

Estimating Calabi-Yau hypersurface and triangulation counts with equation learners

Ross Altman, Jonathan Carifio, James Halverson and Brent D. Nelson

*Department of Physics, Northeastern University,
Boston, MA 02115, U.S.A.*

E-mail: ross@rossealtman.com, j.carifio@neu.edu, j.halverson@neu.edu,
b.nelson@neu.edu

ABSTRACT: We provide the first estimate of the number of fine, regular, star triangulations of the four-dimensional reflexive polytopes, as classified by Kreuzer and Skarke (KS). This provides an upper bound on the number of Calabi-Yau threefold hypersurfaces in toric varieties. The estimate is performed with deep learning, specifically the novel equation learner (EQL) architecture. We demonstrate that EQL networks accurately predict numbers of triangulations far beyond the $h^{1,1}$ training region, allowing for reliable extrapolation. We estimate that number of triangulations in the KS dataset is $10^{10,505}$, dominated by the polytope with the highest $h^{1,1}$ value.

KEYWORDS: Superstring Vacua, Compactification and String Models

ARXIV EPRINT: [1811.06490](https://arxiv.org/abs/1811.06490)

Contents

1	Motivation	1
2	Approach	4
3	Facet classification	5
3.1	Distinguishing unique facets	5
3.2	Performing the classification	6
3.3	Classification results	7
4	Machine learning numbers of triangulations	9
4.1	Training data	10
4.2	Initial attempts	11
4.3	Neural networks	12
4.4	The EQL architecture	14
4.5	Model selection	15
4.6	Estimate for the total upper bound	19
4.7	The $h^{1,1} = 491$ polytope	20
5	Comparison to 3d polytopes	21
5.1	Classification of the 2d facets	21
5.2	Training data	22
5.3	Machine learning and extrapolation	22
6	Conclusion	24

1 Motivation

Triangulations and Calabi-Yau manifolds are objects of intrinsic mathematical interest in combinatorics and algebraic geometry, respectively. In the former case, elementary operations on triangulations such as flips may be used to determine when two triangulations are canonically related to one another, or to populate large ensembles of triangulations. Similarly, elementary topology changing operations on Calabi-Yau manifolds such as flop or conifold transitions may relate two such manifolds to one another by finite distance motion in metric moduli space; these may be utilized to generate large ensembles of geometries. In the case of Calabi-Yau hypersurfaces in toric varieties, Calabi-Yau data is encoded in the structure of a triangulation, and operations on the triangulation often induce topological transitions in the Calabi-Yau hypersurface. The transitions between these objects naturally generate a large network that encodes ensemble structure.

In string theory, Calabi-Yau manifolds are some of the simplest and best-studied backgrounds on which to compactify the extra dimensions of space while preserving supersymmetry in four dimensions [1]. There, flop [2] and conifold [3] transitions induce space-time topology change in a physically consistent manner and may be utilized to generate large ensembles of string compactifications; generalized fluxes [4, 5] also give rise to large ensembles [6]. The associated collection of four-dimensional effective potentials and metastable vacua form the so-called string landscape, which is central to understanding the implications of string theory for particle physics and cosmology.

A natural direction is to understand cosmological dynamics on the landscape and associated mechanisms for vacuum selection. Such dynamics, together with statistical properties of vacua [7], could then lead to concrete statistical predictions for physics in four dimensions. Global structures in the landscape, such as large networks, may play a role in the dynamics. For example, dynamical vacuum selection [8] on a network of string geometries [9] in a well-studied bubble cosmology [10] selects models with large numbers of gauge groups and axions, as well as strong coupling. However, concrete studies of the landscape are difficult due to its enormity [9, 11–15], computational complexity [16–20], and undecidability [17]. It is therefore natural to expect that, in addition to the formal progress that is clearly required, data science techniques such as supervised machine learning will be necessary to understand the landscape; see for initial works [8, 21–23] in this directions and [8, 24–29] for additional promising results.

One concrete goal is to understand the full ensemble of Calabi-Yau threefolds, network-like structures induced by transitions between them, and associated implications for cosmological dynamics and vacuum selection in compactifications of string theory. However, this is far out of reach currently for reasons of enormity and complexity.

Instead, in this paper we take a modest but necessary first step in this direction. We will estimate the number of Calabi-Yau threefold hypersurfaces in toric varieties, which are one of the most-studied ensembles in string theory. Many such manifolds are naturally associated to a triangulation of a 4d reflexive polytope; the latter were classified in a seminal work [30] by Kreuzer and Skarke. With 473,800,776 reflexive polytopes in four dimensions, the number of possibilities already borders on intractable. However, the number grows further since many of these toric varieties contain singularities that must first be resolved in order to obtain a smooth Calabi-Yau hypersurface. The desingularization process takes place via a series of blowups, of which there are typically many possibilities, resulting in significant growth in the number of compactification geometries. The goal in this work is to estimate this growth.

Let us comment further on the desingularization process, since we will be studying the manifolds resulting from this process. In order for the Calabi-Yau threefold hypersurface to be non-singular, its resolved ambient toric variety must have no worse than 1-dimensional singularities. In the language of toric geometry, this corresponds to blowing up all ≥ 2 -dimensional singularities, which in turn corresponds to a fine (i.e. maximal) subdivision of all 2-codimensional cones in the corresponding 4d reflexive polytope. By considering only integral points on the convex hull for this subdivision, we preserve the canonical divisor of the toric variety, and thereby the class of the Calabi-Yau hypersurface itself. Furthermore,

each configuration of simplexes in the maximally subdivided (i.e. fine, star) polytope constitutes a triangulation, and each gives rise to a distinct choice of desingularization of the toric variety. Finally, by considering only regular triangulations such that they can be obtained by projecting down from a higher dimension, we preserve projectivity of the variety.

The number of such fine, regular, star triangulations (FRST) grows exponentially as the complexity of the singular toric variety increases. While this number is partially offset by the necessity to group subsets of triangulation into topologically-equivalent classes, the number of distinct FRSTs gives a concrete upper bound. Our main result is this upper bound

$$n_{\text{FRST}} \simeq 10^{10,505.2 \pm 292.6}, \tag{1.1}$$

where n_{FRST} is the number of FRSTs arising in the Kreuzer-Skarke ensemble. We obtain this estimate with deep learning, specifically a novel neural network architecture known as an equation learner (EQL) [31], which we demonstrate is significantly better at extrapolating triangulation predictions to large $h^{1,1}$ than a standard feed-forward neural network. In particular, by demonstrating accurate extrapolation to high $h^{1,1}$ in the analogous problem for 3d reflexive polytopes, we lend credibility to the 4d prediction, which requires extrapolation significantly beyond the regime in which training data is available. Predicting n_{FRST} provides an estimated upper bound on the number of Calabi-Yau threefold hypersurfaces in toric varieties.

Given these accurate predictions of n_{FRST} , it is interesting to study the interpretability of the predictions made by the neural network. Sometimes referred to as intelligible artificial intelligence, interpretability is a major current goal of machine learning research, and it is one of the reasons for developing the EQL architecture; see [32] for the related idea of conjecture generation, by which interpretable numerical decisions may be turned into rigorous results. In the EQL context, the idea is to mimic what happens in natural sciences such as physics, where a physical phenomenon is often described in terms of an interpretable function that allows for understanding and generalization. Accordingly, by utilizing simpler functions than standard architectures, EQLs in principle increase the likelihood of interpretability.

In the case of n_{FRST} , we found that an EQL that utilizes quadratic functions makes accurate predictions; see sections 4 and 5 for quadratic functions associated to trained EQLs. By studying an associated heat map of coefficients, we demonstrate that some variables and cross-correlations are clearly of more importance than others, but the existence of a large number of warm spots suggest that many variables matter, which makes interpretability difficult. This could be an artifact of having used a quadratic function, which may be suboptimal, but it could also be the case that there is no simple interpretation of n_{FRST} predictions. That is, sufficiently complex phenomena in complex systems may not admit descriptions in terms of simple equations in human-understandable variables.

This paper is organized as follows: in section 2, we give an overview of our approach for estimating the number of FRSTs in the 4d Kreuzer-Skarke database. In section 3 we discuss the method by which we classified the 3d facets of the 4d reflexive polytopes, the results of this classification, and our success in obtaining FRTs of these facets. In section 4

we discuss our unsuccessful initial machine learning attempts, our input features, and our successful EQL model. In section 5, we perform a similar analysis for the 3d reflexive polytopes and show that a model of a similar architecture is able to extrapolate far outside its training range.

2 Approach

Batyrev has shown [33] that a hypersurface in a toric variety can be chosen to be Calabi-Yau if the object underlying the construction of the variety, a lattice polytope, obeys the condition of reflexivity.

A reflexive polytope Δ is defined as the convex hull of a set of points $\{v\} \in \mathbb{Z}^n$ whose dual polytope

$$\Delta^\circ = \{w \in \mathbb{Z}^n \mid w \cdot v \geq -1 \quad \forall v \in \Delta\} \tag{2.1}$$

is itself a lattice polytope. In four dimensions, there are 473,800,776 reflexive polytopes, as classified by Kreuzer and Skarke [30]. The ambient 4d space described by these polytopes is generally a singular toric variety. A proper Calabi-Yau manifold will therefore be a hypersurface in a suitably de-singularized ambient space. The process of de-singularization is equivalent to obtaining a fine, regular, star triangulation (FRST) of the 4d reflexive polytope.

