

Holomorphic feedforward networks

MICHAEL R. DOUGLAS

Abstract: A very popular model in machine learning is the feedforward neural network (FFN). The FFN can approximate general functions and mitigate the curse of dimensionality. Here we introduce FFNs which represent sections of holomorphic line bundles on complex manifolds, and ask some questions about their approximating power. We also explain formal similarities between the standard approach to supervised learning and the problem of finding numerical Ricci flat Kähler metrics, which allow carrying some ideas between the two problems.

I’m delighted to contribute to this volume for Bernie Shiffman, who along with Steve Zelditch initiated me into complex geometry as it is actually done by mathematicians. Our papers on the statistics of flux vacua [17–19] remain important in string theory. They were also my jumping-off point for many other works on complex geometry, in particular work on balanced metrics and the Tian–Yau–Zelditch–Lu expansion with Semyon Klevtsov [14, 15], and on numerical methods for finding Ricci-flat metrics on Calabi–Yau manifolds [6, 7, 12, 13, 16].

More recently, I have been studying machine learning and artificial intelligence. My interest in this goes back to my graduate school days, where I worked with Gerry Sussman [2], a pioneer of AI, and took classes from John Hopfield on his famous model of neural networks. Some reminiscences about these times are in [11]. Although it took many years, machine learning is now at the forefront of applied mathematics and many other fields. Its concepts and technologies are being applied to many problems which at first might not be thought to be related to learning or statistics.

One very active topic is numerical methods for PDE inspired by machine learning, and a few examples are [8, 21, 25]. A good way to explain the relevance of ML is to recall the “curse of dimensionality” [5]. A central part of any numerical method for PDE is to define a finite dimensional space of

Received July 31, 2020.

2010 Mathematics Subject Classification: Primary 32Q25; secondary 65M99, 68T07.

parameterized functions \mathcal{F} which are used to approximate the solutions. The simplest way to do this is to specify the function values at a fixed subset of points. This approximation is controlled by restricting the variation of the target function, for example one can put a uniform bound on its derivatives such as $|f(x) - f(y)| \leq L|x - y|$. While this gets us down to a finite problem which can be put on a computer, on a d -dimensional manifold M we require $\mathcal{O}(L^d)$ parameters. This is a disaster for the high dimensional PDEs of many-body physics and control theory, and problematic even for moderate dimensions such as the $d = 6$ Calabi–Yau manifolds of interest in string theory.

Machine learning also involves approximating functions, often functions on thousands or even millions of variables, such as digital images. Thus the curse of dimensionality is a central problem and there are many ways to deal with it. The simplest class of ML problems is supervised learning, in which one needs to estimate an unknown function $f : M \rightarrow \mathbb{R}$ from data. This includes classification – say, we take M to be the space of maps from pixels to intensities, and we define $f > 0$ (resp. $f < 0$) to mean that the image is a picture of a cat (resp. a dog). What makes this “learning” is that we do not say much about f *a priori*, rather we tell the computer how to infer f from examples. Thus, we give it a dataset \mathcal{D} of input–output pairs $(x_i, f(x_i))$ sampled from some underlying distribution ρ on M . We then postulate a class of functions \mathcal{F} (a model), and an algorithm which looks at the data and produces a $\hat{f} \in \mathcal{F}$ which generalizes to the underlying distribution, *i.e.* $\mathbb{E}_\rho[\hat{f} - f]$ is expected to be small. A straightforward way to try to do this is to numerically minimize the error $\mathbb{E}_\mathcal{D}[\hat{f} - f]$ on the dataset.

In traditional statistics, one postulates a simple \mathcal{F} , say functions which are a linear combination of N fixed basis functions. One then finds a trade-off between accuracy of approximation (which increases with N) and errors arising from fitting the randomness in sampling from ρ (which also increase with N). This is quantified by concepts such as bias-variance tradeoff, and it leads one to prefer simple models, those which use the smallest possible number of parameters N needed to describe the data.

Much of the recent work rejects the idea that \mathcal{F} should be simple. In “deep learning,” one takes \mathcal{F} to be the class of functions which can be realized by a class of feedforward neural networks with a definite number of layers, parameters, *etc.* These functions depend nonlinearly on the parameters in fairly complicated ways, but in return the curse of dimensionality is mitigated. While the nonlinearity of \mathcal{F} would be expected to lead to a bad optimization problem with many local minima, in practice this turns out not to be the case. Indeed, by taking a large number of parameters, one can simplify

the optimization landscape. Given enough parameters (many more than the number of data points), one can find a \hat{f} which precisely fits the data points. For such a \hat{f} , $\mathbb{E}_{\mathcal{D}}[\hat{f} - f] = 0$ and it turns out to be easy to find such minima. However, the dogmas of statistics tell us that such minima are “overfit” and give a poor description of the underlying distribution, *i.e.* $\mathbb{E}_{\rho}[\hat{f} - f]$ is large.

But surprisingly, with the right choices of \mathcal{F} and optimization algorithm, this is not the case. Such models can generalize well, apparently violating this dogma. Understanding how this can be is a very active field of research, of which a few recent works include [3, 4, 29].

Carrying over these advantages from ML to PDE, one could hope that a function class \mathcal{F} produced by a feedforward neural network would have the same advantages in numerical solution of PDE, mitigating the curse of dimensionality and simplifying the optimization problem. These hopes have been borne out in many works such as [8, 21, 25].

The present work is inspired by work on numerical methods to find Ricci flat metrics on hypersurfaces in projective space [10, 13, 22]. Here \mathcal{F} is a class of Kähler metrics, and in previous works this was taken to be the space of all metrics which could be obtained by pullback from a Fubini–Study metric, thus log linear in the parameters. We adapt the feedforward neural network to this problem and define “holomorphic networks” and “bihomogeneous networks,” which represent subsets of these metrics using a nonlinear parameterization. In [16] we study this approach numerically. Here we point out some formal parallels with machine learning, and ask some mathematical questions about the ability of these FFNs to approximate functions.

