be concerned with CICY 3-folds and 4-folds, for which 7,890 and 921,497 configurations were classified, respectively, in [61] and [62]. In particular, we will focus on elliptic fibration structures that those CICYs may admit. As first studied systematically in [23], there is a simple combinatorial method to find elliptic fibrations in terms of their configuration matrices (1). That is, a CICY $\mathcal{M}$ has an obvious elliptic fibration (OEF) if, via row and column permutations, its configuration M can be put in the form,

$$\begin{bmatrix} \mathcal{A}_1 & 0 & \mathcal{F} \\ \mathcal{A}_2 & \mathcal{B} & \mathcal{T} \end{bmatrix}, \tag{2}$$

with the sub-configuration $F := [\mathcal{A}_1 \,|\, \mathcal{F}]$ describing a subvariety of dimension one. Here, $\mathcal{A}_1$ and $\mathcal{A}_2$ are products of m_1 and $m_2 := m - m_1$ projective spaces, respectively, while $\mathcal{F}$, $\mathcal{B}$ and $\mathcal{T}$ are sub-block matrices.

While not every elliptic fibration is necessarily an OEF, upon a classification of fibrations [24], it has been observed that all elliptic CICY 3-folds do admit an OEF. To be specific, 99.33% (all but 53) of the 7,868 CICY 3-fold configurations with genuine $SU(3)$ holonomy admit an OEF structure, where the 53 not only lack an OEF but also cannot be elliptic at all. This classification of elliptic fibrations (EFs) within the CICY3 dataset, as opposed to OEFs, was achieved based on the following topological criteria conjectured in [33]:

A Calabi-Yau n-fold $\mathcal{M}$ is elliptic iff there exists a $(1,1)$-class $D \in H^2(\mathcal{M}, \mathbb{Q})$ such that (a) $D \cdot C \geq 0$ *for every algebraic curve C;* (b) $D^{n-1} \neq 0$; (c) $D^n = 0$.
This has been proven for the 3-fold case subject to the constraints that D is effective or $D \cdot c_2(\mathcal{M}) \neq 0$ [31,32].

Similarly, OEFs of CICY 4-folds were classified [23], resulting in the observation that 99.95% (all but 462) of the 905,684 4-fold configurations with $SU(4)$ holonomy admit an OEF. Although an analogous classification of EFs have not been undertaken (and hence, it is a priori not clear whether or not there exists an elliptic CICY 4-fold lacking an OEF), we will nevertheless use the existence of an OEF as a measure for the ellipticity of a CICY 4-fold when training AI. Table 1 summarizes relevant statistics and characteristic features [23,24] on OEF and EF for the CICY3 and CICY4 cases.

Finally, it is worth noting that not every elliptic fibration admits a section. For instance, the bi-cubic 3-fold, $\begin{bmatrix} \mathbb{P}^2 & 3 \\ \mathbb{P}^2 & 3 \end{bmatrix}$, can be viewed as a fibration of elliptic curves over either of the two $\mathbb{P}^2$ factors, and hence is elliptic. However, it does not admit a section but only a tri-section. Such elliptic fibrations with only a multi-section are sometimes called genus-one fibrations to emphasize the absence of a section. In this letter, however, an elliptic fibration refers to any fibration of elliptic curves over a base manifold, whether there exists a section or not.

3. Machine-learning elliptic fibrations

As already discussed, elliptic fibrations are *typical* within the landscape of known Calabi-Yau manifolds. For instance, all CICY 3-folds with $h^{1,1} > 4$ are elliptic; see Table 1 for relevant statistics for both CICY3 and CICY4 datasets. With small Hodge number, however, there are distinguished ones which are not elliptic. In particular, any cyclic manifold with $h^{1,1} = 1$ (e.g. the Quintic 3-fold) clearly cannot be elliptic.

