# Control in Deep Learning and Neural Networks ${ }^{1}$ 

Enrique Zuazua<br>enrique.zuazua@fau.de

## FAU - AvH

April 13, 2020

[^0]
## What is Machine Learning?



Figure: Source: https://de.mathworks.com/help/stats/machine-learning-in-matlab.html

- Interested in approximating a function $f: \mathbb{R}^{d} \rightarrow \mathbb{R}^{m}$, of some class $\left(C^{0}, L^{1} \ldots\right)$, which we don't know explicitly.
- We have data: its values ${ }^{2}\left\{\vec{y}_{i}\right\}_{i=1}^{S} \in\left(\mathbb{R}^{m}\right)^{S}$ at $S$ distinct points $\left\{\vec{x}_{i}\right\}_{i=1}^{S} \in\left(\mathbb{R}^{d}\right)^{S}$.
- Generally split the $S$ data points into training data $\left\{\vec{x}_{i}, \vec{y}_{i}\right\}_{i=1}^{N}$ and testing data $\left\{\vec{x}_{i}, \vec{y}_{i}\right\}_{i=N+1}^{S}$.
- Machine learning consists in:
(1) Propose a candidate approximation $f_{\Theta}(\cdot): \mathbb{R}^{d} \rightarrow \mathbb{R}^{m}$, depending on parameters $\Theta$ and some hyper-parameters $L \geq 1,\left\{N_{k}\right\}$;
(2) Tune $\Theta$ as to minimize the empirical risk ${ }^{3}$

$$
\sum_{i=1}^{N} \ell\left(f_{\Theta}\left(\vec{x}_{i}\right), \vec{y}_{i}\right),
$$

where $\ell \geq 0, \ell(x, x)=0$ (e.g. $\ell(x, y)=|x-y|^{2}$ ). This is called training.
(3) Check if test error $\sum_{i=N+1}^{S} \ell\left(f_{\Theta}\left(\vec{x}_{i}\right), \vec{y}_{i}\right)$ is small. This is called generalization.

[^1]
## Model validation

Question: After training, how do we know if the optimized neural net is good? Answer: Check if the test error $\sum_{i=N+1}^{S} \ell\left(f_{\Theta}\left(\vec{x}_{i}\right), \vec{y}_{i}\right)$ is small.




Figure: Underfitting, good generalization, and overfitting. We wish to recover the function $f(x)=\cos \left(\frac{3}{2} \pi x\right)$ (blue) on $(0,1)$ from $S=20$ noisy data samples. Constructed approximations using $N=12$ training data, while the remaining 8 samples are used for testing the results. The most complicated model (right) is not necessarily the best (Occam's razor) This is related to the Runge phenomenon.

## Example (Binary classification)

- Assume that all the points in some $\Omega \subset \mathbb{R}^{d}$ are either blue or red in a not too incoherent way.
- Goal: Given an arbitrary blue or red point $\vec{x} \in \Omega \subset \mathbb{R}^{d}$, say what is the color of $\vec{x}$ (for simplicity, 0 is blue and 1 is red).
- Relation point-color represented by a function $f: \mathbb{R}^{d} \rightarrow\{0,1\}$, which we don't know explicitly.
- We know the colors $\left\{\vec{y}_{i}\right\}_{i=1}^{M} \in\{0,1\}^{M}$ of $M$ points $\left\{\vec{x}_{i}\right\}_{i=1}^{M} \in\left(\mathbb{R}^{d}\right)^{M}$. We take $N$ training data.
- Machine learning consists in:
(1) Proposing a candidate approximation $f_{\Theta}(; \cdot): \mathbb{R}^{d} \rightarrow \mathbb{R}$, depending on parameters $\Theta$ and some fixed hyper-parameters $L \geq 0,\left\{N_{k}\right\}_{k=0}^{L}$;
(2) Optimizing $\Theta$ so that $\sum_{i=1}^{N}\left|f_{\Theta}\left(; \vec{x}_{i}\right)-\vec{y}_{i}\right|^{2}$ is small
( Compute optimizer $\widehat{\Theta}$; then given any $\vec{x} \in \Omega$,

$$
F(\vec{x}):=\mathbb{1}_{\left\{f_{L}\left(\hat{\Theta}_{; \cdot}\right) \geq \frac{1}{2}\right\}}(\vec{x})
$$

will yield the color.


Figure: We have the training data points $\left\{\vec{x}_{i}\right\}_{i=1}^{4000} \in\left([0,2]^{2}\right)^{4000}$ on the left, each having respective colors $\left\{y_{i}\right\}_{i=1}^{4000} \in\{0,1\}^{4000}$. The data are arranged in a chess-like pattern. On the right, we plot the level-sets of $F(\vec{x}):=\mathbb{1}_{\left\{f_{L}\left(\widehat{\Theta}^{\prime} \cdot \cdot\right) \geq \frac{1}{2}\right\}}(\vec{x})$ for $\vec{x} \in[-0.5,2.5]^{2}$ which separates the points with $80 \%$ accuracy. We used a neural network with 1 hidden layer and 5 components.

Remark: This classification procedure can be used for classification in image processing (e.g. cat-dog recognition, handwritten digit recognition).