1. Network approximations to Kähler metrics

1.1. Feed-forward networks

A feed-forward network (FFN, also called MLP for multilayer perceptron) is a nonlinear map $F[\vec{W}]$ from a vector space \mathcal{X} to another vector space \mathcal{Y} with parameters \vec{W} . It is built by composing a sequence of maps which alternate between two types, general linear maps W and fixed nonlinear transformations θ , as in

$$(1.1) \quad F[\vec{W}] = W^{(d)} \circ \theta|_{V_{d-1}} \circ W^{(d-1)} \circ \dots \circ \theta|_{V_2} \circ W^{(2)} \circ \theta|_{V_1} \circ W^{(1)}.$$

Each linear map $W^{(i)}$ has as its range a new vector space V_i , so

$$(1.2) \quad W^{(1)} \in \text{Hom}(\mathcal{X}, V_1),$$

$$(1.3) \quad W^{(2)} \in \text{Hom}(V_1, V_2),$$

$$(1.4) \quad \begin{array}{c} \vdots \\ W^{(d)} \end{array} \in \text{Hom}(V_{d-1}, \mathcal{Y})$$

The combination $\theta \circ W$ is called a layer, with the final layer $W^{(d)}$ being an exception in not having θ . The number of layers d is the depth.

Usually one takes the domain $\mathcal{X} \cong \mathbb{K}^D$ and range $\mathcal{Y} \cong \mathbb{K}^{D'}$ with $\mathbb{K} \cong \mathbb{R}$, though $\mathbb{K} \cong \mathbb{C}$ is sometimes used in signal processing applications. We are free to choose the dimensions D_i of the intermediate spaces V_i , the width hyperparameters of the network.¹ By “the width” we will mean $\max D_i$.

Generally one allows the W 's to be arbitrary linear transformations, so the parameters consist of the list of $W^{(i)}$. Thus, defining $D = \dim \mathcal{X}$ and $D' = \dim \mathcal{Y}$, we have

$$(1.5) \quad \vec{W} \in \mathcal{W}_{\vec{D}} \cong \bigoplus_{i=1}^d \text{Hom}(\mathbb{K}^{D_{i-1}}, \mathbb{K}^{D_i}).$$

To define the θ 's, we start with the one dimensional case $\theta|_{\mathbb{K}} : \mathbb{K} \rightarrow \mathbb{K}$, which is called the activation function. This could be any function; two popular choices for $\mathbb{K} = \mathbb{R}$ are $\theta(x) = \tanh x$, and the “rectified linear unit” or ReLU function

$$(1.6) \quad \theta_{ReLU}(x) = \begin{cases} x, & x \geq 0 \\ 0, & x < 0. \end{cases}$$

To define θ_V for a general vector space V , we pick a basis e_i for V and apply $\theta|_{\mathbb{K}}$ componentwise,

$$(1.7) \quad \theta_V \left(\sum_i c_i e_i \right) = \sum_i \theta_{\mathbb{K}}(c_i) e_i.$$

While this depends on the choice of basis, since every θ in Eq. (1.1) appears both prefixed and postfixed by a general linear transformation, the parameterized space of maps is independent of this choice. Thus, for each choice of specific dimensions \vec{D} and activation function $\theta|_{\mathbb{K}}$, we get a space of FFN maps of dimension $DD_1 + D_1D_2 + \dots + D_{d-1}D'$, which we denote $\mathcal{F}_{\vec{D};\theta}$.

This definition was inspired by a simple model in neuroscience (the “perceptron”), in which a neuron is modeled by the computation which takes an

¹The term hyperparameter in machine learning refers to a choice which is not determined automatically by optimization, in contrast to parameters (also called “coefficients” or “weights”) which are generally chosen to minimize an energy (“objective” or “loss”) function.

input vector v and produces a single component of $(\theta \circ W)(v)$. These days neuroscience uses far more sophisticated models of neurons, but FFN’s are widely used in applied math and machine learning as function approximating spaces, analogous to spaces of polynomials or other linear combinations of basis functions. Thus rather than “neuron” one usually uses the term “unit,” and the network $\mathcal{F}[\vec{D}; \theta]$ contains $D_1 + \dots + D_{d-1} + D'$ units.

It has been shown that feed-forward networks can approximate arbitrary real valued functions. This is the case even for $d = 2$ [9], but in this case one can need an exponentially large number of units, as would be the case for simpler methods of interpolation (the “curse of dimensionality”). By using more layers, one can gain many advantages – complicated functions can be represented with many fewer units, and local optimization techniques are much more effective. How exactly this works is not well understood theoretically and there are many interesting observations and hypotheses as to how these advantages arise.

1.2. Multilayer holomorphic embeddings

Our goal is to use the FFN Eq. (1.1) to define a parameterized space of functions from a projective manifold M to a space \mathcal{Y} . To begin, let $\mathcal{L} \rightarrow M$ be a holomorphic line bundle, let $\underline{s} = (s_0, \dots, s_N)$ be a basis of sections of $H^0(\mathcal{L})$, and let ι_s be the corresponding Kodaira embedding of M into $\mathbb{C}\mathbb{P}^N$.

To use Eq. (1.1), we want to regard this embedding as a map into $\mathcal{X} \cong \mathbb{C}^{N+1}$. We can do this in a patch $U \subset M$ by choosing a local frame. If the maps W and θ each have a nice geometric interpretation, the choice of frame should drop out at the end. Taking the V_i ’s to be complex vector spaces, this is evident for W as each $W \in \text{Hom}(V, V')$ corresponds to a holomorphic map $\mathbb{P}V \rightarrow \mathbb{P}V'$.