In the meantime most model-building approaches in string phenomenology have based on manifolds of *small $h^{1,1}$* (or the mirror case of small $h^{2,1}$). This is because there are fewer parameters for them which enter into the effective potential [63–68]. We will also focus especially on small-$h^{1,1}$ manifolds since those with *big enough $h^{1,1}$* are essentially all elliptic; information on (non-)ellipticity can thus be thought of as being learnable mainly through such simpler manifolds. Specifically, the training set will consist of the ellipticity data for CICYs with $h^{1,1}$ up to a certain upper bound, as detailed in the following.

3.1. CICY threefolds

There is a total of 643 CICY 3-folds with $h^{1,1} \leq 4$, 53 of which are not elliptically fibred [24]. Let us test whether standard machine-learning algorithms can distinguish these. Because we have an *unbalanced* dataset, where those which *are* elliptic exceed those which *are not* by one order of magnitude, some enhancement is needed. Luckily, any CICY configuration is equivalent to an arbitrary row and/or column permutation; this gives us a natural way to enhance the number 53. We take 10 random permutations of the rows and the columns each for the 53 configuration matrices and 3 such permutations for the $643 - 53 = 590$. This gives us a total of $10^2 \cdot 53 = 5300$ cases of 0 (non-elliptic) and $3^2 \cdot 590 = 5310$ cases of 1 (elliptic). Furthermore, we note that the maximal number of row and columns are respectively 6 and 7 for these configuration matrices, so we pad with zeros where necessary so that each data-point is of the form

$$M_{6\times 7} \longrightarrow 1 \text{ or } 0 . \qquad (3)$$

In fact, the above counting of enhanced dataset is only an upper bound due to symmetries in the configurations. Nevertheless, we end up with a total labelled dataset $\mathcal{D}$ of size around 10K.

We now have a binary classification problem, for which we perform supervised machine-learning on a training set $\mathcal{T}$ of size $x\%$ and validate on the *unseen* or complementary $(100 - x)\%$ validation set . By validation we mean to construct a 2×2 confusion matrix C where we record the number of false/true positives/negatives of the actual values compared to those predicted by the machine [69]:

		Actual	
		True	False
Predicted	True	True Positive (tp)	False Positive (fp)
Classification	False	False Negative (fn)	True Negative (tn)

From C we have certain "goodness-of-prediction/fit". There is the obvious "precision":

$$\text{Precision} := \frac{tp}{tp + fp} \qquad (4)$$

which is the percentage of correct prediction. However, to avoid assigning too much weight to false positives, the standard accuracy measure to use is the Matthew's correlation coefficient ϕ for C:
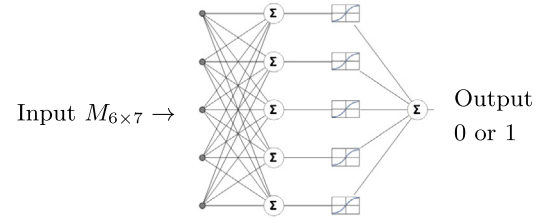

Fig. 1. A graphical schematic of the neural network used (we have combined the element-wise Tanh and Sigmoid layers).
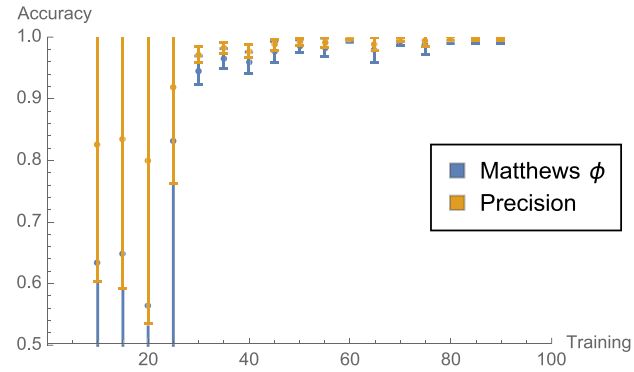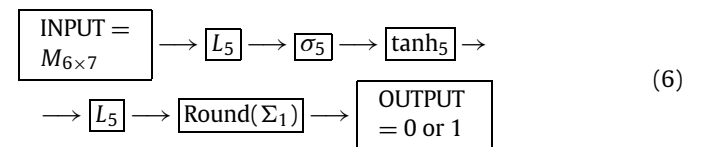