Figure: Classifying handwritten digits from the MNIST dataset using a neural network. The inputs $\left\{\vec{x}_{i}\right\}_{i=1}^{50} \in\left(\mathbb{R}^{28 \times 28}\right)^{50}$ are pixelated images, while the labels $\left\{\vec{y}_{i}\right\} \in\left(\mathbb{R}^{10}\right)^{50}$
correspond to numbers from 0 to 9 , each one identified with an element of the canonical basis $\left\{e_{k}\right\}_{k=1}^{10}$ of $\mathbb{R}^{10}$. We used a network with 2 hidden layers (one with 256 and a second with 64 neurons). Illustrated are the transitions through the layers when the datum is the digit $5 \simeq e_{6}$. The three rightmost plots indicate the three layers of the model, namely the middle 2 are the values of $\sigma\left(A^{0} z^{0}+b^{0}\right) \in[0,1]^{256}$ and $\sigma\left(A^{1} z^{1}+b^{1}\right) \in[0,1]^{64}$ (reshaped in a matrix). The brighter a pixel, the closer the value is to 1

## Example ("Learning" $\sin (x)$ )

- Goal: Given $M=4000$ noisy values of $\sin (x)$ on $(0,6 \pi)$, deduce the original function $\sin (x)$.
- We thus know the values $\left\{y_{i}\right\}_{i=1}^{M} \in[0,1]^{M}$ of $M$ points $\left\{x_{i}\right\}_{i=1}^{M} \in(0,6 \pi)^{M}$. We take $N$ training data.
- Machine learning consists in:
(1) Proposing a candidate approximation $f_{\Theta}(; \cdot): \mathbb{R}^{d} \rightarrow \mathbb{R}$, depending on parameters $\Theta$ and some fixed hyper-parameters $L \geq 0,\left\{N_{k}\right\}_{k=0}^{L}$;
(2) Optimizing $\Theta$ so that $\sum_{i=1}^{N}\left|f_{\Theta}\left(; x_{i}\right)-y_{i}\right|^{2}$ is small
(3) Compute optimizer $\widehat{\Theta}$; then given any $\vec{x} \in(0,1), f_{L}(\widehat{\Theta} ; \cdot)$ should approximate $f(x)=\sin (x)$.
- Remark: The procedure can be used to approximate more complicated functions (including solutions of PDEs, etc.)
- Thus the difference between classification and regression is in the co-domain of the function we wish to approximate. In classification, the co-domain is discrete, contrary to regression, where the function has continuous values.


Figure: The $N$ training data in blue and the trained neural network solution in orange (left), with a zoom in a subinterval (right). We used a neural network with 2 hidden layers consisting of 4 components each.

## Neural networks

## Activation function

Given $\sigma \in C^{0}(\mathbb{R} ; \mathbb{R})$, we denote/define componentwise for $z \in \mathbb{R}^{d}$,

$$
\sigma(z):=\left[\begin{array}{c}
\sigma\left(z_{1}\right)  \tag{1}\\
\sigma\left(z_{2}\right) \\
\vdots \\
\sigma\left(z_{d}\right)
\end{array}\right] \in \mathbb{R}^{d}
$$




## Definition (Neural network)

Let $L \geq 0, d \geq 1, m \geq 1$ and $\left\{N_{k}\right\}_{k=1}^{L} \in \mathbb{N}^{L}$ be given, with $N_{0}:=d$ and $N_{L+1}:=m$.
A neural network with $L$ hidden layers, input dimension $d$, and output dimension $m$, is a parameter-dependent map

$$
\begin{aligned}
f_{\Theta}(; \cdot): \mathbb{R}^{d} & \longrightarrow \mathbb{R}^{m} \\
\vec{x} & \longmapsto\left(\varphi \circ \Lambda_{L+1} \circ \sigma \circ \Lambda_{L} \circ \sigma \circ \ldots \circ \sigma \circ \Lambda_{1}\right)(\vec{x})
\end{aligned}
$$

where $\Lambda_{k+1} \vec{z}:=A^{k} \vec{z}+b^{k}$ for all $k \in\{0, \ldots, L\}$ and $\vec{z} \in \mathbb{R}^{N_{k}}$. The sequence of matrix-vector pairs

$$
\Theta=\left\{\left(A^{k}, b^{k}\right)\right\}_{k=0}^{L}
$$

where

$$
A^{k} \in \mathbb{R}^{N_{k+1} \times N_{k}} \quad \text { and } \quad b^{k} \in \mathbb{R}^{N_{k+1}} \quad \text { for } k \in\{0, \ldots, L\}
$$

are parameters, while $\sigma \in C^{0}(\mathbb{R})$ and $\varphi \in C^{0}(\mathbb{R})$ are two fixed functions.
Comment: Definition of $f_{L}$ hard to read; can be rewritten equivalently as in Slide 15 below.

## Cleaner definition

Consider setting of Slide 14. Given a data point $\vec{x}_{i} \in \mathbb{R}^{d}$ and parameters $\Theta$, a neural network writes as

$$
\begin{equation*}
f_{\Theta}(; \vec{x})=\varphi\left(A^{L} z^{L}+b^{L}\right) \tag{2}
\end{equation*}
$$

where $z^{L}=z_{i}^{L} \in \mathbb{R}^{m}$ being given by the scheme

$$
\left\{\begin{array}{l}
z^{k+1}=\sigma\left(A^{k} z^{k}+b^{k}\right) \quad \text { for } k=0, \ldots, L-1  \tag{3}\\
z^{0}=\vec{x}_{i} \in \mathbb{R}^{d} .
\end{array}\right.
$$

This works for $L \geq 1$. Observe that $z^{k} \in \mathbb{R}^{N_{k}}$ for $k \in\{0, \ldots, L\}$.