Crucially, there are typically very many inequivalent FRSTs for a given polytope. While software packages such as PALP [34] and TOPCOM [35] can, in principle, produce all possible triangulations of a given polytope; in practice, such a request becomes increasingly problematic computationally as the overall size of the polytope grows. A good proxy for this computational load is the value of the topological quantity $h^{1,1}$ associated with the polytope. For example, in [36], all triangulations of all 4d reflexive polytopes with $h^{1,1} \leq 6$ were obtained. The nearly 652,000 unique triangulations required approximately 120,000 core-hours of processing time to obtain. The computational burden can be mitigated somewhat by exploiting the reflexivity property of these polytopes [37]. Even with this assistance, however, obtaining more than a single, canonical triangulation for a given polytope (in a reasonable computational time) becomes difficult for $h^{1,1} \gtrsim 10$ [38]. It would appear, therefore, that an enumeration of the unique triangulations in the KS 4d reflexive polytope dataset must remain unobtainable, barring dramatic advances in computational power or a vastly superior triangulation algorithm. It is this impasse which forms the motivation for the current work.

There is reason to believe that application of machine learning techniques may allow for an estimate of the total number of triangulations in the KS dataset to be achieved. A proof of principle already exists in the KS set of 3d reflexive polytopes, of which there are 4319. In [39], the number of FRSTs for many of the 3d reflexive polytopes was computed, and an estimation of the total number was obtained, using standard triangulation techniques. Soon thereafter, another estimate was obtained using supervised machine learning. Specifically, a decision-tree regression model was trained on known results, and used to estimate the number of FRSTs for the remaining cases [32]. The results were in good agreement with

the original estimate of [39]. The objective of the current work was to obtain such an estimate for the 4d reflexive polytopes via similar techniques.

As was done for the 3d case, the method employed here is to estimate the number of FRSTs of a given polytope as the product of the number of fine, regular triangulations (FRTs) of each of its facets. However, counting the number of vertices in the 4d reflexive dataset (which equals the number of facets by duality) shows that there are over 7.5 billion facets. In order to better get a handle on this set, as well as avoid unnecessary computational repetition, the first step was a classification of these 3d facets. The procedure for doing so is the subject of section 3.

With this classification in hand, we then explicitly computed the number of FRTs for as many of the facets as possible. Due to computational constraints, this consisted of only 1.03% of the total — primarily consisting of facets that first appear in dual polytopes at low values of $h^{1,1}$. The data associated with these explicitly-triangulated facets then became the training data for supervised machine learning. Ultimately, a neural network was employed to construct a model which predicts the natural logarithm of the number of FRTs of each facet. From this we arrive at an estimate of the total number of FRSTs of the 4d reflexive polytopes. The machine learning techniques, and FRST estimate, are the subject of section 4.

Any such machine learning estimate will obviously have limitations, and ours come from the two main issues that we faced. First, as noted above, we were only able to obtain known results for 1.03% of cases, meaning that it was necessary to estimate the value for nearly 99%. Second, as the facets that were able to be triangulated were necessarily ones with fewer triangulations, this problem involved extrapolation to output values beyond the range seen in the training set. However, as we will discuss below, we believe that our estimate is sound. The equation learning (EQL) neural network architecture that we ultimately employed was developed especially for extrapolation outside of a given training set. By withholding the subset of our known results that first appeared at $h^{1,1} \geq 12$ when training, we were able to see that our model extrapolated with stable results over a large number of cases as the number of triangulations increased. We comment further on these limitations, and the potential overcounting associated with our approach, in section 6.

3 Facet classification

In this section we review methods that easily determine when two facets are equivalent, and then apply those methods to the classification of facets in 4d reflexive polytopes.

3.1 Distinguishing unique facets

We must distinguish between three-dimensional (3d) facets of the four-dimensional (4d) reflexive polytopes, since individual facets may appear in multiple polytopes. In particular, we must determine under what conditions two facets should be considered to be equivalent, and then find a method to determine whether the conditions are satisfied.

For their classification of the 3d and 4d reflexive polyhedra [30], Kreuzer and Skarke defined a normal form, with the property that two reflexive polyhedra have the same normal

form if and only if they are related by a $\text{GL}(n, \mathbb{Z})$ transformation. While powerful, this normal form has the restriction that it can only be used on full-dimensional polytopes.

To circumvent this issue, and thereby be able to use the normal form, we employ the method of Grinis and Kasprzyk, as described in section 3.2 of [40]. The origin is interior to every reflexive polytope, and thus we know that none of the facets contain the origin. Thus, for each facet F , we constructed the associated subcone defined by $C_F = \text{conv}(F \cup \{0\})$. The subcone is a full-dimensional polytope, and so its normal form can be computed. Additionally, as the origin is the sole interior lattice point of a reflexive polytope, the origin is the only lattice point in $C_F \setminus F$. This also means that we need not worry about lattice translations when comparing subcones. As the origin is fixed under any $\text{GL}(n, \mathbb{Z})$ transformation, two subcones C_{F_1} and C_{F_2} have the same normal form if and only if their associated facets F_1 and F_2 are related by a $\text{GL}(n, \mathbb{Z})$ transformation.

As an example, consider the following two facets F_1 and F_2 , both of which appear as dual facets to the same $h^{1,1} = 2$ polytope (given by POLYID 21 in the ToricCY database [36]). Each of the facets are the convex hulls of four vertices, as below:

$$F_1 = \text{conv}(\{-1, 0, 0, 0\}, \{-1, 0, 0, 1\}, \{-1, 0, 1, 0\}, \{-1, 1, 0, 0\}) \quad (3.1)$$

$$F_2 = \text{conv}(\{-1, 0, 0, 1\}, \{-1, 0, 1, 0\}, \{1, 0, 0, 0\}, \{2, -1, -1, -1\}). \quad (3.2)$$

Adding the origin to each facet, we obtain the associated subcones

$$C_{F_1} = \text{conv}(\{0, 0, 0, 0\}, \{-1, 0, 0, 0\}, \{-1, 0, 0, 1\}, \{-1, 0, 1, 0\}, \{-1, 1, 0, 0\}) \quad (3.3)$$

$$C_{F_2} = \text{conv}(\{0, 0, 0, 0\}, \{-1, 0, 0, 1\}, \{-1, 0, 1, 0\}, \{1, 0, 0, 0\}, \{2, -1, -1, -1\}) \quad (3.4)$$

Computing the normal form for each subcone, we find that

$$\text{NF}(C_{F_1}) = \text{NF}(C_{F_2}) = \{\{0, 0, 0, 0\}, \{1, 0, 0, 0\}, \{0, 1, 0, 0\}, \{0, 0, 1, 0\}, \{0, 0, 0, 1\}\}. \quad (3.5)$$

We see that the two subcones have the same normal form, and thus the facets are equivalent. Removing the origin, we see that both facets are equivalent to the 3d polytope with vertices $\{\{1, 0, 0, 0\}, \{0, 1, 0, 0\}, \{0, 0, 1, 0\}, \{0, 0, 0, 1\}\}$.

As discussed in section 3.3 below, this particular facet, known as the standard 3-simplex, is the most common facet among the 4d reflexive polytopes. It accounts for 20.45% of all facets, while no other facet accounts for more than 8.57%.

3.2 Performing the classification

The primary practical challenge of the facet classification was the volume of the 4d polytope dataset itself. In order to have a consistent direction for our computation, we worked through the polytopes in order of increasing $h^{1,1}(X)$, where X is an associated Calabi-Yau hypersurface. As it is the dual polytope that is triangulated when constructing this Calabi-Yau, we identified the facets of the dual polytope in each case. Hence, for each $h^{1,1}$ value, we identified the facets that appeared in each dual polytope, keeping both a master list of unique facets as well as recording which facets appeared in each dual polytope (and with which multiplicity).




Figure 1. (Left) The logarithm of the number of new facets at each $h^{1,1}$ value. (Right) The logarithm of the number of reflexive polytopes at each $h^{1,1}$ value.

Computation was done on Northeastern University’s Discovery cluster using the SLURM workload manager. As identifying the dual facets is an independent computation for each polytope, we were able to use distributed computing to decrease our real-world running time by several orders of magnitude. Filtering and removal of duplicate facets was performed after identification had been completed for each $h^{1,1}$ value. The dual and normal form computations were done using a C++ lattice polytope implementation of our own creation which used the PALP source code for many of its underlying calculations. This offered significantly faster performance than the LatticePolytope class in Sage, allowing the classification to be finished on the timespan of a few weeks.

3.3 Classification results

Using the above method for the facets of every 4d reflexive polytope, we found that there are a total of 45,990,557 unique 3d facets, which is an order of magnitude less than the number of polytopes themselves. This is approximately 0.6% of the total number of facets, which is 7,471,985,487. We found that a relatively small number of facets accounted for the majority of the facets that appear across the polytopes, with the most common facet accounting for just over 20% of the total.

As we performed the classification procedure, we recorded the number of facets which appeared for the first time at each $h^{1,1}$ value; i.e., the same facet may appear in many different polytopes with different values of $h^{1,1}$, and we recorded the smallest such value. The distribution of these new facets per $h^{1,1}$ value is shown in the left panel of figure 1, which should be compared to the distribution of the number of polytopes of a given $h^{1,1}$ value, which is given in the right panel of figure 1. The shape of the two distributions are very similar, albeit decreased by around an order of magnitude. We note that the two distributions have different peaks: the number of reflexive polytopes peaks at $h^{1,1} = 27$, while the number of new facets peaks at $h^{1,1} = 33$.




Figure 2. (Left) The number of new facets at each $h^{1,1}$ value, as a fraction of the number of polytopes at that $h^{1,1}$. (Right) The total number of facets found through each $h^{1,1}$ value, as a fraction of the total number of polytopes up to that point.