Fig. 2. The learning curve for the (enhanced) 10K data for $h^{1,1} \leq 4$ CICY 3-folds on distinguishing elliptic fibrations.

$$\phi := \frac{tp \cdot tn - fp \cdot fn}{\sqrt{(tp + fp)(tp + fn)(tn + fp)(tn + fn)}} \qquad (5)$$

which is related to the chi-squared χ^2 of C as $|\phi| = \sqrt{\frac{\chi^2}{n}}$ where n is the total validation size (i.e., sum of entries of C). A value of $\phi = +1$ means completely accurate prediction, -1, complete anti-correlation and 0, random guess [70].

For the actual AI, we choose a neural network which was found to be highly efficient in binary problems in computational geometry in [35,39], viz., a 4-layer-perceptron consisting of (1) a fully connected linear layer L_5 taking the 6×7 matrix input by linear transformation to 5 nodes; (2) an element-wise layer of the sigmoid function $\sigma(z) = (1 - \exp(-z))^{-1}$; (3) an element-wise layer of the hyperbolic tangent function [74]; (4) another fully-connected linear layer L_5 and (5) a summation layer of the 5 nodes into a single real output between 0 and 1, on which we then use integer round to return 0 or 1. Schematically, our multi-layer perceptron [75] looks like

$$\boxed{\begin{array}{c}\text{INPUT} = \\ M_{6\times 7}\end{array}} \longrightarrow \boxed{L_5} \longrightarrow \boxed{\sigma_5} \longrightarrow \boxed{\tanh_5} \rightarrow$$
$$\longrightarrow \boxed{L_5} \longrightarrow \boxed{\text{Round}(\Sigma_1)} \longrightarrow \boxed{\begin{array}{c}\text{OUTPUT} \\ = 0 \text{ or } 1\end{array}} \qquad (6)$$

A graphic representation of the layers of the neural network is given in Fig. 1.

We cross-checked this neural network with other methods such as a classifier with an optimized mixture of decision trees and regressions (implemented in the vanilla version of Classify[ ] function in [71]) and find these also to perform to comparable precision.

We present the learning curve of the accuracy ϕ versus the percentage x of seen training set in Fig. 2. The error bars come from the fact that each training set at a given percentage is randomly chosen 10 times for statistical stability. We see from the figure that the neural network works very well with respect to both of the accuracy measures. Initially, up to about 25% there is a lot of
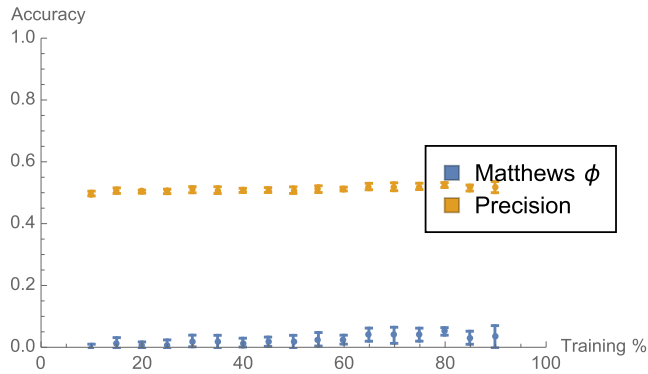
Fig. 3. The learning curve for the (enhanced) 10K data for $h^{1,1} \leq 4$ CICY 3-folds on a control set of a randomly chosen property.