- We can thus define a neural network equivalently by means of (??)-(??).
- (??) is often called the architecture of the neural network
- In principle, we could add different terms in the scheme (??) to obtain other architectures (Slide 49).
- A neural network with $L=1$ is called a one hidden layer (or shallow) network.
- $\sigma \in C^{0}(\mathbb{R})$ generally as Slide 13 (+monotonically increasing)
- $\Lambda_{L+1}$ projects data to arrival space $\mathbb{R}^{m}$, while $\varphi$ can compress it within a desired interval (usually $[0,1]$, to interpret as probabilities).
- In binary classification (i.e. data $\left\{\vec{y}_{i}\right\}_{i=1}^{N} \in\{0,1\}^{N}$ like our red-blue example of Slide 7), we have $\varphi \equiv \sigma$.
In this case, a network with 1 (hidden) layer reads

$$
f(\Theta, \cdot): \vec{x} \mapsto \sigma\left(A^{1} \sigma\left(A^{0} \vec{x}+b^{0}\right)+b^{1}\right)
$$

- Otherwise, depends on the application:
- for regression, we use $\varphi(\vec{x})=\vec{x}$, with $\vec{x} \in \mathbb{R}^{m}$ (no constraints on arrival set).
- for $m>2$ colors we use

$$
\varphi(\vec{x})=\frac{1}{\sum_{j=1}^{m} \exp \left(x_{j}\right)}\left[\begin{array}{c}
\exp \left(x_{1}\right) \\
\vdots \\
\exp \left(x_{m}\right)
\end{array}\right] \in \mathbb{R}^{m} .
$$

Let us elaborate. We want to classify 3 colors (for instance blue, red and green), and we have data of the form $\left\{\vec{x}_{i}\right\}_{i=1}^{N} \in\left(\mathbb{R}^{d}\right)^{N}$. The corresponding labels are $\left\{\vec{y}_{i}\right\}_{i=1}^{N} \in\left(\mathbb{R}^{3}\right)^{N}$ as we identify an element of the canonical basis $\left\{e_{1}, e_{2}, e_{3}\right\}$ of $\mathbb{R}^{3}$ with each of the 3 colors. The function $\varphi$ will take the output $z^{L} \in \mathbb{R}^{3}$ of the last hidden layer of the network, and normalizes it into a probability distribution consisting of 3 probabilities proportional to the exponentials of $\left\{z_{j}^{L}\right\}_{j=1}^{3}$ (i.e. $\left.\varphi\left(z^{L}\right) \in(0,1)^{3}\right)$. Such $\varphi$ is called softmax.


Here $L=2$, input $d=2$ and output $m=1$. Hidden layers are the layers 1 and 2 . The width of the layers $j(0 \leq j \leq 3)$ are $N_{0}=2, N_{1}=4, N_{2}=3, N_{3}=1$, and $z^{j}=\left(z_{1}^{j}, \ldots, z_{N_{j}}^{j}\right)^{\top}$.

Let $L=0$ (zero hidden layers), $d=2$ and $m=1$. Then

$$
f_{\Theta}(, \vec{x})=\varphi\left(\sum_{j=1}^{2} A_{j}^{0} x_{j}+b^{0}\right) \in \mathbb{R}
$$

for $\vec{x}=\left[\begin{array}{l}x_{1} \\ x_{2}\end{array}\right] \in \mathbb{R}^{2}$, where $A^{0}=\left[\begin{array}{ll}A_{1}^{0} & A_{2}^{0}\end{array}\right] \in \mathbb{R}^{1 \times 2}$ (1 row 2 columns) and $b^{0} \in \mathbb{R}$.


Let $L=0$. Then $f_{\Theta}(, \vec{x})=\varphi\left(A^{0} \vec{x}+b^{0}\right)$ for $\vec{x} \in \mathbb{R}^{d}$ with $A^{0} \in \mathbb{R}^{m \times d}$ and $b \in \mathbb{R}^{m}$. Recall by convention that

$$
\varphi\left(A_{0} \vec{x}+b_{0}\right):=\left[\begin{array}{c}
\varphi\left(\sum_{j=1}^{d} A_{1, j}^{0} x_{j}+b_{1}^{0}\right) \\
\varphi\left(\sum_{j=1}^{d} A_{2, j}^{0} x_{j}+b_{2}^{0}\right) \\
\vdots \\
\varphi\left(\sum_{j=1}^{d} A_{m, j}^{0} x_{j}+b_{m}^{0}\right)
\end{array}\right] \in \mathbb{R}^{m} .
$$

The $\mathrm{N}=12$ data points


Figure: We consider 4 blue and 8 red data points.

- Objective: predict the colors of other points in the interval $\longrightarrow$ separate the interval in a set of blue and a set of red points, using the data.
- We consider a (trained) one hidden layer network with one component:

$$
\begin{equation*}
f_{1}(\Theta, x)=\sigma\left(A^{1} \sigma\left(A^{0} x+b^{0}\right)+b^{1}\right) \in \mathbb{R} \quad \text { for } x \in \mathbb{R} \tag{4}
\end{equation*}
$$

where $\left(A^{0}, A^{1}, b^{0}, b^{1}\right) \in \mathbb{R}$ are all scalars.

- In all examples, $\sigma(x)=\frac{1}{1+e^{-x}}$ (see Slide 13).


