# Deep learning <br> 3.4. Multi-Layer Perceptrons 

François Fleuret<br>https://fleuret.org/dlc/

UNIVERSITÉ
DE GENÈVE

A linear classifier of the form

$$
\begin{aligned}
\mathbb{R}^{D} & \rightarrow \mathbb{R} \\
x & \mapsto \sigma(w \cdot x+b),
\end{aligned}
$$

with $w \in \mathbb{R}^{D}, b \in \mathbb{R}$, and $\sigma: \mathbb{R} \rightarrow \mathbb{R}$, can naturally be extended to a multi-dimension output by applying a similar transformation to every output

$$
\begin{aligned}
\mathbb{R}^{D} & \rightarrow \mathbb{R}^{C} \\
x & \mapsto \sigma(w x+b),
\end{aligned}
$$

with $w \in \mathbb{R}^{C \times D}, b \in \mathbb{R}^{C}$, and $\sigma$ is applied component-wise.

Even though it has no practical value implementation-wise, we can represent such a model as a combination of units. More importantly, we can extend it.


Even though it has no practical value implementation-wise, we can represent such a model as a combination of units. More importantly, we can extend it.


Even though it has no practical value implementation-wise, we can represent such a model as a combination of units. More importantly, we can extend it.


This latter structure can be formally defined, with $x^{(0)}=x$,

$$
\forall I=1, \ldots, L, x^{(I)}=\sigma\left(w^{(I)} x^{(I-1)}+b^{(I)}\right)
$$

and $f(x ; w, b)=x^{(L)}$.

This latter structure can be formally defined, with $x^{(0)}=x$,

$$
\forall I=1, \ldots, L, x^{(I)}=\sigma\left(w^{(I)} x^{(I-1)}+b^{(I)}\right)
$$

and $f(x ; w, b)=x^{(L)}$.


Such a model is a Multi-Layer Perceptron (MLP).

Note that if $\sigma$ is an affine transformation, the full MLP is a composition of affine mappings, and itself an affine mapping.

Note that if $\sigma$ is an affine transformation, the full MLP is a composition of affine mappings, and itself an affine mapping.

Consequently:

©
The activation function $\sigma$ should not be affine. Otherwise the resulting MLP would be an affine mapping with a peculiar parametrization.

The two classical activation functions are the hyperbolic tangent

$$
x \mapsto \frac{2}{1+e^{-2 x}}-1
$$



The two classical activation functions are the hyperbolic tangent

$$
x \mapsto \frac{2}{1+e^{-2 x}}-1
$$


and the rectified linear unit (ReLU, Glorot et al., 2011)

$$
x \mapsto \max (0, x)
$$



# Universal approximation 

We can approximate any $\psi \in \mathscr{C}([a, b], \mathbb{R})$ with a linear combination of translated/scaled ReLU functions.


We can approximate any $\psi \in \mathscr{C}([a, b], \mathbb{R})$ with a linear combination of translated/scaled ReLU functions.

$$
f(x)=\sigma\left(w_{1} x+b_{1}\right)
$$



We can approximate any $\psi \in \mathscr{C}([a, b], \mathbb{R})$ with a linear combination of translated/scaled ReLU functions.

$$
f(x)=\sigma\left(w_{1} x+b_{1}\right)+\sigma\left(w_{2} x+b_{2}\right)
$$



We can approximate any $\psi \in \mathscr{C}([a, b], \mathbb{R})$ with a linear combination of translated/scaled ReLU functions.

$$
f(x)=\sigma\left(w_{1} x+b_{1}\right)+\sigma\left(w_{2} x+b_{2}\right)+\sigma\left(w_{3} x+b_{3}\right)
$$



We can approximate any $\psi \in \mathscr{C}([a, b], \mathbb{R})$ with a linear combination of translated/scaled ReLU functions.

$$
f(x)=\sigma\left(w_{1} x+b_{1}\right)+\sigma\left(w_{2} x+b_{2}\right)+\sigma\left(w_{3} x+b_{3}\right)+\ldots
$$



We can approximate any $\psi \in \mathscr{C}([a, b], \mathbb{R})$ with a linear combination of translated/scaled ReLU functions.

$$
f(x)=\sigma\left(w_{1} x+b_{1}\right)+\sigma\left(w_{2} x+b_{2}\right)+\sigma\left(w_{3} x+b_{3}\right)+\ldots
$$



We can approximate any $\psi \in \mathscr{C}([a, b], \mathbb{R})$ with a linear combination of translated/scaled ReLU functions.

$$
f(x)=\sigma\left(w_{1} x+b_{1}\right)+\sigma\left(w_{2} x+b_{2}\right)+\sigma\left(w_{3} x+b_{3}\right)+\ldots
$$



We can approximate any $\psi \in \mathscr{C}([a, b], \mathbb{R})$ with a linear combination of translated/scaled ReLU functions.

$$
f(x)=\sigma\left(w_{1} x+b_{1}\right)+\sigma\left(w_{2} x+b_{2}\right)+\sigma\left(w_{3} x+b_{3}\right)+\ldots
$$