Fig. 4. The learning curve for the (enhanced) 1.54 M data for $h^{1,1} \leq 12$ CICY 4-folds on distinguishing elliptic fibrations.

fluctuation because too few data has been seen to learn anything meaningful. However, starting from around 30%, we see that both naive precision (% of agreement of 0s and 1s) and Matthews' ϕ approach 1. What is remarkable is that even at a relatively low percentage of seen data, we are achieving high accuracy. For example, having seen only about 30% of the total 10K of data of what is an elliptic fibration, the neural network correctly predicts the unseen 70% of cases to over 99% accuracy. Importantly, the computation for each training takes only a few seconds, a vast improvement over any brute-force check of whether a manifold admits an elliptic fibration structure. In summary, determining whether a CICY is elliptically fibred indeed seems to be a machine-learnable problem.

3.2. Control case

It was argued in [35,39] that computational problems in algebraic geometry tend to fall into the machine-learnable category because essentially *any problem* therein boils down to finite steps of finding kernels and cokernels of integer matrices, a task in which the likes of neural networks excel (problems in number theory, in contrast, seem much less amenable to AI). The machine certainly knows nothing about manifolds or bundle fibrations, all it is doing is to catalogue the positive/negative cases together and to manage to spot a pattern not immediately obvious to the human eye (or, for that matter, to the standard methods of algebraic geometry). One might question then whether if one divides any random property into any random subsets, the machine might pick up something, whereby making any specific queries such as elliptic fibration structures rather ineffectual.

It is therefore a good and important thing to test a "control case" as follows. Suppose of the 643 CICY 3-fold configurations, we select 53 *arbitrarily* and *randomly* and assign a property as "0" while all other cases we will call "1" in complete imitation to the above problem and repeat the machine-learning procedure. If elliptic fibration is truly *not a random property*, then the AI should perform poorly in accuracy to this artificial problem.

The learning curve is presented in Fig. 3. It is reassuring that the results are in agreement with complete randomness. The naive precision stays at 50%, consistent with a random guess and subsequently, the ϕ coefficient remains essentially at 0. The control case thus shows that a random assignment of a property does not and should not allow the machine to spot any inherent pattern.

3.3. CICY fourfolds

Since the machine knows nothing about algebraic geometry, there is no reason to stop at CICY 3-folds. The next dataset, which is even more significant to realistic F-theory vacua, is the Calabi-
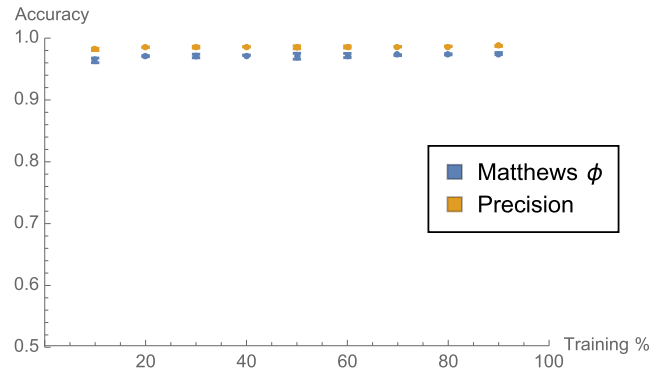
Yau 4-folds arising as complete intersections in products of projective spaces. As reviewed in Section 2, $921,497$ CICY 4-fold configurations were classified in [62] and their OEF structures in [23]; see Table 1 for the relevant statistics.

Up to $h^{1,1} = 12$, there is now a total of 767,642 configurations, 767,180 of which admit elliptic fibrations. Here, everything is a 12×16 integer configuration matrix, upon padding with zeros where necessary as in the 3-fold case. Again, due to the paucity of non-elliptic fibrations, we enhance data by permutations: we randomly permute the non-elliptic ones by row/columns $140^2 = 19600$ times and the elliptic ones by $1^2 = 1$, thereby obtaining a more balanced $\sim 767K$ for 1 (elliptic) and 774K for 0 (non-elliptic) upon removal of redundancies. The learning curve for the CICY 4-folds is presented in Fig. 4. The (small) error bars are due to the fact that each training set at a given percentage $x = 10\%, \ldots, 90\%$ is randomly chosen 5 times for stability. Impressive is the fact that even better than the 3-fold case, the neural network correctly predicts the unseen cases to over 96% accuracy already at $x = 10\%$.