## Example 1: Steps $1 \& 2$ - Stretch, then compress



Figure: Step 1: First apply $\Lambda_{1} x=A^{0} x_{i}+b^{0} \in \mathbb{R}$ (dilating and translating) to each $x_{i}$


Figure: Step 2: Then apply $\sigma$ to Step 1 to compress in $[0,1]$ (first try to obtain a separation)

Example 1: Steps 3 \& 4 - Stretch, compress again


Figure: Step 3: Stretch again by $\Lambda_{1} z^{1}=A^{1} z^{1}+b^{1} \in \mathbb{R}$


Figure: Compress Step 3 by $\sigma$.


Figure: On the right, we have the function which takes any point in the interval and maps to 0 (blue) or 1 (red). On the left, we see its level sets. There was not "enough room" to separate the points.

## Example 2: Same data, 2 components in layer

- 1 component in the hidden layer (i.e. $A^{0} \in \mathbb{R}$ ) doesn't separate well.
- How about 2:

$$
\begin{equation*}
f_{1}(\Theta, x)=\sigma\left(A^{1} \sigma\left(A^{0} x+b^{0}\right)+b^{1}\right) \in \mathbb{R} \quad \text { for } x \in \mathbb{R} \tag{5}
\end{equation*}
$$

but now $A^{0} \in \mathbb{R}^{2 \times 1}, A^{1} \in \mathbb{R}^{1 \times 2}$, and $b^{0} \in \mathbb{R}^{2}, b^{1} \in \mathbb{R}$.

The $\mathrm{N}=12$ data points


Figure: We consider the same 4 blue and 8 red data points.


Figure: Embed the data in $\mathbb{R}^{2}$ via $\Lambda_{1}: \mathbb{R} \rightarrow \mathbb{R}^{2}$ as $A^{0} \in \mathbb{R}^{2 \times 1}$; dilates and translates the points along a diagonal in $\mathbb{R}^{2}$. Applying $\sigma$ "curbs" the points (nonlinearity and monotonicity is important here), making it easier to separate by a line.

## Example 2: Separate the data with a line



Figure: We see how the projection $x \mapsto A^{1} x+b^{1}$ onto $\mathbb{R}$ and compression by $\sigma$ would give separation (see right). Important part was done in the slide before, as working in $\mathbb{R}^{2}$ means that we can separate by a line.


Figure: We project onto $\mathbb{R}$; the angle allows to switch the order.


Figure: Here $\sigma$ is only used as a compression to remain in $[0,1]$ (obtain a probability)


Figure: We see the level sets (left) of the classifier function (right). The function $f_{1}(x)=\sigma\left(A^{1} \sigma\left(A^{0} x+b^{0}\right)+b^{1}\right)$ is with values in $[0,1]$, can be seen as probability distribution; to obtain clear-cut $0-1$ values, we removed the transition zone.

## Example 3: 3 components in layer

We try adding another component in the hidden layer, i.e. take $A^{0} \in \mathbb{R}^{3 \times 1}$.


Figure: We add a component to the hidden layer. Means that the data will be embedded in $\mathbb{R}^{3}$ after the first linear transformation.

## Example 3: Step 1 - Embed in $\mathbb{R}^{3}$

## Linear transformation nb. 1: $A_{1} z^{0}=A^{0} z^{0}+b^{0}$

Linear transformation nb.1: $\Lambda_{1} z^{0}=A^{0} z^{0}+b^{0}$


Figure: The linear transformation stretches the data onto a diagonal line in $\mathbb{R}^{3}$.

## Example 3: Step 2 - Apply $\sigma$ to curb points

Hidden layer nb.1: $z^{1}=\sigma\left(A^{0} z^{0}+b^{0}\right)$
Hidden layer nb.1: $z^{1}=\sigma\left(A^{0} z^{0}+b^{0}\right)$


Figure: We apply the nonlinearity $\sigma$ to each component, to the effect of curbing the data (as well as compressing within $\left.[0,1]^{3}\right)$.

## Example 3: We can separate in $\mathbb{R}^{3}$ with a hyperplane

Hidden layer nb.1: $z^{1}=\sigma\left(A^{0} z^{0}+b^{0}\right)$


Separation before output layer: $z^{2}=\sigma\left(A^{1} z^{1}+b^{1}\right)$
K

Figure: The curbed points may be separated by a hyperplane $\left(\left\{x \in \mathbb{R}^{3}: \sigma\left(A^{1} x+b^{1}\right)=0.5\right\}\right.$ in blue), much like the simpler data of Example 2 can be separated by a line. We see this on the right.

## Example 3: Steps 3 \& 4 - Project and compress



Figure: We see that the data is correctly separated after the projection onto $\mathbb{R}$.


## Example 3: Final classifier




- Augmenting the dimension of the input data is similar to non-intersection phenomena in geometrical optics. Indeed, rays in a plane (2d) starting from different points can intersect, while this does not occur in a cylinder (3d).


Figure: Rays may intercept in $2 d$ (left) but not in $3 d$ (right).

