Kathlén Kohn











feedforward neural networks



. / 14

feedforward neural networks



are parametrized families of functions $\mu : \mathbb{R}^N \longrightarrow \mathcal{M},$ $\theta \longmapsto f_{L,\theta} \circ \ldots \circ f_{1,\theta}$

feedforward neural networks



are parametrized families of functions $\mu : \mathbb{R}^{N} \longrightarrow \mathcal{M},$ $\theta \longmapsto f_{L,\theta} \circ \ldots \circ f_{1,\theta}$ $\mathcal{M} = \text{ function space / neuromanifold, } L = \# \text{ layers}$

training a network

Given training data \mathcal{D}_{r} the goal is to minimize the loss

 $\mathbb{R}^{N} \xrightarrow{\mu} \mathcal{M} \xrightarrow{\ell_{\mathcal{D}}} \mathbb{R}.$

training a network

Given training data \mathcal{D} , the goal is to minimize the loss

 $\mathbb{R}^{N} \xrightarrow{\mu} \mathcal{M} \xrightarrow{\ell_{\mathcal{D}}} \mathbb{R}.$

Geometric questions:

- How does the network architecture affect the geometry of the function space?
- How does the geometry of the function space impact the training of the network?

training a network

Given training data \mathcal{D} , the goal is to minimize the loss

 $\mathbb{R}^N \xrightarrow{\mu} \mathcal{M} \xrightarrow{\ell_{\mathcal{D}}} \mathbb{R}.$

Geometric questions:

- How does the network architecture affect the geometry of the function space?
- How does the geometry of the function space impact the training of the network?

In this talk:

What is the impact of changing from dense layers to convolutional layers?

linear dense networks



In this example:

 $\mu: \mathbb{R}^{2 \times 4} \times \mathbb{R}^{3 \times 2} \longrightarrow \mathbb{R}^{3 \times 4},$ $(W_1, W_2) \longmapsto W_2 W_1.$

linear dense networks



In this example:

 $\begin{array}{c} \mu: \mathbb{R}^{2 \times 4} \times \mathbb{R}^{3 \times 2} \longrightarrow \mathbb{R}^{3 \times 4}, \\ (W_1, W_2) \longmapsto W_2 W_1. \end{array}$

 $\mathcal{M} = \{ \mathcal{W} \in \mathbb{R}^{3 \times 4} \mid \mathrm{rank}(\overline{\mathcal{W}) \leq 2} \}$

linear dense networks



In this example:

 $\mu: \mathbb{R}^{2 \times 4} \times \mathbb{R}^{3 \times 2} \longrightarrow \mathbb{R}^{3 \times 4},$ $(W_1, W_2) \longmapsto W_2 W_1.$

 $\mathcal{M} = \{ \mathcal{W} \in \mathbb{R}^{3 \times 4} \mid \operatorname{rank}(\mathcal{W}) \leq 2 \}$

In general:

$$\frac{\mu: \mathbb{R}^{k_1 \times k_0} \times \mathbb{R}^{k_2 \times k_1} \times \ldots \times \mathbb{R}^{k_L \times k_{L-1}} \longrightarrow \mathbb{R}^{k_L \times k_0},}{(W_1, W_2, \ldots, W_L) \longmapsto W_L \cdots W_2 W_1}.$$

 $\mathcal{M} = \{W \in \mathbb{R}^{k_L \times k_0} \mid \operatorname{rank}(W) \leq \min(k_0, \ldots, k_L)\} \text{ is an algebraic variety and} we know its singularities etc.}$





 $\mu: \mathbb{R}^3 \times \mathbb{R}^2 \longrightarrow \mathbb{R}^5,$ $(u, v) \longmapsto T_{v,1} T_{u,2}, \text{ where }$

 $T_{u,2} = \begin{bmatrix} u_0 & u_1 & u_2 & 0 & 0 \\ 0 & 0 & u_0 & u_1 & u_2 \end{bmatrix}$ $T_{v,1} = \begin{bmatrix} v_0 & v_1 \end{bmatrix}$



 $\mu: \mathbb{R}^3 \times \mathbb{R}^2 \longrightarrow \mathbb{R}^5,$ $(u, v) \longmapsto T_{v,1}T_{u,2}, \text{ where }$