An arbitrary choice of activation function θ will not have a simple geometric interpretation. To find one which does, consider the particular case of $M = \mathbb{C}\mathbb{P}^n$ and $\mathcal{L} = \mathcal{O}(k)$, so that the sections s are degree k polynomials. Clearly we want $\theta(s)$ to also be a homogeneous polynomial, so the natural choice is

$$(1.8) \quad \theta(s) \equiv s^p$$

and its componentwise analog Eq. (1.7). Taking into account the dependence on the frame, the geometric interpretation of this is a map

$$(1.9) \quad \theta : \underbrace{H^0(\mathcal{L}) \oplus \dots \oplus H^0(\mathcal{L})}_{D \text{ times}} \rightarrow \underbrace{H^0(\mathcal{L}^p) \oplus \dots \oplus H^0(\mathcal{L}^p)}_{D \text{ times}}.$$

Another way to relate Eq. (1.1) to geometry is to appeal to the “metatheorem for vector bundles” ([23], Thm. 2.2.3), that any canonical construction in linear algebra gives rise to a geometric version for holomorphic vector bundles. To use this we define θ in terms of a canonical multilinear map on \mathbb{K} , namely the tensor product

$$(1.10) \quad \otimes_p : \mathbb{K}^m \rightarrow \mathbb{K} : (z_1, z_2, \dots, z_k) \rightarrow z_1 z_2 \dots z_p.$$

By evaluating this on p copies of a section, we get a natural map from the line bundle \mathcal{L}^m to its p 'th power \mathcal{L}^{pm} .

This interpretation of Eq. (1.1) was made in [16] in the following concrete form. We took M to be a hypersurface in $\mathbb{C}\mathbb{P}^m$, and $\mathcal{L} = \mathcal{O}(1)$. We then took as our computational representation of M the image in $H^0(\mathcal{O}(1))$ of a set of randomly chosen points in M . We could then use the implementations of the W and θ maps in a standard ML package (Tensorflow).

While we will generally use $p = 2$ in our networks, one could take a different p in each layer, so the network Eq. (1.1) depends on a vector $\vec{p} = (p_1, \dots, p_d)$. As a further generalization, instead of identifying the input sections z_i in Eq. (1.10), one could multiply the outputs of distinct FFN's, or make other geometrically consistent combinations. Each possible “architecture” is naturally associated to a directed graph, in which a particular use of Eq. (1.10) corresponds to a vertex, the inputs z_i correspond to incoming edges, and the output $z_1 \dots z_p$ corresponds to a single outgoing edge. These are examples of tensor networks, about which we will say more below.

The upshot is a construction of parameterized maps from M into $\oplus_D H^0(\mathcal{L}^m)$, defined by Eq. (1.1) and Eq. (1.8) with parameter space Eq. (1.5). Let us denote these maps as $F_{\mathcal{L}; \vec{p}; \vec{D}}^h[\vec{W}]$ with $\vec{W} \in \mathcal{W}_{\vec{D}}$ (or just $F^h[\vec{W}]$ in context). We can regard them as defining embeddings of M into projective space, which for sufficiently large D include all possible embeddings by a complete space of sections.

1.3. Algebraic metrics and holomorphic networks

Next, given an embedding of M into projective space, we can pull back the Fubini–Study metric to get a Kähler metric on M . Let $\mathcal{K}_{\mathcal{L}}^{FS}$ be the space of Fubini–Study metrics on $\mathbb{C}\mathbb{P}^N \cong \mathbb{P}H^0(\mathcal{L})^*$, parameterized by an $(N + 1) \times (N + 1)$ positive definite hermitian matrix $G_{I, \bar{J}}$. The Kähler potentials

$$(1.11) \quad K^{FS}[G] = \log \sum_{I, \bar{J}} G_{I, \bar{J}} s^I \bar{s}^{\bar{J}},$$

then define Kähler metrics on M , the space of which we also denote $\mathcal{K}_{\mathcal{L}}^{FS}$. Replacing \mathcal{L} with the series of line bundles \mathcal{L}^k , we get embeddings into higher dimensional projective spaces, and a sequence of spaces of Fubini–Study metrics $\mathcal{K}_{\mathcal{L}^k}^{FS}$. This gives us the ability to approximate a given Kähler metric to arbitrary accuracy, by taking sufficiently high k . As discussed in [10, 27], this accuracy can go as $k^{-\nu}$ for any ν .

This is a very nice parameterized class of metrics, and we will shortly discuss a method for finding the best numerical approximation to the Ricci flat metric in this class. However, there is a problem with this numerical application. Since $N \sim k^{\dim_{\mathbb{C}} M}$, the number of coefficients will go as $k^{\dim_{\mathbb{R}} M}$, the curse of dimensionality referred to earlier.

Let us try to break this curse by replacing the complete space of sections with the image of $F_{\mathcal{L};\vec{p};\vec{D}}^h[\vec{W}]$, to produce a space of metrics $\mathcal{K}_{\mathcal{L};\vec{p};\vec{D}}^h$. Now the parameter spaces Eq. (1.5) for the construction in §1.2 come in a variety of sizes, obtained by adjusting the intermediate dimensions $\{D_i\}$. And since the degree k is related to the depth as $k = \prod_{i=1}^{d-1} p_i$ (in terms of \vec{p} in Eq. (1.8)), the number of coefficients scales with k only as $\mathcal{O}(D^2 \log k)$.

