Machine learning for complete intersection Calabi-Yau manifolds: A methodological study

Harold Erbin^{*} and Riccardo Finotello[†]

*Dipartimento di Fisica, Università di Torino INFN Sezione di Torino and Arnold–Regge Center,
Via Pietro Giuria 1, I-10125 Torino, Italy*

(Received 21 August 2020; accepted 7 May 2021; published 15 June 2021)

We revisit the question of predicting both Hodge numbers $h^{1,1}$ and $h^{2,1}$ of complete intersection Calabi-Yau (CICY) 3-folds using machine learning (ML), considering both the old and new datasets built respectively by Candelas-Dale-Lutken-Schimmrigk/Green-Hübsch-Lutken and by Anderson-Gao-Gray-Lee. In real-world applications, implementing a ML system rarely reduces to feed the brute data to the algorithm. Instead, the typical workflow starts with an exploratory data analysis (EDA) which aims at understanding better the input data and finding an optimal representation. It is followed by the design of a validation procedure and a baseline model. Finally, several ML models are compared and combined, often involving neural networks with a topology more complicated than the sequential models typically used in physics. By following this procedure, we improve the accuracy of ML computations for Hodge numbers with respect to the existing literature. First, we obtain 97% (resp. 99%) accuracy for $h^{1,1}$ using a neural network inspired by the inception model for the old dataset, using only 30% (resp. 70%) of the data for training. For the new one, a simple linear regression leads to almost 100% accuracy with 30% of the data for training. The computation of $h^{2,1}$ is less successful as we manage to reach only 50% accuracy for both datasets, but this is still better than the 16% obtained with a simple neural network (SVM with Gaussian kernel and feature engineering and sequential convolutional network reach at best 36%). This serves as a proof of concept that neural networks can be valuable to study the properties of geometries appearing in string theory.

DOI: [10.1103/PhysRevD.103.126014](https://doi.org/10.1103/PhysRevD.103.126014)

I. INTRODUCTION

The last few years have seen a major uprising of machine learning (ML), and more particularly of neural networks [1–3]. This technology is extremely efficient at discovering and predicting patterns and now pervades most fields of applied sciences and of the industry. In view of its versatility, it is likely that ML will find its way toward high-energy and theoretical physics (see [4–9] for selected reviews). One of the most critical places where progress can be expected is in understanding the geometries used to describe string compactifications.

String theory is the most developed candidate for a theory of quantum gravity together with the unification of matter and interactions. However, it predicts ten spacetime dimensions: to recover our four-dimensional Universe, it is

necessary to compactify six dimensions. For string theory to be a fundamental theory of reality, a single compactification should describe the current Universe (obviously, other compactifications may enter at early or later stages since spacetime is dynamical). Unfortunately, the number of possibilities—forming the so-called string landscape—is huge (numbers as high as $10^{272\,000}$ have been suggested for some models) [10–19], the mathematical objects entering the compactifications are complex and typical problems are often NP-complete, NP-hard, or even undecidable [9,20,21], making an exhaustive classification impossible. Additionally, there is no single framework to describe all the possible (flux) compactifications. As a consequence, each class of models must be studied with different methods. This has prevented any precise connection to the existing and tested theories (in particular, the Standard Model of particle physics) or the proposal of a sharply defined and doable experiment.

Until recently, the string landscape has been studied using different methods: (1) analytic computations for simple examples, (2) general statistics, (3) random scans, (4) algorithmic enumerations of possibilities. This has been a large endeavor of the string community, and we refer to the reviews [9,22–26] and to references therein for more

^{*}erbin@to.infn.it

[†]riccardo.finotello@to.infn.it

Published by the American Physical Society under the terms of the [Creative Commons Attribution 4.0 International](https://creativecommons.org/licenses/by/4.0/) license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI. Funded by SCOAP³.

details. The main objective of such studies is to understand what are the generic predictions of string theory: even if “the” correct compactification has not been found, this helps to narrow down what to look for experimentally. The first conclusion of these studies is that compactifications giving an effective theory close to the Standard Model are scarce.¹ Each of the four approaches display different limitations: (1) lacks of genericity, (2) is too much general, (3) ignores the structure of the landscape and has few chances to discover rare compactifications, (4) requires too much computational power to move beyond “simple” examples. As a result, no major phenomenological progress has been seen in the last decade and finding a physical compactification looks still as remote. In reaction to these difficulties and starting with the seminal paper [27], new investigations based on ML appeared in the recent years, focusing on different aspects of the string landscape and of the geometries used in compactifications [28–51] (see also [52–65] for related works). For more context and a summary of the state of the art, the reader is referred to the excellent review [9]. ML is extremely adequate when it comes to pattern search, which motivates two main applications to string theory: (1) explore systematically a space of possibilities (if they are not random, ML should be able to find a pattern, even if it is too complicated to be formulated explicitly), (2) obtain approximate results on distributions from which mathematical formulas can be deduced.

We want to address the question of computing the Hodge numbers $h^{1,1}$ and $h^{2,1}$ (positive integers) for *complete intersection Calabi-Yau* (CICY) 3-folds [66] using different machine learning algorithms. A CICY is described by an hypersurface in an ambient space given by a product of projective spaces. The CICY is completely specified by its *configuration matrix* (with entries being positive integers), which encodes the system of polynomial equations representing the hypersurface, and which is the basic input of the algorithms. The CICY 3-folds are the simplest Calabi-Yau and they have been well studied. In particular, they have been completely classified and their topological properties computed [67–69] (see [70–73] for reviews). For these reasons, they provide an excellent sandbox to test ML algorithms in a controlled environment. More particularly, simple tests show that the task is difficult for simple ML algorithms—even neural networks—such that this is an interesting challenge to solve before moving to more difficult problems.

The goal is to predict two positive integers from a matrix of positive integers. This task is complicated by various redundancies in the description (such as an independence in

the permutations of lines and columns). A simple sequential network taking only the matrix as input performs badly, especially for $h^{2,1}$. As a consequence, more advanced methods are needed. While usual physics application of ML reduces to feeding a (big) sequential neural network with raw data, real-world applications are built following a more general workflow [3,74,75]: (1) understanding of the problem, (2) exploratory data analysis (EDA), (3) design of a baseline, (4) definition of a validation strategy, (5) feature engineering and selection, (6) design of ML models, (7) ensembling.

While the first step is straightforward, it is still interesting to notice that computations involved in string geometries (using algebraic topology) are far from standard applications of ML algorithms, which makes the problem even more interesting. EDA aims at understanding better the dataset, in particular, by finding how the variables are distributed, correlated, determining if there are outliers, etc. This analysis naturally leads to designing new features from the existing ones, which is called *feature engineering*. Indeed, putting derived features by hand may make the data more easily understandable by the ML algorithms, for example by emphasizing important properties.² This phase is followed by *feature selection*, where different set of features are chosen according to the need of each algorithm from step (6). In between, one needs to set up a validation strategy to ensure that the predictions appropriately reflect the real values, together with a baseline model, which gives a lower bound on the accuracy together with a working pipeline.³ For instance, we find that a simple linear regression using the configuration matrix as input gives 43.6–48.8% for $h^{1,1}$ and 9.6–10.4% for $h^{2,1}$ using from 20% to 80% of data for training. Hence, any algorithm *must* do better than this to be worth considering. Finally, we can build different models in step (6), in particular, by considering different topologies of neural networks beyond the simplest sequential models. The last optional step consists in combining different models together in order to improve the results. With respect to the whole process, the purpose of this paper is also pedagogical and aims at exemplifying how these steps are performed in an applied ML project.

There is a finite number of 7890 CICY 3-folds. Due to the freedom in representing the configuration matrix, two datasets have been constructed: the “original dataset” [67,68] and the “favorable dataset” [69]. A configuration matrix is said to be favorable if its second cohomology

¹This means that the gauge group is not much bigger than $SU(3) \times SU(2) \times U(1)$ and that there are not too many additional particles. The current bounds on BSM (beyond Standard Model) physics put even stronger restrictions.

²While one could expect ML algorithms to generate these features by themselves, this may complicate the learning process. So in cases where it is straightforward to compute meaningful derived features, it is often worth considering them.

³For example, the original work on this topic [30] did not set up a validation strategy and reported the accuracy over both the training and test data. Correcting this problem leads to an accuracy of 37% [33], which is lower than the linear regression baseline.

descends completely from the second cohomology of the ambient space: this implies that $h^{1,1}$ equals the number of projective spaces in the ambient space [69,76]. In the “favorable” dataset, all configuration matrices are favorable whenever possible (99.1%), whereas in the “original” dataset only 61.8% of the matrices are favorable. Both datasets are described in more details in Sec. II B.

Our analysis continues and generalizes [30,33] at different levels. We compute $h^{2,1}$, which has been ignored in [30,33], where the authors argue that it can be computed from $h^{1,1}$ and from the Euler characteristics (a simple formula exists for the latter). In our case, we want to push the idea of using ML to learn about the physics (or the mathematics) of CY to its very end: we assume that we do not know anything about the mathematics of the CICY, except that the configuration matrix is sufficient to derive all quantities. Moreover, we have already mentioned that ML algorithms have rarely been used to derive data in algebraic topology, which can be a difficult task. For this reason, obtaining also $h^{2,1}$ from ML techniques is an important first step toward using ML for more general problems in string geometries. In particular, this helps to prepare the study of CICY 4-folds (classified in [77]) for which there are four Hodge numbers which are expected to be even more difficult to compute. Finally, regression is also more useful for extrapolating results: a classification approach assumes that we already know all the possible values of the Hodge numbers and has difficulties to predict labels which do not appear in the training set. This is necessary when we move to a dataset for which not all topological quantities have been computed, for instance CY constructed from the Kreuzer–Skarke list of polytopes [78].

In this paper, we compare the performances of the following algorithms: linear regression, support vector machines (SVM) with linear and Gaussian kernels, decision trees and ensemble thereof—random forests and gradient boosting—and deep neural networks. The best results obtained with and without feature engineering are displayed in Fig. 1 for the old dataset. We find that, in all cases except neural networks, using engineered features greatly enhance the performances. The EDA reveals that the number of projective spaces forming the ambient space (equal to the number of rows) is a particularly distinguished feature. In fact, all algorithms yield an accuracy of 99–100% for $h^{1,1}$ in the favorable dataset. For the linear regression, this directly gives the well-known results [69] that $h^{1,1}$ equals the number of projective spaces for favorable configuration matrix. In the case of the original dataset, the best model is a neural network inspired by Google’s inception model [79–81], which allows to reach nearly 100% accuracy. This neural network is further studied in [82]. It is clear that favorable configurations are easier to deal with since the relation is linear: in fact, it is likely that algorithms other than neural networks and SVM get correct results only for favorable configurations when

working in the original dataset (61.8%). The algorithms are not as successful for $h^{2,1}$, with the inception model giving again the best result, close to 50% accuracy—which is still much better than what the baseline or simple models do. We leave improving the computation of $h^{2,1}$ and interpreting what the different algorithms learn for a future work.

The data analysis and ML are programmed in PYTHON using standard open-source packages: pandas [83], matplotlib [84], seaborn [85], scikit-learn [86], scikit-optimize [87], tensorflow [88] (and its high level API Keras [89]). The code and its description are available on Github.

This paper is organized as follows. In Sec. II, we first recall the definition of Calabi-Yau manifolds (Sec. II A) and describe the two existing CICY datasets (Sec. II B). We then engineer new features before performing an EDA for both datasets (Sec. II C), reproducing some well-known figures from the literature. Then, in Sec. III, we implement the different ML algorithms. Our paper culminates in the description of the inceptionlike neural network in Sec. III F 3 where we reach the highest accuracy. Finally, we discuss our results in Sec. IV. Appendix contains details on the different algorithms used in this paper.

II. DATA ANALYSIS

In this section, we introduce *Calabi-Yau* (CY) manifolds before describing the two datasets of CICY manifolds (Sec. II B). Since the CICY have been completely classified, they provide a good opportunity for testing ideas from ML in a controlled setting. In order to select the most appropriate learning algorithm, we perform a preliminary *exploratory data analysis* (EDA) in Sec. II C.

A. Calabi-Yau manifolds

A CY n -fold is a n -dimensional complex manifold X with $SU(n)$ holonomy (they have $2n$ real dimensions). An equivalent definition is the vanishing of its first Chern class. A standard reference for the physicist is [71] (see also [72,73] for useful references).

The most relevant case for superstring compactifications are CY 3-folds. Indeed, superstrings are well defined only in 10 dimensions: in order to recover a 4-dimensional theory, it is necessary to compactify 6 dimensions [71]. Importantly, the compactification on a CY leads to the breaking of a large part of the supersymmetry, which is phenomenologically more realistic.

Calabi-Yau manifolds are characterized by a certain number of topological properties, the most salient being the Hodge numbers $h^{1,1}$ and $h^{2,1}$, counting respectively the Kähler and complex structure deformations, and the Euler characteristics⁴

⁴In full generality, the Hodge numbers $h^{p,q}$ count the numbers of harmonic (p, q) -forms.

FIG. 1. The plots show the best accuracy reached by the models considered in this paper for the old dataset. The models are trained to predict separately $h^{1,1}$ and $h^{2,1}$ and using 30% and 80% of the data for training (original dataset). (a) Accuracy reached when trained using only the configuration matrix (original dataset). (b) Accuracy reached using the best training set for each algorithm.

$$\chi = 2(h^{1,1} - h^{2,1}). \quad (2.1)$$

Interestingly, topological properties of the manifold directly translates into features of the 4-dimensional effective action (in particular, the number of fields, the representations and the gauge symmetry) [71,90].⁵ In particular, the Hodge numbers count the number of chiral multiplets (in heterotic compactifications) and the number of hyper- and vector multiplets (in type II compactifications): these are related to the number of fermion generations (3 in the Standard Model) and is thus an important measure of the distance to the Standard Model.

The simplest CYs are constructed by considering the complete intersection of hypersurfaces in a product $\mathcal{A}$ of projective spaces $\mathbb{P}^{n_i}$ (called the ambient space) [66–69,72,91]:

$$\mathcal{A} = \mathbb{P}^{n_1} \times \dots \times \mathbb{P}^{n_m}. \quad (2.2)$$

Such hypersurfaces are defined by homogeneous polynomial equations: a Calabi-Yau X is described by the solution to the system of equations, i.e., by the intersection of all these surfaces (that the intersection is “complete” means that the hypersurface is nondegenerate).

To gain some intuition, consider the case of a single projective space $\mathbb{P}^n$ with (homogeneous) coordinates Z^I , $I = 0, \dots, n$. In this case, a codimension 1 subspace is obtained by imposing a single homogeneous polynomial equation of degree a on the coordinates

$$\begin{aligned} p_a(Z^0, \dots, Z^n) &= P_{I_1 \dots I_a} Z^{I_1} \dots Z^{I_a} = 0, \\ p_a(\lambda Z^0, \dots, \lambda Z^n) &= \lambda^a p_a(Z^0, \dots, Z^n). \end{aligned} \quad (2.3)$$

Each choice of the polynomial coefficients $P_{I_1 \dots I_a}$ leads to a different manifold. However, it can be shown that the manifolds are (generically) topologically equivalent. Since we are interested only in classifying the CY as topological manifolds and not as complex manifolds, the information about $P_{I_1 \dots I_a}$ can be forgotten and it is sufficient to keep track only on the dimension n of the projective space and of the degree a of the equation. The resulting hypersurface is denoted equivalently as $[\mathbb{P}^n|a] = [n|a]$. Finally, $[\mathbb{P}^n|a]$ is 3-dimensional if $n = 4$ (the equation reduces the dimension by one), and it is a CY (the “quintic”) if $a = n + 1 = 5$ (this is required for the vanishing of its first Chern class). The simplest representative of this class is Fermat’s quintic defined by the equation

⁵Another reason for sticking to topological properties is that there is no CY for which the metric is known. Hence, it is not possible to perform explicitly the Kaluza-Klein reduction in order to derive the 4-dimensional theory.

$$\sum_{I=0}^4 (Z^I)^5 = 0. \quad (2.4)$$

This construction can be generalized to include m projective spaces and k equations, which can mix the coordinates of the different spaces. A CICY 3-fold X as a topological manifold is completely specified by a *configuration matrix* denoted by the same symbol as the manifold:

$$X = \left[\begin{array}{c|ccc} \mathbb{P}^{n_1} & a_1^1 & \dots & a_k^1 \\ \vdots & \vdots & \ddots & \vdots \\ \mathbb{P}^{n_m} & a_1^m & \dots & a_k^m \end{array} \right] \quad (2.5)$$

where the coefficients a_α^r are positive integers and satisfy the following constraints

$$\dim_{\mathbb{C}} X = \sum_{r=1}^m n_r - k = 3, \quad \forall r: n_r + 1 = \sum_{\alpha=1}^k a_\alpha^r. \quad (2.6)$$

The first relation states that the dimension of the ambient space minus the number of equations equals the dimension of the CY 3-fold. The second set of constraints arise from the vanishing of its first Chern class; they imply that the n_i can be recovered from the matrix elements.