 $T_{u,2} = \begin{bmatrix} u_0 & u_1 & u_2 & 0 & 0 \\ 0 & 0 & u_0 & u_1 & u_2 \end{bmatrix}$ $T_{v,1} = \begin{bmatrix} v_0 & v_1 \end{bmatrix}$

In general: $\mu : (w_1, \ldots, w_L) \mapsto T_{w_L, s_L} \cdots T_{w_1, s_1}$, where

/ 14



 $\mu: \mathbb{R}^3 \times \mathbb{R}^2 \longrightarrow \mathbb{R}^5,$ $(u, v) \longmapsto T_{v,1}T_{u,2}, \text{ where }$

 $T_{u,2} = \begin{bmatrix} u_0 & u_1 & u_2 & 0 & 0 \\ 0 & 0 & u_0 & u_1 & u_2 \end{bmatrix}$ $T_{v,1} = \begin{bmatrix} v_0 & v_1 \end{bmatrix}$

In general: $\mu : (w_1, \ldots, w_L) \mapsto T_{w_L, s_L} \cdots T_{w_1, s_1}$, where

is a convolutional matrix of stride *s* with filter *w*

LCNs & sparse polynomial factorization Observation: $\mu(w_1, \ldots, w_L) = T_{w_L, s_L} \cdots T_{w_1, s_1}$ is again a convolutional matrix of stride $s_1 \cdots s_L$.

LCNs & sparse polynomial factorization Observation: $\mu(w_1, \ldots, w_L) = T_{w_L, s_L} \cdots T_{w_1, s_1}$ is again a convolutional matrix of stride $s_1 \cdots s_L$. Its filter can be computed via polynomial multiplication:

LCNs & sparse polynomial factorization Observation: $\mu(w_1, \ldots, w_L) = T_{w_L, s_L} \cdots T_{w_1, s_1}$ is again a convolutional matrix of stride $s_1 \cdots s_L$. Its filter can be computed via polynomial multiplication:

For $S\in\mathbb{Z}_{>0}$, let

 $\pi_{\mathcal{S}}: \mathbb{R}^{k} \longrightarrow \mathbb{R}[x^{\mathcal{S}}]_{\leq k-1},$ $v \longmapsto v_{0}x^{\mathcal{S}(k-1)} + v_{1}x^{\mathcal{S}(k-2)} + \ldots + v_{k-2}x^{\mathcal{S}} + v_{k-1}$

LCNs & sparse polynomial factorization Observation: $\mu(w_1, \ldots, w_L) = T_{w_L, s_L} \cdots T_{w_1, s_1}$ is again a convolutional matrix of stride $s_1 \cdots s_L$. Its filter can be computed via polynomial multiplication:

For $S \in \mathbb{Z}_{>0}$, let

 $\pi_{S} : \mathbb{R}^{k} \longrightarrow \mathbb{R}[x^{S}]_{\leq k-1},$ $v \longmapsto v_{0}x^{S(k-1)} + v_{1}x^{S(k-2)} + \ldots + v_{k-2}x^{S} + v_{k-1}$

and $\pi_S(T_{w,s}) := \pi_S(w)$. Then:

LCNs & sparse polynomial factorization Observation: $\mu(w_1, \ldots, w_L) = T_{w_L, s_L} \cdots T_{w_1, s_1}$ is again a convolutional matrix of stride $s_1 \cdots s_L$. Its filter can be computed via polynomial multiplication:

For $S \in \mathbb{Z}_{>0}$, let

 $\pi_{S} : \mathbb{R}^{k} \longrightarrow \mathbb{R}[x^{S}]_{\leq k-1},$ $v \longmapsto v_{0}x^{S(k-1)} + v_{1}x^{S(k-2)} + \ldots + v_{k-2}x^{S} + v_{k-1}$

and $\pi_S(T_{w,s}) := \pi_S(w)$. Then:

 $\pi_1(\mu(w_1,\ldots,w_L)) = \pi_{S_L}(w_L)\cdots\pi_{S_1}(w_1), \text{ where } S_i := s_1\cdots s_{i-1}.$

LCNs & sparse polynomial factorization Observation: $\mu(w_1, \ldots, w_L) = T_{w_L, s_L} \cdots T_{w_1, s_1}$ is again a convolutional matrix of stride $s_1 \cdots s_L$. Its filter can be computed via polynomial multiplication:

