Quiver moduli from neural networks

M. Reineke

20 May 2021, Forschungsseminar Komplexe Analysis

- joint work in progress with M. Armenta, T. Bruestle, S. Hassoun

Overview

Plan of the talk:

1. Setup: a toy version of neural networks
2. Definition of double-framed quiver moduli
3. Results: qualitative
4. Results: quantitative
5. Methods: representation theory
6. Towards more realistic neural networks

Setup: a toy version of neural networks

$U_1, \ldots, U_m$, $W_1, \ldots, W_n$ vector spaces.
Neural network computes a *network function* $\bigoplus_{i=1}^m U_i \to \bigoplus_{i=1}^n W_i$.
(In our toy version: just a linear function (no activation function)...)

$$\begin{array}{cccc} U_1 & U_2 & \ldots & U_m \\ \downarrow & \downarrow & \ldots & \downarrow \\ \hline \multicolumn{4}{c}{\text{Hidden part}} \\ \hline \downarrow & \downarrow & \ldots & \downarrow \\ W_1 & W_2 & \ldots & W_n \end{array}$$

Setup: a toy version of neural networks

Hidden part is an architecture of *neurons* (vector spaces and linear maps):

$$\begin{array}{ccccc} V_{i_1} & & \ldots & & V_{i_m} \\ & f_1 \searrow & & f_m \swarrow & \\ & & {\color{red}V_j} & & \\ & h_1 \swarrow & & h_n \searrow & \\ V_{k_1} & & \ldots & & W_{i_n} \end{array}$$

computing

$$(v_{i_1}, \ldots, v_{i_m}) \mapsto \left(h_l \left(\sum_k f_k(v_{i_k}) \right) \right)_l .$$

Observation: Invariant under base change ($g \in \mathrm{GL}(V_j)$):

$$(f_1, \ldots, f_m,\ h_1, \ldots, h_n) \mapsto (gf_1, \ldots, gf_m,\ h_1 g^{-1}, \ldots, h_n g^{-1}).$$

$\rightsquigarrow$ **Can be modelled via representations of acyclic quivers**

Example: MLP

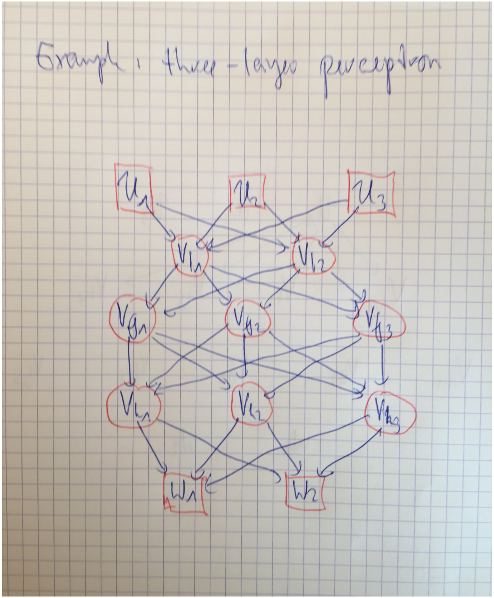

Aim

Aim

Define and study a **moduli space** *encoding the internal states of the network up to base change, and study the behaviour of network function along the moduli space.*

Quiver notation

Q acyclic finite quiver, Q_0 set of vertices, arrows written $\alpha : i \to j$.

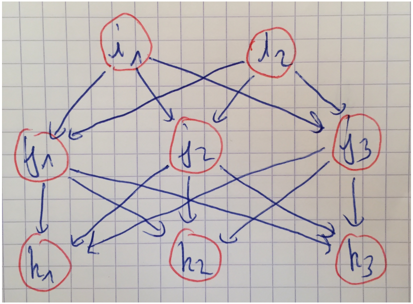

Euler form (non-symmetric bilinear form $\langle \_, \_ \rangle : \mathbb{Z}Q_0 \times \mathbb{Z}Q_0 \to \mathbb{Z}$)

$$\langle \mathbf{d}, \mathbf{e} \rangle = \sum_{i \in Q_0} d_i e_i - \sum_{\alpha : i \to j} d_i e_j$$

for $\mathbf{d}, \mathbf{e} \in \mathbb{Z}Q_0$.

Quiver notation

Fix dimension vector $\mathbf{d} = (d_i)_{i \in Q_0} \in \mathbb{N}Q_0$, fix $\mathbb{C}$-vector spaces V_i of dimension d_i for $i \in Q_0$.
Representation space

$$R_{\mathbf{d}}(Q) = \bigoplus_{\alpha: i \to j} \mathrm{Hom}_k(V_i, V_j)$$

carries action of base change group $G_{\mathbf{d}} = \prod_{i \in Q_0} \mathrm{GL}(V_i)$ via

$$(g_i)_i \cdot (V_\alpha)_\alpha = (g_j f_\alpha g_i^{-1})_{\alpha: i \to j}.$$

Link to representation theory: $G_{\mathbf{d}}$-orbits in $R_{\mathbf{d}}(Q)$ correspond bijectively to isomorphism classes of representations of Q of dimension vector $\mathbf{d}$. Representations of Q of arbitrary dimension vectors form *hereditary* abelian category $\mathrm{rep} Q$, with homological Euler form $\langle \_, \_ \rangle$.