In this case also, two manifolds described by the same configuration matrix but different polynomials are equivalent as real manifold (they are diffeomorphic)—and thus as topological manifolds—but they are different as complex manifolds. Hence, it makes sense to write only the configuration matrix.

A given topological manifold is not described by a unique configuration matrix. First, any permutation of the lines and columns leave the intersection unchanged (it amounts to relabeling the projective spaces and equations). Second, two intersections can define the same manifold. The ambiguity in the line and column permutations is often fixed by imposing some ordering of the coefficients. Moreover, in most cases, there is an optimal representation of the manifold X , called favorable [69]: in such a form, topological properties of X can be more easily derived from the ambient space $\mathcal{A}$.

B. Datasets

Simple arguments [66,67,70] show that the number of CICY is necessarily finite due to the constraints (2.6) together with identities between complete intersection manifolds. The classification of the CICY 3-folds has been tackled in [67], which established a dataset of 7890 CICY.⁶ The topological properties of each of these manifolds have

⁶However, there are redundancies in this set [67,69,92]; this fact will be ignored in this paper.

been computed in [68]. More recently, a new classification has been performed [69] in order to find the favorable representation of each manifold whenever it is possible.

Below we show a list of the CICY properties and of their configuration matrices:

- (i) general properties
 - number of configurations: 7890
 - number of product spaces (block diagonal matrix): 22
 - $h^{1,1} \in [0, 19]$, 18 distinct values [Fig. 2(a)]
 - $h^{2,1} \in [0, 101]$, 65 distinct values [Fig. 2(b)]
 - unique Hodge number combinations: 266
- (ii) “original dataset” [67,68]
 - maximal size of the configuration matrices: 12×15
 - number of favorable matrices (excluding product spaces): 4874 (61.8%)
 - number of nonfavorable matrices (excluding product spaces): 2994
 - number of different ambient spaces: 235
- (iii) “favorable dataset” [69]
 - maximal size of the configuration matrices: 15×18
 - number of favorable matrices (excluding product spaces): 7820 (99.1%)
 - number of nonfavorable matrices (excluding product spaces): 48
 - number of different ambient spaces: 126

The configuration matrix completely encodes the information of the CICY and all topological quantities can be derived from it. However, the computations are involved and there is often no closed-form expression. This situation is typical in algebraic geometry, and it can be even worse for some problems, in the sense that it is not even known how to compute the desired quantity (think to the metric of

CYs). For these reasons, it is interesting to study how we can retrieve these properties using ML algorithms. In the current paper, following [30,33], we focus on the computation of the Hodge numbers with the initial scheme:

$$\begin{aligned} &\text{Input: configuration matrix} \\ &\rightarrow \text{Output: Hodge numbers} \quad (2.7) \end{aligned}$$

To provide a good test case for the use of ML in context where the mathematical theory is not completely understood, we will make no use of known formulas.

C. Exploratory data analysis

A typical ML project does not consist of feeding the raw data—here, the configuration matrix—to the algorithm. It is instead preceded by a phase of exploration in order to better understand the data, which in turn can help to design the learning algorithms. We call *features* properties given as inputs, and *labels* the targets of the predictions. There are several phases in the exploratory data analysis (EDA):

- (1) *feature engineering*: new features are derived from the inputs;
- (2) *feature selection*: the most relevant features are chosen to explain the targets;
- (3) *data augmentation*: new training data is generated from the existing ones;
- (4) *data diminution*: part of the training data is not used.

For pragmatical introductions, the reader is referred to [74,75].

Engineered features are redundant, by definition, but can help the algorithm learn more efficiently by providing an alternative formulation and by drawing attention on salient characteristics. A simple example is the following: given a series of numbers, one can compute different

FIG. 2. Distribution of the Hodge numbers (log scale). (a) $h^{1,1}$. (b) $h^{2,1}$.

FIG. 3. Cumulative retained variance of the principal components of the configuration matrix in the original and favorable dataset. (a) original dataset. (b) favourable dataset.

statistics—median, mean, variance, etc.—and add them to the inputs. It may happen that the initial series becomes then irrelevant once this new information is introduced.

Another approach to improve the learning process is to augment or decrease the number of training samples artificially. For example, one can use invariances of the inputs to generate more training data. This does not help in our case because the entries of the configuration matrices are partially ordered. Another possibility is to remove outliers which can damage the learning process by driving the algorithm far from the best solution. If there are few of them, it is better to ignore them altogether during training since an algorithm which is not robust to outliers will in any case make bad predictions (a standard illustration is given by the Pearson and Spearman correlation coefficients, with the first not being robust to outliers [75]).

Finding good features and selecting those to keep requires trials and errors. In general, it is not necessary to keep track of all steps, but we feel that it is useful to do so in this paper for a pedagogical purpose.

Before starting the EDA, the first step should be to split the data into training and validation sets to avoid biasing the choices of the algorithm and the strategy: the EDA should be performed only on the training set. However, the dataset we consider is complete and quite uniform: a subset of it would display the same characteristics as the entire set. To give a general overview of the properties—which can be useful for the reader interested in understanding the statistics of the CICY and for applications to string compactifications—we work with the full dataset.

1. Engineering

Any transformation of the input data which has some mathematical meaning can be a useful feature. We have

established the following list of possibly useful quantities (most of them are already used to characterize CICY in the literature [71]):

- (i) the number of projective spaces (rows), $m = \text{num_cp}$;
- (ii) the number of equations (columns), $k = \text{num_eqs}$;
- (iii) the number of $\mathbb{P}^1$, $f = \text{num_cp_1}$;
- (iv) the number of $\mathbb{P}^2$, num_cp_2 ;
- (v) the number of $\mathbb{P}^n$ with $n \neq 1$, $F = \text{num_cp_neq1}$;
- (vi) the excess number $N_{ex} = \sum_{r=1}^F (n_r + f + m - 2k) = \text{num_ex}$;
- (vii) the dimension of the cohomology group H^0 of the ambient space, dim_h0_amb ;
- (viii) the Frobenius norm of the matrix, norm_matrix ;
- (ix) the list of the projective space dimensions dim_cp and statistics thereof (min, max, median, mean);
- (x) the list of the equation degrees deg_eqs and statistics thereof (min, max, median, mean);
- (xi) k -means clustering on the components of the configuration matrix (with a number of clusters going from 2 to 15);⁷
- (xii) principal components of the configuration matrix derived using a principal components analysis (PCA) with 99% of the variance retained (see Fig. 3).

2. Selection

a. Correlations To get a first general idea, it is useful to take a look at the correlation matrix of the features and the

⁷The algorithm determines the centroids of conglomerates of data called *clusters* using an iterative process which computes the distance of each sample from the center of the cluster. It then assigns the label of the cluster to the nearest samples. We used the class `cluster.KMeans` in `scikit-learn`.

FIG. 4. Correlations between the engineered scalar features and the labels. (a) original dataset. (b) favourable dataset.

labels.⁸ The correlation matrices for the scalar variables are displayed in Fig. 4 for the original and favorable datasets (this excludes the configuration matrix).

As we can see, some engineered features are strongly correlated, especially in the favorable dataset. In particular $h^{1,1}$ (respectively $h^{2,1}$) correlates (respectively anticorrelates) strongly with the number of projective spaces m and with the norm and rank of the matrix. This gives a first hint that these variables could help improve predictions by feeding them to the algorithm along with the matrix. On the other hand, finer information on the number of projective spaces and equations do not correlate with the Hodge numbers.

From this analysis, in particular from Fig. 4, we find that the values of $h^{1,1}$ and $h^{2,1}$ are also correlated. This motivates the simultaneous learning of both Hodge numbers since it can increase chances for the neural network to learn more universal features. In fact, this is something that often happens in practice: counterintuitively, it has been found that multitasking enhances the ability to generalize [93–97].

b. Feature importance A second option is to sort the features by order of importance. This can be done using a decision tree which is capable to determine the weight of each variable toward making a prediction. One advantage over correlations is that the algorithm is nonlinear and can thus determine subtler relations between the features and labels. To avoid biasing the results using only one decision tree, we trained a random forest of trees (using `ensemble.RandomForestRegressor` in `scikit-learn`). It consists in a large number of decision trees which are trained on different random subsets of the training

dataset and averaged over the outputs (see also Sec. III E and Appendix A 3). The algorithm determines the importance of the different features to make predictions as a by-product of the learning process, because the most relevant features tend to be found at the first branches since they are the most important to make the prediction. The importance of a variable is a number between 0 and 1, and the sum over all of them must be 1. Since a random forest contains many trees, the robustness of the variable ranking usually improves with respect to a single tree (Appendix A 3). Moreover, as the main objective is to obtain a qualitative preliminary understanding of the features, there is no need for fine tuning at this stage and we use the default parameters (in particular, 100 decision trees). We computed feature importance for both datasets and for two different set of variables: one containing the engineered features and the configuration matrix, and one with the engineered features and the PCA components. In the following figures, we show several comparisons of the importance of the features, dividing the figures into scalars, vectors and configuration matrix (or its PCA), and clusters. The sum of importance of all features equals 1.

In Fig. 5, we show the ranking of the scalar features in the two datasets (differences between the set using the configuration matrix and the other using the PCA are marginal and are not shown to avoid redundant plots). As already mentioned, we find again that the number of projective spaces is the most important feature by far. It is followed by the matrix norm in the original dataset, and by the matrix rank for $h^{2,1}$ in the favorable dataset, but in a lesser measure. Finally, it points out that the other features have a negligible impact on the determination of the labels and may as well be ignored during training.

The same analysis can be repeated for the vector features and the configuration matrix component by component. In Fig. 6, we show the cumulative importance of the

⁸The correlation is defined as the ratio between the covariance of two variables $\sigma(x, y) = \sum_i (x_i - \bar{x})(y_i - \bar{y})$ and the product of the standard deviations $\sigma(x)\sigma(y)$ (in this case $\bar{x}$ and $\bar{y}$ are the sample means).

FIG. 5. Importance of the scalar features in the datasets. The same computation involving the PCA of the configuration matrix shows very marginal differences in this case: the importance of the scalar features is mostly unchanged, especially for the higher ranked variables. (a) original dataset. (b) favourable dataset.

features (i.e., the sum of the importance of each component). We can appreciate that the list of the projective space dimensions plays a major role in the determination of the labels in both datasets. In the case of $h^{2,1}$, we also have a large contribution from the dimensions of the cohomology group $\dim_{h^0_amb}$, as can be expected from algebraic topology [71].

In Fig. 7, we finally show the importance associated to the number of clusters used during the EDA: no matter how many clusters we use, their relevance is definitely marginal compared to all other features used in the variable ranking (scalars, vectors, and the configuration matrix or its PCA) for both datasets.

c. Conclusion It seems therefore that the number of projective spaces plays a relevant role in the determination of $h^{1,1}$ and $h^{2,1}$ as well as the list of dimensions of the projective spaces. In order to validate this observation, in Fig. 8 we present a scatter plot of the Hodge number distributions versus the number of projective spaces: it shows that there is indeed a linear dependence in m for $h^{1,1}$, especially in the favorable dataset. In fact, the only exceptions to this pattern in the latter case are the manifolds which do not have a favorable embedding [69]. Hence, a simple data analysis hints naturally toward this mathematical result.

Finally, we found other features which may be relevant and are worth to be included in the algorithm: the matrix rank and norm, the list of projective space dimensions and of the associated cohomology dimensions. However, we want to emphasize one caveat to this analysis: correlations look only for linear relations, and the random forest has not been optimized or could just be not powerful enough to make good predictions. This means that feature selection just gives a hint but it may be necessary to adapt.

3. Removing outliers

The Hodge number distributions (Figs. 2 and 9) display few outliers which lie outside the tail of the main distributions. Such outliers may negatively impact the learning process and drive down the accuracy: it makes sense to remove them from the training set.

It is easy to see that the 22 outlying manifolds with $h^{1,1} = h^{2,1} = 0$ are product spaces, recognisable from their block-diagonal matrix. Moreover, we will also remove outliers with $h^{1,1} = 19$ and $h^{2,1} > 86$, which represent 15 and 2 samples. In total, this represents 39 samples, or 0.49% of the total data.

To simplify the overall presentation and because the dataset is complete, we will mainly focus on the pruned subset of the data obtained by removing outliers, even from the test set.⁹ This implies that Hodge numbers lie in the ranges $1 \leq h^{1,1} \leq 16$ and $15 \leq h^{2,1} \leq 86$. Except when stated otherwise, accuracy is indicated for this pruned dataset. Obviously, the very small percentage of outliers makes the effect of removing them from the test set negligible when stating accuracy.

III. MACHINE LEARNING ANALYSIS

In this section, we compare the performances of different ML algorithms: linear regression, SVM, random forests, gradient boosted trees and neural networks. Before reporting the results for each algorithm, we detail the

⁹There is no obligation to use a ML algorithm to label outliers in the training set, it is perfectly fine to decide which data to include or not, even based on targets. However, for a real-world application, outliers in the test set should be labeled by some process based only on the input features. Flagging possible outliers may improve the predictions by helping the machine understand that such samples require more caution.

FIG. 6. Importance of the vector features the configuration matrix (or its principal components) in the datasets: notice how the PCA plays a much more important role with respect to the full configuration matrix. (a) original dataset. (b) favourable dataset.

feature selection (Sec. III A) and the evaluation strategy (Sec. III B). We obtain the best results in Sec. III F 3 where we present a neural network inspired by the inception model [79–81]. We provide some details on the different algorithms in Appendix and refer the reader to the literature [1–3,5,7,9,74,75] for more details.

A. Feature extraction

In Sec. II, the EDA showed that several engineered features are promising for predicting the Hodge numbers. In what follows, we will compare the performances of various algorithms using different subsets of features:

- (i) only the configuration matrix (no feature engineering);
- (ii) only the number of projective spaces m ;
- (iii) only a subset of engineered features and not the configuration matrix nor its PCA;

- (iv) a subset of engineered features and the PCA of the matrix.

Following the EDA and feature engineering, we finally select the features we use in the analysis by choosing the highest ranked features. We will therefore keep the number of projective spaces (`num_cp` in the dataset) and the list of the dimension of the projective spaces (`dim_cp`) for both $h^{1,1}$ and $h^{2,1}$. We will also include the dimension of the cohomology group of the ambient space `dim_h0_amb` but only for $h^{2,1}$.

B. Analysis strategy

For the ML analysis, we split the dataset into training and test sets: we fit the algorithms on the first and then show the predictions on the test set, which will not be touched until the algorithms are ready.

FIG. 7. Incidence of the numbers of clusters on the variable ranking. Also in this case the difference between using the configuration matrix or its PCA is marginal and actually the clusters have even lower importance when using the latter. We therefore avoid presenting nugatory information and show only the importance of clusters when using the configuration matrix. (a) original dataset. (b) favourable dataset.

FIG. 8. Distribution of the labels with respect to the number of projective spaces. (a) original dataset. (b) original dataset. (c) favourable dataset. (d) favourable dataset.

1. Test split and validation

The training set is made of 90% of the samples for training, which leaves the remaining 10% in the test set (i.e., 785 manifolds out of the 7851 in the set).¹⁰

For most algorithms, we use *leave-one-out* cross-validation on the training set as evaluation of the algorithm: we subdivide the training set in 9 subsets, each of them containing 10% of the *total* amount of samples, then, we train the algorithm on 8 of them and evaluate it on the 9th. We then repeat the procedure changing the evaluation fold until the algorithm has been trained and evaluated on all of them. The performance measure in validation is given by the average over all the left out folds.

When training neural networks, we will however use a single *holdout validation* set made of 10% of the *total* samples.

2. Predictions and metrics

Since we are interested in predicting exactly the Hodge numbers, the appropriate metric measuring the success of the predictions is the accuracy (for each Hodge number separately):

$$\text{accuracy} = \frac{1}{N} \sum_{i=1}^N \delta(y_i^{\text{true}} - y_i^{\text{pred}}), \quad (3.1)$$

where N is the number of samples. In the paper, accuracy of the predictions on the test set is rounded to the nearest integer.