Along with the cumulative distribution, we also show these distributions as fractions of the relevant numbers of polytopes in figure 2. In the left panel, we show the number of new facets that appear at a given $h^{1,1}$ value, as a fraction of the total number of polytopes at that $h^{1,1}$ value. The scatter that emerges for $h^{1,1} \gtrsim 250$ represents the relatively small number of polytopes at these rather large $h^{1,1}$ values. Coupled with the information contained in figure 1, it is clear that the overwhelming majority of facets are encountered well before $h^{1,1} \simeq 250$, though (as we will see), the number of triangulations are dominated by the outliers in the very large $h^{1,1}$ bins.

Finally, the right panel of figure 2 shows the number of new facets at a given $h^{1,1}$ value, normalized by the *cumulative* number of polytopes to that particular $h^{1,1}$ value. We see that after the peak in the right panel of figure 1, the ratio quickly saturates to a value of approximately 0.1. This is nothing more than the ratio of total unique facets found (47×10^6) to the number of 4d reflexive polytopes in the KS database (470×10^6).

In understanding the reliability of our estimation methods, it is important to ask how often the unique facets that we have enumerated actually appear in the 4d polytopes in the KS database. This is the subject of figure 3, in which we show a histogram of the number of appearances of a given facet.

One can see from the left panel of figure 3 that a small number of facets dominate the database, with the frequency counts decreasing linearly on a logarithmic scale (the right panel of figure 3). The 100 most common facets account for 74% of the total number. As mentioned in section 1, we were able to achieve all possible FRTs for a very small set of the total number of 3d facets, but these cases represent over 88% of facets by appearance. This is indicated by the blue shading in both panels of figure 3.




Figure 3. Histograms of the frequencies of facet appearances, with bins spaced linearly (*left panel*), and logarithmically (*right panel*). The orange bars represent the total number of facets, while the blue bars show the amount for which the number of FRTs is explicitly computed, and thus known.

The most common facet is the 3d polyhedron known as the standard 3-simplex (S3S), which has normal form vertices

$$\{\{1, 0, 0, 0\}, \{0, 1, 0, 0\}, \{0, 0, 1, 0\}, \{0, 0, 0, 1\}\}. \tag{3.6}$$

The S3S appears a total of 1,528,150,671 times in the database, accounting for 20.45% of all facets, and appears in 87.8% of the polytopes. It first appears as a dual facet at $h^{1,1} = 1$ and maintains a consistent presence through the database, as evidenced by the graphs in figure 4. However, the S3S, being itself a simplex with no interior points, has only one FRT, and thus has essentially no effect on the total number of polytope FRSTs.

4 Machine learning numbers of triangulations

In supervised machine learning (often called simply *supervised learning*), the output of the machine learning algorithm is a function, commonly called a *model*, which takes a specified set of inputs and produces a unique output value. This function contains numerical parameters, or *weights*, which are adjusted by the machine learning algorithm. The algorithm is trained on a set of input \rightarrow output pairs, and attempts to minimize a predetermined loss function by adjusting the weights of the function.

In this section we will estimate the number of FRSTs of 4d reflexive polytopes by using supervised learning to predict the number of FRTs of the facets. For this application, the input is a set of features which describe a facet, and the desired output is the number of FRTs for that facet. We describe in this section the features which we used as input and the structure of the model that we obtained via supervised learning.




Figure 4. (Left) The percentage of dual polytopes that contain S3S at each $h^{1,1}$ value. (Right) The same graph, truncated at $h^{1,1} \leq 120$. We note that the erratic portion of the figure 4 for $h^{1,1} \gtrsim 120$ accounts for fewer than 1 million polytopes. The truncated graph (figure 4b) shows more clearly the prevalence of S3S.

$h^{1,1}$	Facets	Triangulated	% Triangulated
1 – 11	142,257	142,257	100%
12	92,178	92,162	99.983%
13	132,153	108,494	82.097%
14	180,034	124,700	69.625%
15	236,476	3,907	1.652%
> 15	45,207,459	1,360	0.003%
Total	45,990,557	472,896	1.028%

Table 1. Dual facet FRT numbers obtained, binned by the first $h^{1,1}$ value at which they appear.

4.1 Training data

The first step in our machine learning process was to generate data on which to train a model. In this case, this meant it was necessary to triangulate as many of the 3d facets as possible, in order to obtain the best possible training set.

Due to computational restrictions, we were only able to find the actual number of FRTs for 472,880 of the facets, which represents 1.03% of the total. These consist almost entirely of facets which first appear in the dual polytope at relatively low $h^{1,1}$ values. Table 1 shows our progress. As one can see, $h^{1,1} = 14$ was the last value at which we were able to obtain an appreciable fraction of the real FRT values. While the triangulated set may constitute only a small fraction of the unique facets, they account for over 88% of all facet appearances.

To obtain the FRT values, we first computed the 2-skeleton of each facet using our C++ code. The n -skeleton of a lattice polytope consists of all points which are not interior to any $(n+1)$ -dimensional face. We then passed the 2-skeleton to the TOPCOM executable `points2nfinetriangs` in order to obtain the FRT number. As with the classification process, we used Northeastern’s Discovery cluster to run simultaneous calculations for several hundred facets at a time.

When training with this data, and throughout our process of finding a good model, we organized the facets by the first $h^{1,1}$ at which they appeared as a dual facet, which we will refer to from here on as the $h^{1,1}$ value of the facet. Our reason for this was that there is a rough correlation between this $h^{1,1}$ value and the number of triangulations. This is reflected in table 1: as $h^{1,1}$ increased, fewer cases were able to complete as the facets became more complex. For this reason, the training set for each model only contained facets with $h^{1,1}$ up to some maximum value (typically 11). This allowed us to see how well the model would extrapolate to FRT numbers higher than it had seen, something that would prove to be an issue throughout our attempts.

4.2 Initial attempts

Our first attempt at obtaining a model was to follow the same pattern used for the 3d polytopes in section 3 of [32]. As was done in the 3d case, for each polytope we constructed the 4-tuple (n_p, n_i, n_b, n_v) , consisting of the number of points, interior points, boundary points, and vertices, respectively. Our metric for determining the model’s success was the mean absolute percent error (MAPE), which is defined as

$$\text{MAPE} = \frac{100}{n} \times \sum_{i=1}^n \left| \frac{A_i - P_i}{A_i} \right|, \tag{4.1}$$

where n is the number of data points, and P_i and A_i are the predicted and actual values for the output, which here is $\ln(N_{\text{FRT}})$ for the i^{th} facet.

We trained models using the same four algorithms as discussed in [32]: Linear Discriminant Analysis (LDA), k-Nearest Neighbors (KNNR), Classification and Regression Trees (CART), and Naive Bayes (NB). Of these, the CART algorithm gave impressive performance on the training set, as well as on the test data for the $h^{1,1}$ values that it had trained on. However, its ability to extrapolate to higher $h^{1,1}$ values was poor. As the majority of the facets lie at higher $h^{1,1}$ values, the inability to extrapolate well was unacceptable.

As an example, we show here the results from one of these attempts. The model is an `ExtraTreesRegressor` model from the `scikit-learn` Python package, using 35 estimators. We trained the model on a data set consisting of 60% of the known values from the range $5 \leq h^{1,1} \leq 10$. On the training set, the model achieved a MAPE of 5.723. It performed similarly on the test set consisting of the other 40% of these values, with a MAPE of 5.821. If we could be confident of a similar accuracy for the rest of the facets, this could be an acceptable model. However, when we used the model to predict on facets from higher $h^{1,1}$ values whose FRT numbers are known, we obtained the results shown in table 2.

As one can see, the MAPE value consistently increases with the $h^{1,1}$ value, indicating that the model is less accurate farther away from its training region. Further, the rightmost

$h^{1,1}$	MAPE	Actual mean	Predicted mean
11	6.566	9.582	9.189
12	9.065	10.882	9.903
13	11.566	11.755	10.067
14	17.403	12.638	10.179

Table 2. Prediction results for $\ln(N_{\text{FRT}})$, using the ExtraTreesRegressor model, for $h^{1,1}$ values outside of its training region.

two columns illustrate the problem with this model: it is under-predicting on facets with a higher number of triangulations. The model never predicts a $\ln(N_{\text{FRT}})$ value greater than 12.467, which is similar to 12.595, the highest value in its training set.

This example is emblematic of the problems we faced with traditional machine learning. Models either fit the training data poorly, or had trouble extrapolating outside of their training regions. The primary difference between the 3d case (where this simple approach worked), and the 4d case, is the inability to generate a representative training set. In the 3d case, the number of FRTs for all facets that appear up through $h^{1,1} = 22$ were triangulated, in a set where the maximum value of $h^{1,1}$ is 35. Conversely, in the 4d case, the number of FRTs was only obtained for all facets up through $h^{1,1} = 11$, and for a majority of cases up to $h^{1,1} = 14$, in a data set where the maximum $h^{1,1}$ value is 491.

4.3 Neural networks

Given the poor performance of traditional supervised machine learning as described above, we shifted our focus to artificial neural networks. In particular, a feed-forward network seemed the most suited to our purposes.

We recall briefly here the definitions of a neuron and a neural network. A *neuron* is a function $f(\sum_i^k w_i \cdot x_i + b)$ whose argument x_i is the input. The value of f is called the *output*. The parameters w_i are called *weights*, b is called the *bias* (or *offset*), and \cdot represents the appropriate tensor contraction. The function f itself is called the *activation function* and the choice of f is one of the characteristics that define the neuron.