## Universal approximation theorem(s)

- Theorems generally of the form: The class of neural networks is dense with respect to some topology in some function class $\mathcal{C}$.
- Interpretation: given $f \in \mathcal{C}$ and $\varepsilon>0$, there exists a neural network $f_{L(\varepsilon)}$ with $L(\varepsilon)>0$ hidden layers, each layer consisting of $N_{k}(\varepsilon)$ components, as well as parameters $\Theta(\varepsilon)$, in a way that $\left\|f-f_{\Theta}((\varepsilon))\right\|_{\mathcal{C}}<\varepsilon$;
- Sometimes, $L>0$ can be fixed too (see Slide 42). Then, more and more components in the hidden layers are needed to approximate $f$;
- The theorem does not say what these parameters are, or how to find them. It can basically be seen as an existence result for neural networks.


## Theorem (Cybenko '89, MCSS)

Let $\sigma: \mathbb{R} \rightarrow \mathbb{R}$ be a nonconstant, bounded and continuous function.
Let $d \geq 1$ and $L=1$.
For any $\varepsilon>0$ and any $f \in C^{0}\left([0,1]^{d}\right)$, there exists $N_{1} \in \mathbb{N}$, coefficients $A^{0} \in \mathbb{R}^{N_{1} \times d}$, $A^{1} \in \mathbb{R}^{1 \times N_{1}}$ and $b^{0} \in \mathbb{R}^{N_{1}}$ such that

$$
f_{L}(\vec{x})=A^{1} \sigma\left(A^{0} \vec{x}+b^{0}\right)
$$

satisfies

$$
\sup _{\vec{x} \in[0,1]^{d}}\left|f(\vec{x})-f_{L}(\vec{x})\right|<\varepsilon
$$

- This theorem corresponds to $L=1$ (one hidden layer) and $\varphi(x)=x$ (regression) in Definition Slide 14-15.
- Ingredients of the proof: Contradiction argument + Hahn Banach + Riesz-Representation + properties of $\sigma$.
- $\sigma(x)=\frac{1}{1+e^{-x}}$ fits the assumptions for instance.

We actually also have the following improved result.

## Theorem (Poggio et al. '17)

Let $\sigma \in C^{\infty}(\mathbb{R})$ be bounded and not a polynomial.
Let $d \geq 1$ and $L=1$.
For any $\varepsilon>0$ and any $f \in C^{\kappa}\left([0,1]^{d}\right)$ with $\kappa \geq 1$ there exist coefficients $A^{0} \in \mathbb{R}^{N_{1} \times d}, A^{1} \in \mathbb{R}^{1 \times N_{1}}$ and $b^{0} \in \mathbb{R}^{N_{1}}$ with

$$
N_{1}=\mathcal{O}\left(\varepsilon^{-d / \kappa}\right)
$$

such that

$$
f_{L}(\vec{x})=A^{1} \sigma\left(A^{0} \vec{x}+b^{0}\right)
$$

satisfies

$$
\sup _{\vec{x} \in[0,1]^{d}}\left|f(\vec{x})-f_{L}(\vec{x})\right|<\varepsilon .
$$

- Theorem implies that one can use a network with one hidden layer for approximating $C^{\kappa}$ functions. However, in general the number $N_{1}$ of neurons needed for a fixed approximation accuracy $\varepsilon>0$ grows exponen- tially in $d$, and so does the number of parameters in $A^{0}, A^{1}, b^{0}$. This means that the storage requirement as well as the effort to determine $A^{0}, A^{1}, b^{0}$ easily exeeds all reasonable bounds already for moderate dimensions $d$. Hence, this approach suffers from the curse of dimensionality.

More recent results for $\sigma(x)=\max (x, 0)$ include

## Theorem (Hanin '17)

Let $d \geq 1$ and let $f:[0,1]^{d} \rightarrow \mathbb{R}$ be a positive and continuous function with $\|f\|_{\infty}=1$. Then for any $\varepsilon>0$, there exists a neural network $f_{L}$ with $\operatorname{ReLU}$ activation of depth

$$
L=\frac{2 d!}{w_{f}(\varepsilon)^{d}}
$$

and width $\max _{k} N_{k} \leq d+3$ such that

$$
\left\|f-f_{L}\right\|_{\infty} \leq \varepsilon
$$

Here $w_{f}: \delta \mapsto \sup \{|f(x)-f(y)|:|x-y| \leq \delta\}$ denotes the modulus of continuity of $f$.

- Improved results for functions $f: \mathbb{R}^{d} \rightarrow \mathbb{R}^{m}$ can be found in Müller '20.
- Results for other function classes $\left(L^{p}, W^{k, p}\right)$ exist as well.
- It is true that any (say continuous) function can be approximated arbitrarily close both by a neural network and a polynomial. True for a lot of constructs (Galerkin, etc.).
- In machine learning one doesn't really want a function that fits through the data perfectly (overfitting). Rather, one wants something that fits well, but also probably works for points that have not been seen yet.
- Neural networks are (empirically) best-in-breed at solving a very specific kind of problem: computing a function $f: X \rightarrow Y$ given the values of $f$ on a large but finite subset $A \subset X$. Typically $Y$ (the space of "labels") is finite and small relative to $A$.
- As it turns out, neural networks can "learn" $f$ so well that they can produce points in $f^{-1}(y)$ 's near a prescribed $x \in X$, which is why they are good at image / language generation ${ }^{4}$.