Since the Hodge numbers are integers, the problem of predicting them looks like a classification task. However, as argued in the Introduction, we prefer to use a regression approach. Indeed, regression does not require to specify the data boundaries and allows to extrapolate beyond them, contrary to a classification approach where the categories are fixed at the beginning.¹¹

Most algorithms need a differentiable loss function since the optimization of parameters (such as neural networks weights) uses some variant of gradient descent. For this reason, the accuracy cannot be used and the models are trained by minimizing the mean squared error (MSE), which is simply the squared ℓ_2 -norm between of the difference between the predictions and the real values. There will however be also a restricted number of cases in which we will use either the mean absolute error (MAE),

which is the ℓ_1 -norm of the same difference, or a weighted linear combination of MSE and MAE (also known as *Huber* loss): we will point them out at the right time. When predicting both Hodge numbers together, the total loss is the sum of each individual loss with equal weight: $h^{1,1}$ is simpler to learn so it is useful to put emphasis on learning $h^{2,1}$, but the magnitudes of the latter are higher, such that the associated loss is naturally bigger (since we did not normalize the data).

Since predictions are real numbers, we need to turn them into integers. In general, rounding to the nearest integer gives the best result, but we found algorithms (such as linear regression) for which flooring to the integer below works better. The optimal choice of the integer function is found for each algorithm as part of the hyperparameter optimization (described below). The accuracy is computed after the rounding stage.

Learning curves for some models are displayed. They show how the performances of a model improves by using more training data, for fixed hyperparameters. To obtain it, we train models using from 10% to 90% of all the data (“training ratio”) and evaluate the accuracy on the remaining data.¹²

To avoid redundant information and to avoid cluttering the paper with graphs, the results for models predicting separately the Hodge numbers for the test set are reported in tables, while the results for the models predicting both numbers together are reported in the learning curves. For the same reason, the latter are not displayed for the favorable dataset.

3. Visualization of the performance

Complementary to the predictions and the accuracy results, we also provide different visualizations of the performance of the models in the form of univariate plots (histograms) and multivariate distributions (scatter plots).

The usual assumption behind the statistical inference of a distribution is that the difference between the observed data and the predicted values can be modelled by a random variable called *residual* [98,99].¹³ As such we expect that its values can be sampled from a normal distribution with a constant variance (i.e., constant width), since it should not depend on specific observations, and centered around zero, since the regression algorithm tries to minimize the squared difference between observed and predicted values.

¹⁰Remember that we have removed outliers, see Sec. II C 3. Scores quoted in this paper are slightly different from [82] because, in that paper, outliers are kept in the test set.

¹¹A natural way to transform the problem in a regression task is to *normalize* the Hodge numbers, for example by shifting by the mean value and diving by the standard deviation. Under this transformation, the Hodge numbers are mapped to real numbers. While normalizing often improve ML algorithms, we found that the impact was mild or even negative.

¹²Statistics are not provided due to the limitations of our available computational resources. However, we check manually on few examples that the reported results are typical.

¹³The difference between the non observable *true* value of the model and the observed data is known as *statistical error*. The difference between residuals and errors is subtle but the two definitions have different interpretations in the context of the regression analysis: in a sense, residuals are an estimate of the errors.

FIG. 9. Summary of the statistics for the distributions of both Hodge numbers. The colored box shows the three quartiles of the distributions, with the internal horizontal line corresponding to the median. The “whiskers” cover the interquartile range, i.e., 1.5 times the distance between the first and third quartiles, from the lower and upper limits of the boxes. Isolated points show the remaining outliers, which we however choose to keep to avoid excessively pruning the dataset. (a) $h^{1,1}$. (b) $h^{2,1}$.

Histograms of the residual errors should therefore exhibit such properties graphically.

Another interesting kind of visual realization of the residuals is to show their distribution against the variables used for the regression model: in the case of a simple regression model in one variable, it is customary to plot the residuals as a function of the independent variable, but in a multivariable regression analysis (such as the case at hand) the choice usually falls on the values predicted by the fit (not the observed data). We shall therefore plot the residuals as functions of the predicted values.¹⁴ Given the assumption of the random distribution of the residuals, they should not present strong correlations with the predictions and should not exhibit trends. In general the presence of correlated residuals is an indication of an incomplete or incorrect model which cannot explain the variance of the predicted data, meaning that the model is either not suitable for predictions or that we should add information (that is, add features) to it.

4. Hyperparameter optimization

One of the key steps in a ML analysis is the optimization of the *hyperparameters* of the algorithm. These are internal parameters of each estimator (such as the number of trees in a random forest or the amount of regularization in a linear model): they are not modified during the training of the model, but they directly influence it in terms of performance and outcome.

Hyperparameter optimization is performed by training many models with different hyperparameters, and keeping

those which perform best according to some metric on the validation set(s). As it does not need to be differentiable, we use the accuracy as a scoring function to evaluate the models. There is however a subtle issue because it is not clear how to combine the accuracy of $h^{1,1}$ and $h^{2,1}$ to get a single metric. For this reason, we will perform the analysis on both Hodge numbers separately. Then, we can design a single model computing both Hodge numbers simultaneously by making a compromise by hand between the hyperparameters found for the two models computing the Hodge numbers separately.

The optimization is implemented using the API from `scikit-learn`, using the function `metrics.make_scorer` and the accuracy as a custom scoring function. There are several approaches to perform this search automatically, in particular: grid search, random search, genetic evolution, and Bayes optimization.

Grid and random search are natively implemented in `scikit-learn`. The first takes a list of possible discrete values of the hyperparameters and will evaluate the algorithm over all possible combinations. The second samples values in both discrete sets and continuous intervals according to some probability distributions, repeating the process a fixed number of times. The grid search method is particularly useful for discrete hyperparameters, less refined searches or for a small number of combinations, while the second method can be used to explore the hyperparameter space on a larger scale [100]. Genetic algorithms are based on improving the choice of hyperparameters over *generations* that successively select only the most promising values: in general, they require a lot of tuning and are easily influenced by the fact that the replication process can also lead to worse results totally at random [101]. They are however effective when dealing with very deep or complex neural networks.

¹⁴We will use the same strategy also for the fit using just the number of projective spaces in order to provide a way to compare the plots across different models.

TABLE I. Hyperparameter choices of the ℓ_1 regression model used. In addition to the known hyperparameters α and `fit_intercept`, we also include the `normalize` parameter which indicates whether the samples have been centered and scaled by their ℓ_2 norm before the fit: it is ignored when the intercept is ignored.

		matrix		num_cp		eng. feat.		PCA	
		<i>old</i>	<i>fav.</i>	<i>old</i>	<i>fav.</i>	<i>old</i>	<i>fav.</i>	<i>old</i>	<i>fav.</i>
α	$h^{1,1}$	2.0×10^{-6}	3.0×10^{-5}	0.10	2.0×10^{-6}	0.05	0.05	0.07	0.08
	$h^{2,1}$	1.0×10^{-6}	1.0×10^{-5}	0.1	1.0×10^{-6}	3.0×10^{-4}	1.2×10^{-3}	2.0×10^{-6}	1.2×10^{-3}
<code>fit_intercept</code>	$h^{1,1}$	False	False	True	False	True	True	False	True
	$h^{2,1}$	True	True	True	True	True	False	True	False
<code>normalize</code>	$h^{1,1}$			False		False	False		False
	$h^{2,1}$	False	True	False	False	False		True	

Bayes optimization [102,103] is a very well-established mathematical procedure to find the stationary points of a function without knowing its analytical form [104]. It relies on assigning a *prior* probability to a given parameter and then multiply it by the probability distribution (or *likelihood*) of the scoring function to compute the probability of finding a better results given a set of hyperparameters. This has proven to be very effective in our case and we adopted this solution as it does not require fine tuning and leads to better results for models which are not deep neural networks. We choose to use `scikit-optimize` [105] whose method `BayesSearchCV` has a very well implemented PYTHON interface compatible with `scikit-learn`. We will in general perform 50 iterations of the Bayes search algorithm, unless otherwise specified.

C. Linear models

Linear models attempt to describe the labels as a linear combinations of the input features while keeping the coefficients at order one (Appendix A 1). However, non-linearity can still be introduced by engineering features which are nonlinear in terms of the original data.

From the results of Sec. II C, we made a hypothesis on the linear dependence of $h^{1,1}$ on the number of projective spaces m . As a first approach, we can try to fit a linear model to the data as a baseline computation and to test whether there is actual linear correlation between the two quantities. We will consider different linear models, including their regularized versions.

1. Parameters

The linear regression is performed with the class `linear_model.ElasticNet` from `scikit-learn`. The hyperparameters involved in this case are the amount of regularization α , the relative ratio (`l1_ratio`) between the ℓ_1 and ℓ_2 regularization losses, and the fit of the intercept.

By performing the hyperparameter optimization, we found that ℓ_2 regularization has a minor impact and can

be removed, which corresponds to setting the relative ratio to 1 (this is equivalent to using `linear_model.Lasso`).

In Table I we show the choices of the hyperparameters for the different models we built using the ℓ_1 regularized linear regression.

For the original dataset, we floored the predictions to the integers below, while in the favorable we rounded to the next integer. This choice for the original dataset makes sense: the majority of the samples lie on the line $h^{1,1} = m$, but there are still many samples with $h^{1,1} > m$ (see Fig. 8). As a consequence, the ML prediction pulls the line up, which can only damage the accuracy. Choosing the floor function is a way to counteract this effect. Note that accuracy for $h^{2,1}$ is only slightly affected by the choice of rounding, so we just choose the same one as $h^{1,1}$ for simplification.

2. Results

In Table II, we show the accuracy for the best hyperparameters. For $h^{1,1}$, the most precise predictions are given by the number of projective spaces which actually confirms the hypothesis of a strong linear dependence of $h^{1,1}$ on the number of projective spaces. In fact, this gives close to 100% accuracy for the favorable dataset, which shows that there is no need for more advanced ML algorithms. Moreover, adding more engineered features *decreases* the accuracy in most cases where regularization is not appropriate. The accuracy for $h^{2,1}$ remains low but including engineered features definitely improves it.

TABLE II. Best accuracy of the linear model using ℓ_1 regularization on the test split.

		matrix	num_cp	eng. feat.	PCA
<i>original</i>	$h^{1,1}$	51%	63%	63%	64%
	$h^{2,1}$	11%	8%	21%	21%
<i>favorable</i>	$h^{1,1}$	95%	100%	100%	100%
	$h^{2,1}$	14%	15%	19%	19%

FIG. 10. Plots of the residual error for the ℓ_1 regularized linear model: rows show the different scenarios (fit with only the matrix, with only the number of projective spaces, with the engineered features, with the engineered features and the PCA). Plots refer to the test split of the original dataset.

In Fig. 10, we show the plots of the residual errors of the model on the original dataset. For the ℓ_1 regularized linear model, the univariate plots show that the errors seem to follow normal distributions peaked at 0 as they generally should: in the case of $h^{1,1}$, the width is also quite contained. The scatter plots instead show that, in general, there is no correlation between a particular sector of the predictions and the error made by the model, thus the variance of the residuals is in general randomly distributed over the predictions. Only the case of the fit of the number of projective spaces seems to show a slight correlation for $h^{2,1}$, signaling that the model using only one feature might be actually incomplete: in fact it is better to include also other engineered features.

The learning curves (Fig. 11) clearly shows that the model underfits. Moreover, we also noticed that the models

are only marginally affected by the number of samples used for training. In particular, this provides a very strong baseline for $h^{1,1}$. For comparison, we also give the learning curve for the favorable dataset in Fig. 12: this shows that a linear regression is completely sufficient to determine $h^{1,1}$ in that case.

D. Support vector machines

Support vector machines (SVM) are a family of algorithms which use a *kernel trick* to map the space of input data vectors into a higher dimensional space where samples can be accurately separated and fitted to an appropriate curve (Appendix A 2).

In this analysis, we show two such kernels, namely a linear kernel (also known as *no kernel* since no

FIG. 11. Learning curves for the linear regression (original dataset), including outliers and using a single model for both Hodge numbers. (a) input: matrix, $\alpha = 2 \times 10^{-4}$. (b) input: num_cp, $\alpha = 1$.

transformations are involved) and a Gaussian kernel (known as `rbf` in ML literature, from *radial basis function*).

1. Linear kernel

For this model we use the class `svm.LinearSVR` in `scikit-learn`.

a. Parameters In Table III we show the choices of the hyperparameters used for the model. As we show in Appendix A 2 parameters C and ϵ are related to the penalty assigned to the samples lying outside the no-penalty boundary (the loss in this case is computed according to the ℓ_1 or ℓ_2 norm of the distance from the boundary as specified by the `loss` hyperparameter). Other parameters are related to the use of the intercept to improve the prediction.

We rounded the predictions to the floor for the original dataset and to the next integer for the favorable dataset.

FIG. 12. Learning curves for the linear regression (favorable dataset), including outliers and using a single model for both Hodge numbers. (a) input: num_cp, $\alpha = 1$.

b. Results In Table IV, we show the accuracy on the test set for the linear kernel. As we can see the performance of the algorithm strongly resembles a linear model in terms of the accuracy reached.

It is interesting to notice that the contributions of the PCA do not improve the predictions using just the engineered features: it seems that the latter work better than using the configuration matrix or its principal components.

The residual plots in Fig. 13 confirm what we already said about the linear models with regularization: the model with only the number of projective spaces shows a tendency to heteroscedasticity¹⁵ which can be balanced by adding more engineered feature, also helping in having more precise predictions (translated into peaked univariate distributions). In all cases, we notice that the model slightly overestimates the real values (residuals are computed as the difference between the prediction and the real value) as the second, small peaks in the histograms for $h^{1,1}$ suggest: this may also explain why flooring the predictions produces the highest accuracy.

As in general for linear models, the influence of the number of samples used for training is marginal also in this case: we only noticed a decrease in accuracy when also including the PCA or directly the matrix.

2. Gaussian kernel

We then consider SVM using a Gaussian function as kernel. The choice of the function can heavily influence the outcome of the predictions since they map the samples into a much higher dimensional space and create highly non-linear combinations of the features before fitting the algorithm. In general, this can help in the presence of

¹⁵That is the tendency to have a correlation between the predictions and the residuals: theoretically speaking, there should not be any, since we suppose the residuals to be independent on the model and normally sampled.

TABLE III. Hyperparameter choices of the linear SVR regression. The parameter `intercept_scaling` is clearly only relevant when the intercept is used. The different losses used simply distinguish between the ℓ_1 norm of the ϵ -dependent boundary where no penalty is assigned and its ℓ_2 norm.

		matrix		num_cp		eng. feat.		PCA	
		old	fav.	old	fav.	old	fav.	old	fav.
C	$h^{1,1}$	0.13	24	0.001	0.0010	0.13	0.001	0.007	0.4
	$h^{2,1}$	0.30	100	0.05	0.0016	0.5	0.4	1.5	0.4
ϵ	$h^{1,1}$	0.7	0.3	0.4	0.00	0.9	0.0	0.5	0.0
	$h^{2,1}$	0.0	0.0	10	0.03	0.0	0.0	0.0	0.6
fit_intercept	$h^{1,1}$	True	False	True	False	True	False	False	False
	$h^{2,1}$	True	False	True	True	True	True	True	False
intercept_scaling	$h^{1,1}$	0.13		100		0.01			
	$h^{2,1}$	100		13	92	100	0.01	100	
loss	$h^{1,1}$	$ \epsilon $	$ \epsilon $	$ \epsilon $	$ \epsilon ^2$	$ \epsilon $	$ \epsilon $	$ \epsilon $	$ \epsilon $
	$h^{2,1}$	$ \epsilon $	$ \epsilon $	$ \epsilon ^2$	$ \epsilon $	$ \epsilon $	$ \epsilon $	$ \epsilon $	$ \epsilon $

“obscure” features which badly correlate one another. In our case, we can hope to leverage the already good correlations we found in the EDA with the kernel trick. The implementation is done with the class `svm.SVR` from `scikit-learn`.

a. Parameters As we show in Appendix A 2, this particular choice of kernel leads to profoundly different behavior with respect to linear models: we will round the predictions to the next integer in both datasets since the loss function strongly penalizes unaligned samples.

In Table V, we show the choices of the hyperparameters for the models using the Gaussian kernel. As usual the hyperparameter C is connected to the penalty assigned to the samples outside the soft margin boundary (see Appendix A 2) delimited by the ϵ . Given the presence of a non linear kernel we have to introduce an additional hyperparameter γ which controls the width of the Gaussian function used for the support vectors.

b. Results In Table VI, we show the accuracy of the predictions on the test sets. In the favorable dataset, we can immediately appreciate the strong linear dependence of $h^{1,1}$ on the number of projective spaces: even though there are a few non favorable embeddings in the dataset, the kernel trick is able to map them in a better representation and improve the accuracy. The predictions for the original

dataset have also improved and are the best results we found using shallow learning. The predictions using only the configuration matrix matches [33], but we can slightly improve the accuracy by using a combination of engineered features and PCA.

In Fig. 14, we show the residual plots and their histograms for the original dataset: residuals follow peaked distributions which, in this case, do not present a second smaller peak (thus we need to round to the next integer the predictions) and good variate distribution over the predictions.