We performed a similar control case for the CICY 4-folds, i.e., taking a random sample of $767642 - 767180 = 462$ cases and declaring them to have property "0" and the remainder, property "1". After data enhancement by permutation, more on the "0" than the "1", we find a training curve identical to Fig. 3. That is, when there is no pattern, the neural network will give a completely random guess with Matthew's phi equalling 0.

Interestingly, if we train, on the CICY 3-fold set, on which ones are elliptically fibred and attempt to predict which CICY 4-folds are elliptic, the network consistently (over repeated trials) predicts that all CICY 4-folds are elliptically fibred. This is actually not too far from the truth since non-elliptic fibration is comparatively rare. However, working within the *same* dataset (using 3-folds to predict about 3-folds, and 4-folds, about 4-folds) actually picks out which of the specific rare cases that are not elliptic.

4. Conclusions and outlook

In this letter, we have demonstrated that distinguishing whether a complete intersection Calabi-Yau manifold is elliptic is a problem well suited for machine-learning. The fact that we have used a large dataset of specific Calabi-Yau manifolds is only for the purposes of being concrete; it is expected that the general problem of identifying fibration structure for arbitrary varieties should also be addressable by the likes of neural networks.

On another matter, distinguishing the elliptic fibrations with a section from those without one is in general a complicated task in algebraic geometry [72]. It will be interesting to see if machine-learning techniques can be employed into such a task. Yet another

interesting aspect of elliptic CICYs is that they in general admit more than one fibration structures [19], which was fully analyzed for CICY 3-folds [24]. Machine-learning aspects of such an enumeration problem, we leave to future investigation.

In some sense, the key question of [35,39] is "what class of problems in mathematics and computation can be machine-learned". In particular, it was found that typical problems in algebraic geometry, such as cohomology calculation or toric variety combinatorics, seem to be better suited for a neural network than for traditional methods. After all, as far as the AI is concerned, it only knows a trainable set of inputs (configuration matrices of manifolds or bundles, etc.) and outputs (ranks of cohomology groups, dimensions of loci, yes/no answer of existence of fibrations), perfectly adapted for supervised learning. From this, the AI finds some optimal arrangement, by regression, classifiers, neural networks, decision trees, etc., in order to guess the answer in general.

This is at once mysterious and invigorating. As is well-known, computational problems in algebraic geometry ultimately suffers from the exponentially expensive steps of finding Groebner bases and syzygies, because of the high degree multi-variate polynomials involved. Very quickly, one would find the desired problem intractable even with the most advanced algorithms and computers. That the neural network is able to find the correct answer in a matter of seconds (precisely because it knows nothing of the mathematics) suggests that there is some underlying method/structure which is yet to be discovered. In this sense, our "control experiment", not emphasized in previous studies, is important. We found that for a random assignment of fictitious property, the AI gives completely random guesses ($\phi \simeq 0$, precision $\simeq 50\%$). This suggests that there truly exists an underlying pattern to meaningful properties such as whether a manifold is an elliptic fibration.