For $S \in \mathbb{Z}_{>0}$, let

 $\pi_{S} : \mathbb{R}^{k} \longrightarrow \mathbb{R}[x^{S}]_{\leq k-1},$ $v \longmapsto v_{0}x^{S(k-1)} + v_{1}x^{S(k-2)} + \ldots + v_{k-2}x^{S} + v_{k-1}$

and $\pi_S(T_{w,s}) := \pi_S(w)$. Then:

 $\pi_1(\mu(w_1, \ldots, w_L)) = \pi_{S_L}(w_L) \cdots \pi_{S_1}(w_1), \text{ where } S_i := s_1 \cdots s_{i-1}.$

Hence, we reinterpret μ as

$$\mu: \mathbb{R}[x^{S_1}]_{\leq d_1} \times \ldots \times \mathbb{R}[x^{S_L}]_{\leq d_L} \longrightarrow \mathbb{R}[x]_{\leq d_1 S_1 + \ldots + d_L S_L},$$
$$(P_1, \ldots, P_L) \longmapsto P_L \cdots P_1$$

 $\mu: \mathbb{R}[x^{S_1}]_{\leq d_1} \times \ldots \times \mathbb{R}[x^{S_L}]_{\leq d_l} \longrightarrow \mathbb{R}[x]_{\leq d}$, where $d:=\sum_i d_i S_i$ $(P_1,\ldots,P_L)\longmapsto P_L\cdots P_1,$

 $\mu: \mathbb{R}[x^{S_1}]_{\leq d_1} \times \ldots \times \mathbb{R}[x^{S_L}]_{\leq d_L} \longrightarrow \mathbb{R}[x]_{\leq d}, \text{ where } d := \sum_i d_i S_i$ $(P_1, \ldots, P_L) \longmapsto P_L \cdots P_1,$

Theorem: The function space $\mathcal{M}_{d,S} = \operatorname{im}(\mu)$ is a semi-algebraic, Euclidean-closed subset of $\mathbb{R}[x]_{\leq d}$ of dimension $d_1 + \ldots + d_L + 1$.



 $\mu: \mathbb{R}[x]_{\leq 2} \times \mathbb{R}[x^2]_{\leq 1} \to \mathbb{R}[x]_{\leq 4}$

 $\mu: \mathbb{R}[x]_{\leq 1} \times \mathbb{R}[x]_{\leq 1} \times \mathbb{R}[x^2]_{\leq 1} \to \mathbb{R}[x]_{\leq 4}$

 $\mu: \mathbb{R}[x^{S_1}]_{\leq d_1} \times \ldots \times \mathbb{R}[x^{S_L}]_{\leq d_L} \longrightarrow \mathbb{R}[x]_{\leq d}, \text{ where } d := \sum_i d_i S_i$ $(P_1, \ldots, P_L) \longmapsto P_L \cdots P_1,$

Theorem: The function space $\mathcal{M}_{d,S} = \operatorname{im}(\mu)$ is a semi-algebraic, Euclidean-closed subset of $\mathbb{R}[x]_{\leq d}$ of dimension $d_1 + \ldots + d_L + 1$.

Corollary: $\mathcal{M}_{d,S}$ is full-dimensional in $\mathbb{R}[x]_{\leq d}$ if and only if all strides $s_i = 1$.



 $\mu: \mathbb{R}[x^{S_1}]_{\leq d_1} \times \overline{\ldots \times \mathbb{R}[x^{S_L}]_{\leq d_L}} \longrightarrow \mathbb{R}[x]_{\leq d}, \text{ where } d := \sum_i d_i S_i$ $(P_1, \ldots, P_L) \longmapsto P_L \cdots P_1,$

Example: s = (1, 1) and d = (1, 1). $\Rightarrow \mathcal{M}_{d,S} = \operatorname{im}(\mu) =$

 $\mu: \mathbb{R}[x^{S_1}]_{\leq d_1} \times \ldots \times \mathbb{R}[x^{S_L}]_{\leq d_L} \longrightarrow \mathbb{R}[x]_{\leq d}, \text{ where } d := \sum_i d_i S_i$ $(P_1, \ldots, P_L) \longmapsto P_L \cdots P_1,$

Example: s = (1, 1) and d = (1, 1). $\Rightarrow M_{d,s} = im(\mu) = \{ \text{ quadratic polynomials with 2 real roots } \}$