The Gaussian kernel is also more influenced by the size of the training set. Using 50% of the samples as training set we witnessed a drop in accuracy of 3% while using engineered features and the PCA, and around 1% to 2% in all other cases. The learning curves (Fig. 15) show that the accuracy improves by using more data. Interestingly, it shows that using all engineered features leads to an overfit on the training data since both Hodge numbers reach almost 100%, while this is not the case for $h^{2,1}$. For comparison, we also display in Fig. 16 the learning curve for the favorable dataset: this shows that predicting $h^{1,1}$ accurately works out-of-the-box.

E. Decision trees

We now consider two algorithms based on decision trees: random forests and gradient boosted trees. Decision trees are powerful algorithms which implement a simple decision rule (in the style of an *if... then... else...* statement) to classify or assign a value to the predictions. However, they have a tendency to adapt too well to the training set and to not be robust enough against small changes in the training data. We consider a generalization of this algorithm used for *ensemble learning*: this is a technique in ML which uses multiple estimators (they can be the same or different) to improve the performances.

TABLE IV. Accuracy of the linear SVM on the test split.

		matrix	num_cp	eng. feat.	PCA
original	$h^{1,1}$	61%	63%	65%	62%
	$h^{2,1}$	11%	9%	21%	20%
favorable	$h^{1,1}$	96%	100%	100%	100%
	$h^{2,1}$	14%	14%	19%	20%

FIG. 13. Plots of the residual errors for the SVM with linear kernel.

We will present the results of *random forests* of trees which increase the bias compared to a single decision tree, and *gradient boosted* decision trees, which can use smaller trees to decrease the variance and learn better

representations of the input data by iterating their decision functions and use information on the previous runs to improve (see Appendix A 3 for a more in-depth description).

TABLE V. Hyperparameter choices of the SVR regression with Gaussian kernel.

		matrix		num_cp		eng. feat.		PCA	
		<i>old</i>	<i>fav.</i>	<i>old</i>	<i>fav.</i>	<i>old</i>	<i>fav.</i>	<i>old</i>	<i>fav.</i>
C	$h^{1,1}$	14	1000	170	36	3	40	1.0	1000
	$h^{2,1}$	40	1000	1.0	1.0	84	62	45	40
ϵ	$h^{1,1}$	0.01	0.01	0.45	0.03	0.05	0.3	0.02	0.01
	$h^{2,1}$	0.01	0.01	0.01	0.09	0.29	0.10	0.20	0.09
γ	$h^{1,1}$	0.03	0.002	0.110	0.009	0.07	0.003	0.02	0.001
	$h^{2,1}$	0.06	0.100	0.013	1000	0.016	0.005	0.013	0.006

TABLE VI. Accuracy of the Gaussian SVM on the test split.

		matrix	num_cp	eng. feat.	PCA
<i>original</i>	$h^{1,1}$	70%	63%	66%	72%
	$h^{2,1}$	22%	10%	36%	34%
<i>favorable</i>	$h^{1,1}$	99%	100%	100%	100%
	$h^{2,1}$	22%	17%	32%	33%

1. Random forests

The random forest algorithm is implemented with Scikit's `ensemble.RandomForestRegressor`.

a. Parameters Hyperparameter tuning for decision trees can in general be quite challenging. From the general theory on random forests (Appendix A 3), we can try and

look for particular shapes of the trees: this ensemble learning technique usually prefers a small number of fully grown trees. We performed only 25 iterations of the optimization process due to the very long time taken to train all the decision trees.

In Table VII, we show the hyperparameters used for the predictions. As we can see from `n_estimator`, random forests are usually built with a small number of fully grown (specified by `max_depth` and `max_leaf_nodes`) trees (not always the case, though). In order to avoid overfit we also tried to increase the number of samples necessary to split a branch or create a leaf node using `min_samples_leaf` and `min_samples_split` (introducing also a weight on the samples in the leaf nodes specified by `min_weight_fraction_leaf` to balance the tree). Finally the criterion chosen by the optimization reflects the choice of the trees to measure the impurity

FIG. 14. Plots of the residual errors for the SVM with Gaussian kernel.

FIG. 15. Learning curves for the SVM with Gaussian kernel (original dataset), using a single model for both Hodge numbers. (a) input: matrix, $C = 15$, $\gamma = 0.03$, $\epsilon = 0.1$. (b) input: all, $C = 10$, $\gamma = 0.03$, $\epsilon = 0.1$.

FIG. 16. Learning curves for the SVM with Gaussian kernel (favorable dataset), using a single model for both Hodge numbers. (a) input: matrix, $C = 20$, $\gamma = \text{scale}$, $\epsilon = 0.1$. (b) input: all, $C = 20$, $\gamma = \text{scale}$, $\epsilon = 0.1$.

TABLE VII. Hyperparameter choices of the random forest regression.

		matrix		num_cp		eng. feat.		PCA	
		old	fav.	old	fav.	old	fav.	old	fav.
criterion	$h^{1,1}$	mse	mse	mae	mae	mae	mse	mae	mae
	$h^{2,1}$	mae	mae	mae	mae	mae	mae	mae	mae
max_depth	$h^{1,1}$	100	100	100	30	90	30	30	60
	$h^{2,1}$	90	100	90	75	100	100	100	60
max_leaf_nodes	$h^{1,1}$	100	80	90	20	20	35	90	90
	$h^{2,1}$	90	100	100	75	100	60	100	100
min_samples_leaf	$h^{1,1}$	1	1	1	15	1	15	1	1
	$h^{2,1}$	3	1	4	70	1	70	30	1
min_samples_split	$h^{1,1}$	2	30	20	35	10	10	100	100
	$h^{2,1}$	30	2	50	45	2	100	2	100
min_weight_fraction_leaf	$h^{1,1}$	0.0	0.0	0.0	1.7×10^{-3}	0.0	0.009	0.0	0.0
	$h^{2,1}$	3.0×10^{-4}	0.0	1.0×10^{-4}	0.13	0.0	0.0	0.0	0.0
n_estimators	$h^{1,1}$	10	100	45	120	155	300	10	300
	$h^{2,1}$	190	10	160	300	10	10	10	300

TABLE VIII. Accuracy of the random forests on the test split.

		matrix	num_cp	eng. feat.	PCA
<i>original</i>	$h^{1,1}$	55%	63%	66%	64%
	$h^{2,1}$	12%	9%	17%	18%
<i>favorable</i>	$h^{1,1}$	89%	99%	98%	98%
	$h^{2,1}$	14%	17%	22%	27%

of the predictions using either the mean squared error (mse) or the mean absolute error (mae) of the predictions (see Appendix A 3).

b. Results In Table VIII, we summarize the accuracy reached using random forests of decision trees as estimators. As we already expected, the contribution of the number of projective spaces helps the algorithm to generate

better predictions. In general, it seems that the engineered features alone can already provide a good basis for predictions. In the case of $h^{2,1}$, the introduction of the principal components of the configuration matrix also increases the prediction capabilities. As in most other cases, we used the floor function for the predictions on the original dataset and the rounding to next integer for the favorable one.

As usual, in Fig. 17 we show the histograms of the distribution of the residual errors and the scatter plots of the residuals. While the distributions of the errors are slightly wider than the SVM algorithms, the scatter plots of the residual show a strong heteroscedasticity in the case of the fit using the number of projective spaces: though quite accurate, the model is strongly incomplete. The inclusion of the other engineered features definitely helps and also leads to better predictions. Learning curves are displayed in Fig. 18.

FIG. 17. Plots of the residual errors for the random forests.

FIG. 18. Learning curves for the random forest (original dataset), including outliers and using a single model for both Hodge numbers. (a) input: matrix, default parameters. (b) input: all, default parameters.

2. Gradient boosted trees

We used the class ensemble.GradientBoostingRegressor from Scikit in order to implement the gradient boosted trees.

a. Parameters Hyperparameter optimization has been performed using 25 iterations of the Bayes search algorithm since by comparison the gradient boosting algorithms took the longest learning time. We show the chosen hyperparameters in Table IX.

With respect to the random forests, for the gradient boosting we also need to introduce the `learning_rate` (or *shrinking parameter*) which controls the gradient descent of the optimization which is driven by the choice of the `loss` parameters (`ls` is the ordinary least squares loss, `lad` is the least absolute deviation and `huber` is a combination of the previous two losses weighted by the hyperparameter α). We also introduce the subsample hyperparameter which chooses a fraction of the samples to be fed into the algorithm at each iteration. This procedure has both a regularization effect on the trees, which should

TABLE IX. Hyperparameter choices of the gradient boosted decision trees.

		matrix		num_cp		eng. feat.		PCA	
		old	fav.	old	fav.	old	fav.	old	fav.
α	$h^{1,1}$	0.4							
	$h^{2,1}$		0.11			0.99			
criterion	$h^{1,1}$	mae	mae	friedman_mse	mae	friedman_mse	friedman_mse	mae	mae
	$h^{2,1}$	mae	mae	friedman_mse	mae	mae	mae	mae	mae
learning_rate	$h^{1,1}$	0.3	0.04	0.6	0.03	0.15	0.5	0.04	0.03
	$h^{2,1}$	0.6	0.5	0.3	0.5	0.04	0.02	0.03	0.07
loss	$h^{1,1}$	huber	ls	lad	ls	ls	lad	ls	ls
	$h^{2,1}$	ls	huber	ls	ls	huber	ls	ls	lad
max_depth	$h^{1,1}$	100	100	15	60	2	100	55	2
	$h^{2,1}$	85	100	100	30	35	60	15	2
min_samples_split	$h^{1,1}$	2	30	20	35	10	10	100	100
	$h^{2,1}$	30	2	50	45	2	100	2	100
min_weight_fraction_leaf	$h^{1,1}$	0.03	0.0	0.0	0.2	0.2	0.0	0.06	0.0
	$h^{2,1}$	0.0	0.0	0.16	0.004	0.0	0.0	0.0	0.0
n_estimators	$h^{1,1}$	90	240	120	220	100	130	180	290
	$h^{2,1}$	100	300	10	20	200	300	300	300
subsample	$h^{1,1}$	0.8	0.8	0.9	0.6	0.1	0.1	1.0	0.9
	$h^{2,1}$	0.7	1.0	0.1	0.9	0.1	0.9	0.1	0.2

TABLE X. Accuracy of the gradient boosting on the test split.

		matrix	num_cp	eng. feat.	PCA
original	$h^{1,1}$	50%	63%	61%	58%
	$h^{2,1}$	14%	9%	23%	21%
favorable	$h^{1,1}$	97%	100%	99%	99%
	$h^{2,1}$	17%	16%	35%	22%

not adapt too much to the training set, and speeds up the training (at least by a very small amount).

b. Results We show the results of gradient boosting in Table X. As usual, the linear dependence of $h^{1,1}$ on the number of projective spaces is evident and in this case also produces the best accuracy result (using the floor function for the original dataset and rounding to the next integer for

the favorable dataset) for $h^{1,1}$. $h^{2,1}$ is once again strongly helped by the presence of the redundant features.

In Fig. 19, we finally show the histograms and the scatter plots of the residual errors for the original dataset showing that also in this case the choice of the floor function can be justified and that the addition of the engineered features certainly improves the overall variance of the residuals.

F. Neural networks

In this section we approach the problem of predicting the Hodge numbers using artificial neural networks (ANN), which we briefly review in Appendix A 4. We use Google's *Tensorflow* framework and *Keras*, its high-level API, to implement the architectures and train the networks [88,89]. We explore different architectures and discuss the results.

Differently from the previous algorithms, we do not perform a cross-validation scoring but we simply retain

FIG. 19. Plots of the residual errors for the gradient boosted trees.

10% of the total set as a holdout validation set (also referred to as *development* set) due to the computation power available. Thus, we use 80% of the samples for training, 10% for evaluation and 10% as a test set. For the same reason, the optimization of the algorithm has been performed manually.

We always use the Adam optimizer with default learning rate 10^{-3} to perform the gradient descent and a fix batch size of 32. The network is trained for a large number of epochs to avoid missing possible local optima. In order to avoid overshooting the minimum, we dynamically reduce the learning rate both using the *Adam* optimizer, which implements learning rate decay, and through the callback `callbacks.ReduceLROnPlateau` in Keras, which scales the learning rate by a given factor when the monitored quantity (e.g., the validation loss) does not decrease): we choose to reduce it by 0.3 when the validation loss does not improve for at least 75 epochs. Moreover, we stop training when the validation loss does not improve during 200 epochs. Clearly, we then keep only the weights of the neural networks which gave the best results. Batch normalization layers are used with a momentum of 0.99.

Training and evaluation were performed on a NVidia GeForce 940MX laptop GPU with 2 GB of RAM memory.

1. Fully connected network

First, we reproduce the analysis from [33] for the prediction of $h^{1,1}$.

a. Model The neural network presented in [33] for the regression task contains 5 hidden layers with 876, 461, 437, 929 and 404 units (Fig. 20). All layers (including the output

FIG. 20. Architecture of the fully connected network to predict $h^{1,1}$. For simplicity we do not draw the dropout and batch normalization layers present after every FC layer.

TABLE XI. Accuracy (approximate) for $h^{1,1}$ obtained in [33], Fig. 1].

	Training data				
	10%	30%	50%	70%	90%
Regression	58%	68%	72%	75%	75%
Classification	68%	78%	82%	85%	88%

TABLE XII. Accuracy using *rint* rounding on the predictions of the ANNs on $h^{1,1}$ and $h^{2,1}$ on the test set.

	DenseNet		classic ConvNet		inception ConvNet	
	<i>old</i>	<i>fav.</i>	<i>old</i>	<i>fav.</i>	<i>old</i>	<i>fav.</i>
$h^{1,1}$	77%	97%	94%	99%	99%	99%
$h^{2,1}$			36%	31%	50%	48%

layer) are followed by a ReLU activation and by a dropout layer with a rate of 0.2072. This network contains roughly 1.58×10^6 parameters.

The other hyperparameters (like the optimizer, batch size, number of epochs, regularization, etc.) are not mentioned. In order to reproduce the results, we have filled the gap as follows:

- (i) Adam optimizer with batch size of 32;
- (ii) maximal number epochs of 2000 without early stopping;¹⁶
- (iii) we implement learning rate reduction by 0.3 after 75 epochs without improvement of the validation loss;
- (iv) no ℓ_1 or ℓ_2 regularization;
- (v) a batch normalization layer [106] after each fully connected layer.

b. Results We have first reproduced the results from [33], which are summarized in Table XI. The training process was very quick and the loss function is reported in Fig. 21. We obtain an accuracy of 77% both on the development and the test set of the original dataset with 80% of training data (see Table XII). Using the same network, we also achieved 97% of accuracy in the favorable dataset.

2. Convolutional network

We then present a new purely convolutional network to predict $h^{1,1}$ and $h^{2,1}$, separately or together. The advantage of such networks is that it requires a smaller number of parameters and is insensitive to the size of the inputs. The latter point can be helpful to work without padding the matrices (of the same or different representations), but the use of a flatten layer removes this benefit.

¹⁶It took around 20 minutes to train the model.

FIG. 21. Loss function of the FC network in the original dataset.

a. Model The neural network has 4 convolutional layers. They are connected to the output layer with a intermediate flatten layer. After each convolutional layer, we use the ReLU activation function and a batch normalization layer (with momentum 0.99). Convolutional layers use the padding option same and a kernel of size (5,5) to be able to extract more meaningful representations of the input, treating the configuration matrix somewhat similarly to an object segmentation task [107]. The output layer is also followed by a ReLU activation in order to force the prediction to be a positive number. We use a dropout layer only after the convolutional network (before the flatten layer), but we introduced a combination of ℓ_2 and ℓ_1 regularization to reduce the variance. The dropout rate is 0.2 in the original dataset and 0.4 for the favorable dataset, while ℓ_1 and ℓ_2 regularization are set to 10^{-5} . We train the model using the *Adam* optimizer with a starting learning rate of 10^{-3} and a minibatch size of 32.

The architecture is more similar in style to the old *LeNet* presented for the first time in 1998 by Y. LeCun during the ImageNet competition. In our implementation, however, we do not include the pooling operations and swap the usual order of batch normalization and activation function by first putting the ReLU activation.

In Fig. 22, we show the model architecture in the case of the original dataset and of predicting $h^{1,1}$ alone. The convolution layers have 180, 100, 40, and 20 units each.

b. Results With this setup, we were able to achieve an accuracy of 94% on both the development and the test sets for the “old” database and 99% for the favorable dataset in both validation and test sets (results are briefly summarized in Table XII). We thus improved the results of the densely connected network and proved that convolutional networks can be valuable assets when dealing with the extraction of a

good representation of the input data: not only are CNNs very good at recognizing patterns and rotationally invariant objects inside pictures or general matrices of data, but deep architectures are also capable of transforming the input using nonlinear transformations [108] to create new patterns which can then be used for predictions.

