



**HAL**  
open science

# Inception neural network for complete intersection Calabi-Yau 3-folds

Harold Erbin, Riccardo Finotello

► **To cite this version:**

Harold Erbin, Riccardo Finotello. Inception neural network for complete intersection Calabi-Yau 3-folds. *Machine Learning: Science and Technology*, 2021, 2 (2), pp.02LT03. 10.1088/2632-2153/abda61 . cea-04082047

**HAL Id: cea-04082047**

**<https://cea.hal.science/cea-04082047>**

Submitted on 26 Apr 2023

**HAL** is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

# Inception Neural Network for Complete Intersection Calabi-Yau 3-folds

**H. Erbin, R. Finotello**

Dipartimento di Fisica, Università di Torino and I.N.F.N. – sezione di Torino,  
via P. Giuria 1, I-10125 Torino, Italy

E-mail: [erbin@to.infn.it](mailto:erbin@to.infn.it), [riccardo.finotello@to.infn.it](mailto:riccardo.finotello@to.infn.it)

**Abstract.** We introduce a neural network inspired by Google’s Inception model to compute the Hodge number  $h^{1,1}$  of complete intersection Calabi-Yau (CICY) 3-folds. This architecture improves largely the accuracy of the predictions over existing results, giving already 97% of accuracy with just 30% of the data for training. Accuracy climbs to 99% when using 80% of the data for training. This proves that neural networks are a valuable resource to study geometric aspects in both pure mathematics and string theory.


*Keywords:* deep learning, algebraic geometry, string theory

## 1. Introduction

The last few years witnessed the uprising of *deep learning* as a very efficient method to elaborate, process and learn patterns in data [1]. While the underlying ideas behind neural networks are not recent [2, 3], larger databases, and computational capabilities together with new techniques led deep learning to pervade most fields of scientific research and industrial development.

Understanding geometrical structures is an emerging application of machine learning, referred to as *geometric deep learning* [4, 5] when neural networks are used. This is an important problem for different fields: for example in the industry (e.g. for 3d modelling of objects), computer science (e.g. for gradient optimisation [6]), pure mathematics, and theoretical physics. For this reason it is crucial to adapt existing techniques or to design new ones if needed.

In this paper, we focus on the computation of the Hodge number  $h^{1,1}$  for complete intersection Calabi-Yau (CICY) 3-folds [7]. This is a challenging mathematical problem *per se* because traditional methods from algebraic topology lead to complicated algorithms, without closed-form expressions in most cases. Machine learning techniques give the possibility to speed up computations and to obtain hints to better understand the mathematical structures. Moreover, Calabi-Yau manifolds, beyond being important mathematical objects, also have a distinguished role in string theory as they are needed to describe the compactified dimensions [8]. In particular the general properties of the 4-dimensional effective field theory are completely determined by the topology. Given the complexity of the space of string vacua, developing faster and efficient computational techniques is essential in the search of the Standard Model (or an extension compatible with experiments) within string theory at low energy. Finally, this type of objects is quite remote from typical data considered in machine



**Figure 1.** Accuracy reached by different models. The percentage in parenthesis indicates the ratio of training data. “He” refers to [12], “Bull et al.” to [13]. Each model except the Inception one keeps outliers in the training set (the effects is marginal in linear regression and SVM).

learning, which calls for an evaluation of existing techniques in this context and, if they are not sufficient, the development of new approaches.

The CICY 3-folds are appropriate for this task: since they have been completely classified [9, 10, 11], they provide a simple playground where it is possible to test different machine learning techniques. The goal of this paper is to continue the study started in [12, 13], which used machine learning techniques to compute  $h^{1,1}$  (see also [14, 15, 16] for other papers on CICY 3-folds). Related applications on the study of cohomology groups are [17, 18, 19]. For an introduction to machine learning and its applications to string theory, we refer to the excellent review [20].