References

[1] P. Candelas, G. Horowitz, A. Strominger, E. Witten, Nucl. Phys. B 258 (1985) 46.
[2] C. Vafa, arXiv:hep-th/0509212.
[3] S.J. Lee, W. Lerche, T. Weigand, J. High Energy Phys. 1810 (2018) 164.
[4] S.J. Lee, W. Lerche, T. Weigand, Nucl. Phys. B 938 (2019) 321.
[5] S.J. Lee, W. Lerche, T. Weigand, arXiv:1901.08065 [hep-th].
[6] S.J. Lee, W. Lerche, T. Weigand, arXiv:1904.06344 [hep-th].
[7] T.W. Grimm, E. Palti, I. Valenzuela, J. High Energy Phys. 1808 (2018) 143.
[8] T.W. Grimm, C. Li, E. Palti, J. High Energy Phys. 1903 (2019) 016.
[9] P. Corvilain, T.W. Grimm, I. Valenzuela, arXiv:1812.07548 [hep-th].
[10] F. Marchesano, M. Wiesner, arXiv:1904.04848 [hep-th].
[11] A. Font, A. Herraez, L.E. Ibanez, arXiv:1904.05379 [hep-th].
[12] C. Vafa, Nucl. Phys. B 469 (1996) 403.
[13] V. Braun, Y.H. He, B.A. Ovrut, T. Pantev, Phys. Lett. B 618 (2005) 252.
[14] V. Braun, Y.H. He, B.A. Ovrut, T. Pantev, J. High Energy Phys. 0605 (2006) 043.
[15] V. Bouchard, R. Donagi, Phys. Lett. B 633 (2006) 783.
[16] V. Braun, Y.H. He, B.A. Ovrut, J. High Energy Phys. 0606 (2006) 032.
[17] K. Hori, et al., Mirror Symmetry, AMS, 2003.
[18] D.A. Cox, S. Katz, Mirror Symmetry and Algebraic Geometry, Math. Surv. & Mono., vol. 68, AMS, 1999.
[19] L.B. Anderson, X. Gao, J. Gray, S.J. Lee, J. High Energy Phys. 1610 (2016) 105.
[20] M. Gross, Duke Math. J. 74 (2) (1994) 271.
[21] A. Grassi, Math. Ann. 290 (1991) 287.
[22] G. Di Cerbo, R. Svaldi, arXiv:1608.02997 [math.AG].
[23] J. Gray, A.S. Haupt, A. Lukas, J. High Energy Phys. 1409 (2014) 093.
[24] L.B. Anderson, X. Gao, J. Gray, S.J. Lee, J. High Energy Phys. 1710 (2017) 077.
[25] L.B. Anderson, J. Gray, B. Hammack, J. High Energy Phys. 1808 (2018) 128.
[26] Y.C. Huang, W. Taylor, J. High Energy Phys. 1903 (2019) 014.
[27] F. Rohsiepe, arXiv:hep-th/0502138.
[28] S.B. Johnson, W. Taylor, J. High Energy Phys. 1410 (2014) 23.
[29] S.B. Johnson, W. Taylor, Fortschr. Phys. 64 (2016) 581.
[30] P. Candelas, A. Constantin, H. Skarke, Commun. Math. Phys. 324 (2013) 937.
[31] K. Oguiso, Int. J. Math. 3 (1993) 439.
[32] P.M.H. Wilson, Math. Ann. 300 (1994) 693.
[33] J. Kollár, arXiv:1206.5721 [math.AG].
[34] Non-Linear Algebra: ICERM Semester Programme, https://icerm.brown.edu/programs/sp-f18/.
[35] Y.H. He, arXiv:1706.02714 [hep-th].
[36] D. Krefl, R.K. Seong, Phys. Rev. D 96 (6) (2017) 066014.
[37] F. Ruehle, J. High Energy Phys. 1708 (2017) 038.
[38] J. Carifio, J. Halverson, D. Krioukov, B.D. Nelson, J. High Energy Phys. 1709 (2017) 157.
[39] Y.H. He, Phys. Lett. B 774 (2017) 564.
[40] V. Jejjala, A. Kar, O. Parrikar, arXiv:1902.05547 [hep-th].
[41] A. Cole, G. Shiu, arXiv:1812.06960 [hep-th].
[42] A. Mutter, E. Parr, P.K.S. Vaudrevange, Nucl. Phys. B 940 (2019) 113.
[43] T. Rudelius, J. Cosmol. Astropart. Phys. 1902 (2019) 044.