We now write this out explicitly for $d = 2$ and $p = 2$. A slightly simpler expression is obtained by removing the final weight matrix $W^{(2)}$ from Eq. (1.1), which is redundant if we can vary G in Eq. (1.11). We then have

$$(1.12) \quad K^h[W^{(1)}, G] = \log \sum_{1 \leq I, \bar{J} \leq D_1} G_{I, \bar{J}} \left(\sum_i W_i^{(1), I} s^i \right)^2 \left(\sum_j W_j^{(1)*, \bar{J}} \bar{s}^{\bar{j}} \right)^2.$$

The nonlinear dependence on the weights W is a general feature of FFN’s.

We studied this construction computationally along lines we describe below, and found that it is not very good at approximating Ricci flat metrics. This may be because these metrics are pullbacks of degenerate FS metrics on the embedding space, called partial Fubini–Study metrics (see for example [28]).

1.4. Bihomogeneous networks

A variation on the previous construction is to first make sesquilinear combinations of the holomorphic sections and then apply the nonlinear network construction. Thus, we could take

$$(1.13) \quad \mathcal{X} \equiv H^0(\mathcal{L}) \otimes H^0(\bar{\mathcal{L}}),$$

embed $M \hookrightarrow \mathcal{X}$ by taking the outer product $s^I \times \bar{s}^{\bar{J}}$ of the basis of sections with its complex conjugate, and use the network Eq. (1.1) to define parameterized

functions $F[\vec{W}] : \mathcal{X} \rightarrow \mathbb{R}$. We then replace the Kähler potential Eq. (1.11) with

$$(1.14) \quad K_{\mathcal{L}; \vec{p}; \vec{D}}^b[\vec{W}] = \log F_b$$

$$(1.15) \quad F^b = W^{(d)} \circ \theta \circ W^{(d-1)} \circ \dots \circ \theta \circ W^{(1)} \cdot (\operatorname{Re} s^I \bar{s}^{\bar{J}}, \operatorname{Im} s^I \bar{s}^{\bar{J}}).$$

In terms of a concrete embedding by polynomials, this amounts to taking the linear spaces V_i to be subspaces of the space of bihomogeneous polynomials of degree (n_i, n_i) , while θ is still Eq. (1.8).

We will denote a particular Kähler potential of this type as $K^b[\vec{W}]$ and the space of these as $\mathcal{K}_{\mathcal{L}; \vec{p}; \vec{D}}^b$. Here $D = D_0 = b^0(\mathcal{L})^2$ and $D' = D_d = 1$, with the intermediate D_i adjustable. Note that \mathcal{X} has a real structure, and we can take the intermediate values and weights in Eq. (1.1) to be real. Thus this network has $DD_1 + \dots + D_{d-1}$ real parameters. For sufficiently large D_i , this reproduces the complete space $\mathcal{K}_{\mathcal{L}^{p_1 \dots p_{d-1}}}^{FS}$ of FS Kähler potentials, though with a more complicated parameterization.

Here is a depth 2 bihomogeneous network with $p = 2$ (the analog of Eq. (1.12)), and a depth 3 network:

$$(1.16) \quad K_{\mathcal{L}; 2; D_1}^b = \log \sum_{1 \leq I \leq D_1} W_I^{(2)} \left(\sum_i W_{i, \bar{j}}^{(1), I} s^i \bar{s}^{\bar{j}} \right)^2$$

$$(1.17) \quad K_{\mathcal{L}; \vec{2}; D_1, D_2}^b = \log \sum_{1 \leq I \leq D_2} W_I^{(3)} \left(\sum_{1 \leq J \leq D_1} W_J^{(2), I} \left(\sum_i W_{i, \bar{j}}^{(1), J} s^i \bar{s}^{\bar{j}} \right)^2 \right)^2.$$

These networks are better behaved than the holomorphic networks for small intermediate widths. In particular, the function class with d layers is simply contained in the function class with $d + 1$ layers. To be precise, take $p = 2$, then we apply a rescaling (denoted $2 \times$) to match the Kähler classes from the two constructions,

$$(1.18) \quad 2 \times \mathcal{K}[\mathcal{L}; \vec{2}; \vec{D}] \subset \mathcal{K}[\mathcal{L}; \vec{2}; \vec{D} \oplus D_d].$$

For \mathcal{K}^h one needs $D_d \geq h^0(\mathcal{L}^{2^{d-1}})$. But for \mathcal{K}^b , this can be accomplished with $D_d = 1$ (for example, $D_2 = 1$ in Eq. (1.17)).

In particular one can start with the 1 layer network $F = |W \cdot z|^2$ and repeatedly apply $\theta \cdot \mathbf{id}$ to get $F = |W \cdot z|^{2^k}$, a function sharply peaked at $z \propto W$. One can then add this to another network F' , increasing its widths by one and adding a feature with size $1/k$. Thus one can describe structure on the same short length scales as Eq. (1.11) but using many fewer parameters.

In [16] we give numerical evidence that these metrics are quite good at representing the Ricci flat metrics. Again, there are variations of this construction labeled by directed graphs. We could also take differences of these potentials to define classes of real valued functions.

1.5. Matrix product states

These were originally proposed to describe spin chains in quantum physics, and have been applied to ML problems in [1, 26] and other works. We will use them to make a third definition of a parameterized subset of the Fubini–Study metrics.

Without going into all the details, a spin is a quantum system whose wave function is a point in $\mathbb{P}V^*$ with $V \cong \mathbb{C}^D$, while a chain of N spins has a wave function Ψ in $\mathbb{P}\mathcal{X}$ where $\mathcal{X} \equiv \otimes^N V^*$. Since N might be the number of atoms in a macroscopic object, the curse of dimensionality is a dominant aspect of these problems. Thus the need for parameterized low dimensional subsets of \mathcal{X} is even greater.