Even though the convolution operation is very time consuming, another advantage of CNN is the extremely reduced number of parameters with respect to FC networks.¹⁷ The architectures we used were in fact made of approximately 5.8×10^5 parameters: way less than half the number of parameters used in the FC network. Ultimately, this leads to a smaller number of training epochs necessary to achieve good predictions (see Fig. 23).

Using this classic setup, we tried different architectures. The network for the original dataset seems to work best in the presence of larger kernels, dropping by roughly 5% in accuracy when a more “classical” 3×3 kernel is used. We also tried to use to set the padding to *valid*, reducing the input from a 12×15 matrix to a 1×1 feature map over the course of 5 layers with 180, 100, 75, 40, and 20 filters. The advantage is the reduction of the number of parameters (namely $\sim 4.9 \times 10^5$) mainly due to the small FC network at the end, but accuracy dropped to 87%. The favorable dataset seems instead to be more independent of the specific architecture, retaining accuracy also with smaller kernels.

The analysis for $h^{2,1}$ follows the same prescriptions. For both the original and favorable dataset, we opted for 4 convolutional layers with 250, 150, 100, and 50 filters and no FC network for a total amount of 2.1×10^6 parameters.

In this scenario we were able to achieve 36% of accuracy in the development set and 40% on the test set for $h^{2,1}$ in the “old” dataset and 31% in both development and test sets in the favorable set (see Table XII).

The learning curves for both Hodge numbers are given in Fig. 24. This model uses the same architecture as the one for predicting $h^{1,1}$ only, which explains why it is less accurate as it needs to also adapt to compute $h^{2,1}$ —a difficult task, as we have seen (see for example Fig. 28).

3. Inception-like neural network

In the effort to find a better architecture, we took inspiration from Google’s winning CNN in the annual *ImageNet challenge* in 2014 [79–81]. The architecture presented uses *inception* modules in which separate 3×3 , 5×5 convolutions are performed side by side (together with *max pooling* operations) before recombining the outputs. The modules are then repeated until the output layer is reached. This has two evident advantages: users can avoid taking a completely arbitrary decision on the type of convolution to use since the network will take care of it

¹⁷It took around 4 hours of training (and no optimization) for each Hodge number in each dataset.

FIG. 22. Pure convolutional neural network for predicting $h^{1,1}$. It is made of 4 modules composed by convolutional layer, ReLU activation, batch normalization (in this order), followed by a dropout layer, a flatten layer and the output layer (in this order).

FIG. 23. Loss function of the networks for the prediction of $h^{1,1}$ and $h^{2,1}$. We can see that the validation loss flattens out while the training loss keeps decreasing: we took care of the overfit by using the weights of the network when the validation loss reached its minimum. The use of minibatch gradient descent also completely spoils the monotonicity of the loss functions which can therefore increase moving from one epoch to the other, while keeping the descending trend for most of its evolution. (a) Loss function of $h^{1,1}$. (b) Loss function of $h^{2,1}$.

tuning the weights, and the number of parameters is extremely restricted as the network can learn complicated functions using fewer layers. As a consequence the architecture of such models can be made very deep while keeping the number of parameters contained, thus being able to learn very difficult representations of the input and producing accurate predictions. Moreover, while the training phase might become very long due to the complicated convolutional operations, the small number of parameters is such that predictions can be generated in a very small amount of time, making inceptionlike models extremely appropriate whenever quick predictions are necessary.

Another advantage of the architecture is the presence of different kernel sizes inside each module: the network automatically learns features at different scales and different positions, thus leveraging the advantages of a deep architecture with the ability to learn different representations at the same time and compare them.

a. Model In Fig. 25, we show a schematic of our implementation. Differently from the image classification task, we drop the pooling operation and implement two side-by-side convolution over rows (12×1 kernel for the original dataset, 15×1 for the favorable) and one over

FIG. 24. Learning curves for the classic convolutional neural network (original dataset), using a single model for both Hodge numbers.

columns (1×15 and 1×18 respectively).¹⁸ We use same as padding option. The output of the convolutions are then concatenated in the filter dimensions before repeating the “inception” module. The results from the last module are directly connected to the output layer through a flatten layer. In both datasets, we use batch normalization layers (with momentum 0.99) after each concatenation layer and a dropout layer (with rate 0.2) before the FC network.¹⁹

For both $h^{1,1}$ and $h^{2,1}$ (in both datasets), we used 3 modules made by 32, 64 and 32 filters for the first Hodge number, and 128, 128 and 64 filters for the second. We also included ℓ_1 and ℓ_2 regularization of magnitude 10^{-4} in all cases. The number of parameters was thus restricted to 2.3×10^5 parameters for $h^{1,1}$ in the original dataset and 2.9×10^5 in the favorable set, and 1.1×10^6 parameters for $h^{2,1}$ in the original dataset and 1.4×10^6 in the favorable dataset. In all cases, the number of parameters has decreased by a significant amount: in the case of $h^{1,1}$ they are roughly $\frac{1}{3}$ of the parameters used in the classical CNN and around $\frac{1}{6}$ of those used in the FC network.

For training we used the *Adam* gradient descent with an initial learning rate of 10^{-3} and a batch size of 32. The callbacks helped to contain the training time (without

¹⁸Pooling operations are used to shrink the size of the input. Similar to convolutions, they use a window of a given size to scan the input and select particular values inside. For instance, we could select the average value inside the small portion selected, performing an *average pooling* operation, or the maximum value, a *max pooling* operation. This usually improves image classification and object detection tasks as it can be used to sharpen edges and borders.

¹⁹The position of the batch normalization is extremely important as the parameters computed by such layer directly influence the following batch. We however opted to wait for the scan over rows and columns to finish before normalising the outcome to avoid biasing the resulting activation function.

optimization) under 5 hours for each Hodge number in each dataset. Training curves are given in Fig. 26.

b. Results With these architectures, we were able to achieve more than 99% of accuracy for $h^{1,1}$ in the test set (same for the development set) and 50% of accuracy for $h^{2,1}$ (a slightly smaller value for the development set). We report the results in Table XII.

We therefore increased the accuracy for both Hodge numbers (especially $h^{2,1}$) compared to what can achieve a simple sequential network, while at the same time reducing significantly the number of parameters of the network.²⁰ This increases the robustness of the method and its generalization properties.

In Fig. 27, we show the distribution of the residuals and their scatter plot, showing that the distribution of the errors does not present pathological behavior and the variance of the residuals is well distributed over the predictions.

In fact, this neural network is much more powerful than the previous networks we considered, as can be seen by studying the learning curves (Fig. 28). When predicting only $h^{1,1}$, it surpasses 97% accuracy using only 30% of the data for training. While it seems that the predictions suffer when using a single network for both Hodge numbers, this remains much better than any other algorithm. It may seem counterintuitive that convolutions work well on this data since they are not translation or rotation invariant, but only permutation invariant. However, convolution alone is not sufficient to ensure invariances under these transformations but it must be supplemented with pooling operations [1], which we do not use. Moreover, convolution layers do more than just taking translation properties into account: they allow to make highly complicated combinations of the inputs and to share weights among components, which allow to find subtler patterns than standard fully connected layers. This network is more studied in more details in [82].

4. Boosting the inceptionlike model

To improve further the accuracy of $h^{2,1}$, we have tried to modify the network by adding engineered features as auxiliary inputs. This can be done by adding inputs to the inception neural network and merging the different branches at different stages. There are two possibilities to train such a network: (1) train the full network directly, or (2) train the inception network alone, then freeze its weights and connect it to the additional inputs, training

²⁰In an attempt to improve the results for $h^{2,1}$ even further, we also considered to first predict $\ln(1 + h^{2,1})$ and then transform it back. However, the predictions dropped by almost 10% in accuracy even using the “inception” network: the network seems to be able to approximate quite well the results (not better nor worse than simply $h^{2,1}$) but the subsequent exponentiation is taking apart predictions and true values. Choosing a correct rounding strategy then becomes almost impossible.

FIG. 25. In each concatenation module (here shown for the “old” dataset) we operate with separate convolution operations over rows and columns, then concatenate the results. The overall architecture is composed of 3 “inception” modules made by two separate convolutions, a concatenation layer and a batch normalization layer (strictly in this order), followed by a dropout layer, a flatten layer and the output layer with ReLU activation (in this order).

FIG. 26. The loss functions of “inception” network for $h^{1,1}$ and $h^{2,1}$ in the original dataset show that the number of epochs required for training is definitely larger than for simpler architectures, despite the reduced number of parameters. (a) Loss of $h^{1,1}$. (b) Loss of $h^{2,1}$.

only the new layer. We found that the architectures we tried did not improve the accuracy, but we briefly describe our attempts for completeness.

We focused in particular on the number of projective spaces, the vector of dimensions of the projective spaces and the vector of dimensions of the principal cohomology

group) and predicting $h^{1,1}$ and $h^{2,1}$ at the same time. The core of the neural network is the inception network described in Sec. III F 3. Then, the engineered features are processed using fully connected layers and merged to the predictions from the inception branch using a concatenation layer. Obviously, output layers for $h^{1,1}$ and $h^{2,1}$ can

FIG. 27. Histograms of the residual errors and residual plots of the Inception network. (a) Residuals of $h^{1,1}$. (b) Residuals of $h^{2,1}$.

FIG. 28. Learning curves for the inception neural network (original dataset). (a) predicting both $h^{1,1}$ and $h^{2,1}$. (b) predicting $h^{1,1}$ only.

be located on different branches, which allow for different processing of the features.

As mentioned earlier, a possible approach is to first train the inception branch alone, before freezing its weights and connecting it to the rest of the network. This can prevent spoiling the already good predictions and speed up the new learning process. This is a common technique called *transfer learning*: we can use a model previously trained on a slightly different task and use its weights as part of the new architecture.

Our trials involved shallow fully connected layers (1–3 layers with 10 to 150 units) between the engineered features and after the concatenation layer. Since the EDA analysis (Sec. II C) shows a correlation between both Hodge numbers, we tried architectures where the result for $h^{1,1}$ is used to predict $h^{2,1}$.

For the training phase, we also tried an alternative to the canonical choice of optimising the sum of the losses. We first train the network and stop the process when the validation loss for $h^{1,1}$ does not longer improve, load back the best weights and save the results, keep training and stop when the loss for $h^{2,1}$ reaches a plateau.

With this setup we were able to slightly improve the predictions of $h^{1,1}$ in the original dataset, reaching almost 100% of accuracy in the predictions, while the favorable dataset stayed at around 99% of accuracy. The only few missed predictions (4 manifolds out of 786 in the test set) are in very peculiar regions of the distribution of the Hodge number. For $h^{2,1}$ no improvement has been noticed.

G. Ensemble learning: Stacking

We conclude the ML analysis by describing a method very popular in ML competitions [74]: ensembling. This consists in taking several ML algorithms and combining together the predictions of each individual model to obtain a more precise predictions. Using this technique it is possible to decrease the variance and improve generalization by compensating weaknesses of algorithms with strengths of others. Indeed, the idea is to put together algorithms which perform best in different zones of the label distribution in order to combine them to build an algorithm better than any individual component.

The simplest such algorithm is *stacking* whose principle is summarized in Fig. 29. First, the original training set is split in two parts (not necessarily even). Second, a certain number of *first-level learners* is trained over the first split and used to generate predictions over the second split. Third, a “meta learner” is trained of the second split to combine the predictions from the first-level learners. Predictions for the test set are obtained by applying both level of models one after the other.

We have selected the following models for the first level: linear regression, SVR with the Gaussian kernel, the random forest and the “inception” neural network. The meta-learner is a simple linear regression with ℓ_1 regularization (Lasso).

FIG. 29. Stacking ensemble learning with two level learning. The original training set is split into two training folds and the first level learners are trained on the first. The trained models are then used to generate a new training set (here the “1st level labels”) using the second split as input features. The same also applies to the test set. Finally a “meta-learner” uses the newly generated training set to produce the final predictions on the test set.

The motivations for the first-level algorithms is that stacking works best with a group of algorithms which work in the most diverse way among them.

Also in this case, we use a cross-validation strategy with 5 splits for each level of the training: from 90% of total training set, we split into two halves containing each 45% of the total samples and then use 5 splits to grade the algorithm, thus using 9% of each split for cross correlation at each iteration) and the Bayes optimization for all algorithms but the ANN (50 iterations for elastic net, SVR and lasso and 25 for the random forests). The ANN was trained using a holdout validation set containing the same number of samples as each cross-validation fold, namely 9% of the total set. The accuracy is then computed as usual using `numpy rint` for SVR, neural networks, the meta learner and $h^{1,1}$ in the original dataset in general, and `numpy floor` in the other cases.

In Table XIII, we show the accuracy of the ensemble learning. We notice that accuracy improves slightly only for $h^{2,1}$ (original dataset) compared to the first-level learners. However, this is much lower than what has been achieved in Sec. III F 3. The reason is that the learning suffers from the reduced size of the training set. Another reason is that the different algorithms may perform similarly well in the same regions.

IV. DISCUSSION

In this paper, we have proved how a proper data analysis can lead to improvements in predictions of Hodge numbers

TABLE XIII. Accuracy of the first and second level predictions of the stacking ensemble for elastic net regression (EN), support vector with `rbf` kernel (SVR), random forest (RF) and the artificial neural network (ANN) as first level learners and lasso regression as meta learner.

		$h^{1,1}$		$h^{2,1}$	
		<i>old</i>	<i>fav.</i>	<i>old</i>	<i>fav.</i>
<i>1st level</i>	EN	65%	100%	19%	19%
	SVR	70%	100%	30%	34%
	RF	61%	98%	18%	24%
	ANN	98%	98%	33%	30%
<i>2nd level</i>	Lasso	98%	98%	36%	33%

$h^{1,1}$ and $h^{2,1}$ for CICY 3-folds. Moreover, considering more complex neural networks—in particular, architectures inspired by the inception model [79–81]—allowed us to reach close to 100% accuracy for $h^{1,1}$ with much less data and less parameters than in previous works.

While our analysis improved the accuracy for $h^{2,1}$ over what can be expected from a simple sequential neural network, we barely reached 50%. Hence, it would be interesting to push further our study to improve the accuracy. Possible solutions would be to use a deeper inception network, find a better architecture including engineered features, and refine the ensembling (for example using StackNet [109]).

Another interesting question to probe is related to representation learning, i.e., finding a better description of the Calabi-Yau. Indeed, one of the main difficulty in making predictions is the redundancy of the possible descriptions of a single manifold. For example, one could try to set up a map from any matrix to its favorable representation (if it exists). Or, on the contrary, one could generate more matrices for the same manifold in order to increase the size of the training set. Another possibility is to use the graph representation of the configuration matrix to which is automatically invariant under permutations [71] (another graph representation has been decisive in [48] to get a good accuracy). Techniques such as (variational) autoencoder [110–112], cycle GAN [113], invertible neural networks [114], graph neural networks [115,116] or more generally techniques from geometric deep learning [117] could be helpful. Related to this point, one could focus on nonfavorable configurations only (possibly, converting matrices in the original dataset which are already favorable) to understand better how this affects the learning. It could also be interesting to understand better how the algorithms perform on outliers when they are included in the test set but not in the training set.

Finally, our techniques apply directly to CICY 4-folds [76,77]. However, there are much more manifolds in this case, such that one can expect to reach a better accuracy for the different Hodge numbers (the different learning curves

for the 3-folds indicate that the model training would benefit from more data). We hope to report soon on these issues. Another interesting class of manifolds to explore with our techniques are generalized CICY 3-folds [118].

We leave these questions to future explorations.

ACKNOWLEDGMENTS

We are grateful to Sven Krippendorf and Fabian Rühle for useful discussions on the applications of machine learning to string theory and for comments on the draft. The work of H.E. has been conducted under a Carl Friedrich von Siemens Research Fellowship of the Alexander von Humboldt Foundation for postdoctoral researchers during part of this project. H.E. and R.F. are partially supported by the MIUR PRIN Contract No. 2015MP2CX4 “Non-perturbative Aspects of Gauge Theories and Strings”.

APPENDIX A: MACHINE LEARNING ALGORITHMS

1. Linear regression

Considering a set of F features $\{x_n\}$ where $n = 1, \dots, F$, a linear model learns a function

$$f(x_n) = \sum_{n=1}^F w_n x_n + b, \quad (\text{A1})$$

where w and b are the *weights* and *intercept* of the fit.