[^2]
## Training

- We are given data $\left\{\left(\vec{x}_{i}, \vec{y}_{i}\right)\right\}_{j=1}^{N} \in\left(\mathbb{R}^{d}\right)^{N} \times\left(\mathbb{R}^{m}\right)^{N}$;
- Training consists in solving the optimization problem:

$$
\begin{equation*}
\min _{\Theta=\left\{\left(A^{k}, b^{k}\right)\right\}_{k=0}^{L}} \sum_{i=1}^{N}\left|\vec{y}_{i}-f_{\Theta}\left(; \vec{x}_{i}\right)\right|^{2}+\epsilon \mathcal{R}(\Theta) ; \tag{6}
\end{equation*}
$$

$\epsilon>0$ is a penalization/regularisation parameter, $\mathcal{R}$ convex;

- Non-convex optimization problem because of $f_{L}$;
- Existence of a minimizer may be shown by a direct method ( $\sigma \in C^{0}$ );

Once training is done:

- Minimizer $\widehat{\Theta}$; if $\left\{\vec{y}_{i}\right\}_{i=1}^{N} \in\{0,1\}^{N}$ (classification), we set

$$
F(\vec{x}):=\mathbb{1}_{\left\{f_{L}\left(\widehat{\Theta}_{; \cdot}\right) \geq \frac{1}{2}\right\}}(\vec{x}) \quad \forall \vec{x} \in \mathbb{R}^{d}
$$

and we are done. Otherwise $F(\vec{x}):=f_{L}(\widehat{\Theta}, \vec{x})$ (regression).

The functional to be minimized is of the form

$$
\begin{equation*}
J(\Theta)=\sum_{i=1}^{N} J_{i}(\Theta) \tag{7}
\end{equation*}
$$

We could do gradient descent:

$$
\Theta^{n+1}:=\Theta^{n}-\eta \nabla J\left(\Theta^{n}\right)
$$

$\eta$ is step-size. But often $N \gg 1$ ( $N=10^{3}$ or more).

- Stochastic gradient descent (Robbins-Monro 50s, Bottou et al, SIREV '18):
(1) pick $i \in\{1, \ldots, N\}$ uniformly at random
(3) $\Theta^{n+1}:=\Theta^{n}-\eta \nabla J_{i}\left(\Theta^{n}\right)$
- Mini-batch GD can also be considered (pick a subset of data instead of just one point)
- Use chain rule and adjoints to compute these gradients ("backpropagation")
- Issues: might not converge to global minimizer; also how does one initialize the weights in the iteration (usually done at random)?

Recall that for a set of parameters $\Theta=\left\{\left(A^{k}, b^{k}\right)\right\}_{k=0}^{L}$ the neural network writes as

$$
f_{\Theta}(; \vec{x}):=\varphi\left(\Lambda_{L+1} \circ \sigma \circ \Lambda_{L} \circ \sigma \ldots \circ \sigma \circ \Lambda_{1}\right)(\vec{x})=\varphi\left(A^{L} z^{L}+b^{L}\right)
$$

where for $\vec{x} \in \mathbb{R}^{d}$,

$$
\begin{cases}z^{k+1} & =\sigma\left(A^{k} z^{k}+b^{k}\right) \quad \text { for } k=0, \ldots, L-1  \tag{8}\\ z^{0} & =\vec{x}\end{cases}
$$

- We recognize a discrete-time dynamical system;
- Training can thus be rewritten as a constrained optimization problem:

$$
\min _{\Theta=\left\{\left(A^{k}, b^{k}\right)\right\}_{k=0}^{L}} \sum_{i=1}^{N}\left|\vec{y}_{i}-\varphi\left(A^{L} z^{L}+b^{L}\right)\right|^{2}+\epsilon \mathcal{R}(\Theta)
$$

with $z^{L}=z^{L}\left(\Theta, \vec{x}_{i}\right)$, subject to (??) with $z^{0}=\vec{x}_{i}$.

- In this case, deep learning may be seen as discretized optimal control (the parameters $\Theta$ play the role of controls).

Back to neural networks. Let $L \geq 1$ (at least 1 hidden layer).
A different architecture (Weinan E et al. '18): given any datum $\vec{x}_{i} \in \mathbb{R}^{d}$ for some $i=1, \ldots, N$,

$$
\begin{cases}z^{k+1} & =z^{k}+\Delta t \sigma\left(A^{k} z^{k}+b^{k}\right) \quad \text { for } k=0, \ldots L-1  \tag{9}\\ z^{0} & =\vec{x}_{i} \in \mathbb{R}^{d}\end{cases}
$$

with $\Delta t=\frac{T}{L-1}$ and $T>0$ given time horizon.

- requires same widths $N_{k}$ at each layer $k \in\{0, \ldots, L\}$, but we can extend the data by 0 .
- Recognize explicit Euler scheme for ODE

$$
\left\{\begin{array}{l}
z^{\prime}(t)=\sigma(A(t) z(t)+b(t)) \quad \text { for } t \in(0, T)  \tag{10}\\
z(0)=\vec{x}_{i} \in \mathbb{R}^{d}
\end{array}\right.
$$

- $z(t)=z_{i}(t)$, but $A, b$ should be independent of the data.
- Universal approximation theorems exist also for ResNet architectures (see Lin \& Jegelka '18)
- Can be advantageous to consider the continuous-time optimal control problem:

$$
\begin{equation*}
\inf _{u(t) \in U,(\alpha, \beta)} \sum_{i=1}^{N}\left|\vec{y}_{i}-\varphi(\alpha z(T)+\beta)\right|^{2}+\epsilon \mathcal{R}(u) \tag{11}
\end{equation*}
$$

where $z=z_{i}$ solves

$$
\left\{\begin{array}{l}
z^{\prime}(t)=F(u(t), z(t)) \quad \text { in }(0, T)  \tag{12}\\
z(0)=\vec{x}_{i} \in \mathbb{R}^{d}
\end{array}\right.
$$