A *neural network* is a (finite) directed graph, each node of which is a neuron. For each arrow in the graph, the output of the neuron at the tail is used as the input for the neuron at the head. Conceptually, a neural network is organized into *layers*, such that there are no connections between nodes in the same layer. The set of nodes whose arguments explicitly involve the original input data is known as the *input layer*, while the set whose output involves the actual output data is called the *output layer*. All other layers are referred to as *hidden layers*.

In the general case, a neural network may contain cycles, allowing for complicated connections between layers. In a *feed-forward* neural network, there are no cycles, and so information moves only in one direction. Cycles in a neural network allow for the current output to be influenced by the previous output. As the number of triangulations of a given facet is independent of the number for other facets, we can restrict ourselves to the feed-forward case.

We initially tried applying a feed-forward network to our 4-tuples of data, but despite trying a variety of architectures faced similar issues as in the previous subsection, with models consistently underpredicting outside of the training regime. To aid our networks in their prediction attempts, additional input data was generated. In addition to the numbers of points, interior points, boundary points, and vertices, the following quantities were computed for each facet to be used as inputs to the neural network:

- The number of points in the 1- and 2-skeletons
- The first $h^{1,1}$ value at which the facet appears in a dual polytope
- The number of faces
- The number of edges
- The number of flips of a seed triangulation of the 2-skeleton. Two triangulations differ by a flip if one can be obtained from the other by removing one edge and inserting another.
- Several quantities obtained from a single fine, regular triangulation of the facet:
 - The total numbers of 1-, 2-, and 3-simplices in the triangulation
 - The numbers of 1- and 2-simplices in the triangulation, without accounting for redundancy between higher-dimensional simplices
 - The numbers of 1- and 2-simplices shared between N 2- and 3-simplices, respectively, for N up to 5

As an example we consider the facet F with normal form is given by

$$\begin{aligned} & \{ \{0, 0, 0, 0\}, \{1, 0, 0, 0\}, \{1, 2, 0, 0\}, \{0, 0, 1, 0\}, \{0, 1, 1, 0\}, \\ & \{0, 0, 0, 1\}, \{3, 3, -1, -1\}, \{3, 2, -2, 0\}, \{0, 1, 0, 1\} \}. \end{aligned} \tag{4.2}$$

This facet has 11 lattice points, of which 10 lie on the boundary (leaving 1 in the interior). Of these 10 boundary points, 8 are vertices. The 1- and 2-skeletons of this facet contain 9 and 10 points, respectively. It first appears in a dual polytope with $h^{1,1} = 8$, and has 10 faces and 16 edges. Its flip graph has 13 nodes. Obtaining one FRT for this facet from TOPCOM, we find that it contains 72 1-simplices, 48 2-simplices, and 12 3-simplices. Of the 1-simplices, there were 29 unique ones, while there were 32 unique 2-simplices. There were 13 1-simplices shared between two 2-simplices, 6 shared between three, 3 shared between 4, and 1 shared between 5. There were 16 2-simplices shared between two 3-simplices, with none shared between more than two. Hence the input vector describing this facet was

$$(1, 10, 8, 9, 10, 8, 10, 16, 13, 72, 48, 12, 29, 32, 13, 6, 3, 1, 16, 0, 0, 0). \tag{4.3}$$

Adding this additional data to our original inputs improved results, but not to the point of satisfaction. As an example of the struggles faced by a traditional neural network, we show here the results from one such model. This neural network has two hidden layers,

$h^{1,1}$	MAPE	Mean value	Predicted mean
12	5.904	10.882	10.324
13	6.550	11.755	10.753
14	10.915	12.638	11.094

Table 3. Prediction results for $\ln(N_{\text{FRT}})$, using the traditional neural network, for $h^{1,1}$ values outside of its training region.




Figure 5. Histograms of the percent error of the feed-forward neural network’s predictions in the extrapolation region. Compare to extrapolations in the EQL network in figure 7.

each with 30 nodes. The first layer has a sigmoid activation function, while the second has a tanh activation function. The final layer is a rectified linear unit (ReLU), meaning that its activation function f_{act} is equal to the positive part of its argument:

$$f_{\text{act}}(x) = \max(0, x). \tag{4.4}$$

A ReLU unit was chosen for the final layer as we want $\ln(N_{\text{FRT}})$ to be positive.

The model was trained on a data set consisting of an equal number of randomly chosen points from $6 \leq h^{1,1} \leq 11$ to avoid biasing the model towards higher $h^{1,1}$ values. The model performed well on the test set of values from the same $h^{1,1}$, with a MAPE of 6.304. However, when we evaluated the model’s progress on higher $h^{1,1}$ values, our results were as shown in table 3. This model performs much better than the ExtraTreesRegressor outside of the training region, with a MAPE of 10.915 at $h^{1,1} = 14$. However, the MAPE values are increasing with $h^{1,1}$, and the mean predicted values are falling behind the true mean values, leading one to believe that the model is underpredicting at higher $h^{1,1}$. This belief is confirmed by examining the histograms of the percent error values in the extrapolation range, shown in figure 5. As $h^{1,1}$ increases, the percent error distribution skews to the left, indicating that the model is not keeping up with the increasing values. Hence this model is not suitable for extrapolation to the higher $h^{1,1}$ values which make up the bulk of our data set.

4.4 The EQL architecture

In this section we will instead utilize an equation learner (EQL) architecture, which was first introduced in [31], and will find that it has significantly better extrapolation ability.




Figure 6. A representation of a simple EQL layer with $n_m = 4$ and $n_o = 3$ sandwiched between two fully-connected layers. The first two elements of the intermediate representation are each acted on by activation functions f_i , while the remaining two elements are multiplied together.

Each layer in an EQL network consists of two stages: a standard linear stage followed by a non-linear stage. The non-linear stage essentially replaces the activation function in a standard neural network layer, but differs significantly in that it changes the shape of the linear stage’s output tensor. Like any feed-forward neural network layer, an EQL layer accepts some number of input values n_i , and outputs some number n_o of output values. The linear stage maps the n_i -dimensional input x to an intermediate n_m -dimensional representation z via an affine transformation. That is,

$$z = Wx + b \tag{4.5}$$

for some weight matrix $W \in \mathbb{R}^{n_m \times n_i}$ and some bias vector $b \in \mathbb{R}^{n_m}$.

The second stage takes this intermediate representation z to the final n_o -dimensional output y via a non-linear transformation. For this transformation, the elements of z are divided into two parts which are acted upon differently. To u of the elements, an activation function is applied. These nodes are called unary units. In principle, a different activation function can be applied to each unary unit. The other v elements are pairwise multiplied, giving $\frac{v}{2}$ output values. These nodes are called binary units. This pairwise multiplication step is the key aspect of the layer, as it allows for nonlinear interactions between the nodes of the network. More generally, nonlinear interactions between the nodes other than pairwise multiplication could also be utilized. A diagrammatic representation of an EQL layer is shown in figure 6.

4.5 Model selection

Our final neural network model is simple, with only one hidden EQL layer between the input and output layers. The input layer consists of 22 nodes as previously described. The linear stage of the EQL layer contains 45 nodes. Of these, 15 are unary units, with the other 30 being the binary units. The output from this layer thus consists of $15 + \frac{30}{2} = 30$

			Extrapolation MAPE		
$h_{\min}^{1,1}$	$h_{\max}^{1,1}$	Test MAPE	$h^{1,1} = 12$	$h^{1,1} = 13$	$h^{1,1} = 14$
6	10	7.297	6.647	6.699	6.598
7	10	6.001	7.512	7.626	7.469
8	10	7.184	5.048	5.172	5.834
6	11	5.643	4.393	4.490	4.416
7	11	6.967	7.512	7.626	7.469
8	11	5.551	4.444	4.463	4.934

Table 4. Results of training our model on various $h^{1,1}$ ranges. The model with $h_{\min}^{1,1} = 6$, $h_{\max}^{1,1} = 11$ performs well on the test set and the best on extrapolation to higher $h^{1,1}$ values.

nodes. The unary activation which was found to be the most successful was to square each of the unary units. The output stage consisted of one node, whose value represents the natural logarithm of the number of FRTs. This final output had a ReLU activation function to ensure that the output was nonnegative.

The Adam optimizer was used, with hyperparameters $\beta_1 = 0.9$, $\beta_2 = 0.99$, and $\epsilon = 1 \times 10^{-8}$. At each layer, an L1 regularizer with $\lambda = 0.001$ was used on all weights in order to select for the most important features. In addition, a dropout rate of 0.1 was used for the EQL hidden layer to help select the most optimal neuron configuration.

To select a model, we trained across multiple ranges $h_{\min}^{1,1} \leq h^{1,1} \leq h_{\max}^{1,1}$ and examined the results. As the facets at $h^{1,1} \leq 5$ have few triangulations and constitute relatively few data points, we chose $h_{\min}^{1,1} \geq 6$ in each case. In order to have an adequate range on which to test our model’s extrapolation, we also chose $h_{\max}^{1,1} \leq 11$.

The training data sets consisted of an equal number of randomly chosen data points from each $h^{1,1}$ value. This was done rather than taking a percentage of the data to avoid biasing the model towards fitting to the higher $h^{1,1}$ values; there are, for example, over 37 times as many facets at $h^{1,1} = 11$ as at $h^{1,1} = 6$. We then tested each model on the full set of values across its training range to see how it performed. Our metric of choice was the mean absolute percentage error (MAPE), as this gives a size-independent measurement of how well the model is performing. The results of our training are shown in table 4.