One of the key assumptions behind a linear fit is the independence of the residual error between the predicted point and the value of the model, which can therefore be assumed to be sampled from a normal distribution peaked at the average value [98,99]. The parameters of the fit are then chosen to maximize their *likelihood* function, or conversely to minimize its logarithm with a reversed sign (the χ^2 function). A related task is to minimize the mean squared error, without assuming a statistical distribution of the residual error: ML for regression usually implements this as loss function of the estimators. In this sense, loss functions for regression are more general than a likelihood approach, but are nonetheless related. For plain linear regression, the associated loss is

$$\mathcal{L}(w, b) = \frac{1}{2N} \sum_{i=1}^N \sum_{n=1}^F (y^{(i)} - (w_n x_n^{(i)} + b))^2, \quad (\text{A2})$$

where N is the number of samples and $x_n^{(i)}$ the n th feature of the i th sample. The values of the parameters will therefore be

$$(w, b) = \underset{w, b}{\operatorname{argmin}} \mathcal{L}(w, b). \quad (\text{A3})$$

This usually requires looping over all samples and all features, thus the *least squares* method has a time complexity of $O(F \times N)$: while the increase of the number of samples might be an issue, the number of engineered features and matrix components usually does not change and does not represent a huge effort in terms of rescaling the algorithm.

There are however different versions of possible regularization which we might add to constrain the parameters of the fit and avoid adapting too well to the training set. In particular we may be interested in adding a ℓ_1 regularization:

$$\mathcal{L}_1(w) = \sqrt{\sum_{n=1}^F w_n^2}, \quad (\text{A4})$$

or the ℓ_2 version:

$$\mathcal{L}_2(w) = \sum_{n=1}^F w_n^2. \quad (\text{A5})$$

Notice that in general we do not regularize the intercept. These terms can be added to the plain loss function to try and avoid large parameters to influence the predictions and to keep better generalization properties:

(i) add both ℓ_1 and ℓ_2 regularization (this is called *elastic net*):

$$\begin{aligned} \mathcal{L}_{en}(w, b; \alpha_{en}, L) = & \mathcal{L}(w, b) + \alpha_{en} \cdot L \cdot \mathcal{L}_1(w) \\ & + \frac{\alpha_{en}}{2} \cdot (1 - L) \cdot \mathcal{L}_2(w), \end{aligned} \quad (\text{A6})$$

(ii) keep only ℓ_1 regularization (i.e., the *lasso* regression):

$$\mathcal{L}_{lss}(w, b; \alpha_{lss}) = \mathcal{L}(w, b) + \alpha_{lss} \cdot \mathcal{L}_1(w), \quad (\text{A7})$$

(iii) keep only ℓ_2 regularization (*ridge* regression):

$$\mathcal{L}_{rdg}(w, b; \alpha_{rdg}) = \mathcal{L}(w, b) + \alpha_{rdg} \cdot \mathcal{L}_2(w). \quad (\text{A8})$$

The role of the hyperparameter L is to balance the contribution of the additional terms. For larger values of the hyperparameter α , w (and b) assume smaller values and adapt less to the particular training set.

2. Support vector machines for regression

This family of supervised ML algorithms were created with classification tasks in mind [119] but have proven to be effective also for regression problems [120]. Differently from the linear regression, instead of minimizing the squared distance of each sample, the algorithm assigns a penalty to predictions of samples $x^{(i)} \in \mathbb{R}^F$ (for $i = 1, 2, \dots, N$) which are further away than a certain hyperparameter ε from their true value y , allowing however a *soft margin* of tolerance represented by the penalties ζ above and ξ below. This is achieved by minimizing w , b , ζ , and ξ in the function²¹

$$\begin{aligned} \mathcal{L}(w, b, \zeta, \xi) = & \frac{1}{2} \sum_{n=1}^{F'} w_n^2 + C \sum_{i=1}^N (\zeta^{(i)} + \xi^{(i)}) \\ & + \sum_{i=1}^N \sum_{n=1}^{F'} \alpha^{(i)} (y^{(i)} - w_n \phi_n(x^{(i)}) - b - \varepsilon - \zeta^{(i)}) \\ & + \sum_{i=1}^N \sum_{n=1}^{F'} \beta^{(i)} (w_n \phi_n(x^{(i)}) + b - y^{(i)} - \varepsilon - \xi^{(i)}) \\ & - \sum_{i=1}^N (\rho^{(i)} \zeta^{(i)} + \sigma^{(i)} \xi^{(i)}) \end{aligned} \quad (\text{A9})$$

where $\alpha^{(i)}, \beta^{(i)}, \rho^{(i)}, \sigma^{(i)} \geq 0$ such that the previous expression encodes the constraints

$$\begin{cases} y^{(i)} - \sum_{n=1}^{F'} w_n \phi_n(x^{(i)}) - b \leq \varepsilon + \zeta^{(i)}, & \varepsilon \geq 0, \quad \zeta^{(i)} \geq 0, \quad i = 1, 2, \dots, N \\ \sum_{n=1}^{F'} w_n \phi_n(x^{(i)}) + b - y^{(i)} \leq \varepsilon + \xi^{(i)}, & \varepsilon \geq 0, \quad \xi^{(i)} \geq 0, \quad i = 1, 2, \dots, N \end{cases} \quad (\text{A10})$$

and where $\phi(x^{(i)}) \in \mathbb{R}^{F'}$ is a function mapping the feature vector $x^{(i)} \in \mathbb{R}^F$ in a higher dimensional space ($F' > F$),

whose interpretation will become clear in an instant. The minimization problem leads to

$$\begin{cases} w_n - \sum_{i=1}^N (\alpha^{(i)} - \beta^{(i)}) \phi_n(x^{(i)}) = 0 \\ \sum_{i=1}^N (\alpha^{(i)} - \beta^{(i)}) = 0 \\ \sum_{i=1}^N (\alpha^{(i)} + \rho^{(i)}) = \sum_{i=1}^N (\beta^{(i)} + \sigma^{(i)}) = C \end{cases} \quad (\text{A11})$$

such that $0 \leq \alpha^{(i)}, \beta^{(i)} \leq C, \forall i = 1, 2, \dots, N$. This can be reformulated as a *dual* problem in finding the extrema of $\alpha^{(i)}$ and $\beta^{(i)}$ in

²¹In a classification task the training objective would be the minimization of the opposite of the log-likelihood function of predicting a positive class, that is $y^{(i)}(w_n \phi_n(x^{(i)}) + b)$, which should equal the unity for good predictions (we can consider $\varepsilon = 1$), instead of the regression objective $y^{(i)} - w_n \phi_n(x^{(i)}) - b$. The differences between SVR for classification purposes and regression follow as shown.

$$W(\alpha, \beta) = \frac{1}{2} \sum_{i,j=1}^N \theta^{(i)} \theta^{(j)} K(x^{(i)}, x^{(j)}) - \varepsilon \sum_{i=1}^N (\alpha^{(i)} + \beta^{(i)}) + \sum_{i=1}^N y^{(i)} \theta^{(i)}, \quad (\text{A12})$$

where $\theta = \alpha - \beta$ are called *dual coefficients* (accessible through the attribute `dual_coef_` of `svm.SVR` in `scikit-learn`) and $K(x^{(i)}, x^{(j)}) = \sum_{n=1}^{F'} \phi_n(x^{(i)}) \phi_n(x^{(j)})$ is the *kernel* function. Notice that the Lagrange multipliers $\alpha^{(i)}$ and $\beta^{(i)}$ are nonvanishing only for particular sets of vectors $l^{(i)}$ which lie outside the ε dependent bounds of (A10) and operate as landmarks for the others. They are called *support vectors* (accessible using the attribute `support_vectors_` in `svm.SVR`), hence the name of the algorithm. There can be at most N when $\varepsilon \rightarrow 0^+$. As a consequence any sum involving $\alpha^{(i)}$ or $\beta^{(i)}$ can be restricted to the subset of support vectors. Using the kernel notation, the predictions will therefore be

$$y_{\text{pred}}^{(i)} = y_{\text{pred}}(x^{(i)}) = \sum_{n=1}^{F'} w_n \phi_n(x^{(i)}) + b = \sum_{a \in A} \theta^{(a)} K(x^{(i)}, l^{(a)}) + b, \quad (\text{A13})$$

where $A \subset \{1, 2, \dots, N\}$ is the subset of labels of the support vectors.

In Sec. III D we consider two different implementations of the SVM algorithm:

- (i) the *linear kernel*, namely the case when $K \equiv \text{id}$ and the loss, in the `scikit-learn` implementation of `svm.LinearSVR`, can be simplified to

$$\mathcal{L}(w, b) = C \sum_{i=1}^N \sum_{n=1}^{F'} \max(0, |y^{(i)} - w_n \phi_n(x^{(i)}) - b| - \varepsilon) + \frac{1}{2} \sum_{n=1}^{F'} w_n^2, \quad (\text{A14})$$

without resolving to the dual formulation of the problem.

- (ii) the Gaussian kernel (called `rbf`, from *radial basis function*) in which

$$K(x^{(i)}, l^{(a)}) = \exp\left(-\gamma \sum_{n=1}^F (x_n^{(i)} - l_n^{(a)})^2\right). \quad (\text{A15})$$

From the definition of the loss function (A9) and the kernels, we can appreciate the role of the main hyperparameters of the algorithm. While the interpretation of ε is straightforward as the margin allowed without penalty for the prediction, γ represents the width of the normal

distribution used to map the features in the higher dimensional space. Furthermore, C plays a similar role to the l_2 additional term in (A8) by controlling the entity of the penalty for samples outside the ε -dependent bound, however its relation to the linear regularization is $\alpha_{\text{ridge}} = C^{-1}$, thus $C > 0$ by definition.

Given the nature of the algorithm, SVMs are powerful tools which usually grant better results in both classification and regression tasks with respect to logistic and linear regression, but they scale poorly with the number of samples used during training. In particular the time complexity is at worst²² $O(F \times N^3)$ due to the quadratic nature of (A12) and the computation of the kernel function for all samples: for large datasets ($N \gtrsim 10^4$) they are usually outperformed by ANNs.

3. Decision trees, random forests, and gradient boosting

Decision trees are supervised ML algorithms which model simple decision rules based on the input data [121,122]. They are informally referred to with the acronym CART (from *classification and regression trees*) and their name descends from the binary tree structure coming from such decision functions separating the input data at each iteration (*node*), thus creating a bifurcating structure with *branches* (the different paths, or decisions made) and *leaves* (the samples in each branch): the basic idea behind them is an *if... then... else* structure. In `scikit-learn` this is implemented in the classes `tree.DecisionTreeClassifier` and `tree.DecisionTreeRegressor`.

The idea behind it is to take input samples $x^{(i)} \in \mathbb{R}^F$ (for $i = 1, 2, \dots, N$) and partition the space in such a way that data with the same label $y^{(i)} \in \mathbb{R}$ is on the same subset of samples (while for classification this may be natural to visualize, for regression this amounts to approximate the input data with a step function whose value is constant inside the partition). Let in fact $j = 1, 2, \dots, F$ be a feature and $x_j^{(i)}$ the corresponding value for the sample i , at each node n of the tree we partition the set of input data $\mathcal{M}_n$ into two subsets:

$$\begin{aligned} \mathcal{M}_n^{[1]}(t_{j,n}) &= \{(x^{(i)}, y^{(i)}) \in \mathbb{R}^F \times \mathbb{R} | x_j^{(i)} < t_{j,n} \forall i \in A_n\}, \\ \mathcal{M}_n^{[2]}(t_{j,n}) &= \mathcal{M}_n \setminus \mathcal{M}_n^{[1]}(t_{j,n}), \end{aligned} \quad (\text{A16})$$

where A_n is the full set of labels of the data samples in the node n and $t_{j,n} \in \mathbb{R}$ is a threshold value for the feature j at node n .

²²In general it is plausible that the time complexity is $O(F \times N^2)$ based on good implementations of caching in the algorithm.

The measure of the ability of the split to reach the objective (classifying or creating a regression model to predict the labels) is modeled through an *impurity* function (i.e., the measure of how often a random data point would be badly classified or how much it would be badly predicted). Common choices in classification tasks are the Gini impurity, a special quadratic case of the Tsallis entropy (which in turn is a generalization of the Boltzmann-Gibbs entropy, recovered as the first power of the Tsallis entropy) and the information theoretic definition of the entropy. In regression tasks it is usually given by the l_1 and l_2 norms of the deviation from different estimators (mean and median) for each node n :

(i) *mean absolute error*

$$H_n^{[l]}(x; t_{j,n}) = \frac{1}{|\mathcal{M}_n^{[l]}(t_{j,n})|} \sum_{i \in A_n^{[l]}} |y^{(i)} - \tilde{y}_{\text{pred},n}^{[l]}(x)|, \quad (x^{(i)}, y^{(i)}) \in \mathcal{M}_n(t_{j,n}), \quad (\text{A17})$$

(ii) *mean squared error*:

$$H_n^{[l]}(x; t_{j,n}) = \frac{1}{|\mathcal{M}_n^{[l]}(t_{j,n})|} \sum_{i \in A_n^{[l]}} (y^{(i)} - \tilde{y}_{\text{pred},n}^{[l]}(x))^2, \quad (x^{(i)}, y^{(i)}) \in \mathcal{M}_n(t_{j,n}), \quad (\text{A18})$$

where $|\mathcal{M}_n^{[l]}(t_{j,n})|$ is the cardinality of the set $\mathcal{M}_n^{[l]}(t_{j,n})$ for $l = 1, 2$ and

$$\begin{aligned} \tilde{y}_{\text{pred},n}^{[l]}(x) &= \text{median } i \in A_n^{[l]} y_{\text{pred}}(x^{(i)}), \\ \tilde{y}_{\text{pred},n}^{[l]}(x) &= \frac{1}{|A_n^{[l]}|} \sum_{i \in A_n^{[l]}} y_{\text{pred}}(x^{(i)}), \end{aligned} \quad (\text{A19})$$

where $A_n^{[l]} \subset A_n$ are the subset of labels in the left and right splits ($l = 1$ and $l = 2$, that is) of the node n .

The full measure of the impurity of the node n and for a feature j is then:

$$G_{j,n}(\mathcal{M}; t_{j,n}) = \frac{|\mathcal{M}_n^{[1]}(t_{j,n})|}{|\mathcal{M}_n|} H_n^{[1]}(x; t_{j,n}) + \frac{|\mathcal{M}_n^{[2]}(t_{j,n})|}{|\mathcal{M}_n|} H_n^{[2]}(x; t_{j,n}), \quad (\text{A20})$$

from which we select the parameters

$$\hat{t}_{j,n} = \text{argmin } t_{j,n} G_n(\mathcal{M}_n; t_{j,n}). \quad (\text{A21})$$

We then recurse over all $\mathcal{M}_n^{[l]}(\hat{t}_{j,n})$ (for $l = 1, 2$) until we reach the maximum allowed depth of the tree (at most $|\mathcal{M}_n| = 1$).

Other than just predicting a class or a numeric value, decision trees provide a criterion to assign the importance of each feature appearing in the nodes. The implementation of the procedure can however vary between different libraries: in `scikit-learn` the importance of a feature is computed by the total reduction in the objective function due to the presence of the feature, normalized over all nodes. Namely it is defined as the difference between the total impurity normalized by the total amount of samples in the node and the sum of the separate impurities of the left and right split normalized over the number of samples in the respective splits, summed over all the nodes. Thus features with a high *variable ranking* (or *variable importance*) are those with a higher impact in reducing the loss of the algorithm and can be expected to be seen in the initial branches of the tree. A measure of the variable importance is in general extremely useful for feature engineering and feature selection since it gives a natural way to pick features with a higher chance to provide a good prediction of the labels.

By nature decision trees have a query time complexity of $O(\log(N))$ as most binary search algorithms. However their definition requires running over all F features to find the best split for each sample thus increasing the time complexity to $O(F \times N \log(N))$. Summing over all samples in the whole node structure leads to the worst case scenario of a time complexity $O(F \times N^2 \log(N))$. Well balanced trees (that is, nodes are approximately symmetric with the same amount of data samples inside) can usually reduce that time by a factor N , but it may not always be the case.

Decision trees have the advantage to be very good at classifying or creating regression relations in the presence of “well separable” data samples and they usually provide very good predictions in a reasonable amount of time (especially when balanced). However if F is very large, a small variation of the data will almost always lead to a huge change in the decision thresholds and they are usually prone to overfit. There are however smart ways to compensate this behavior based on *ensemble* learning such as *bagging*²³ and *boosting* as well as *pruning* methods such as limiting the depth of the tree or the number of splits and introducing a dropout parameter to remove certain nodes of the tree. Also random forests of trees provide a variable ranking system by averaging the importance of each feature across all base estimators in the bagging aggregator.

²³The term *bagging* comes from the contraction of *bootstrap* and *aggregating*: predictions are in fact made over randomly sampled partitions of the training set with substitution (i.e., samples can appear in different partitions, known as *bootstrap* approach) and then averaged together (*aggregating*). Random forests are an improvement to this simple idea and work best for decision trees: while it is possible to bag simple trees and take their predictions, using the random subsampling as described usually leads to better performance and results.