A matrix product state (MPS) is defined by specifying a series of auxiliary linear spaces L_i , and a series T_i of multilinear maps

$$(1.19) \quad T_i : V \rightarrow \text{Hom}(L_i, L_{i+1}).$$

Typically one takes all of the L_i to be isomorphic, and then writes

$$(1.20) \quad \Psi_T = \text{Tr } T_1 T_2 \dots T_N.$$

This gives us an $N(\dim L)^2 - 1$ dimensional subset of the full $D^N - 1$ dimensional space of wave functions.

To adapt this to the problem at hand, we could take $V \equiv H^0(\mathcal{L})$ to get a parameterized subset of holomorphic sections of $\text{Sym}^N V$. While we have not yet tried this construction or its many variations, one which seems particularly natural is a bihomogeneous version with

$$(1.21) \quad K^T \equiv \log \text{Tr } (T_1 \cdot z)(T_1 \cdot z)^\dagger (T_2 \cdot z)(T_2 \cdot z)^\dagger \dots (T_d \cdot z)(T_d \cdot z)^\dagger.$$

1.6. Numerical methods

Suppose we want to approximate a given Kähler potential K with one from these classes, say \mathcal{K}^b . This is a problem in interpolation of functions, which is closely related to supervised machine learning. To get a dataset, we sample

points x_i from \mathcal{X} and evaluate $y_i = K(x_i)$, to get “data” (x_i, y_i) . We then postulate a “loss function,” a measure of the distance between the dataset and the model $(x_i, K_b(x_i))$. This is a function of the parameters of K^b , which we then minimize using computational methods, usually gradient descent.

Now, the problem of finding a Ricci flat Kähler metric can be phrased as an approximation problem. Since the first Chern class vanishes, the canonical bundle has a unique section up to overall scale. This is a nonvanishing holomorphic d -form, call it Ω . Then the Ricci flat Kähler form ω satisfies

$$(1.22) \quad \omega^d = \mathcal{N}\Omega \wedge \bar{\Omega}$$

where the constant \mathcal{N} can be determined by integrating both sides. One could then use the L^p norm of the difference as a loss function, as in [22].

In [16] we carried this out to find numerical Ricci flat metrics on quintic hypersurfaces of dimension 3, following the general approach developed in [13, 22]. To sample from a hypersurface in $\mathbb{C}\mathbb{P}^{d+1}$, we sample pairs of points using some measure μ on $\mathbb{C}\mathbb{P}^{d+1}$. We then find the intersection of the corresponding line with the hypersurface. By a result of Shiffman and Zelditch ([24], lemma 3.1), the resulting set of points is distributed by the pullback of μ . Then, since one can compute the right hand side of Eq. (1.22), one can treat this as a problem in function interpolation, choosing the weights in Eq. (1.14) to best fit this given function. This is a task well suited to machine learning software.

2. Accuracy of approximation

What determines the accuracy of the FFN description Eq. (1.14) of a canonical Kähler metric? Is it the depth of the network or the number of its parameters? Here is a precise form of the question and some speculations.

2.1. Approximation by algebraic metrics

To begin, recall the discussion of approximation of metrics from [10, 27]. Given any Kähler metric ω on M , we can find a sequence of algebraic metrics (pullbacks of degree k Fubini–Study metrics Eq. (1.11) to M) such that

$$(2.1) \quad \|\omega - \omega_k\| = o(k^{-\nu})$$

for any ν . This can be shown by using the Tian–Yau–Zelditch–Lu expansion for the density of states on M ,

$$(2.2) \quad \rho_h(z) \equiv \sum_I |s^I(z)|_h^2$$

$$(2.3) \quad = 1 + k^{-1}a_1(\omega) + k^{-2}a_2(\omega) + \dots,$$

where h is a Hermitian metric on L^k with curvature $-2\pi i k \omega_h$, and the a_i are local invariants constructed from ω . To define ρ we introduce a Hermitian metric on $H^0(L^k)$,

$$(2.4) \quad \|s\|_{H^{ilb}(h)}^2 = R \int_M d\mu_h |s|_h^2,$$

where $d\mu_h = \omega_h^n/n!$. The sum in Eq. (2.3) is taken over an orthonormal basis with respect to Eq. (2.4). Thus it depends on h both through the inner product Eq. (2.4) and explicitly in $|\dots|_h$.

Now, starting with h such that $\omega_h = \omega$, we can find a \tilde{h} whose curvature is

$$(2.5) \quad \omega_{\tilde{h}} = \omega_h + k^{-1} i \partial \bar{\partial} \log(\rho_h).$$

One way to see this is to regard the sum in Eq. (2.2) as a bihomogeneous polynomial in the sense of §1.4. As such it can be directly interpreted as \tilde{h} . Taking $\partial \bar{\partial} \log$ of Eq. (2.3) then gives us Eq. (2.5).

We can now define the balanced metrics. These have $\tilde{h} = h$ in Eq. (2.5), and thus

$$(2.6) \quad \rho_{\tilde{h}}(z) = 1.$$

When a balanced metric exists, by Eq. (2.3) it is a canonical approximation to a Ricci flat or constant scalar curvature metric.

The ideas above can be combined to show Eq. (2.1). One can use Eq. (2.3) to compute the leading term in $\rho_{\tilde{h}}$ in an expansion in $1/k$, and then Eq. (2.5) to define a corrected \tilde{h}' for which this term agrees with the one in ρ_h . This procedure can be carried out to arbitrary order. The problem is that one needs k^{2n} coefficients to describe this sequence. This fits with the expectation from “curse of dimensionality” for a basis of functions which is localized on a length scale $L \sim 1/k$.