As the table shows, the model trained on the largest range of $h^{1,1}$ values, from 6 to 11, performed the best on both the test set and when extrapolating to higher values. This is the model that we used to generate predictions for the rest of the data set. We note that the extrapolation range contains a large number of points: there are 92162, 108494, and 124700 triangulated facets at $h^{1,1} = 12, 13$, and 14, respectively. We also note that, regardless of the training range, the models have stable MAPE values in the extrapolation region.

One might worry that, despite the promising MAPE values, this model is persistently underpredicting like the previous model. However, examining the mean predicted values (table 5) and the distributions of percent errors (figure 7), we can see that this is not the case. The mean predicted values stay close to the true means, and the percent errors stay centered around zero as $h^{1,1}$ increases.

$h^{1,1}$	Mean value	Predicted mean
12	10.733	10.722
13	11.755	11.591
14	12.638	12.492

Table 5. The true mean values and the mean predicted by our model in the extrapolation region.




Figure 7. Histograms of the percent error of our chosen model’s predictions in the extrapolation region. Compare to extrapolations in the EQL network in figure 5.

On the other hand, the extrapolation here is only to three $h^{1,1}$ values that are not in the training set, which provides some cause for concern since we will be interested in the polytope with maximal $h^{1,1}$, which has $h^{1,1} = 491$. For this reason we will perform an analogous analysis of $3d$ polytopes in section 5, and will demonstrate that in that case the EQL accurately predicts the number of FRTs for $h^{1,1}$ values significantly beyond that in the training set; e.g., training up to $h^{1,1} = 11$, we will accurately predict numbers of FRTs for all available facets, which go up to $h^{1,1} = 25$. This suggests that the EQL may also make accurate predictions in the $4d$ polytope case at values of $h^{1,1}$ significantly beyond the training set.

For completeness, we present in figure 8 the formula learned by our EQL model. The variables x_i are in the same order as in the example in section 4.3. This formula does not seem to lend itself to any intuitive interpretation. Its interpretability is also hindered by the fact that its variables will generally be of different scales. To attempt to adjust for this, we normalized each variable so that it has an expectation value of 1. This was done using the mean value of each variable in the training data. We then calculated $\log_{10}(|c|)$ for each coefficient c in the normalized expression to determine their relative importance on the value. The results of this are shown in figure 9. We can see that there are two brightest squares, which correspond to the interactions of the number of unique 1-simplices in the seed triangulation with both the number of 3-simplices, and the number of 1-simplices without accounting for redundancy. However, while these terms have the highest values, there are many others at a similar order of magnitude. Thus, while we can say that some terms are more important than others, the complicated form of the function, along with a lack of knowledge of how these variables will change at larger $h^{1,1}$, makes any stronger statements difficult.

$$\begin{aligned}
\ln(N_{\text{FRT}}) = & -0.425x_0^2 + 0.29447x_0x_1 - 0.2304x_1^2 + 0.02462x_0x_2 - 0.17529x_1x_2 - 0.3368x_2^2 + 0.72012x_0x_3 + 0.0707x_1x_3 + \\
& 0.00583x_2x_3 - 0.40825x_3^2 - 0.01146x_0x_4 - 0.00008x_1x_4 - 0.0789x_2x_4 + 0.00599x_3x_4 - 0.02246x_4^2 + 0.35742x_0x_5 + \\
& 0.00696x_1x_5 + 0.39255x_2x_5 - 0.35135x_3x_5 + 0.09x_4x_5 - 0.2482x_5^2 - 0.19063x_0x_6 - 0.02357x_1x_6 - 0.08904x_2x_6 + \\
& 0.20651x_3x_6 - 0.04321x_4x_6 + 0.21098x_5x_6 - 0.0609x_6^2 + 0.20861x_0x_7 + 0.05763x_1x_7 + 0.5381x_2x_7 - 0.19461x_3x_7 + \\
& 0.00197x_4x_7 - 0.41043x_5x_7 + 0.12497x_6x_7 - 0.15195x_7^2 + 0.02359x_0x_8 + 0.0095x_1x_8 + 0.05497x_2x_8 - 0.03016x_3x_8 + \\
& 0.00692x_4x_8 - 0.10282x_5x_8 + 0.06078x_6x_8 + 0.00153x_7x_8 + 0.00301x_8^2 - 0.29528x_0x_9 - 0.00072x_1x_9 - 0.20046x_2x_9 + \\
& 0.2274x_3x_9 - 0.01136x_4x_9 + 0.25023x_5x_9 - 0.13467x_6x_9 + 0.23766x_7x_9 + 0.08246x_8x_9 - 0.11071x_9^2 + 0.02683x_0x_{10} - \\
& 0.00254x_1x_{10} + 0.02553x_2x_{10} + 0.01278x_3x_{10} + 0.00554x_4x_{10} - 0.04584x_5x_{10} + 0.03714x_6x_{10} + 0.01352x_7x_{10} + \\
& 0.00105x_8x_{10} + 0.20665x_9x_{10} - 0.00099x_{10}^2 - 0.34428x_0x_{11} + 0.05335x_1x_{11} - 0.00715x_2x_{11} + 0.33024x_3x_{11} - \\
& 0.01736x_4x_{11} + 0.16713x_5x_{11} - 0.08848x_6x_{11} + 0.08456x_7x_{11} + 0.03144x_8x_{11} - 0.11092x_9x_{11} - 0.00151x_{10}x_{11} - \\
& 0.07556x_{11}^2 + 0.02349x_0x_{12} + 0.0031x_1x_{12} - 0.01155x_2x_{12} - 0.03125x_3x_{12} - 0.00223x_4x_{12} + 0.01197x_5x_{12} - \\
& 0.00373x_6x_{12} - 0.03727x_7x_{12} + 0.02864x_8x_{12} - 0.07238x_9x_{12} + 0.01227x_{10}x_{12} + 0.0125x_{11}x_{12} - 0.00375x_{12}^2 - \\
& 0.00624x_0x_{13} + 0.00116x_1x_{13} + 0.02883x_2x_{13} - 0.00045x_3x_{13} - 0.00025x_4x_{13} - 0.02332x_5x_{13} + 0.00198x_6x_{13} - \\
& 0.0065x_7x_{13} - 0.07052x_8x_{13} + 0.04967x_9x_{13} - 0.04768x_{10}x_{13} - 0.00442x_{11}x_{13} + 0.01206x_{12}x_{13} + 0.00333x_{13}^2 - \\
& 0.0373x_0x_{14} - 0.01724x_1x_{14} + 0.00677x_2x_{14} + 0.08978x_3x_{14} + 0.02568x_4x_{14} + 0.03319x_5x_{14} - 0.01481x_6x_{14} + \\
& 0.00051x_7x_{14} - 0.00392x_8x_{14} - 0.06852x_9x_{14} - 0.05699x_{10}x_{14} - 0.0266x_{11}x_{14} + 0.02735x_{12}x_{14} - 0.00365x_{13}x_{14} + \\
& 0.00757x_{14}^2 - 0.07426x_0x_{15} - 0.07933x_1x_{15} - 0.00701x_2x_{15} + 0.11551x_3x_{15} + 0.01621x_4x_{15} + 0.0459x_5x_{15} - \\
& 0.03323x_6x_{15} + 0.01862x_7x_{15} + 0.00938x_8x_{15} - 0.03075x_9x_{15} - 0.00161x_{10}x_{15} - 0.03737x_{11}x_{15} + 0.00556x_{12}x_{15} - \\
& 0.00167x_{13}x_{15} - 0.01574x_{14}x_{15} - 0.01246x_{15}^2 - 0.07638x_0x_{16} + 0.02935x_1x_{16} + 0.08977x_2x_{16} + 0.01429x_3x_{16} + \\
& 0.02235x_4x_{16} - 0.06464x_5x_{16} + 0.03652x_6x_{16} + 0.03147x_7x_{16} + 0.00244x_8x_{16} + 0.00924x_9x_{16} - 0.01379x_{10}x_{16} - \\
& 0.00638x_{11}x_{16} + 0.02373x_{12}x_{16} - 0.03973x_{13}x_{16} - 0.00033x_{14}x_{16} + 0.00165x_{15}x_{16} - 0.0025x_{16}^2 + 0.16209x_0x_{17} + \\
& 0.00283x_1x_{17} + 0.00744x_2x_{17} - 0.18455x_3x_{17} + 0.00375x_4x_{17} - 0.08754x_5x_{17} + 0.05244x_6x_{17} - 0.04801x_7x_{17} - \\
& 0.02091x_8x_{17} + 0.10988x_9x_{17} - 0.00019x_{10}x_{17} + 0.05868x_{11}x_{17} - 0.00919x_{12}x_{17} + 0.00468x_{13}x_{17} - 0.00073x_{14}x_{17} + \\
& 0.02233x_{15}x_{17} - 0.01548x_{16}x_{17} - 0.02151x_{17}^2 + 0.21655x_0x_{18} + 0.00257x_1x_{18} - 0.00186x_2x_{18} - 0.16816x_3x_{18} - \\
& 0.0052x_4x_{18} - 0.06696x_5x_{18} + 0.05542x_6x_{18} - 0.07719x_7x_{18} - 0.01278x_8x_{18} + 0.04658x_9x_{18} - 0.11095x_{10}x_{18} + \\
& 0.07151x_{11}x_{18} + 0.02977x_{12}x_{18} + 0.0495x_{14}x_{18} + 0.02117x_{15}x_{18} + 0.0208x_{16}x_{18} - 0.07068x_{17}x_{18} - 0.01711x_{18}^2 - \\
& 0.04685x_0x_{19} + 0.12461x_1x_{19} - 0.00152x_2x_{19} + 0.00549x_3x_{19} - 0.0037x_4x_{19} + 0.02334x_5x_{19} - 0.00204x_6x_{19} + \\
& 0.01064x_7x_{19} - 0.00468x_8x_{19} - 0.01531x_9x_{19} - 0.00025x_{10}x_{19} - 0.02014x_{11}x_{19} + 0.00253x_{12}x_{19} + 0.00039x_{13}x_{19} + \\
& 0.01247x_{14}x_{19} + 0.00623x_{15}x_{19} - 0.01978x_{16}x_{19} + 0.01116x_{17}x_{19} + 0.00997x_{18}x_{19} - 0.00125x_{19}^2 - 0.24571x_0x_{20} + \\
& 0.01527x_1x_{20} + 0.01513x_2x_{20} + 0.21444x_3x_{20} - 0.00042x_4x_{20} + 0.01284x_5x_{20} + 0.01556x_6x_{20} + 0.10175x_7x_{20} + \\
& 0.01378x_8x_{20} + 0.01132x_9x_{20} - 0.00014x_{10}x_{20} - 0.08197x_{11}x_{20} + 0.05976x_{12}x_{20} - 0.09782x_{13}x_{20} - 0.0311x_{14}x_{20} - \\
& 0.02498x_{15}x_{20} - 0.01261x_{16}x_{20} + 0.04463x_{17}x_{20} + 0.03959x_{18}x_{20} - 0.01115x_{19}x_{20} - 0.02539x_{20}^2 - 0.00024x_0x_{21} + \\
& 0.01386x_1x_{21} + 0.02516x_2x_{21} + 0.00222x_3x_{21} + 0.00004x_4x_{21} - 0.06183x_5x_{21} + 0.04606x_6x_{21} + 0.01789x_7x_{21} - \\
& 0.01325x_8x_{21} + 0.02409x_9x_{21} + 0.00023x_{10}x_{21} + 0.00006x_{11}x_{21} - 0.00317x_{12}x_{21} + 0.00541x_{13}x_{21} - 0.00012x_{14}x_{21} + \\
& 0.00012x_{15}x_{21} - 0.00592x_{16}x_{21} + 0.02185x_{17}x_{21} - 0.00001x_{18}x_{21} - 0.00003x_{20}x_{21} + 0.0257x_{21}^2 - 1.548990x_0 + \\
& 3.819380x_1 + 4.156280x_2 - 1.006350x_3 + 1.110140x_4 - 2.591030x_5 + 1.540610x_6 - 2.829660x_7 - 1.5943x_8 + 1.383880x_9 - \\
& 2.551x_{10} + 0.18922x_{11} + 0.72398x_{12} + 0.04856x_{13} + 0.48053x_{14} + 0.6916x_{15} - 1.253460x_{16} - 0.70274x_{17} + 0.44341x_{18} - \\
& 1.167680x_{19} - 1.8572x_{20} - 1.170990x_{21} - 23.70712
\end{aligned}$$