[44] J. Halverson, F. Ruehle, Phys. Rev. D 99 (4) (2019) 046015.
[45] D. Klaewer, L. Schlechter, Phys. Lett. B 789 (2019) 438.
[46] M. Demirtas, C. Long, L. McAllister, M. Stillman, arXiv:1808.01282 [hep-th].
[47] K. Bull, Y.H. He, V. Jejjala, C. Mishra, Phys. Lett. B 785 (2018) 65.
[48] K. Bull, Y.H. He, V. Jejjala, C. Mishra, arXiv:1903.03113 [hep-th].
[49] R. Jinno, arXiv:1805.12153 [hep-th].
[50] W.J. Cunningham, arXiv:1805.04463 [gr-qc].
[51] Y.N. Wang, Z. Zhang, J. High Energy Phys. 1808 (2018) 009.
[52] G.K. Demir, N. Sönmez, H. Dogan, arXiv:1803.01550 [hep-ph].
[53] K. Hashimoto, S. Sugishita, A. Tanaka, A. Tomiya, Phys. Rev. D 98 (4) (2018) 046019.
[54] J. Carifio, W.J. Cunningham, J. Halverson, D. Krioukov, C. Long, B.D. Nelson, Phys. Rev. Lett. 121 (10) (2018) 101602.
[55] J. Liu, J. High Energy Phys. 1712 (2017) 149.
[56] T. Cohen, M. Freytsis, B. Ostdiek, J. High Energy Phys. 1802 (2018) 034.
[57] J. Halverson, B. Nelson, F. Ruehle, arXiv:1903.11616 [hep-th].
[58] Y.H. He, arXiv:1812.02893 [hep-th].
[59] L.B. Anderson, F. Apruzzi, X. Gao, J. Gray, S.J. Lee, Nucl. Phys. B 906 (2016) 441.
[60] L.B. Anderson, F. Apruzzi, X. Gao, J. Gray, S.J. Lee, Phys. Rev. D 93 (8) (2016) 086001.
[61] P. Candelas, A.M. Dale, C.A. Lutken, R. Schimmrigk, Nucl. Phys. B 298 (1988) 493;
P. Candelas, C.A. Lutken, R. Schimmrigk, Nucl. Phys. B 306 (1988) 113;
M. Gagnon, Q. Ho-Kim, Mod. Phys. Lett. A 9 (1994) 2235.
[62] J. Gray, A.S. Haupt, A. Lukas, J. High Energy Phys. 1307 (2013) 070.
[63] R. Altman, J. Gray, Y.H. He, V. Jejjala, B.D. Nelson, J. High Energy Phys. 1502 (2015) 158.
[64] P. Candelas, A. Constantin, C. Mishra, Fortschr. Phys. 66 (2018) 1800029.
[65] P. Candelas, R. Davies, Fortschr. Phys. 58 (2010) 383;
P. Candelas, A. Constantin, Fortschr. Phys. 60 (2012) 345–369;
P. Candelas, A. Constantin, C. Mishra, Fortschr. Phys. 64 (6–7) (2016) 463.
[66] P. Candelas, X. de la Ossa, Y.H. He, B. Szendroi, Adv. Theor. Math. Phys. 12 (2) (2008) 429.
[67] A. Constantin, J. Gray, A. Lukas, J. High Energy Phys. 01 (2017) 001.
[68] L.B. Anderson, Y.H. He, A. Lukas, J. High Energy Phys. 0707 (2007) 049.
[69] K-M. Ting, Encyclopedia of Machine Learning, Springer, ISBN 978-0-387-30164-8, 2011.
[70] B. Matthews, Biochim. Biophys. Acta, Protein Struct. 405 (2) (1975) 442–451.
[71] Wolfram Research, Inc., Mathematica, Version 11.3, Champaign, IL, 2018.
[72] L.B. Anderson, X. Gao, J. Gray, S.J. Lee, J. High Energy Phys. 1611 (2016) 004.
[73] D. Yarotsky, A. Zhevnerchuk, arXiv:1906.09477 [cs.NE].
[74] With a simplified network lacking the tanh layer we still find a compelling result but it gets slightly better when this layer is included; see [73] for a relevant study in the computer science literature.
[75] As for the hyperparameters, we have used 20 epochs, batch size 64, learning rate 0.001, and ADAM optimizer.