- Recall that often $\varphi \equiv \sigma$ (classification) or $\varphi(x)=x$ (regression).
- $u(t)=(A(t), b(t)) \in\left(\mathbb{R}^{d \times d} \times \mathbb{R}^{d}\right)^{2}$ and $F(u(t), z(t))=\sigma(A(t) z(t)+b(t))$.
- $\sigma \in \operatorname{Lip}(\mathbb{R})$; existence of a minimizer see Trélat '05; here
$U \subset L^{\infty}\left(0, T ;\left(\mathbb{R}^{d \times d} \times \mathbb{R}^{d}\right)^{2}\right)$
- Easier to write optimality system (Weinan E et al. '18); we can use different algorithms for training (shooting method);
- Other schemes (Runge-Kutta) for time-discretization to obtain new architectures


Figure: We see that the trajectories of the dynamical system $z_{i}^{\prime}(t)=\sigma\left(A(t) z_{i}(t)+b(t)\right)$ with $z_{i}(0)=\left[x_{i}, 0\right]$ separate the points $x_{i}$.


Figure: We see that the trajectories of the dynamical system $z_{i}^{\prime}(t)=\sigma\left(A(t) z_{i}(t)+b(t)\right)$ with $z_{i}(0)=\left[x_{i}, 0\right]$ separate the points $x_{i}$. Here we ran the optimisation simulation on the same dataset.





Figure: We ran the simulation for solving the optimisation problem 10 times on the same dataset as previous slide. The left figures are the losses in each run, with the mean shown in bold. On the right, we see the mean in bold, and the deviation is shaded. In blue, we have the ResNet neural network, while in green we solve the continuous ODEs with Runge-Kutta.

Recall again that for parameters $\Theta=\left\{\left(A^{k}, b^{k}\right)\right\}_{k=0}^{L}$ the neural network writes ${ }^{5}$

$$
f_{\Theta}(; \vec{x}):=\varphi\left(\Lambda_{L+1} \circ \sigma \circ \Lambda_{L} \circ \sigma \ldots \circ \sigma \circ \Lambda_{2} \circ \sigma \circ \Lambda_{1}\right)(\vec{x}) .
$$

We can also (in addition to Slide 49) write $f_{\Theta}(; \vec{x})=\varphi\left(z^{L}\right)$, where now

$$
\begin{cases}z^{k+1}=A^{k+1} \sigma\left(z^{k}\right)+b^{k+1} & \text { for } k=0, \ldots, L-1  \tag{13}\\ z^{0}=A^{0} \vec{x}+b^{0} & \end{cases}
$$

When $L \gg 1$, WLOG ${ }^{6}$ assume $b^{0}=0$ and $A^{0}=E: \mathbb{R}^{d} \rightarrow \mathbb{R}^{N_{1}}$ fixed embedding map. Then relabel the parameters $\Theta$ by shifting the indexes $k \leftarrow k+1$ and write

$$
\left\{\begin{array}{l}
z^{k+1}=A^{k} \sigma\left(z^{k}\right)+b^{k} \quad \text { for } k=0, \ldots, L-1  \tag{14}\\
z^{0}=E \vec{x}
\end{array}\right.
$$

[^3]Let $\left\{\left(\vec{x}_{i}, \vec{y}_{i}\right)\right\}_{i=1}^{N} \in\left(\mathbb{R}^{d}\right)^{N} \times\left(\mathbb{R}^{m}\right)^{N}$ be given training data. Let $m>1$. Now assume that $N_{k}=m$ for all $k \in\{1, \ldots, L+1\}$. We can consider

$$
\left\{\begin{array}{l}
z^{k+1}=z^{k}+\Delta t\left(A^{k} \sigma\left(z^{k}\right)+b^{k}\right) \quad \text { for } k=0, \ldots, L-1  \tag{15}\\
z^{0}=E \vec{x}_{i} \in \mathbb{R}^{m}
\end{array}\right.
$$

for all $i$, where $\Delta t=\frac{T}{L-1}$ with $T>0$ fixed.

- Can be seen as discretization of ODE: $z=z_{i}(t)$ solves

$$
\left\{\begin{array}{l}
z^{\prime}(t)=A(t) \sigma(z(t))+b(t) \quad \text { in }(0, T)  \tag{16}\\
z(0)=E \vec{x}_{i}
\end{array}\right.
$$

for all $i \in\{1, \ldots, N\}$. Here $A(t) \in \mathbb{R}^{m \times m}$ and $b(t) \in \mathbb{R}^{m}$ are the controls.

- Compared to (??), (??) is an affine-control system (perhaps easier for theoretical purposes).
- "Controllability": for $N \in \mathbb{N}$, find $A(t), b(t)$ such that $\varphi\left(z_{i}(T)\right)=\vec{y}_{i}$ for all $i \in\{1, \ldots, N\}$, where $z_{i}$ solves (??).
- If $\varphi(x)=x$, we can try to apply Chow-Rashevski theorem (Iterated Lie brackets).