Figure 8. The formula learned by our model for the number of FRTs of a facet. The meanings of the x_i are described in section 4.3.




Figure 9. A heatmap showing $\log_{10}(|c|)$ for each coefficient c in the formula, after rescaling each variable to have expectation value 1. The top row corresponds to the constant term (top left square) and the linear terms.

4.6 Estimate for the total upper bound

Having obtained a satisfactory model, we generated the necessary input data for all of the facets. We then fed the data into the neural network to obtain the estimated number of FRTs for each facet.

Naively, for a given reflexive polytope, its FRSTs would be given by the set of all possible combinations of its facets' FRTs. But in reality, the number of FRSTs is reduced by two considerations. The first is that given two 3d facets F_1 and F_2 , the triangulation of each induces a triangulation on the intersection $F_1 \cap F_2$, and these induced triangulations may not overlap. The second consideration is that even if the induced triangulations do overlap, the aggregated triangulation of the reflexive polytope may fail to be regular even though the individual facet FRTs are regular. Thus, using the estimated number of FRTs for each facet, we are only able to estimate an upper bound for each 4d reflexive polytope. So, for a given reflexive polytope Δ with facets F_i , we know that

$$N_{\text{FRST}}(\Delta) \leq \prod_i N_{\text{FRT}}(F_i), \tag{4.6}$$

where N_{FRST} and N_{FRT} are the number of FRSTs and FRTs of the 2-skeleton of a given reflexive polytope and facet, respectively. Via this inequality, the upper bound is given by the product of the facet FRTs. One can then calculate the estimate for each 4d polytope and sum the results. Determining the degree to which the product of FRTs overestimates the number of FRSTs is left to a future work.

In practice, since our neural network predicted the natural logarithm of the number of FRSTs, we calculated the estimate for each polytope as

$$N_{\text{FRST}}(\Delta) = \exp\left(\sum_i \ln(N_{\text{FRT}}(F_i))\right). \tag{4.7}$$

In these calculations, we used the actual value for $\ln(N_{\text{FRT}}(F_i))$ whenever it was known, and the neural network prediction otherwise. As noted in section 3.3, the triangulated facets account for over 88% of all facets.

Performing this calculation for each polytope, it was found that the single polytope whose dual has $h^{1,1} = 491$ (the greatest $h^{1,1}$ value) dominated the count. Our estimation method predicts that it has $1.465 \times 10^{10,505}$ FRSTs. This is more than any other polytope by over 1300 orders of magnitude, and so this value effectively serves as the total triangulation number for the entire set. Considering that this polytope has 680 integral points, 40 more than any other, it is unsurprising that our method has identified this polytope as possessing the most triangulations.

4.7 The $h^{1,1} = 491$ polytope

The polytope whose FRST count dominates the database is the polytope dual to the single $h^{1,1} = 491$ polytope, which we will call Δ_{491}° . This polytope has 680 integral points and five facets, of which only four are unique. The polytope and its four facets F_i are given by the convex hulls

$$\begin{aligned}
 \Delta_{491}^\circ &= \text{conv}(\{\{1, 0, 0, 0\}, \{0, 1, 0, 0\}, \{0, 0, 1, 0\}, \{21, 28, 36, 42\}, \{-63, -56, -48, -42\}\}) \\
 F_1 &= \text{conv}(\{\{1, 0, 0, 0\}, \{0, 1, 0, 0\}, \{0, 0, 1, 0\}, \{21, 28, 36, 42\}\}) \\
 F_2 &= \text{conv}(\{\{1, 0, 0, 0\}, \{0, 1, 0, 0\}, \{3, 4, 6, 0\}, \{3, 4, 6, 84\}\}) \\
 F_3 &= \text{conv}(\{\{1, 0, 0, 0\}, \{0, 1, 0, 0\}, \{7, 8, 14, 0\}, \{7, 8, 14, 84\}\}) \\
 F_4 &= \text{conv}(\{\{1, 0, 0, 0\}, \{0, 1, 0, 0\}, \{7, 15, 21, 0\}, \{7, 15, 21, 84\}\}).
 \end{aligned} \tag{4.8}$$

The facet F_1 first appears as a dual facet at $h^{1,1} = 23$, and appears twice in Δ_{491}° . The facets F_2, F_3 and F_4 each appear once in Δ_{491}° and nowhere else in the database.

Our model predicts that F_1 will have $e^{29.32} = 5.41 \times 10^{12}$ FRTs. As it appears twice, it will contribute $e^{2 \times 29.32} = 2.93 \times 10^{25}$ to the triangulation number for this polytope, which is a subdominant contribution. F_2 has a larger contribution, as our model predicts that it will have $e^{2391.5} = 4 \times 10^{1038}$ FRTs. However, the triangulation number for the polytope is dominated by F_3 and F_4 , which our model predicts to have $e^{10,753.0} = 1 \times 10^{4670}$ and $e^{10,985.9} = 1.25 \times 10^{4771}$ FRTs, respectively. Combining these, we get the estimate of the number of FRSTs for this polytope:

$$N_{\text{FRST}}(\Delta_{491}^\circ) = (2.93 \times 10^{25})(4 \times 10^{1038})(1 \times 10^{4670})(1.25 \times 10^{4771}) = 1.5 \times 10^{10,505}. \tag{4.9}$$

Finally, we note that according to our model, the facets F_4 and F_3 have the most and second-most FRTs among all facets.

To estimate a range of uncertainty for this value, we can use the extrapolated MAPE values from our model in table 4. For $h^{1,1} = 12, 13, 14$, our model had MAPE values of 4.393, 4.490, and 4.416. Using the average of these values, 4.433, as the percent error for

each $\ln(N_{\text{FRT}})$ estimate in Δ_{491}° , we have

$$\begin{aligned} F_1 : \quad \ln(N_{\text{FRT}}) &= 29.32 \pm 1.30, \\ F_2 : \quad \ln(N_{\text{FRT}}) &= 2391.5 \pm 106.0, \\ F_3 : \quad \ln(N_{\text{FRT}}) &= 10753.0 \pm 476.7, \\ F_4 : \quad \ln(N_{\text{FRT}}) &= 10985.9 \pm 487.0. \end{aligned}$$

Propagating these errors, we find that for Δ_{491}° we have $\log_{10}(N_{\text{FRT}}) = 10,505.2 \pm 292.6$. This gives a range for our estimate of $10^{10,505.2 \pm 292.6} = [10^{10,212.6}, 10^{10,797.8}]$.

5 Comparison to 3d polytopes

The biggest point of uncertainty involving our analysis of the 4d reflexive polytopes is the inability to test our model on facets that first appear beyond $h^{1,1} = 15$. As all evidence suggests that these facets — and in particular those that first appear at high $h^{1,1}$ values — will be responsible for the dominant contribution to the total number of FRSTs, it would be preferable to have some confidence in our model’s accuracy in this region.