Most breakthroughs in AI and industrial applications of deep learning usually followed the discovery of a new network model. This is particularly true in computer vision where convolutional, Inception and residual networks [3, 21, 22, 23, 24] have been major cornerstones. In this work, we introduce an alternative version of Google’s *Inception* network [21, 22, 23] (see [20] for a review) to predict  $h^{1,1}$  from the configuration matrix of CICY 3-folds. Using 30% of training data, we reach close to 97% accuracy on the predictions, improving by a large measure previous results [12, 13] with much less training data and parameters ( $\approx 234\,000$ ). Using 80% for the data for training we obtain 99% accuracy.

This must be compared with the following accuracies: 37% (regression, fully connected network,  $\approx 280\,000$  parameters, 63% training data) in [12], 75% (regression, fully connected network,  $\approx 1\,580\,000$  parameters, 70% training data) and 85% (classification, convolutional network, 70% training data) in [13] (Figure 1). More generally, we found that the Inception-like network performs much better than any other machine learning algorithm, even after feature engineering [25]: the best algorithm after neural networks is SVM with an RBF kernel, which reaches 68%

accuracy with 80 % of training data [13, 25]. This shows that neural networks are able to make accurate predictions for Hodge numbers, as long as the correct architecture is found. This opens the door to new applications to theoretical physics and mathematics which may lead to even further progress.

The code is written in Python and relies on the following packages: `scikit-learn` [26], `tensorflow` [27] (and its high level API, `keras` [28]) and the `scipy` ecosystem for visualisation and computations [29].

## 2. General Setup

The dataset [9, 10] is made of 7890 CICY 3-folds, described by their configuration matrices and their topological properties, including the Hodge numbers  $h^{1,1}$  and  $h^{2,1}$ . We focus on predicting the Hodge number  $h^{1,1} \in \mathbb{N}$ , which lies in the closed interval  $[0, 19]$  with 18 distinct values (with  $h^{1,1} = 17, 18$  not present), from the configuration matrix:

$$\left[ \begin{array}{c|ccc} \mathbb{P}^{n_1} & a_1^1 & \cdots & a_k^1 \\ \vdots & \vdots & \ddots & \vdots \\ \mathbb{P}^{n_m} & a_1^m & \cdots & a_k^m \end{array} \right], \quad a_\alpha^r \in \mathbb{N} \quad \longrightarrow \quad h^{1,1} \in \mathbb{N}. \quad (1)$$

The configuration matrix describes the CICY as the intersection of  $k$  hypersurfaces, characterised by a system of homogeneous polynomial equations, inside the ambient space  $\mathbb{P}^{n_1} \times \cdots \times \mathbb{P}^{n_m}$ , where  $m$  denotes the number of complex projective spaces. The coefficients  $a_\alpha^r$  of the matrix denote the power of the coordinates of each projective space entering each polynomial equation. This data is sufficient to characterise the topology. For more information on CICY we refer the reader to the literature [30, 31, 9, 10, 32, 33].

We consider the problem as a regression task and not as a classification task even if the outputs are integers. Indeed, the latter requires knowledge of all possible Hodge numbers which can appear and prevents any extrapolation, which is not desirable in the current context. Since regression algorithms output a real number, it is necessary to map predictions to integers before comparing with the real values.

The dataset is split into three subsets: one for training (used to learn the optimal model weights with gradient descent), one for validation (hyperparameter tuning and early stopping in neural networks), and one for testing.

In the following section, we discuss a few properties of the dataset which play an important role in the training of the neural network introduced in the next section.

### 2.1. Exploratory Data Analysis

The first step before writing the neural network is to better understand the data. Displaying the distribution of the Hodge numbers (Figure 2) and the whisker plot in the left side of Figure 3, one finds the presence of outliers at small and high Hodge numbers. Outliers can strongly impede the learning process of most algorithms and they must be handled with care. In this paper we obtained the best accuracy by simply removing them from the training data (but keeping them in the test set).