2.2. Approximation by bihomogeneous metrics

Could Eq. (2.1) hold with ω_k taken from some restricted class of metrics? In particular, consider the subsets of metrics $\mathcal{K}_{\mathcal{L}; \vec{p}; \vec{D}}^b$ defined by the bihomogeneous networks Eq. (1.14). For sufficiently large widths \vec{D} , these include the general Fubini–Study metric, and as we decrease the widths we get subsets which use many fewer parameters, of order width $\max D_i$ times depth d . Furthermore, as we explained these metrics can contain structure on the same

length scale $L \sim 1/k$ as the general Fubini–Study metric. Could these networks break the curse of dimensionality?

To sharpen the question, define a polynomial width sequence of networks to be a sequence of networks with $d \in \mathbb{N}$, such that the width $\max D_i$ grows at most polynomially in d . Take $p = 2$ for definiteness, then these networks define sequences of metrics ω_k with $k = 2^{d-1}$. Define ω to be of polynomial complexity if a sequence exists satisfying Eq. (2.1) for some $\nu > 0$. We would like to know whether there exist metrics ω which are not of polynomial complexity.

Let us try to bound the maximum distance from the algebraic metrics to an approximate metric, as

$$(2.7) \quad \forall K \in \mathcal{K}_{\mathcal{L}^k}, \exists \tilde{K} \in \mathcal{K}_{\mathcal{L}; \vec{p}; \vec{D}}^b \text{ s.t. } \|\omega_K - \omega_{\tilde{K}}\| \leq B_{\mathcal{L}^k; \vec{p}; \vec{D}}.$$

For concreteness we can take $M = \mathbb{C}\mathbb{P}^n$, but our arguments here will only be qualitative. Let us consider a two layer network with $p = 2$, so the inputs are the complete set of degree $(k/2, k/2)$ polynomials, and $\theta(z) = z^2$. We then want to understand the dependence $B_{\mathcal{L}^k; p=2; D_1}$ of the upper bound.

The simplest hypothesis is that $B = 0$ once we match the counting of parameters between $\mathcal{K}_{\mathcal{L}^k}$ and $\mathcal{K}_{\mathcal{L}^{k/2}; 2; D_1}^b$, in other words when the width D_1 satisfies

$$(2.8) \quad \binom{k+n}{k}^2 = D_1 \binom{\frac{k}{2} + n}{\frac{k}{2}}^2.$$

This has the asymptotic behaviors

$$(2.9) \quad D_1 \sim \begin{cases} 2^{2n} & \text{for } n \ll k, \\ \left(\frac{n}{k}\right)^k & \text{for } n \gg k. \end{cases}$$

The intuition is that while we need to combine general quadratic functions of the inputs, these are only independent for sufficiently high dimension n ; in low dimension they are linearly dependent.

The qualitative behavior Eq. (2.9) can be checked by simplifying even more, to take the inputs to be vectors X^I in \mathbb{R}^D . For $n = 1$, it follows because one can write any degree $4m$ real univariate polynomial as a difference of squares of two such polynomials.² For the $n \gg k$ limit, let X^I be a basis for

²Write $p = g^2 - h^2 = (g-h)(g+h)$ and then factor p into terms of degree 1 and 2.

the degree $k/2$ polynomials of dimension D , then we are trying to minimize an error of the form

$$(2.10) \quad \mathcal{E} = \left\| \sum_{I,J} f_{IJ} X^I X^J - \sum_{\alpha=1}^{D_1} \left(\sum_I W_{\alpha,I}^{(1)} X^I \right)^2 \right\|.$$

If we use the Frobenius norm, this is minimized by regarding f_{IJ} as a matrix and diagonalizing it. The optimal rank D_1 solution is to write $f = c^t w c$ and keep the top K (in magnitude) values of w . The worst case for approximation is all eigenvalues equal, in which case we need $D_1 = D$ reproducing the second case in Eq. (2.9). In this case one also sees that taking $D' < D$ inputs has the same effect.

In the $n \gg k$ limit we even see that the best rank D_1 approximation to a quadratic on \mathbb{R}^D has worst case relative L_2 error $1 - D_1/D$. If we boldly extrapolate this to general $B_{n,k;D_1}$ by replacing D with the critical value of D_1 satisfying Eq. (2.8), we would conjecture that

$$(2.11) \quad B_{n,k;2;D_1} = \max 0, B_{n,k;1} \left(1 - \frac{D_1 - 1}{D_{match} - 1} \right); \quad D_{match} \equiv \frac{\binom{k+n}{k}^2}{D_0},$$

where $D_0 = \binom{k/2+n}{k/2}^2$ and $B_{n,k;1}$ is the worst case error bound for approximating a degree (k, k) polynomial with a single degree $(k/2, k/2)$ polynomial. In other words, the improvement in approximation ability for a two layer network is linear in the ratio of the number of available parameters to the minimal number of parameters for a universal network, meaning one which can exactly reproduce any degree k polynomial.³

A strong form of the conjecture is that this is even the case if we allow the choice of inputs to depend on the function we are approximating. Thus, if we take the inputs to be a subspace of dimension D_0 but which can depend on K in Eq. (2.7), the worst case error would still be Eq. (2.11). This is not at all obvious, but the idea is motivated by the effect of taking $D' < D$ in Eq. (2.10).

This strong form of the conjecture could be generalized to a multilayer network, and suggests that in low dimensions n , the growth k^{2n} is inevitable. Let us consider a d layer network. By the above, the final layer can approximate a general function in terms of 2^{2n} outputs of independent $d - 1$ -layer networks. For example, we can modify Eq. (1.17) by replacing the layer 1

³The linearity corresponds to the worst case in which f in Eq. (2.10) is the identity. Constraints on the spectrum of f could lead to better bounds.

weights $W_{i,\bar{j}}^{(1),J}$ with independent weights $W_{i,\bar{j}}^{(1),IJ}$ for each choice of subnetwork indexed by I , to get

(2.12)

$$K_{\mathcal{L};\bar{2};D_1,D_2}^b = \log \sum_{1 \leq I \leq D_2} W_I^{(3)} \left(\sum_{1 \leq J \leq D_1} W_J^{(2),I} \left(\sum_i W_{i,\bar{j}}^{(1),IJ} s^i \bar{s}^{\bar{j}} \right)^2 \right)^2.$$

We can repeat this process for each successive layer, producing a tree structured network (an idea introduced in [20]) whose total number of parameters would be $N \sim 2^{2dn} \sim k^{2n}$. Since we did not change the assumptions leading to Eq. (2.8), the asymptotic parameter count is still the same as that for a general polynomial.