Though obtaining the true number of FRTs for these largest facets is currently infeasible, we hope to gain confidence in our model by showing that we can achieve the same goal for the 2d facets of the 3d reflexive polytopes. As these facets are smaller, the number of FRTs for the majority of the facets can be triangulated. Therefore a model’s extrapolation ability can be more thoroughly tested. Specifically, we will demonstrate that EQLs trained to predict $\ln(N_{\text{FRT}})$ predictions on data with relatively low $h^{1,1}$ nevertheless accurately extrapolate to much higher $h^{1,1}$ values, lending some credence to our above assumption the 4d case.

5.1 Classification of the 2d facets

As in section 2, we began by classifying the 2d facets of the 3d reflexive polytopes to avoid redundancy in our calculations. The same classification method as in section 2 — for each facet, we obtained a 3d polytope by attaching the origin, calculated the normal form, and removed the origin afterwards.

In section 2 we organized the facets by the first $h^{1,1}$ value at which they appeared. As $h^{1,1}$, as obtained by Batyrev’s formula, is equal to 20 for every 3d reflexive polytope, we instead organized the facets by the number of integral points $|\Delta|$ in the smallest polytope Δ in which they appeared. We note that this quantity is equal to $h^{1,1}(B) + 4$, where B is the smooth weak Fano toric threefold associated to Δ . For the remainder of this section, by $h^{1,1}$ we will mean $h^{1,1}(B) = |\Delta| - 4$.

As in section 2, the calculations were performed using our C++ lattice polytope implementation. Due to the small size of the 3d reflexive dataset, the classification was completed in just over 20 seconds.

Performing our classification, we found that the 4,319 3d reflexive polytopes contain 344 unique 2d facets. The numbers of facets and polytopes that first appear at each $h^{1,1}$ value are shown in figure 10.




Figure 10. (Left) The number of new facets at each $h^{1,1}$ value. (Right) The number of reflexive polytopes at each $h^{1,1}$ value.

5.2 Training data

Having classified the 344 unique facets, we set out to triangulate as many as possible in order to test our model. We again used TOPCOM to obtain the FRT count for each facet. We were able to successfully obtain the number of FRTs for 322 of the 344 facets, with all facets that appear up to $h^{1,1} = 25$ being completed.

Additionally, for 6 of the remaining 12 facets, as well as all of the other 322, we were able to obtain the number of fine triangulations (FTs). For the 322 facets for which the number of FRTs was found, there is very close agreement between the numbers of FRTs and FTs. More specifically, the number of FTs was at most 1.026 times the number of FRTs. Thus, we used the number of FTs as an approximation for the number of FRTs for the six additional facets. This gave us values for all facets that appear up through $h^{1,1} = 30$.

5.3 Machine learning and extrapolation

Having triangulated as many facets as possible, we trained models with an EQL hidden layer. Our input data was simpler than for the 3d facets, and consisted of

- The number of integral points
- The number of boundary points
- The number of interior points
- The number of vertices
- The length of the longest side
- The length of the shortest side
- The average side length

$h^{1,1}$	Facets	MAPE	MSE
6	5	7.865	0.032
7	13	16.583	0.061
8	15	8.805	0.055
9	12	5.851	0.075
10	19	5.808	0.087
11	19	10.678	0.213
12	15	8.754	0.334
13	18	9.128	0.330
14	21	10.200	0.722
15	22	8.756	0.538
16	14	9.103	0.755
17	22	9.071	0.610
18	13	7.850	0.619
19	21	10.491	1.314
20	13	10.962	2.050
21	7	9.259	1.158
22	23	9.167	0.798
23	14	14.333	6.504
24	7	7.894	1.075
25	3	2.649	0.373
26	6	12.112	2.228
27	5	3.985	0.644
28	3	6.900	0.380
29	1	5.258	0.295
30	1	2.074	0.146

Table 6. Performance of our model across $h^{1,1}$ of the 3d reflexive polytopes.

As before, we had the neural network train on $\ln(N_{\text{FRT}})$. As the goal of this exercise was to mimic the 4d case, where we only have data for the smallest facets, we restricted our training data to facets that first appear at $4 \leq h^{1,1} \leq 11$. There are 243 facets that first appear at $h^{1,1} \geq 12$, meaning that our extrapolation region contains 70.6% of the facets.

Our trained model contains a single EQL layer, which used a combination of linear and quadratic unary activation functions. The performance of the model is shown in table 6. Recall that $h^{1,1} = 25$ is the highest value for which all of the FRT values are known; for $26 \leq h^{1,1} \leq 30$ we approximate the number of FRT by the number of FTs for some facets.

As we can see, the model performs well through the extrapolation region of $17 \leq h^{1,1} \leq 25$. With our trained model, we predicted the number of FRTs for all 344 facets, and then estimated the number of FRSTs for any given polytope as the product of the predictions for each facet. We also computed these estimates using the real FRT numbers for polytopes where all facets have been triangulated. Taking the mean value at each value of $h^{1,1}$, we obtain the graph shown in figure 11. From this graph we see that




Figure 11. The mean value of $\log_{10}(N_{\text{FRST}})$ at $h^{1,1}$, using predicted facet FRT values (*blue*) and known facet FRT and FT values (*red*).

$$\begin{aligned} \ln(N_{\text{FRT}}) = & 0.01418x_0^2 - 0.03435x_0x_1 - 0.02165x_1^2 + 0.11134x_0x_2 + 0.00201x_1x_2 + 0.00206x_2^2 - \\ & 0.03566x_0x_3 - 0.02813x_1x_3 + 0.00993x_2x_3 + 0.05023x_3^2 - 0.03399x_0x_4 - 0.00929x_1x_4 - \\ & 0.01405x_2x_4 + 0.11072x_3x_4 + 0.0694x_4^2 + 0.04551x_0x_5 - 0.04939x_1x_5 + 0.04087x_2x_5 - \\ & 0.00532x_3x_5 + 0.00719x_4x_5 - 0.00774x_5^2 - 0.07105x_0x_6 + 0.04438x_1x_6 - 0.11917x_2x_6 - \\ & 0.07082x_3x_6 - 0.14734x_4x_6 - 0.007x_5x_6 + 0.14731x_6^2 - 0.28707x_0 + 0.46716x_1 - 0.59766x_2 - \\ & 0.4975x_3 - 0.35609x_4 - 0.49381x_5 + 1.354040x_6 + 5.530190 \end{aligned}$$

Figure 12. The formula learned by our model for the number of FRTs of a facet. The meanings of the x_i are described earlier in this section.

the estimates from the model predictions are in good agreement with the estimate made from known values well outside of the training region of $4 \leq h^{1,1} \leq 11$. In particular, we note that at the highest $h^{1,1}$ value available to us, $h^{1,1} = 30$, our model is in nearly perfect agreement with the true result.

We give below in figure 12 the formula learned by our model, with the x_i in the order of the list at the beginning of this section. As in the 4d case, we also rescaled the variables to have expectation value 1, with the results being show in figure 13. Again, we see many warm spots, making further interpretation of the formula difficult.

6 Conclusion

In this paper we have provided the first concrete estimate of the number of fine, regular, star triangulations of 4d reflexive polytopes,

$$n_{\text{FRST}} \simeq 10^{10,505.2 \pm 292.6}. \tag{6.1}$$




Figure 13. A heatmap showing $\log_{10}(|c|)$ for each coefficient c in the formula, after rescaling each variable to have expectation value 1. The top row corresponds to the constant term (top left square) and the linear terms.

This provides an upper bound on the number of topologically distinct Calabi-Yau threefold hypersurfaces in toric varieties.

We attempted a variety of supervised learning techniques for predicting n_{FRST} and found that a neural network with the equation learning (EQL) architecture performed best. The estimation was computed by taking products of numbers of FRTs of facets, where the EQL was trained on 4d polytopes up to $h^{1,1} = 11$ and made accurate predictions up to $h^{1,1} = 14$, which was the highest $h^{1,1}$ at which we were able to compute FRTs of facets for the sake of validating the trained EQL. While encouraging that the EQL extrapolated to higher $h^{1,1}$, it is so much smaller than the maximum $h^{1,1} = 491$ that it is hard to trust such a high extrapolation without further evidence. For that reason, we performed an analogous analysis in the case of 3d polytopes and found that the EQL was able to extrapolate accurately to $h^{1,1} = 30$, where the maximum possible value is $h^{1,1} = 35$, despite the fact that it was only trained up through $h^{1,1} = 11$. This provides some evidence that the corresponding extrapolation in the 4d case may be trustworthy, despite needing to extrapolate far beyond the training region. Indeed, such extrapolations were part of the motivation for the EQL architecture [31] in the first place.

After training successful models, we extracted the equations that they learned. The variables were rescaled to have an expectation value of 1 in an attempt to interpret these equations. We found that some variables and terms were more important than others, but the majority of variables made significant contributions, making interpretation of the equation difficult. This may be due to the functional form utilized in our EQLs.

In the process of making the prediction, we have demonstrated the overall utility of deep neural networks in this context, in particular their ability to extrapolate to regions of higher topological complexity. This provides motivation for further studies of triangulations and Calabi-Yau manifolds using techniques from data science.

Though this result is a modest prediction of a single number, it is a necessary step in understanding the ensemble of Calabi-Yau threefold hypersurfaces as a whole. We wish to turn to refined structures in the ensemble in the future, and in particular their implications for cosmology. For instance, studying axion-like particles (ALPs) in this context is particularly well-motivated since the ensemble is strongly dominated by the polytope that gives rise to the large number of ALPs.