The outliers fall into two classes. First, the product spaces are recognisable by having vanishing Hodge numbers  $h^{1,1} = h^{2,1} = 0$  and a block-diagonal configuration




Figure 2. Distribution of  $h^{1,1}$  (log scale).





Figure 3. Whisker plot for the distribution of  $h^{1,1}$  before (*full*) and after (*clean*) removing the outliers. The coloured boxes highlight the interval between the first and third quartiles while the internal horizontal line represents the median value. The “whiskers” delimit the interquartile range, while isolated points mark the remaining outliers.

matrix.‡ Second we deal with manifolds with high Hodge numbers. We keep only manifolds such that  $h^{1,1} \in [1, 16]$  and  $h^{2,1} \in [15, 86]$  in the training data. Over the full dataset only 39 samples are excluded, or 0.49%. Hence training samples are taken as a subset of the distribution given in the right side of Figure 3. We expect systematical errors on test samples among outliers but they are too few to drastically impact the accuracy.

‡ Note that  $h^{1,1} = 0$  is not the the actual value of  $h^{1,1}$  but indicates merely that the CICY is factorizable into products of tori and K3 surfaces.



**Figure 4.** Scatter plots of  $h^{1,1}$  versus  $m$ . Manifolds on the diagonal are favourable.

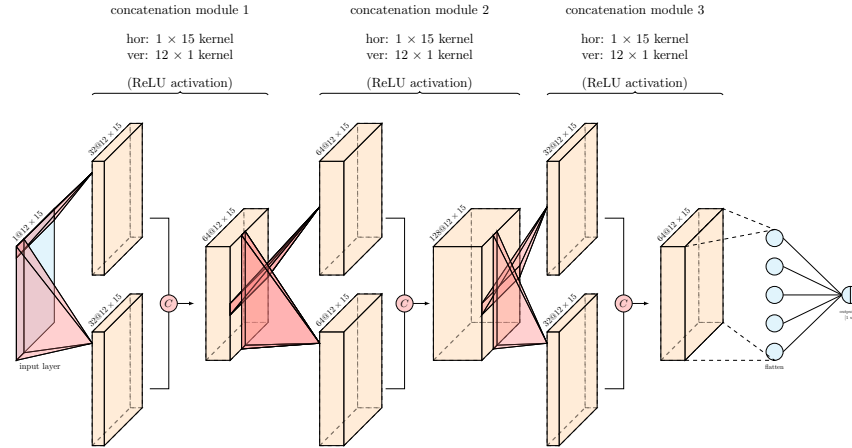
## 2.2. Baseline

It is important to design a simple baseline model to quantify the gain of using a neural network. Here, we consider a linear regression with  $\ell_1$  regularisation with parameter  $2 \times 10^{-4}$  and without intercept. Integers are obtained by flooring the predictions to the next lower integers. We obtain 47% to 51% accuracy using 20% to 80% of the data for training.

Moreover a simple analysis [25] shows that the number of projective spaces  $m$  (number of rows of the matrix) is an important feature. Performing a linear regression with  $\ell_1$  weight of 1.0, we obtain 63% of accuracy. This is related to a known mathematical result [32] stating that the so-called favourable matrices have  $h^{1,1} = m$  (in the dataset from [9, 10], there are 4874 favourable matrices, visible on the diagonal of the scatter plot in Figure 4). If it had not been known, the linear regression could have led to conjecture that this formula – indeed, conjecture generation is another distinguished use of machine learning techniques for theoretical physics [34, 19]. Note that SVM with RBF kernel is the best ML algorithm outside neural networks but improves only marginally over linear regression (Figure 1) [25].