 $\mu: \mathbb{R}[x^{S_1}]_{\leq d_1} \times \ldots \times \mathbb{R}[x^{S_L}]_{\leq d_L} \longrightarrow \mathbb{R}[x]_{\leq d}, \text{ where } d := \sum_i d_i S_i$ $(P_1, \ldots, P_L) \longmapsto P_L \cdots P_1,$

Example: s = (1, 1) and d = (1, 1). $\Rightarrow M_{d,S} = im(\mu) = \{ \text{ quadratic polynomials with 2 real roots } \}$

Example: s = (1, 1) and d = (1, 2) $\Rightarrow \mathcal{M}_{d,s} = \operatorname{im}(\mu) =$

 $\mu: \mathbb{R}[x^{S_1}]_{\leq d_1} \times \ldots \times \mathbb{R}[x^{S_L}]_{\leq d_L} \longrightarrow \mathbb{R}[x]_{\leq d}, \text{ where } d := \sum_i d_i S_i$ $(P_1, \ldots, P_L) \longmapsto P_L \cdots P_1,$

Example: s = (1, 1) and d = (1, 1). $\Rightarrow \mathcal{M}_{d,S} = \operatorname{im}(\mu) = \{ \text{ quadratic polynomials with 2 real roots } \}$

Example: s = (1, 1) and d = (1, 2) $\Rightarrow M_{d,S} = im(\mu) = \{ \text{ all cubic polynomials } \}$

 $\mu: \mathbb{R}[x^{S_1}]_{\leq d_1} \times \ldots \times \mathbb{R}[x^{S_L}]_{\leq d_L} \longrightarrow \mathbb{R}[x]_{\leq d}, \text{ where } d := \sum_i d_i S_i$ $(P_1, \ldots, P_L) \longmapsto P_L \cdots P_1,$

Example: s = (1, 1) and d = (1, 1). $\Rightarrow M_{d,s} = im(\mu) = \{ \text{ quadratic polynomials with 2 real roots } \}$

Example: s = (1, 1) and d = (1, 2) $\Rightarrow M_{d,s} = im(\mu) = \{ \text{ all cubic polynomials } \}$

Proposition: s = (1, 1, ..., 1) and $d = (d_1, d_2, ..., d_L)$. $\Rightarrow \mathcal{M}_{d,S} = \operatorname{im}(\mu) = \mathbb{R}[x]_{\leq d}$ if and only if

 $\mu: \mathbb{R}[x^{S_1}]_{\leq d_1} \times \ldots \times \mathbb{R}[x^{S_L}]_{\leq d_L} \longrightarrow \mathbb{R}[x]_{\leq d}, \text{ where } d := \sum_i d_i S_i$ $(P_1, \ldots, P_L) \longmapsto P_L \cdots P_1,$

Example: s = (1, 1) and d = (1, 1). $\Rightarrow M_{d,s} = im(\mu) = \{ \text{ quadratic polynomials with 2 real roots } \}$

Example: s = (1, 1) and d = (1, 2) $\Rightarrow M_{d,s} = im(\mu) = \{ \text{ all cubic polynomials } \}$

Proposition: s = (1, 1, ..., 1) and $d = (d_1, d_2, ..., d_L)$. $\Rightarrow \mathcal{M}_{d,S} = \operatorname{im}(\mu) = \mathbb{R}[x]_{< d}$ if and only if at most one d_i is odd.

$\begin{array}{c} \mathsf{LCN \ function \ spaces} \\ \mu: \mathbb{R}[x^{S_1}]_{\leq d_1} \times \ldots \times \mathbb{R}[x^{S_L}]_{\leq d_L} \longrightarrow \mathbb{R}[x]_{\leq d}, \text{where } d := \sum_i d_i S_i \\ (P_1, \ldots, P_L) \longmapsto P_L \cdots P_1, \end{array}$



 $\mu : \mathbb{R}[x]_{\leq 2} \times \mathbb{R}[x^2]_{\leq 1} \to \mathbb{R}[x]_{\leq 4} \qquad \mu : \mathbb{R}[x]_{\leq 1} \times \mathbb{R}[x]_{\leq 1} \times \mathbb{R}[x^2]_{\leq 1} \to \mathbb{R}[x]_{\leq 4}$ $\uparrow \mathcal{M}_{\boldsymbol{d},\boldsymbol{S}} = \{(a + bx + c^2)(d + ex^2) \mid a, \dots, e \in \mathbb{R}\}$