Acknowledgments

We thank Cody Long and Jiahua Tian for useful discussions. B.D.N. and J.C. are supported by NSF Grant PHY-1620575. J.H. is supported by NSF grant PHY-1620526.

Open Access. This article is distributed under the terms of the Creative Commons Attribution License ([CC-BY 4.0](https://creativecommons.org/licenses/by/4.0/)), which permits any use, distribution and reproduction in any medium, provided the original author(s) and source are credited.

References

- [1] P. Candelas, G.T. Horowitz, A. Strominger and E. Witten, *Vacuum Configurations for Superstrings*, *Nucl. Phys. B* **258** (1985) 46 [[INSPIRE](#)].
- [2] P.S. Aspinwall, B.R. Greene and D.R. Morrison, *Calabi-Yau moduli space, mirror manifolds and space-time topology change in string theory*, *Nucl. Phys. B* **416** (1994) 414 [[hep-th/9309097](#)] [[INSPIRE](#)].
- [3] A. Strominger, *Massless black holes and conifolds in string theory*, *Nucl. Phys. B* **451** (1995) 96 [[hep-th/9504090](#)] [[INSPIRE](#)].
- [4] K. Dasgupta, G. Rajesh and S. Sethi, *M theory, orientifolds and G-flux*, *JHEP* **08** (1999) 023 [[hep-th/9908088](#)] [[INSPIRE](#)].
- [5] S.B. Giddings, S. Kachru and J. Polchinski, *Hierarchies from fluxes in string compactifications*, *Phys. Rev. D* **66** (2002) 106006 [[hep-th/0105097](#)] [[INSPIRE](#)].
- [6] R. Bousso and J. Polchinski, *Quantization of four form fluxes and dynamical neutralization of the cosmological constant*, *JHEP* **06** (2000) 006 [[hep-th/0004134](#)] [[INSPIRE](#)].
- [7] M.R. Douglas, *The Statistics of string/M theory vacua*, *JHEP* **05** (2003) 046 [[hep-th/0303194](#)] [[INSPIRE](#)].
- [8] J. Carifio, W.J. Cunningham, J. Halverson, D. Krioukov, C. Long and B.D. Nelson, *Vacuum Selection from Cosmology on Networks of String Geometries*, *Phys. Rev. Lett.* **121** (2018) 101602 [[arXiv:1711.06685](#)] [[INSPIRE](#)].
- [9] J. Halverson, C. Long and B. Sung, *Algorithmic universality in F-theory compactifications*, *Phys. Rev. D* **96** (2017) 126006 [[arXiv:1706.02299](#)] [[INSPIRE](#)].
- [10] J. Garriga, D. Schwartz-Perlov, A. Vilenkin and S. Winitzki, *Probabilities in the inflationary multiverse*, *JCAP* **01** (2006) 017 [[hep-th/0509184](#)] [[INSPIRE](#)].
- [11] S. Ashok and M.R. Douglas, *Counting flux vacua*, *JHEP* **01** (2004) 060 [[hep-th/0307049](#)] [[INSPIRE](#)].

- [12] F. Denef and M.R. Douglas, *Distributions of flux vacua*, *JHEP* **05** (2004) 072 [[hep-th/0404116](#)] [[INSPIRE](#)].
- [13] W. Taylor and Y.-N. Wang, *The F-theory geometry with most flux vacua*, *JHEP* **12** (2015) 164 [[arXiv:1511.03209](#)] [[INSPIRE](#)].
- [14] W. Taylor and Y.-N. Wang, *A Monte Carlo exploration of threefold base geometries for 4d F-theory vacua*, *JHEP* **01** (2016) 137 [[arXiv:1510.04978](#)] [[INSPIRE](#)].
- [15] W. Taylor and Y.-N. Wang, *Scanning the skeleton of the 4D F-theory landscape*, *JHEP* **01** (2018) 111 [[arXiv:1710.11235](#)] [[INSPIRE](#)].
- [16] F. Denef and M.R. Douglas, *Computational complexity of the landscape. I.*, *Annals Phys.* **322** (2007) 1096 [[hep-th/0602072](#)] [[INSPIRE](#)].
- [17] M. Cvetič, I. Garcia-Etxebarria and J. Halverson, *On the computation of non-perturbative effective potentials in the string theory landscape: IIB/F-theory perspective*, *Fortsch. Phys.* **59** (2011) 243 [[arXiv:1009.5386](#)] [[INSPIRE](#)].
- [18] N. Bao, R. Bousso, S. Jordan and B. Lackey, *Fast optimization algorithms and the cosmological constant*, *Phys. Rev. D* **96** (2017) 103512 [[arXiv:1706.08503](#)] [[INSPIRE](#)].
- [19] F. Denef, M.R. Douglas, B. Greene and C. Zukowski, *Computational complexity of the landscape II — Cosmological considerations*, *Annals Phys.* **392** (2018) 93 [[arXiv:1706.06430](#)] [[INSPIRE](#)].
- [20] J. Halverson and F. Ruehle, *Computational Complexity of Vacua and Near-Vacua in Field and String Theory*, *Phys. Rev. D* **99** (2019) 046015 [[arXiv:1809.08279](#)] [[INSPIRE](#)].
- [21] Y.-H. He, *Deep-Learning the Landscape*, [arXiv:1706.02714](#) [[INSPIRE](#)].
- [22] F. Ruehle, *Evolving neural networks with genetic algorithms to study the String Landscape*, *JHEP* **08** (2017) 038 [[arXiv:1706.07024](#)] [[INSPIRE](#)].
- [23] D. Krefl and R.-K. Seong, *Machine Learning of Calabi-Yau Volumes*, *Phys. Rev. D* **96** (2017) 066014 [[arXiv:1706.03346](#)] [[INSPIRE](#)].
- [24] D. Klaewer and L. Schlechter, *Machine Learning Line Bundle Cohomologies of Hypersurfaces in Toric Varieties*, *Phys. Lett. B* **789** (2019) 438 [[arXiv:1809.02547](#)] [[INSPIRE](#)].
- [25] K. Bull, Y.-H. He, V. Jejjala and C. Mishra, *Machine Learning CICY Threefolds*, *Phys. Lett. B* **785** (2018) 65 [[arXiv:1806.03121](#)] [[INSPIRE](#)].
- [26] Y.-N. Wang and Z. Zhang, *Learning non-Higgsable gauge groups in 4D F-theory*, *JHEP* **08** (2018) 009 [[arXiv:1804.07296](#)] [[INSPIRE](#)].
- [27] K. Hashimoto, S. Sugishita, A. Tanaka and A. Tomiya, *Deep learning and the AdS/CFT correspondence*, *Phys. Rev. D* **98** (2018) 046019 [[arXiv:1802.08313](#)] [[INSPIRE](#)].
- [28] J. Liu, *Artificial Neural Network in Cosmic Landscape*, *JHEP* **12** (2017) 149 [[arXiv:1707.02800](#)] [[INSPIRE](#)].
- [29] R. Jinno, *Machine learning for bounce calculation*, [arXiv:1805.12153](#) [[INSPIRE](#)].
- [30] M. Kreuzer and H. Skarke, *Complete classification of reflexive polyhedra in four-dimensions*, *Adv. Theor. Math. Phys.* **4** (2002) 1209 [[hep-th/0002240](#)] [[INSPIRE](#)].
- [31] G. Martius and C.H. Lampert, *Extrapolation and learning equations*, [arXiv:1610.02995](#).
- [32] J. Carifio, J. Halverson, D. Krioukov and B.D. Nelson, *Machine Learning in the String Landscape*, *JHEP* **09** (2017) 157 [[arXiv:1707.00655](#)] [[INSPIRE](#)].

- [33] V.V. Batyrev, *Dual polyhedra and mirror symmetry for Calabi-Yau hypersurfaces in toric varieties*, *J. Alg. Geom.* **3** (1994) 493 [[alg-geom/9310003](#)] [[INSPIRE](#)].
- [34] M. Kreuzer and H. Skarke, *PALP: A Package for analyzing lattice polytopes with applications to toric geometry*, *Comput. Phys. Commun.* **157** (2004) 87 [[math/0204356](#)] [[INSPIRE](#)].
- [35] J. Rambau, *TOPCOM: Triangulations of point configurations and oriented matroids*, in *Mathematical Software — ICMS 2002*, A.M. Cohen, X.-S. Gao and N. Takayama eds., pp. 330–340, World Scientific (2002) [<http://www.zib.de/PaperWeb/abstracts/ZR-02-17>].
- [36] R. Altman, J. Gray, Y.-H. He, V. Jejjala and B.D. Nelson, *A Calabi-Yau Database: Threefolds Constructed from the Kreuzer-Skarke List*, *JHEP* **02** (2015) 158 [[arXiv:1411.1418](#)] [[INSPIRE](#)].
- [37] C. Long, L. McAllister and P. McGuirk, *Heavy Tails in Calabi-Yau Moduli Spaces*, *JHEP* **10** (2014) 187 [[arXiv:1407.0709](#)] [[INSPIRE](#)].
- [38] M. Demirtas, C. Long, L. McAllister and M. Stillman, *The Kreuzer-Skarke Axiverse*, [arXiv:1808.01282](#) [[INSPIRE](#)].
- [39] J. Halverson and J. Tian, *Cost of seven-brane gauge symmetry in a quadrillion F-theory compactifications*, *Phys. Rev. D* **95** (2017) 026005 [[arXiv:1610.08864](#)] [[INSPIRE](#)].
- [40] R. Grinis and A. Kasprzyk, *Normal forms of convex lattice polytopes*, [arXiv:1301.6641](#).