## 3. Inception Neural Network

In this section, we introduce a new deep learning architecture capable of predicting accurately  $h^{1,1}$  from the configuration matrix of the CICY manifolds. Though different both in purpose and in definition, the model is inspired by Google’s *Inception Network* [21, 22, 23]. This deep neural network uses inception modules performing different concurrent convolutional operations to enhance, process, and rearrange its input (in Google’s case, images to be classified over 1000 classes in the *ImageNet* repository). This architecture encountered great success as it obtained results much



**Figure 5.** Schematic figure of the inception model used to predict  $h^{1,1}$  from configuration matrices.

better than any other machine learning algorithm until then. Modifications of the original model brought even higher accuracy and enhancement of computer vision capabilities. We refer the reader to [20] for a review of Inception networks.

We arrived at this network by going through neural network architectures used in computer vision. Indeed, the configuration matrix being a matrix of integers, it resembles an image with one channel. Since a sequential convolutional network does not reach a sufficient accuracy and needs a lot of training data, the Inception is the next natural step. Its structure has been guided by the form of the configuration matrix (see subsection 3.4 for more details).

Adapting this network to our problem, we obtain close to 100% accuracy already by training with only 30% of the data, which is much higher than existing results [12, 13]. A more general machine learning analysis of this problem will appear in [25].

### 3.1. Architecture

The architecture is schematically depicted in Figure 5: it is divided into three *inception* modules followed by an output layer with a single unit for the prediction of the Hodge number.

The first layer takes the configuration matrices as input, which are represented as tensors of shape  $(12, 15, 1)$  (matrices with a single channel). Next two parallel convolutions (shown in red in Figure 5) are performed: one over the rows ( $12 \times 1$  kernel, processing each projective space at a time) and one over the columns ( $1 \times 15$  kernel, processing each equation of the polynomial system at a time). The outputs of both layers are concatenated together over the channel dimension. These two steps form an inception module, which is repeated 3 times in total, with respectively 32, 64, and 32 filters. All convolutional layers and the final layer are followed by a ReLU activation function and each concatenation by a batch normalisation with momentum 0.99. A dropout layer with a rate 0.2 after the last inception module and before flattening the results, to connect it to the final output layer. Finally, all layers have  $\ell_1$  and  $\ell_2$  regularisation, respectively with weights  $10^{-4}$  and  $10^{-3}$ .

	layer	shape	parameters
<i>input</i>	Input	(12, 15, 1)	0
<i>module 1</i>	Conv2D@12x1	(12, 15, 32)	416
	Conv2D@1x15	(12, 15, 32)	512
	concatenate	(12, 15, 64)	0
	BatchNorm	(12, 15, 64)	256
<i>module 2</i>	Conv2D@12x1	(12, 15, 64)	49 216
	Conv2D@1x15	(12, 15, 64)	61 504
	concatenate	(12, 15, 128)	0
	BatchNorm	(12, 15, 128)	512
<i>module 3</i>	Conv2D@12x1	(12, 15, 32)	49 184
	Conv2D@1x15	(12, 15, 32)	61 472
	concatenate	(12, 15, 64)	0
	BatchNorm	(12, 15, 64)	256
<i>dropout</i>	Dropout	(12, 15, 64)	0
<i>fully connected</i>	Flatten	(11520,)	0
	Dense	(1,)	11 521
<i>total parameters</i>			234 849
<i>trainable parameters</i>			234 337
<i>non trainable parameters</i>			512

**Table 1.** Summary of the model with the number of parameters for each layer.

Table 1 summarises the network and the number of parameters in each layer. The network has  $\approx 234\,000$  parameters, which is less than previous proposals [12, 13]. This is achieved by using only convolutional layers with relatively small kernels.

Note that there are no pooling layers. Convolutions use `same` as padding value which allows us to keep the same size (12, 15) as the input. The output layer is followed by a ReLU activation function which forces the result to be positive, as it should be for Hodge numbers.


This architecture has two evident advantages over a fully connected (FC) network or even a more classical convolutional structure. First, the network concurrently learns different representations and automatically combines them in more complex representations. Second, the number of parameters is extremely restricted.