 $\mu : \mathbb{R}[x]_{\leq 2} \times \mathbb{R}[x^2]_{\leq 1} \to \mathbb{R}[x]_{\leq 4} \qquad \mu : \mathbb{R}[x]_{\leq 1} \times \mathbb{R}[x]_{\leq 1} \times \mathbb{R}[x^2]_{\leq 1} \to \mathbb{R}[x]_{\leq 4}$ $\uparrow \mathcal{M}_{\boldsymbol{d},\boldsymbol{S}} = \{(\boldsymbol{a} + \boldsymbol{b}x + \boldsymbol{c}^2)(\boldsymbol{d} + \boldsymbol{e}x^2) \mid \boldsymbol{a}, \dots, \boldsymbol{e} \in \mathbb{R}\}$ $= \{A + Bx + Cx^2 + Dx^3 + Ex^4 \mid AD^2 + B^2E = BCD, \ C^2 \geq 4AE\}$

A convolution on D-dimensional signals is defined using a filter w of format

 $k^{(1)} \times \cdots \times \overline{k^{(D)}}$

and a stride tuple $s = (s^{(1)}, \ldots, s^{(D)}) \in \mathbb{Z}_{>0}^D$.

A convolution on D-dimensional signals is defined using a filter w of format

 $k^{(1)} \times \cdots \times k^{(D)}$

and a stride tuple $s = (s^{(1)}, \dots, s^{(D)}) \in \mathbb{Z}^D_{\geq 0}.$

The convolution is a linear map α_w that produces an output tensor of format $d^{(1)} \times \cdots \times d^{(D)}$ from an input tensor X of format $(s^{(1)}(d^{(1)}-1)+k^{(1)}) \times \cdots \times (s^{(D)}(d^{(D)}-1)+k^{(D)})$:

$$(\alpha_{w,s}(X))_{i_1,\ldots,i_D} = \sum_{j_1=0}^{k^{(1)}-1} \cdots \sum_{j_D=0}^{k^{(D)}-1} w_{j_1,\ldots,j_D} X_{i_1s^{(1)}+j_1,\ldots,i_Ds^{(D)}+j_D}$$

for $i_m = 0, 1, \ldots, d^{(m)} - 1$.

A convolution on D-dimensional signals is defined using a filter w of format

 $k^{(1)} imes \cdots imes k^{(D)}$

and a stride tuple $s = \overline{(s^{(1)}, \dots, s^{(D)})} \in \mathbb{Z}^D_{\geq 0}.$

The convolution is a linear map α_w that produces an output tensor of format $d^{(1)} \times \cdots \times d^{(D)}$ from an input tensor X of format $(s^{(1)}(d^{(1)}-1)+k^{(1)}) \times \cdots \times (s^{(D)}(d^{(D)}-1)+k^{(D)})$:

$$(\alpha_{w,s}(X))_{i_1,\ldots,i_D} = \sum_{j_1=0}^{k^{(1)}-1} \cdots \sum_{j_D=0}^{k^{(D)}-1} w_{j_1,\ldots,j_D} X_{i_1s^{(1)}+j_1,\ldots,i_Ds^{(D)}+j_D}$$

for $i_m = 0, 1, \ldots, d^{(m)} - 1$.

This linear map is represented by a 2*D*-dimensional convolutional tensor T of format

$$d^{(1)} imes \cdots imes d^{(D)} imes (s^{(1)}(d^{(1)}-1)+k^{(1)}) imes \cdots imes (s^{(D)}(d^{(D)}-1)+k^{(D)})$$

Example: D = 2, s = (1, 1), filter w of format 2×2 : The convolution of the input

$$X = \begin{bmatrix} * & * \\ * \circ & * \circ \\ \circ & \circ \end{bmatrix}$$

with filter w yields an output of size 2×1 . The associated convolutional tensor $T \in \mathbb{R}^{2 \times 1 \times 3 \times 2}$ has the slices:

$$T_{00::} = \begin{bmatrix} w_{00} & w_{01} \\ w_{10} & w_{11} \\ 0 & 0 \end{bmatrix} \text{ and } T_{10::} = \begin{bmatrix} 0 & 0 \\ w_{00} & w_{01} \\ w_{10} & w_{11} \end{bmatrix}$$

LCNs

A linear convolutional network composes L such convolutions, yielding another convolution!