We can approximate any $\psi \in \mathscr{C}([a, b], \mathbb{R})$ with a linear combination of translated/scaled ReLU functions.

$$
f(x)=\sigma\left(w_{1} x+b_{1}\right)+\sigma\left(w_{2} x+b_{2}\right)+\sigma\left(w_{3} x+b_{3}\right)+\ldots
$$



We can approximate any $\psi \in \mathscr{C}([a, b], \mathbb{R})$ with a linear combination of translated/scaled ReLU functions.

$$
f(x)=\sigma\left(w_{1} x+b_{1}\right)+\sigma\left(w_{2} x+b_{2}\right)+\sigma\left(w_{3} x+b_{3}\right)+\ldots
$$



We can approximate any $\psi \in \mathscr{C}([a, b], \mathbb{R})$ with a linear combination of translated/scaled ReLU functions.

$$
f(x)=\sigma\left(w_{1} x+b_{1}\right)+\sigma\left(w_{2} x+b_{2}\right)+\sigma\left(w_{3} x+b_{3}\right)+\ldots
$$



We can approximate any $\psi \in \mathscr{C}([a, b], \mathbb{R})$ with a linear combination of translated/scaled ReLU functions.

$$
f(x)=\sigma\left(w_{1} x+b_{1}\right)+\sigma\left(w_{2} x+b_{2}\right)+\sigma\left(w_{3} x+b_{3}\right)+\ldots
$$



We can approximate any $\psi \in \mathscr{C}([a, b], \mathbb{R})$ with a linear combination of translated/scaled ReLU functions.

$$
f(x)=\sigma\left(w_{1} x+b_{1}\right)+\sigma\left(w_{2} x+b_{2}\right)+\sigma\left(w_{3} x+b_{3}\right)+\ldots
$$



We can approximate any $\psi \in \mathscr{C}([a, b], \mathbb{R})$ with a linear combination of translated/scaled ReLU functions.

$$
f(x)=\sigma\left(w_{1} x+b_{1}\right)+\sigma\left(w_{2} x+b_{2}\right)+\sigma\left(w_{3} x+b_{3}\right)+\ldots
$$



We can approximate any $\psi \in \mathscr{C}([a, b], \mathbb{R})$ with a linear combination of translated/scaled ReLU functions.

$$
f(x)=\sigma\left(w_{1} x+b_{1}\right)+\sigma\left(w_{2} x+b_{2}\right)+\sigma\left(w_{3} x+b_{3}\right)+\ldots
$$



We can approximate any $\psi \in \mathscr{C}([a, b], \mathbb{R})$ with a linear combination of translated/scaled ReLU functions.

$$
f(x)=\sigma\left(w_{1} x+b_{1}\right)+\sigma\left(w_{2} x+b_{2}\right)+\sigma\left(w_{3} x+b_{3}\right)+\ldots
$$



This is true for other activation functions under mild assumptions.

Extending this result to any $\psi \in \mathscr{C}\left([0,1]^{D}, \mathbb{R}\right)$ requires a bit of work.
We can approximate the sin function with the previous scheme, and use the density of Fourier series to get the final result:

$$
\begin{aligned}
\forall \epsilon>0, \exists K, w \in \mathbb{R}^{K \times D}, b \in \mathbb{R}^{K}, \omega \in \mathbb{R}^{K}, \text { s.t. } \\
\max _{x \in[0,1]^{D}}|\psi(x)-\omega \cdot \sigma(w x+b)| \leq \epsilon
\end{aligned}
$$

So we can approximate any continuous function

$$
\psi:[0,1]^{D} \rightarrow \mathbb{R}
$$

with a one hidden layer perceptron

$$
x \mapsto \omega \cdot \sigma(w x+b)
$$

where $b \in \mathbb{R}^{K}, w \in \mathbb{R}^{K \times D}$, and $\omega \in \mathbb{R}^{K}$.


So we can approximate any continuous function

$$
\psi:[0,1]^{D} \rightarrow \mathbb{R}
$$

with a one hidden layer perceptron

$$
x \mapsto \omega \cdot \sigma(w x+b)
$$

where $b \in \mathbb{R}^{K}, w \in \mathbb{R}^{K \times D}$, and $\omega \in \mathbb{R}^{K}$.


This is the universal approximation theorem.

©
A better approximation requires a larger hidden layer (larger $K$ ), and this theorem says nothing about the relation between the two.

So this results states that we can make the training error as low as we want by using a larger hidden layer. It states nothing about the test error.

4
A better approximation requires a larger hidden layer (larger $K$ ), and this theorem says nothing about the relation between the two.

So this results states that we can make the training error as low as we want by using a larger hidden layer. It states nothing about the test error.

Deploying MLP in practice is often a balancing act between under-fitting and over-fitting.

The end

## References

X. Glorot, A. Bordes, and Y. Bengio. Deep sparse rectifier neural networks. In International Conference on Artificial Intelligence and Statistics (AISTATS), 2011.