### 3.2. Training and validation strategy

We use a *holdout* validation strategy: the dataset is divided into three subsets for training (gradient descent to optimize the neural network’s weights), validation (early stopping, and hyperparameter tuning) and testing purposes (final assessment of our model). We retain respectively 80% of all samples for training, 10% for validation, and 10% for testing.

Before feeding the configuration matrix to the neural network, we first remove the outliers as discussed previously. We have tried to rescale the matrix by dividing by the highest entry (5), but this does not bring any significant improvement.





**Figure 6.** On the left we show the evolution of the loss evaluated on the training (blue curve) and evaluation (red curve) sets during training with 80 % training data (the colored area denotes the  $1\sigma$  region). On the right we present the distribution of the real and predicted values for the Inception network (80 % training data).

We did not use any data augmentation. Adding matrices with permutations of rows and columns seem to decrease the performance of the neural network: one possible explanation is that matrix components are ordered lexicographically [9]. Moreover, we did not generate more matrices using mathematical equivalences [9] since the final accuracy is high enough.

Hyperparameter tuning (number of inception modules and filters, dropout rate, etc.) has been performed by hand by evaluating several models on the validation set. After finding the appropriate architecture, described in the previous subsection, we have also evaluated the accuracy by training with 30 % and 50 % of the data (keeping always 10 % for the validation set, necessary for early stopping).

The neural network is trained using the *Adam* [35] optimiser with default parameters, initial learning rate  $10^{-3}$  and a batch size of 32. We use the mean squared error of the predictions as a loss function. The learning rate is reduced by a factor of 0.3 when the validation loss does not decrease during 75 epochs. We also use early stopping: the network is trained until the validation loss does not decrease for 200 epochs, restoring the weights associated with the lowest validation loss.

Predictions are obtained by averaging the results of 5 neural networks (bagging), which allows us to reduce the variance and obtain the standard deviation of the results. Since predictions are real numbers at this point, they are rounded to the closest integers before comparing them with the real value. The performance of the model is measured by the accuracy, which is the ratio of predictions matching exactly the real values.

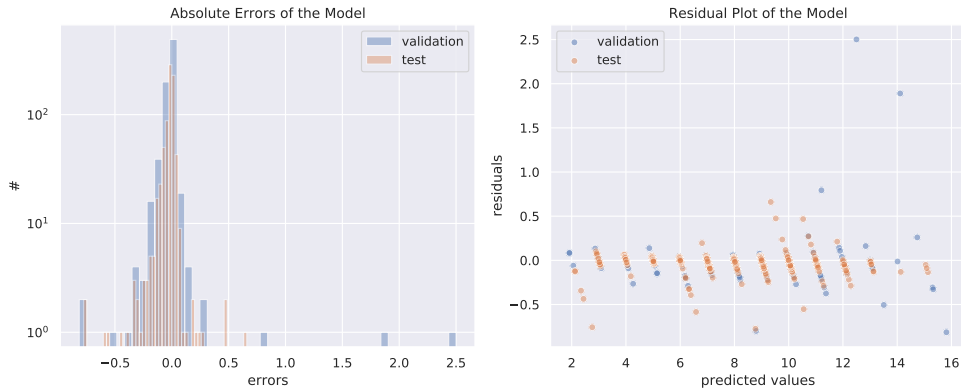
Finally, we will also provide learning curves for the neural network described in the previous section. For this, we split the dataset into training and validation subsets with different relative ratios and we compute the accuracy on both sets after training. In each case, we keep 10 % of the training data for early stopping. Exception made for this, the rest of the setup is the same.

### 3.3. Results

In the left plot in Figure 6, we show the evolution of the training and validation loss (mean squared error) during training. Curiously the mean absolute error is smaller for the validation set.