Thus the question is, can we improve on this by sharing intermediate results. Since the intermediate width of the tree structured network at layer l is $2^{2n(d-l)}$, in principle there is a lot of scope for this. And in the early layers, the number of independent polynomials is $N \sim (2^l)^{2n} \sim 2^{2nl} \ll 2^{2n(d-l)}$, so for $l < d/2$ one might save on parameters by replacing the many subnetworks with a single network which constructs a complete basis. However without further assumptions this does not buy much. Rather than write the analog of Eq. (2.12), let us just give the modification at layer $d/2$. Replacing the subnetworks with the general sections of degree \sqrt{k} , we take

(2.13)
$$W_{I_{d/2-1}}^{(d/2),I_d I_{d-1} \dots I_{d/2}} \Rightarrow W_{i_1 \dots i_{d/2-1}, \bar{j}_1 \dots \bar{j}_{d/2-1}}^{(d/2),I_d I_{d-1} \dots I_{d/2}}.$$

The 2^{dn} subnetworks each have $\mathcal{O}(2^{dn})$ parameters, so the count is still $\mathcal{O}(k^{2n})$.

Now, it still might be that the subnetworks do not need all of these inputs and one could use fewer parameters. To study this systematically we would want generalizations of Eq. (2.7) to bound the error for jointly approximating $D_i > 1$ functions, given a subspace of the possible inputs. This second dependence (on the subspace) is crucial as we are trying to drastically reduce the number of inputs to the intermediate layers. But according to the strong form of Eq. (2.11), the error would not depend on this choice; it would still be governed by the ratio of the number of parameters to the number for a universal network. Any error bound which decreases with k would require a number of parameters which grows as the same power of k .

Thus, assuming a very strong conjecture, we have argued that to obtain Eq. (2.1) for a general metric we need a number of coefficients growing as k^{2n} . Whether or not we believe this argument, still this would not be too surprising in the worst case. Could the situation be better in special cases, such as the Ricci flat and balanced metrics?

Is there a useful definition of a balanced bihomogeneous metric? Since these are a subset of the embedding metrics, one could try to define a balanced bihomogeneous metric by finding a functional on embeddings whose minimum is the balanced metric. One could then take its minimum over the subset of bihomogeneous metrics.

Indeed, such a functional exists [10],

$$(2.14) \quad \psi(G) = \int_M \nu(z) K_G(s(z)) - \frac{1}{N+1} \operatorname{tr} \log G,$$

where $\nu(z)$ is a measure on M , and K and G are as in Eq. (1.11). To get the balanced metric defined by Eq. (2.6) one must take ν to be the volume form for this metric. But the definition makes sense for a general ν , and on a Calabi–Yau manifold we can use the other canonical volume form (the r.h.s. of Eq. (1.22)) to get a simpler definition which also converges to the Ricci flat metric.

The problem with this is that, as a function of the weights W , Eq. (2.14) is probably nonconvex. If one substitutes in Eq. (1.14) to get an explicit expression for it (analogous to Eq. (2.10)), this will be degree 2^{d-i} in the weights $W^{(i)}$. Thus it probably does not have a unique minimum. This problem is familiar from the neural network literature, but surprisingly enough turns out not to be a serious problem in the applications to machine learning. Still, for present purposes an unambiguous definition would be better. Perhaps there are more geometric definitions of holomorphic network which avoid this problem.

Acknowledgements

We thank Steve Zelditch for suggestions which helped us improve the manuscript.

References

- [1] ANIMASHREE ANANDKUMAR, RONG GE, DANIEL HSU, SHAM M. KAKADE, and MATUS TELGARSKY, *Tensor decompositions for learning latent variable models*, Journal of Machine Learning Research **15** (2014), 2773–2832. [MR3270750](#)
- [2] JAMES H. APPLGATE, MICHAEL R. DOUGLAS, YEKTA GÜRSEL, PETER HUNTER, CHARLES L. SEITZ, and GERALD JAY SUSSMAN, *A digital orrery*, The Use of Supercomputers in Stellar Dynamics, Springer, 1986, pp. 86–95.