As a reference, *random forests* of decision trees (ensemble.RandomForestRegressor in scikit-learn) are ensemble learning algorithms based on fully grown (deep) decision trees. They were created to overcome the issues related to overfitting and variability of the input data and are based on random sampling of the training data [123]. The idea is to take K random partitions of the training data and train a different decision tree for each of them and combine the results: for a classification task this would resort to averaging the *a posteriori* (or conditional) probability of predicting the class c given an input x (i.e., the Bayesian probability $P(c|x)$) over the K trees, while for regression this amounts to averaging the predictions of the trees $y_{\text{pred},\hat{n}}^{(i)\{k\}}$ where $k = 1, 2, \dots, K$ and $\hat{n}$ is the final node (i.e., the node containing the final predictions). This defines what has been called a *random forest* of trees which can usually help in improving the predictions by reducing the variance due to trees adapting too much to training sets.

Boosting methods are another implementation of ensemble learning algorithms in which more *weak learners*, in this case shallow decision trees, are trained over the training dataset [124,125]. In general parameters $\hat{t}_{j,n}$ in (A21) can be approximated by an expansion

$$t_{j,n}(x) = \sum_{m=0}^M t_{j,n}^{\{m\}}(x) = \sum_{m=0}^M \beta_{j,n}^{\{m\}} g(x; a_{j,n}^{\{m\}}), \quad (\text{A22})$$

where $g(x; a_{j,n}^{\{m\}})$ are called *base learners* and M is the number of iterations.²⁴ The values of $a_{j,n}^{\{m\}}$ and $\beta_{j,n}^{\{m\}}$ are enough to specify the value of $t_{j,n}(x)$ and can be computed by iterating (A21):

$$(a_{j,n}^{\{m\}}, \beta_{j,n}^{\{m\}}) = \underset{\{a_{j,n}, \beta_{j,n}\}}{\text{argmin}} G_{j,n}(\mathcal{M}_n; t_{j,n}^{\{m-1\}}(x) + \beta_{j,n} g(x; a_{j,n})). \quad (\text{A23})$$

The specific case of boosted trees is simpler since the base learner predicts a constant value $g(x; a_{j,n}^{\{m\}})$, thus (A23) simplifies to

$$\gamma_{j,n}^{\{m\}} = \underset{\gamma_{j,n}}{\text{argmin}} G_{j,n}(\mathcal{M}_n; t_{j,n}^{\{m-1\}}(x) + \gamma_{j,n}). \quad (\text{A24})$$

Ultimately, the value of the parameters in (A22) are updated using gradient descent as

²⁴Different implementations of the algorithm refer to the number of iterations in different way. For instance scikit-learn calls them `n_estimators` in the class `ensemble.GradientBoostingRegressor` in analogy to the random forest where the same name is given to the number of trained decision trees, while XGBoost prefers `num_boost_rounds` and `num_parallel_tree` to name the number of boosting rounds (the iterations) and the number of trees trained in parallel in a forest.

$$t_{j,n}^{\{m\}}(x) = t_{j,n}^{\{m-1\}}(x) + \nu \gamma_{j,n}^{\{m\}}, \quad (\text{A25})$$

where $0 \leq \nu \leq 1$ is the *learning rate* which controls the magnitude of the update. Through this procedure, boosted trees can usually vastly improve the predictions of very small decision trees by increasing variance over bias. Another way to prevent overfitting the training set is to randomly *subsample* the features vector by taking a subset of them (in scikit-learn it is represented as a percentage of the total number of features). Moreover scikit-learn introduces various ways to control the loss of gradient boosting: apart from the aforementioned *least squares* and *least absolute deviation*, we can have hybrid versions of these such as the *huber loss* which combines the two previous losses with an additional hyperparameter α [126]. While more implementations are present, also the boosted trees provide a way to measure the importance of the variables as any decision tree algorithm.

4. Artificial neural networks

ANNs are a state of the art algorithm in ML. They usually outperform any other algorithm in very large datasets (the size of our dataset is roughly at the threshold) and can learn very complicated decision boundaries and functions.²⁵ In the main text we used two types of neural networks: *fully connected* (FC) networks and *convolutional neural networks* (CNN). They both rely on being built in a layered structure, starting from the input layers (e.g., the configuration matrix of CY manifolds or an RGB image or several engineered features) going toward the output layers (e.g., the Hodge numbers or the classification class of the image).

In FC networks the input of layer l is a feature vector $a^{(i)\{l\}} \in \mathbb{R}^{n_l}$ (for $i = 1, 2, \dots, N$) and, as shown in Fig. 20, each layer is densely connected to the following.²⁶ In other words, each entry of the vectors $a_j^{(i)\{l\}}$ (for $j = 1, 2, \dots, n_l$) is mapped through a function ψ to all the components of the following layer $a^{\{l+1\}} \in \mathbb{R}^{n_{l+1}}$:

$$\psi: \mathbb{R}^{n_l} \rightarrow \mathbb{R}^{n_{l+1}} \\ a^{(i)\{l\}} \mapsto a^{(i)\{l+1\}} = \psi_j(a^{(i)\{l\}}), \quad (\text{A26})$$

such that

²⁵Despite their fame in the face of the general public, even small networks can prove to be extremely good at learning complicated functions in a small amount of time.

²⁶Clearly the input vector $x \in \mathbb{R}^F$ is equivalent to the vector $a^{\{0\}}$ and $n_0 = F$. Inputs to each layer are here represented as a matrix $a^{\{l\}}$ whose columns are made by samples and whose rows are filled with the values of the features.

$$a_j^{(i)\{l+1\}} = \psi_j(a^{(i)\{l\}}) = \phi\left(\sum_{k=1}^{n_l} a_k^{(i)\{l\}} W_{kj}^{\{l\}} + b^{\{l\}} \mathbb{I}_j\right), \quad (\text{A27})$$

where $\mathbb{I} \in \mathbb{R}^{n_{l+1}}$ is an identity vector. The matrix $W^{\{l\}}$ is *weight matrix* and $b^{\{l\}}$ is the *bias* term. The function ϕ is a non linear function and plays a fundamental role: without it the successive application of the linear map $a^{\{l\}} \cdot W^{\{l\}} + b \mathbb{I}$ would prevent the network from learning more complicated decision boundaries or functions as the ANN would only be capable of reproducing linear relations. ϕ is known as *activation function* and can assume different forms, as long as its non linearity is preserved (e.g., a *sigmoid* function in the output layer of a network squeezes the results in the interval $[0, 1]$ thus reproducing the probabilities of a classification). A common choice is the *rectified linear unit* (ReLU) function

$$\phi(z) = \text{ReLU}(z) = \max(0, z), \quad (\text{A28})$$

which has been proven to be better at training deep learning architectures [127], or its modified version LeakyReLU(z) = $\max(\alpha z, z)$ which introduces a slope $\alpha > 0$ to improve the computational performance near the nondifferentiable point in the origin.

CNN architectures were born in the context of computer vision and object localization [128]. As one can suspect looking at Fig. 22 for instance, the fundamental difference with FC networks is that they use a convolution operation $K^{\{l\}} * a^{(i)\{l\}}$ instead of a linear map to transform the output of the layers, before applying the activation function.²⁷ This way the network is no longer densely connected, as the results of the convolution (*feature map*) depends only on a restricted neighborhood of the original feature, depending on the size of the *kernel* window $K^{\{l\}}$ used and the shape of the input $a^{(i)\{l\}}$, which is no longer limited to flattened vectors. In turn, its size influences the convolution operator which we can compute: one way to see this is to visualize an image being scanned by a smaller window function over all pixels or by skipping some a certain number of them (the length of the *stride* of the kernel). In general the output will therefore be different than the input, unless the latter is *padded* (with zeros usually) before the convolution. The size of the output is therefore:

²⁷In general the input of each layer can be a generic tensor with an arbitrary number of axis. For instance, an RGB image can be represented by a three dimensional tensor with indices representing the width of the image, its height and the number of filters (in this case 3, one for each color channel).

$$O_n = \frac{I_n - k_n + 2p_n}{S_n} + 1, \quad n = 1, 2, \dots, \quad (\text{A29})$$

where O is the output size, I the input size, k the size of the kernel used, p the amount of padding (symmetric at the start and end of the axis considered) and S the stride. In the formula, n runs over the number of components of the input tensor. While any padding is possible, we are usually interested in two kinds of possible convolutions:

- (i) “same” convolutions for which $O_n = I_n$, thus $p_n = \frac{I_n(S_n-1) - S_n + k_n}{2}$,
- (ii) “valid” convolutions for which $O_n < I_n$ and $p_n = 0$.

In both cases the learning process aims to minimize the loss function defined for the task: in our regression implementation of the architecture we used the mean squared error of the predictions. The objective is to find best possible values of weight and bias terms $W^{\{l\}}$ and $b^{\{l\}}$ or to build the best filter kernel $K^{\{l\}}$ through *backpropagation* [129], that is by reconstructing the gradient of the loss function climbing back the network from the output layer to the input and then using the usual gradient descent procedure to select the optimal parameters. For instance, in the case of FC networks we need to find

$$(\hat{W}^{\{l\}}, \hat{b}^{\{l\}}) = \underset{W^{\{l\}}, b^{\{l\}}}{\operatorname{argmin}} \frac{1}{2N} \sum_{i=1}^N (y^{(i)} - a^{(i)\{L\}})^2$$

$$\forall l = 1, 2, \dots, L, \quad (\text{A30})$$

where L is the total number of layers in the network. A similar relation holds in the case of CNN architectures. In the main text we use the *Adam* [130] implementation of gradient descent and add batch normalization layers to improve the convergence of the algorithm.

As we can see from their definition, ANNs are capable of learning very complex structures at the cost of having a large number of parameters to tune. The risk of overfitting the training set is therefore quite evident. There are in general several techniques to counteract the tendency to adapt the training set, one of them being the introduction of regularization (l_2 and l_1) in the same fashion of a linear model (we show it in Appendix A 1). Another successful way is to introduce *dropout* layers [131] where connections are randomly switched off according to a certain retention probability (or its complementary, the *dropout rate*): this regularization technique allows to keep good generalization properties since the prediction can rely in a less incisive way on the particular architecture since which is randomly modified during training (dropout layers however act as the identity during predictions to avoid producing random results).

- [1] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning* (The MIT Press, Cambridge, Massachusetts, USA, 2016).
- [2] F. Chollet, *Deep Learning with Python*, 1st ed. (Manning Publications, Shelter Island, New York, USA, 2017).
- [3] A. Géron, *Hands-On Machine Learning with Scikit-Learn, Keras, and TensorFlow*, 2nd ed. (O'Reilly Media, USA, 2019).
- [4] K. Albertsson *et al.*, Machine learning in high energy physics community white paper, in *Proceedings, 18th International Workshop on Advanced Computing and Analysis Techniques in Physics Research (ACAT 2017)* (J. Phys. Conf. Ser., Seattle, WA, USA, 2018), Vol. 1085, p. 022008 [[arXiv:1807.02876](#)].
- [5] P. Mehta, M. Bukov, C.-H. Wang, A. G. R. Day, C. Richardson, C. K. Fisher, and D. J. Schwab, A high-bias, low-variance introduction to machine learning for physicists, *Phys. Rep.* **810**, 1 (2019).
- [6] M. Ntampaka *et al.*, The role of machine learning in the next decade of cosmology, [arXiv:1902.10159](#).
- [7] G. Carleo, I. Cirac, K. Cranmer, L. Daudet, M. Schuld, N. Tishby, L. Vogt-Maranto, and L. Zdeborová, Machine learning and the physical sciences, *Rev. Mod. Phys.* **91**, 045002 (2019).
- [8] M. Buchanan, The power of machine learning, *Nat. Phys.* **15**, 1208 (2019).
- [9] F. Ruehle, Data science applications to string theory, *Phys. Rep. Data Sci. Appl. String Theory*, **839**, 1 (2020).
- [10] W. Lerche, D. Lüst, and A. N. Schellekens, Chiral four-dimensional heterotic strings from self-dual lattices, in *Current Physics—Sources and Comments*, Superstring Construction, Vol. 4, edited by B. Schellekens (Elsevier, New York, 1989), pp. 252–282.
- [11] M. R. Douglas, The statistics of string/M theory vacua, *J. High Energy Phys.* **03** (2003) 046.
- [12] S. Ashok and M. R. Douglas, Counting flux vacua, *J. High Energy Phys.* **01** (2004) 060.
- [13] M. R. Douglas, Basic results in vacuum statistics, *C.R. Phys.* **5**, 965 (2004).
- [14] M. R. Douglas and S. Kachru, Flux compactification, *Rev. Mod. Phys.* **79**, 733 (2007).
- [15] W. Taylor and Y.-N. Wang, The F-theory geometry with most flux vacua, *J. High Energy Phys.* **12** (2015) 164.
- [16] A. N. Schellekens, Big numbers in string theory, [arXiv:1601.02462](#).
- [17] J. Halverson, C. Long, and B. Sung, On algorithmic universality in F-theory compactifications, *Phys. Rev. D* **96**, 126006 (2017).
- [18] W. Taylor and Y.-N. Wang, Scanning the skeleton of the 4D F-theory landscape, *J. High Energy Phys.* **01** (2018) 111.
- [19] A. Constantin, Y.-H. He, and A. Lukas, Counting string theory Standard Models, *Phys. Lett. B* **792**, 258 (2019).
- [20] F. Denef and M. R. Douglas, Computational complexity of the landscape I, *Ann. Phys. (Amsterdam)* **322**, 1096 (2007).
- [21] J. Halverson and F. Ruehle, Computational complexity of vacua and near-vacua in field and string theory, *Phys. Rev. D* **99**, 046015 (2019).
- [22] M. Graña, Flux compactifications in string theory: A comprehensive review, *Phys. Rep.* **423**, 91 (2006).
- [23] D. Lüst, Seeing through the string landscape—A string Hunter's companion in particle physics and cosmology, *J. High Energy Phys.* **03** (2009) 149.
- [24] L. E. Ibáñez and A. M. Uranga, *String Theory and Particle Physics: An Introduction to String Phenomenology*, 1st ed. (Cambridge University Press, Cambridge, England, 2012).
- [25] T. D. Brennan, F. Carta, and C. Vafa, The string landscape, the swampland, and the missing corner, in *Proceedings of Theoretical Advanced Study Institute Summer School 2017; Physics at the Fundamental Frontier;—PoS(TASI2017)* (SISSA Medialab, Trieste, Italy, 2018), Vol. 305, p. 015 [[arXiv:1711.00864](#)].
- [26] J. Halverson and P. Langacker, TASI lectures on remnants from the string landscape, *Proc. Sci. TASI2017* (2018) 019 [[arXiv:1801.03503](#)].
- [27] S. Abel and J. Rizos, Genetic algorithms and the search for viable string vacua, *J. High Energy Phys.* **08** (2014) 010.
- [28] D. Krefl and R.-K. Seong, Machine learning of Calabi-Yau volumes, *Phys. Rev. D* **96**, 066014 (2017).
- [29] F. Ruehle, Evolving neural networks with genetic algorithms to study the string landscape, *J. High Energy Phys.* **08** (2017) 038.
- [30] Y.-H. He, Machine-learning the string landscape, *Phys. Lett. B* **774**, 564 (2017).
- [31] J. Carifio, J. Halverson, D. Krioukov, and B. D. Nelson, Machine learning in the string landscape, *J. High Energy Phys.* **09** (2017) 157.
- [32] R. Altman, J. Carifio, J. Halverson, and B. D. Nelson, Estimating Calabi-Yau hypersurface and triangulation counts with equation learners, *J. High Energy Phys.* **03** (2019) 186.
- [33] K. Bull, Y.-H. He, V. Jejjala, and C. Mishra, Machine learning CICY threefolds, *Phys. Lett. B* **785**, 65 (2018).
- [34] A. Cole and G. Shiu, Topological data analysis for the string landscape, *J. High Energy Phys.* **03** (2019) 054.
- [35] D. Klaewer and L. Schlechter, Machine learning line bundle cohomologies of hypersurfaces in toric varieties, *Phys. Lett. B* **789**, 438 (2019).
- [36] A. Mütter, E. Parr, and P. K. S. Vaudrevange, Deep learning in the heterotic orbifold landscape, *Nucl. Phys. B* **940**, 113 (2019).
- [37] Y.-N. Wang and Z. Zhang, Learning Non-higgsable gauge groups in 4D F-theory, *J. High Energy Phys.* **08** (2018) 009.
- [38] A. Ashmore, Y.-H. He, and B. Ovrut, Machine learning Calabi-Yau metrics, *Fortschr. Phys.* **68**, 2000068 (2020).
- [39] C. R. Brodie, A. Constantin, R. Deen, and A. Lukas, Machine learning line bundle cohomology, *Fortschr. Phys.* **68**, 1900087 (2020).
- [40] K. Bull, Y.-H. He, V. Jejjala, and C. Mishra, Getting CICY high, *Phys. Lett. B* **795**, 700 (2019).
- [41] A. Cole, A. Schachner, and G. Shiu, Searching the landscape of flux vacua with genetic algorithms, *J. High Energy Phys.* **11** (2019) 045.
- [42] A. E. Faraggi, G. Harries, B. Percival, and J. Rizos, Towards machine learning in the classification of $Z_2 \times Z_2$