<i>training data</i>	Fully connected	Convolution	Inception
80 %	≈77 %	92.5 %	98.7 %
50 %	≈74 %	84.9 %	98.3 %
30 %	≈68 %	78.5 %	97.6 %


**Table 2.** Accuracy for the Inception neural network for different sizes of the training dataset, with standard deviations between 0.1% to 0.5%. Results obtained for other models are added for comparison: fully connected network [13] (read from Figure 1), convolutional network [25]. See also Figure 1.



**Figure 7.** Difference between the true values of  $h^{1,1}$  and their predictions seen as a univariate distribution (on the left) and as a function of the predicted value (on the right): the histogram shows the distribution and extension in values of the difference between true values and predictions, while the scatter plot exhibits the constant variance of the residual error (the network performs equally well over the entire range of  $h^{1,1}$ ).

The agreement between the predictions and real values is excellent on the test fold. The distributions are displayed on the right of Figure 6. The results at different ratios of training data are given in Table 2, where we also display the accuracy for other regression models: the fully connected network from [13] and an improved sequential convolutional network described in [25] (see also the introduction for more details). Even though the sequential model can already achieve very high accuracy, the Inception network performs even better with fewer parameters and much less training data. The learning curve is given in Figure 8: it does not show signs of overfitting and clearly demonstrates the quick convergence to almost 100% accuracy.

As presented in Figure 7, the network performs equally well over the entire range of  $h^{1,1}$  both in the validation and test sets: the variance of the difference between the observed values of the Hodge number and its predictions (i.e. the residuals) is constant as shown by the scatter plot. Moreover, the histogram of the residuals shows that the distribution is peaked around 0 and very few predictions lie far from the central value: the variance is in fact very small.



**Figure 8.** Learning curve of the Inception network. The colored area denotes the  $1\sigma$  region.

### 3.4. Ablation study

We can now study in detail the relative impact of each improvement introduced in our paper. The three points of comparison are 1) parallel vs sequential convolution layers, 2) using  $1d$  kernels  $12 \times 1$  and  $1 \times 15$  or  $2d$  kernels  $3 \times 3$  and  $5 \times 5$  (without changing the number of layers), 3) including or removing outliers from the training data. A comparison of the accuracy achieved by different models is displayed in Figure 9.


First, we want to measure the benefit of using parallel instead of sequential convolutions. In [25] we have built a convolutional network (*convnet* in Figure 9) made of 4 layers with 180, 100, 40 and 20 units, all with a  $5 \times 5$  kernel and  $\ell_1$  and  $\ell_2$  regularisation  $10^{-4}$  and  $10^{-3}$  ( $\approx 580\,000$  parameters). The accuracies of this network at a few training ratios are given in Table 2 and we refer the reader to [25] for more details. While this network performs better than earlier models (compare Figures 1 and 9), its accuracy is below the Inception model.

Second, we wish to uncover the effect of using  $1d$  kernels  $12 \times 1$  and  $1 \times 15$  instead of  $2d$  kernels. For this, we have trained a new version of the Inception model with the  $1d$  kernels replaced by concurrent  $3 \times 3$  and  $5 \times 5$  kernels (typical in computer vision tasks), leaving all other hyperparameters identical ( $\approx 290\,000$  parameters). From Figure 9, we find that this network performs even less well than the sequential convolutional network. One possible explanation is that the two  $1d$  convolutional windows process separately the information of each single projective spaces (columns) or polynomial equation (rows), scanning all of them one after the other. This could explain why it is necessary to have two  $1d$  kernels: one for the projective spaces, one for the equations.

Third, we have argued that removing outliers from the training and validation sets helps the network to learn better. The effect is not as important as the previous two points, but still noticeable (Figure 9).

Finally, we compared the difference between regression and classification. We have one-hot encoded the Hodge numbers, replaced the last layer of the network described in subsection 3.1 by a softmax, and used the cross-entropy loss for optimization. We find that classification is less efficient than regression (Figure 9). Adding two additional Inception modules brings the accuracy to 96%, still below the result from the regression network.