Associated ResNet: assume that $N_{k}=m$ for all $k \in\{1, \ldots, L+1\}$. We can consider

$$
\left\{\begin{array}{l}
z^{k+1}=z^{k}+\Delta t\left(A^{k} \sigma\left(z^{k}\right)+b^{k}\right) \quad \text { for } k=0, \ldots, L-1  \tag{17}\\
z^{0}=E \vec{x}_{i} \in \mathbb{R}^{m}
\end{array}\right.
$$

for all $i$, where $\Delta t=\frac{T}{L-1}$ with $T>0$ fixed.
Now

- Can be seen as discretization of ODE: $z=z_{i}(t)$ solves

$$
\left\{\begin{array}{l}
z^{\prime}(t)=A(t) \sigma(z(t))+b(t) \quad \text { in }(0, T) \\
z(0)=E \vec{x}_{i}
\end{array}\right.
$$

for all $i \in\{1, \ldots, N\}$. Here $A(t) \in \mathbb{R}^{m \times m}$ and $b(t) \in \mathbb{R}^{m}$ are the controls.

- Cuchiero, Larsson, Teichman '19 consider a simplified system

$$
\left\{\begin{array}{l}
z^{\prime}(t)=\sum_{j=1}^{5} u^{j}(t) v^{j}(z(t)) \quad \text { in }(0, T)  \tag{18}\\
z(0)=E \vec{x}_{i} \in \Omega \subset \mathbb{R}^{m}
\end{array}\right.
$$

- $v^{j}(\cdot) \in C^{\infty}\left(\mathbb{R}^{m}, \mathbb{R}^{m}\right)$ are polynomial fields at most of order 2 ; controls $u^{j}(t)$ are scalars.
- They prove Objective for (??) using (Ch-Ra) for stacked system satisfied by $\left\{z_{i}\right\}_{i=1}^{N}$


## Theorem (Chow-Rashevski)

Let $\Omega \subset \mathbb{R}^{m}$ be bounded domain, and assume $v^{1}, \ldots, v^{5}$ satisfy the Hörmander condition

$$
\operatorname{Lie}\left(v^{1}, \ldots, v^{5}\right)(x)=\mathbb{R}^{m}
$$

for all $x \in \Omega$. Then, (??) is exactly controllable (in usual sense).

- Lie algebra evaluated at $x$ :
$\operatorname{Lie}\left(v^{1}, \ldots, v^{5}\right)(x):=\left\{W(x): W \in \operatorname{span}\left\{v^{1}, \ldots v^{5}\right.\right.$ and all iterated Lie brackets $\left.\}\right\}$
- Lie bracket: $[u, v](x)=v^{\prime} u-u^{\prime} v$ for $u, v \in C^{\infty}\left(\mathbb{R}^{m}, \mathbb{R}^{m}\right)\left(v^{\prime}\right.$ and $u^{\prime}$ Jacobians).


## How to quantify generalization?

- Let training data $\left\{\overrightarrow{x_{i}}, \overrightarrow{y_{i}}\right\}_{i=1}^{N}$ be samples of some unknown probability distribution $p ; f_{\Theta}$ trained neural net
- Empirical risk (training error): $\mathcal{R}_{N}\left[f_{\Theta}\right]=\frac{1}{N} \sum_{i=1}^{N} \ell\left(f_{\Theta}\left(\vec{x}_{i}\right), \vec{y}_{i}\right)$
- Expected risk (non-sampled error):

$$
\mathcal{R}\left[f_{\Theta}\right]=\mathbb{E}_{x, y \sim p}\left[\ell\left(f_{\Theta}(x), y\right)\right]
$$

(non-computable since $p$ unknown)

- Goal: minimize expected risk. Possible approach: bound

$$
\text { generalization gap }:=\mathcal{R}_{N}\left[f_{\Theta}\right]-\mathcal{R}\left[f_{\Theta}\right]
$$

by $\mathcal{O}\left((L / N)^{\alpha}\right), \alpha>0, L$ is number of hidden layers.
ne Understanding deep learning requires rethinking generalizationattributes small generalization error either to properties of the model family, or to the ...
is 74 Cited by 1461 Related articles All 20 versions $\theta$
Generalization is not well understood for deep neural nets.

## Questions and perspectives

Many questions persist:

- How can one quantify/describe the stability of the deep learning process with respect to perturbations in the data?
- What about the choice of the activation function $\sigma$ ?
- How does one best choose the widths $N_{k}$ of the neural network w.r.t. the data?
- What mathematical control results does one transfer to deep learning (e.g. Lie brackets)?
- Other architectures were not presented (e.g. Convolutional Neural Networks)


[^0]:    ${ }^{1}$ This material was provided by Borjan Geshovski and is part of his PhD Thesis planned for 2021 at UAM-Madrid

[^1]:    ${ }^{2}$ Possibly noisy, so $f\left(\vec{x}_{i}\right)=\vec{y}_{i}+\varepsilon_{i}$ with $\varepsilon_{i} \sim \mathcal{N}\left(0, \varsigma^{2}\right)$ for instance.
    ${ }^{3}$ Also called training error

[^2]:    ${ }^{4}$ Related to GAN (Generative adversarial networks), see Goodfellow et al. '16

[^3]:    ${ }^{\mathbf{5}}$ Recall that $\Lambda_{k+1} z=A^{k} z+b^{k} \in \mathbb{R}^{N_{k+1}}$ for $z \in \mathbb{R}^{N_{k}}$ and $k \in\{0, \ldots, L\}$.
    ${ }^{6}$ If $L \gg 1$, we can choose to optimize the parameters $A^{k}$ at a later hidden layer, not necessarily at the first one.