- orbifold compactifications, *J. Phys. Conf. Ser.* **1586**, 012032 (2020).
- [43] J. Halverson, B. Nelson, and F. Ruehle, Branes with brains: Exploring string vacua with deep reinforcement learning, *J. High Energy Phys.* **06** (2019) 003.
- [44] Y.-H. He and S.-J. Lee, Distinguishing elliptic fibrations with AI, *Phys. Lett. B* **798**, 134889 (2019).
- [45] M. Bies, M. Cvetič, R. Donagi, L. Lin, M. Liu, and F. Ruehle, Machine learning and algebraic approaches towards complete matter spectra in 4d F-theory, *J. High Energy Phys.* **01** (2021) 196.
- [46] N. C. Bizet, C. Damian, O. Loaiza-Brito, D. K. M. Peña, and J. A. Montañez-Barrera, Testing swampland conjectures with machine learning, *Eur. Phys. J. C* **80**, 766 (2020).
- [47] J. Halverson and C. Long, Statistical predictions in string theory and deep generative models, *Fortschr. Phys.* **68**, 2000005 (2020).
- [48] S. Krippendorff and M. Syvaeri, Detecting symmetries with neural networks, [arXiv:2003.13679](https://arxiv.org/abs/2003.13679).
- [49] H. Otsuka and K. Takemoto, Deep learning and K-means clustering in heterotic string vacua with line bundles, *J. High Energy Phys.* **05** (2020) 047.
- [50] E. Parr and P. K. S. Vaudrevange, Contrast data mining for the MSSM from strings, *Nucl. Phys. B* **952**, 114922 (2020).
- [51] E. Parr, P. K. S. Vaudrevange, and M. Wimmer, Predicting the orbifold origin of the MSSM, *Fortschr. Phys.* **68**, 2000032 (2020).
- [52] H. Erbin and S. Krippendorff, GANs for generating EFT models, *Phys. Lett. B* **810**, 135798 (2020).
- [53] P. Betzler and S. Krippendorff, Connecting dualities and machine learning, *Fortschr. Phys.* **68**, 2000022 (2020).
- [54] H.-Y. Chen, Y.-H. He, S. Lal, and M. Z. Zaz, Machine learning etudes in conformal field theories, [arXiv:2006.16114](https://arxiv.org/abs/2006.16114).
- [55] W.-C. Gan and F.-W. Shu, Holography as deep learning, *Int. J. Mod. Phys. D* **26**, 1743020 (2017).
- [56] K. Hashimoto, S. Sugishita, A. Tanaka, and A. Tomiya, Deep learning and AdS/CFT, *Phys. Rev. D* **98**, 046019 (2018).
- [57] K. Hashimoto, S. Sugishita, A. Tanaka, and A. Tomiya, Deep learning and holographic QCD, *Phys. Rev. D* **98**, 106014 (2018).
- [58] K. Hashimoto, AdS/CFT as a deep Boltzmann machine, *Phys. Rev. D* **99**, 106017 (2019).
- [59] J. Tan and C.-B. Chen, Deep learning the holographic black hole with charge, *Int. J. Mod. Phys. D* **28**, 1950153 (2019).
- [60] T. Akutagawa, K. Hashimoto, and T. Sumimoto, Deep learning and AdS/QCD, *Phys. Rev. D* **102**, 026020 (2020).
- [61] Y.-K. Yan, S.-F. Wu, X.-H. Ge, and Y. Tian, Deep learning black hole metrics from shear viscosity, *Phys. Rev. D* **102**, 101902 (2020).
- [62] I. M. Comsa, M. Firsching, and T. Fischbacher, SO(8) supergravity and the magic of machine learning, *J. High Energy Phys.* **08** (2019) 057.
- [63] N. Bobev, T. Fischbacher, F. F. Gautason, and K. Pilch, A cornucopia of AdS₅ vacua, *J. High Energy Phys.* **07** (2020) 240.
- [64] N. Bobev, T. Fischbacher, and K. Pilch, A new $N = 1$ AdS₄ vacuum of maximal supergravity, *J. High Energy Phys.* **01** (2020) 099.
- [65] C. Krishnan, V. Mohan, and S. Ray, Machine learning $\mathcal{N} = 8, D = 5$ gauged supergravity, *Fortschr. Phys.* **68**, 2000027 (2020).
- [66] P. Green and T. Hübsch, Calabi-Yau manifolds as complete intersections in products of complex projective spaces, *Commun. Math. Phys.* **109**, 99 (1987).
- [67] P. Candelas, A. M. Dale, C. A. Lütken, and R. Schimmrigk, Complete intersection Calabi-Yau manifolds, *Nucl. Phys. B* **298**, 493 (1988).
- [68] P. S. Green, T. Hübsch, and C. A. Lütken, All the hodge numbers for all Calabi-Yau complete intersections, *Classical Quantum Gravity* **6**, 105 (1989).
- [69] L. B. Anderson, X. Gao, J. Gray, and S.-J. Lee, Fibrations in CICY threefolds, *J. High Energy Phys.* **10** (2017) 077.
- [70] C. A. Lütken, Recent progress in Calabi-Yauology, *Nucl. Phys. B, Proc. Suppl.* **5**, 175 (1988).
- [71] T. Hübsch, *Calabi-Yau Manifolds: A Bestiary For Physicists* (Wspc, Singapore, 1992).
- [72] L. B. Anderson and M. Karkheiran, TASI lectures on geometric tools for string compactifications, *Proc. Sci. TASI2017* (2018) 013 [[arXiv:1804.08792](https://arxiv.org/abs/1804.08792)].
- [73] Y.-H. He, Calabi-Yau spaces in the string landscape, [arXiv:2006.16623](https://arxiv.org/abs/2006.16623).
- [74] D. Altukhov, A. Guschin, M. Michailidis, M. Trofimov, and D. Ulyanov, How to win a data science competition: Learn from Top Kagglers.
- [75] S. S. Skiena, *The Data Science Design Manual* (Springer, New York, 2017).
- [76] J. Gray, A. S. Haupt, and A. Lukas, Topological invariants and fibration structure of complete intersection Calabi-Yau four-folds, *J. High Energy Phys.* **09** (2014) 093.
- [77] J. Gray, A. S. Haupt, and A. Lukas, All complete intersection Calabi-Yau four-folds, *J. High Energy Phys.* **07** (2013) 070.
- [78] M. Kreuzer and H. Skarke, Complete classification of reflexive polyhedra in four dimensions, *Adv. Theor. Math. Phys.* **4**, 1209 (2000).
- [79] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich, Going deeper with convolutions, [arXiv:1409.4842](https://arxiv.org/abs/1409.4842).
- [80] C. Szegedy, V. Vanhoucke, S. Ioffe, J. Shlens, and Z. Wojna, Rethinking the inception architecture for computer vision, [arXiv:1512.00567](https://arxiv.org/abs/1512.00567).
- [81] C. Szegedy, S. Ioffe, V. Vanhoucke, and A. Alemi, Inception-v4, inception-resnet and the impact of residual connections on learning, [arXiv:1602.07261](https://arxiv.org/abs/1602.07261).
- [82] H. Erbin and R. Finotello, Inception neural network for complete intersection Calabi-Yau 3-folds, *Mach. Learn. Sci. Technol.* **2**, 02LT03 (2021).
- [83] W. McKinney, Data structures for statistical computing in Python, in *Proceedings of the 9th Python in Science Conference*, edited by S. van der Walt and J. Millman (2010), pp. 56–61.
- [84] J. D. Hunter, Matplotlib: A 2d graphics environment, *Comput. Sci. Eng.* **9**, 90 (2007).
- [85] M. Waskom *et al.*, *mwaskom/seaborn*: v0.8.1 (2017).

- [86] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and É. Duchesnay, Scikit-learn: Machine learning in python, *J. Mach. Learn. Res.* **12**, 2825 (2011), <https://dl.acm.org/doi/10.5555/1953048.2078195>.
- [87] T. Head *et al.*, scikit-optimize/scikit-optimize: v0.5.2 (2018).
- [88] M. Abadi *et al.*, Tensorflow: Large-scale machine learning on heterogeneous systems (2015).
- [89] F. Chollet, Keras, <https://keras.io>, 2015.
- [90] K. Becker, M. Becker, and J. H. Schwarz, *String Theory and M-Theory: A Modern Introduction*, 1st ed. (Cambridge University Press, Cambridge, England, 2006).
- [91] P. Green and T. Hübsch, Polynomial deformations and cohomology of Calabi-Yau manifolds, *Commun. Math. Phys.* **113**, 505 (1987).
- [92] L. B. Anderson, Y.-H. He, and A. Lukas, Monad bundles in heterotic string compactifications, *J. High Energy Phys.* **08** (2008) 104.
- [93] S. Thrun, Is learning the n-th thing any easier than learning the first?, in *Proceedings of the 8th International Conference on Neural Information Processing Systems*, NIPS'95 (MIT Press, 1995), pp. 640–646.
- [94] R. Caruana, Multitask learning, *Mach. Learn.* **28**, 41 (1997).
- [95] J. Baxter, A model of inductive bias learning, *J. Artif. Intell. Res.* **12**, 149 (2000).
- [96] A. Maurer, M. Pontil, and B. Romera-Paredes, The benefit of multitask representation learning, *J. Mach. Learn. Res.* **17**, 2853 (2016), <https://dl.acm.org/doi/10.5555/2946645.3007034>.
- [97] A. Ndirango and T. Lee, Generalization in multitask deep neural classifiers: A statistical physics approach, in *Advances in Neural Information Processing Systems* (2019), pp. 15862–15871 [[arXiv:1910.13593](https://arxiv.org/abs/1910.13593)].
- [98] L. Lista, *Statistical Methods for Data Analysis in Particle Physics*, Vol. 941 (Springer, New York, 2017).
- [99] B. Caffo, J. Leek, and R. D. Peng, Data science specialization.
- [100] J. Bergstra and Y. Bengio, Random search for hyperparameter optimization, *J. Mach. Learn. Res.* **13**, 281 (2012), <https://dl.acm.org/doi/10.5555/2188385.2188395>.
- [101] G. Rudolph, Convergence analysis of canonical genetic algorithms, *IEEE Trans. Neural Networks* **5**, 96 (1994).
- [102] J. Snoek, H. Larochelle, and R. P. Adams, Practical bayesian optimization of machine learning algorithms, in *Proceedings of the 25th International Conference on Neural Information Processing Systems—Volume 2*, NIPS'12 (Curran Associates Inc., Red Hook, NY, USA, 2012), pp. 2951–2959 [[arXiv:1206.2944](https://arxiv.org/abs/1206.2944)].
- [103] B. Shahriari, K. Swersky, Z. Wang, R. P. Adams, and N. de Freitas, Taking the human out of the loop: A review of bayesian optimization, *Proc. IEEE* **104**, 148 (2016).
- [104] J. Moćkus, On Bayesian methods for seeking the extremum, in *Optimization Techniques IFIP Technical Conference Novosibirsk, 1974*, Lecture Notes in Computer Science, edited by G. I. Marchuk (Springer, New York, 1975), pp. 400–404.
- [105] T. Head *et al.*, Scikit-optimize.
- [106] S. Ioffe and C. Szegedy, Batch normalization: Accelerating deep network training by reducing internal covariate shift, [arXiv:1502.03167](https://arxiv.org/abs/1502.03167).
- [107] C. Peng, X. Zhang, G. Yu, G. Luo, and J. Sun, Large kernel matters—improve semantic segmentation by global convolutional network, in *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Los Alamitos, CA, USA (2017), pp. 1743–1751 [[arXiv:1703.02719](https://arxiv.org/abs/1703.02719)].
- [108] S. Mallat, Understanding deep convolutional networks, *Phil. Trans. R. Soc. A* **374**, 20150203 (2016).
- [109] M. Michailidis, Stacknet (2020).
- [110] D. P. Kingma and M. Welling, Auto-encoding variational Bayes, [arXiv:1312.6114](https://arxiv.org/abs/1312.6114).
- [111] D. J. Rezende, S. Mohamed, and D. Wierstra, Stochastic backpropagation and approximate inference in deep generative models (2014), 1401.4082.
- [112] T. Salimans, D. P. Kingma, and M. Welling, Markov chain Monte Carlo and variational inference: Bridging the gap, [arXiv:1410.6460](https://arxiv.org/abs/1410.6460).
- [113] J.-Y. Zhu, T. Park, P. Isola, and A. A. Efros, Unpaired image-to-image translation using cycle-consistent adversarial networks, [arXiv:1703.10593](https://arxiv.org/abs/1703.10593).
- [114] L. Ardizzone, J. Kruse, S. Wirkert, D. Rahner, E. W. Pellegrini, R. S. Klessen, L. Maier-Hein, C. Rother, and U. Köthe, Analyzing inverse problems with invertible neural networks, [arXiv:1808.04730](https://arxiv.org/abs/1808.04730).
- [115] M. Gori, G. Monfardini, and F. Scarselli, A new model for learning in graph domains, in *Proceedings of the International Joint Conference on Neural Networks*, Montreal, QC, Canada, Vol. 2 (2005), pp. 729–734.
- [116] F. Scarselli, A. C. Tsoi, M. Gori, and M. Hagenbuchner, Graphical-based learning environments for pattern recognition, in *Structural, Syntactic, and Statistical Pattern Recognition*, Lecture Notes in Computer Science, edited by A. Fred, T. M. Caelli, R. P. W. Duin, A. C. Campilho, and D. de Ridder (Springer, New York, 2004), pp. 42–56.
- [117] M. M. Bronstein, J. Bruna, Y. LeCun, A. Szlam, and P. Vandergheynst, Geometric deep learning: Going beyond euclidean data, *IEEE Signal Process. Mag.* **34**, 18 (2017).
- [118] L. B. Anderson, F. Apruzzi, X. Gao, J. Gray, and S.-J. Lee, A new construction of Calabi-Yau manifolds: Generalized CICYs, *Nucl. Phys.* **B906**, 441 (2016).
- [119] C. Cortes and V. Vapnik, Support-vector networks, *Mach. Learn.* **20**, 273 (1995).
- [120] H. Drucker, C. J. Burges, L. Kaufman, A. J. Smola, and V. Vapnik, Support vector regression machines, in *Advances in Neural Information Processing Systems* (MIT Press, Cambridge, 1997), pp. 155–161.
- [121] J. R. Quinlan, Induction of decision trees, *Mach. Learn.* **1**, 81 (1986).
- [122] K. Wittkowski, Classification and regression trees, in *Metrika*, edited by L. Breiman, J. H. Friedman, R. A. Olshen, and C. J. Stone (1986), Vol. 33, p. 128.
- [123] Tin Kam Ho, Random decision forests, in *Proceedings of 3rd International Conference on Document Analysis and Recognition*, Los Alamitos, CA, USA (1995), Vol. 1, pp. 278–282.

- [124] J.H. Friedman, Greedy function approximation: A gradient boosting machine, *Ann. Stat.* **29**, 1189 (2001).
- [125] J.H. Friedman, Stochastic gradient boosting, *Computational Statistics and Data Analysis* **38**, 367 (2002).
- [126] T. Fawcett, Using rule sets to maximize roc performance, in *Proceedings 2001 IEEE International Conference on Data Mining, Los Alamitos, CA, USA* (2001), pp. 131–138.
- [127] X. Glorot, A. Bordes, and Y. Bengio, Deep sparse rectifier neural networks, in *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics, Fort Lauderdale, FL, USA* (2011), pp. 315–323.
- [128] J. Tompson, R. Goroshin, A. Jain, Y. LeCun, and C. Bregler, Efficient object localization using convolutional networks, in *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Los Alamitos, CA, USA* (2015).
- [129] D. E. Rumelhart, G. E. Hinton, and R. J. Williams, Learning representations by back-propagating errors, *Nature (London)* **323**, 533 (1986).
- [130] D. P. Kingma and J. Ba, Adam: A method for stochastic optimization, [arXiv:1412.6980](https://arxiv.org/abs/1412.6980).
- [131] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov, Dropout: A simple way to prevent neural networks from overfitting, *J. Mach. Learn. Res.* **15**, 1929 (2014).