In conclusion, we see that convolutional layers working in parallel are responsible for a large part of the performance boost. That convolution is useful for CICY



**Figure 9.** Comparison of accuracy for different properties. The label “no outliers” means that outliers are excluded from the training and validation set. The labels  $1d$  and  $2d$  refer respectively to the kernels  $(12 \times 1, 1 \times 15)$  and  $(3 \times 3, 5 \times 5)$ .

may seem counter-intuitive [13] since the configuration matrices are not rotation nor translation invariant but only permutation invariant. However, we first note that convolution alone is only equivariant to global translation: it is not invariant to rotation nor translation (even locally), both of which require the addition of pooling layers (which we do not have) [1]. Moreover, convolution layers can be understood more generally as a way to spot different patterns in data by sharing weights, storing them in multiple channels, and recombining them in more complicated representations in subsequent layers. For instance, the original Inception models [21, 22, 23] include layers with  $1 \times 1$  kernel, which clearly do not exploit invariance properties. Another motivation for using convolution layers is parameter sharing: the same operations are applied at different locations of the input. Parameter sharing with the  $1d$  shape of the kernels implies that the same formulas are applied to each equation and each projective space, as can be expected for a geometric object.

#### 4. Conclusion

We have introduced a new type of neural network to compute the Hodge number  $h^{1,1}$  of complete intersection Calabi–Yau 3-folds. This neural network inspired by Google’s Inception model gets near-perfect accuracy, using fewer data and parameters than existing models. This improves largely the prediction power of the network and proves that deep learning is perfectly adapted for computations in algebraic topology. Hence, this network should definitely be explored at length to exploit its potential, which seems to be as promising for theoretical physics and mathematics as it has been in computer vision.

The next step consists in predicting also the Hodge number  $h^{2,1}$ . A preliminary analysis shows that the task is harder and the Inception network reaches only 50% accuracy – but it is higher than all other models, the best of which reach at most 35% (for SVM with Gaussian kernel and sequential convolutional network) [25]. One

solution is to use a better representation of the data. A first possibility is to use the favourable representation from [11], but this does not help [25]. Another more promising avenue is to use the graph representation introduced in [16]. It will also be interesting to extend our analysis to other topological objects useful for string theory. A last open question is to understand what the neural network learned and if it is possible to extract any interesting information from the weights. We leave these questions for the future.

### Acknowledgements

We are grateful to Sven Krippendorf and Fabian Rühle for discussions, and to Siavash Golkar 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 2015MP2CX4 “Non-perturbative Aspects of Gauge Theories and Strings”.