The end-to-end convolution has:

• Stride: $(s_1^{(1)} \cdots s_L^{(1)}, \dots, s_1^{(D)} \cdots s_L^{(D)})$

LCNs

A linear convolutional network composes L such convolutions, yielding another convolution!

The end-to-end convolution has:

• Stride: $(s_1^{(1)} \cdots s_L^{(1)}, \dots, s_1^{(D)} \cdots s_L^{(D)})$

Filter computed via polynomial multiplication:

For

•
$$x = (x_1, \dots, y_D),$$

• $S = (S^{(1)}, \dots, S^{(D)}),$
• $k = (k^{(1)}, \dots, k^{(D)}),$

we denote by

$$\mathbb{R}[x^S]_{\leq k-1}$$

the vector space of polynomials of degree at most $k^{(m)} - 1$ in $x_m^{S^{(m)}}$.

LCNs & sparse polynomial factorization

We identify the space $\mathbb{R}[x^S]_{\leq k-1}$ with the tensor space $\mathbb{R}^{k^{(1)} \times \cdots \times k^{(D)}}$:

$$\pi_{S}: \mathbf{v} \mapsto \sum_{i_{1}=0}^{k^{(1)}-1} \cdots \sum_{i_{D}=0}^{k^{(D)}-1} \mathbf{v}_{i_{1},...,i_{D}} x_{1}^{S^{(1)}i_{1}} \cdots x_{D}^{S^{(D)}i_{D}}$$

LCNs & sparse polynomial factorization

We identify the space $\mathbb{R}[x^S]_{\leq k-1}$ with the tensor space $\mathbb{R}^{k^{(1)} \times \cdots \times k^{(D)}}$:

$$\pi_{S}: \mathbf{v} \mapsto \sum_{i_{1}=0}^{k^{(1)}-1} \cdots \sum_{i_{D}=0}^{k^{(D)}-1} \mathbf{v}_{i_{1},\dots,i_{D}} x_{1}^{S^{(1)}i_{1}} \cdots x_{D}^{S^{(D)}i_{D}}$$

Proposition: $\pi_{(1,...,1)}$ of the end-to-end filter is

 $\pi_{S_L}(w_L)\cdots\pi_{S_1}(w_1),$

where $S_{i}^{(m)} := s_{1}^{(m)} \cdots s_{i-1}^{(m)}$.

LCNs & sparse polynomial factorization

We identify the space $\mathbb{R}[x^S]_{\leq k-1}$ with the tensor space $\mathbb{R}^{k^{(1)} \times \cdots \times k^{(D)}}$:

$$\pi_{S}: \mathbf{v} \mapsto \sum_{i_{1}=0}^{k^{(1)}-1} \cdots \sum_{i_{D}=0}^{k^{(D)}-1} \mathbf{v}_{i_{1},\dots,i_{D}} x_{1}^{S^{(1)}i_{1}} \cdots x_{D}^{S^{(D)}i_{D}}$$

Proposition: $\pi_{(1,...,1)}$ of the end-to-end filter is

 $\pi_{S_L}(w_L)\cdots\pi_{S_1}(w_1),$

where $S_i^{(m)} := s_1^{(m)} \cdots s_{i-1}^{(m)}$.

This is factorization is typically unique!

https://doi.org/10.1007/s44426-025-00002-2

ARTICLE



Algebraic complexity and neurovariety of linear convolutional networks

Vahid Shahverdi¹

Received: 10 April 2025 / Accepted: 28 April 2025 © The Author(s) 2025



In this paper, we study linear convolutional networks with one-dimensional filters and arbitrary strides. The neuromanifold of such a network is a semialgebraic set, represented by a space of polynomials admitting specific factorizations. Introducing a recursive algorithm, we generate polynomial equations whose common zero locus corresponds to the Zariski closure of the corresponding neuromanifold. Furthermore, we explore the algebraic complexity of training these networks employing tools from metric algebraic geometry. Our findings reveal that the number of all complex critical points in the optimization of such a network is equal to the generic Euclidean distance degree of a Segre variety. Notably, this count significantly surpasses the number of critical points encountered in the training of a fully connected linear network with the same number of parameters.

Keywords Function space description of neural networks \cdot Linear networks \cdot Euclidean distance degree \cdot Critical points

What about higher dimensions ?