- [3] PETER L. BARTLETT, ANDREA MONTANARI, and ALEXANDER RAKHLIN, *Deep learning: a statistical viewpoint*, arXiv:2103.09177 [cs, math, stat] (2021). [MR4295218](#)
- [4] MIKHAIL BELKIN, DANIEL HSU, SIYUAN MA, and SOUMIK MANDAL, *Reconciling modern machine-learning practice and the classical bias–variance trade-off*, Proceedings of the National Academy of Sciences **116** (2019), no. 32, 15849–15854. [MR3997901](#)
- [5] R. BELLMAN, Rand Corporation, and Karreman Mathematics Research Collection, *Dynamic Programming*, Rand Corporation research study, Princeton University Press, 1957.
- [6] VOLKER BRAUN, TAMAZ BRELIDZE, MICHAEL R. DOUGLAS, and BURT A. OVRUT, *Calabi–Yau metrics for quotients and complete intersections*, Journal of High Energy Physics **2008** (2008), no. 05, 080. [MR2411279](#)
- [7] VOLKER BRAUN, TAMAZ BRELIDZE, MICHAEL R. DOUGLAS, and BURT A. OVRUT, *Eigenvalues and eigenfunctions of the scalar Laplace operator on Calabi–Yau manifolds*, Journal of High Energy Physics **2008** (2008), no. 07, 120. [MR2430064](#)
- [8] GIUSEPPE CARLEO and MATTHIAS TROYER, *Solving the Quantum Many-Body Problem with Artificial Neural Networks*, Science **355** (2017), no. 6325, 602–606, arXiv:1606.02318. [MR3642415](#)
- [9] GEORGE CYBENKO, *Approximation by superpositions of a sigmoidal function*, Mathematics of Control, Signals and Systems **2** (1989), no. 4, 303–314. [MR1015670](#)
- [10] S. K. DONALDSON, *Some numerical results in complex differential geometry*, Pure and Applied Mathematics Quarterly **5** (2009), no. 2, 571–618. [MR2508897](#)
- [11] MICHAEL R. DOUGLAS, *From algebraic geometry to machine learning*, Pure and Applied Mathematics Quarterly, to appear (2020). [MR4257595](#)
- [12] MICHAEL R. DOUGLAS, ROBERT L. KARP, SERGIO LUKIC, and RENÉ REINBACHER, *Numerical solution to the Hermitian Yang–Mills equation on the fermat quintic*, Journal of High Energy Physics **2007** (2007), no. 12, 083. [MR2366473](#)
- [13] MICHAEL R. DOUGLAS, ROBERT L. KARP, SERGIO LUKIC, and RENÉ REINBACHER, *Numerical Calabi–Yau metrics*, Journal of Mathematical Physics **49** (2008), no. 3, 032302. [MR2406784](#)

- [14] MICHAEL R. DOUGLAS and SEMYON KLEVTSOV, *Black holes and balanced metrics*, arXiv:[0811.0367](#) [hep-th] (2008). [MR2714117](#)
- [15] MICHAEL R. DOUGLAS and SEMYON KLEVTSOV, *Bergman kernel from path integral*, *Communications in Mathematical Physics* **293** (2009), no. 1, 205–230. [MR2563804](#)
- [16] MICHAEL R. DOUGLAS, SUBRAMANIAN LAKSHMINARASIMHAN, and YIDI QI, *Numerical Calabi–Yau metrics from holomorphic networks*, To appear in the proceedings of MSML 2021 (2020), arXiv:[2012.04797](#).
- [17] MICHAEL R. DOUGLAS, BERNARD SHIFFMAN, and STEVE ZELDITCH, *Critical Points and Supersymmetric Vacua I*, *Communications in Mathematical Physics* **252** (2004), no. 1–3, 325–358 (English). [MR2104882](#)
- [18] MICHAEL R. DOUGLAS, BERNARD SHIFFMAN, and STEVE ZELDITCH, *Critical points and supersymmetric vacua, II: Asymptotics and extremal metrics*, arXiv:math.CV (2004). [MR2219939](#)
- [19] MICHAEL R. DOUGLAS, BERNARD SHIFFMAN, and STEVE ZELDITCH, *Critical Points and Supersymmetric Vacua, III: String/M Models*, *Communications in Mathematical Physics* **265** (2006), no. 3, 617–671 (English). [MR2231684](#)
- [20] WEINAN E and STEPHAN WOJTOWYTSCH, *On the Banach spaces associated with multi-layer ReLU networks: Function representation, approximation theory and gradient descent dynamics*, arXiv:[2007.15623](#) [cs, math, stat] (2020).
- [21] JIEQUN HAN, ARNULF JENTZEN, and WEINAN E, *Solving high-dimensional partial differential equations using deep learning*, *Proceedings of the National Academy of Sciences* **115** (2018), no. 34, 8505–8510. [MR3847747](#)
- [22] MATTHEW HEADRICK and ALI NASSAR, *Energy functionals for Calabi–Yau metrics*, *Advances in Theoretical and Mathematical Physics* **17** (2013), no. 5, 867–902. [MR3251826](#)
- [23] DANIEL HUYBRECHTS, *Complex geometry: an introduction*, Springer Science & Business Media, 2005. [MR2093043](#)
- [24] BERNARD SHIFFMAN and STEVE ZELDITCH, *Distribution of zeros of random and quantum chaotic sections of positive line bundles*, *Communications in Mathematical Physics* **200** (1999), no. 3, 661–683, arXiv:math/9803052. [MR1675133](#)

- [25] JUSTIN SIRIGNANO and KONSTANTINOS SPILIOPOULOS, *DGM: A deep learning algorithm for solving partial differential equations*, Journal of Computational Physics **375** (2018), 1339–1364, arXiv:[1708.07469](#). [MR3874585](#)
- [26] JAMES STOKES and JOHN TERILLA, *Probabilistic modeling with matrix product states*, Entropy **21** (2019), no. 12, 1236. [MR4055725](#)
- [27] GANG TIAN et al., *On a set of polarized kähler metrics on algebraic manifolds*, Journal of Differential Geometry **32** (1990), no. 1, 99–130. [MR1064867](#)
- [28] STEVE ZELDITCH and PENG ZHOU, *Pointwise weyl law for partial bergman kernels*, Algebraic and analytic microlocal analysis, Springer, 2013, pp. 589–634. [MR3903328](#)
- [29] CHIYUAN ZHANG, SAMY BENGIO, MORITZ HARDT, BENJAMIN RECHT, and ORIOL VINYALS, *Understanding deep learning requires rethinking generalization* (2016), arXiv:[1611.03530](#).

Michael R. Douglas

CMSA

Harvard University

Cambridge 02138

USA

E-mail: mdouglas@cmsa.fas.harvard.edu