### References

- [1] I. Goodfellow, Y. Bengio, and A. Courville. *Deep learning*. MIT press, 2016.
- [2] F. Rosenblatt. The perceptron: A probabilistic model for information storage and organization in the brain. *Psychological Review*, pages 65–386, 1958.
- [3] Y. Lecun, L. Bottou, Y. Bengio, and P. Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.
- [4] M. M. Bronstein et al. Geometric Deep Learning: Going beyond Euclidean Data. *IEEE Signal Processing Magazine*, 34(4):18–42, July 2017.
- [5] Geometric deep learning. <http://geometricdeeplearning.com/>.
- [6] N. Lei et al. A geometric understanding of deep learning. *Engineering*, 2020.
- [7] P. Green and T. Hübsch. Calabi-yau Manifolds as Complete Intersections in Products of Complex Projective Spaces. *Commun. Math. Phys.*, 109:99, 1987.
- [8] L. E. Ibáñez and A. M. Uranga. *String theory and particle physics: An introduction to string phenomenology*. Cambridge University Press, 2012.
- [9] P. Candelas, A. M. Dale, C. A. Lütken, and R. Schimmrigk. Complete intersection calabi-yau manifolds. *Nuclear Physics B*, 298(3):493 – 525, 1989.
- [10] P. S. Green, T. Hübsch, and C. A. Lütken. All the hodge numbers for all calabi-yau complete intersections. *Classical and Quantum Gravity*, 6(2):105–124, feb 1989.
- [11] L. B. Anderson, X. Gao, J. Gray, and S.-J. Lee. Fibrations in CICY Threefolds. *JHEP*, 10:077, 2017.
- [12] Y.-H. He. Deep-Learning the Landscape. 6 2017.
- [13] K. Bull, Y.-H. He, V. Jejjala, and C. Mishra. Machine Learning CICY Threefolds. *Phys. Lett. B*, 785:65–72, 2018.
- [14] K. Bull, Y.-H. He, V. Jejjala, and C. Mishra. Getting CICY High. *Phys. Lett. B*, 795:700–706, 2019.
- [15] Y.-H. He and S.-J. Lee. Distinguishing elliptic fibrations with AI. *Phys. Lett. B*, 798:134889, 2019.
- [16] Sven Krippendorf and Marc Syvaeri. Detecting symmetries with neural networks. *Machine Learning: Science and Technology*, 2(1):015010, 2020.
- [17] F. Ruehle. Evolving neural networks with genetic algorithms to study the String Landscape. *JHEP*, 08:038, 2017.
- [18] D. Klaewer and L. Schlechter. Machine Learning Line Bundle Cohomologies of Hypersurfaces in Toric Varieties. *Phys. Lett. B*, 789:438–443, 2019.
- [19] C. R. Brodie, A. Constantin, R. Deen, and A. Lukas. Machine Learning Line Bundle Cohomology. *Fortsch. Phys.*, 68(1):1900087, 2020.
- [20] F. Ruehle. Data science applications to string theory. *Phys. Rept.*, 839:1–117, 2020.
- [21] C. Szegedy et al. Going Deeper with Convolutions. *arXiv e-prints*, page arXiv:1409.4842, September 2014.

- [22] C. Szegedy et al. Rethinking the inception architecture for computer vision. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2818–2826, 2016.
- [23] C. Szegedy, S. Ioffe, V. Vanhoucke, and A. A. Alemi. Inception-v4, inception-resnet and the impact of residual connections on learning. In *Thirty-first AAAI conference on artificial intelligence*, 2017.
- [24] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [25] Harold Erbin and Riccardo Finotello. Machine learning for complete intersection Calabi-Yau manifolds: a methodological study. 7 2020.
- [26] F. Pedregosa et al. Scikit-learn: Machine learning in python. *Journal of Machine Learning Research*, 12(85):2825–2830, 2011.
- [27] Martín Abadi et al. TensorFlow: Large-scale machine learning on heterogeneous systems, 2015. Software available from tensorflow.org.
- [28] F. Chollet. Keras, 2015.
- [29] P. Virtanen et al. SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods*, 17:261–272, 2020.
- [30] P. S. Green, T. Hübsch, and C. A. Lütken. All Hodge Numbers of All Complete Intersection Calabi-Yau Manifolds. *Class. Quant. Grav.*, 6:105–124, 1989.
- [31] P. S. Green and T. Hübsch. Polynomial Deformations and Cohomology of Calabi-yau Manifolds. *Commun. Math. Phys.*, 113:505, 1987.
- [32] L. B. Anderson, X. Gao, J. Gray, and S.-J. Lee. Fibrations in CICY Threefolds. *JHEP*, 10:077, 2017.
- [33] L. B. Anderson and M. Karkheiran. TASI Lectures on Geometric Tools for String Compactifications. *PoS*, TASI2017:013, 2018.
- [34] J. Carifio, J. Halverson, D. Krioukov, and B. D. Nelson. Machine Learning in the String Landscape. *JHEP*, 09:157, 2017.
- [35] D. P. Kingma and J. Ba. Adam: A Method for Stochastic Optimization. *arXiv e-prints*, page arXiv:1412.6980, December 2014.