Double-framed quiver moduli

Fix additional dimension vectors $\mathbf{p}, \mathbf{q} \in \mathbb{N}Q_0$, fix $\mathbb{C}$-vector space U_i and W_i of dimension p_i and q_i, respectively, for all $i \in Q_0$.

Definition

Double-framed representation space $R_{\mathbf{d},\mathbf{p},\mathbf{q}}(Q) =$

$$= \underbrace{\bigoplus_{\alpha:i\to j} \mathrm{Hom}_k(V_i, V_j)}_{\text{hidden part}} \times \underbrace{\bigoplus_{i\in Q_0} \mathrm{Hom}_k(U_i, V_i)}_{\text{input part}} \times \underbrace{\bigoplus_{i\in Q_0} \mathrm{Hom}_k(V_i, W_i)}_{\text{output part}}$$

with points $(V, f, h) = ((V_\alpha)_\alpha, (f_i)_i, (h_i)_i)$. Carries $G_\mathbf{d}$-action

$$(g_i)_i \cdot ((V_\alpha)_\alpha, (f_i)_i, (h_i)_i) = ((g_j V_\alpha g_i^{-1})_{\alpha:i\to j}, (g_i f_i)_i, (h_i g_i^{-1})_i).$$

Double-framed quiver moduli

Definition

A double-framed representation (V, f, h) is called

- Θ^+-*stable* if the only subrepresentation containing $(\mathrm{Im}(f_i))_i$ is V itself,
- Θ^--*stable* if the only subrepresentation contained in $(\mathrm{Ker}(h_i))_i$ is 0,
- *simple* if it is both Θ^+-stable and Θ^--stable.

Definition

For fixed Q, $\mathbf{d}$, $\mathbf{p}$, $\mathbf{q}$, define moduli spaces as quotients

- $\mathcal{M} = R_{\mathbf{d},\mathbf{p},\mathbf{q}}(Q)//G_{\mathbf{d}}$,
- $\mathcal{M}^{\Theta^+} = R_{\mathbf{d},\mathbf{p},\mathbf{q}}(Q)^{\Theta^+-\mathrm{st}}/G_{\mathbf{d}}$,
- $\mathcal{M}^{\Theta^-} = R_{\mathbf{d},\mathbf{p},\mathbf{q}}(Q)^{\Theta^--\mathrm{st}}/G_{\mathbf{d}}$,
- $\mathcal{M}^{\mathrm{simp}} = R_{\mathbf{d},\mathbf{p},\mathbf{q}}(Q)^{\mathrm{simp}}/G_{\mathbf{d}}$.

Results I: Qualitative

Theorem

Assume $p_i \geq \langle \mathbf{d}, \mathbf{i} \rangle$, $q_i \geq \langle \mathbf{i}, \mathbf{d} \rangle$ *for all* $i \in Q_0$.

- $\mathcal{M}$ *is an irreducible, normal, affine, rational variety of dimension* $(\mathbf{p} + \mathbf{q}) \cdot \mathbf{d} - \langle \mathbf{d}, \mathbf{d} \rangle$.
- $\mathcal{M}^{\mathrm{simp}} \subset \mathcal{M}$ *is open and smooth.*
- $\exists\, \mathcal{M}^{\Theta^{\pm}} \to \mathcal{M}$ *proper, isomorphism over* $\mathcal{M}^{\mathrm{simp}}$, *desingularization, fibres admit affine pavings.*
- $\mathcal{M}^{\Theta^{\pm}}$ *isomorphic to vector bundle over tower of Grassmann bundles.*

Example: the trivial quiver – classical invariant theory

Given $p, q \geq d$, consider

$$U \xrightarrow{f} V \xrightarrow{h} W$$

up to base change $(f, h) \mapsto (gf, hg^{-1})$.
Quotient map $(f, h) \mapsto hf$.

$$\mathcal{M} \simeq \{q \in \mathrm{Hom}(U, W) \,|\, \mathrm{rk}(q) \leq d\}$$

$$\mathcal{M}^{\mathrm{simp}} \simeq \{q \in \mathrm{Hom}(U, W) \,|\, \mathrm{rk}(q) = d\}$$

$$\mathcal{M}^{\Theta^+} \simeq \{(U, q) \in \mathrm{Gr}^d(U) \times \mathrm{Hom}(U, W) \,|\, U \subset \mathrm{Ker}(q)\}$$

$$\mathcal{M}^{\Theta^-} \simeq \{(q, U') \in \mathrm{Hom}(U, W) \times \mathrm{Gr}_d(W) \,|\, \mathrm{Im}(q) \subset U'\}.$$

Results II: Quantitative

Path $\omega : i \rightsquigarrow j$ in Q given by $\omega : i = i_0 \stackrel{\alpha_1}{\to} i_1 \stackrel{\alpha_2}{\to} \ldots \stackrel{\alpha_s}{\to} i_s = j$.
Define $V_\omega := V_{\alpha_s} \circ \ldots \circ V_{\alpha_1} : V_i \to V_j$.

Theorem

$\mathcal{M}$ admits closed embedding into $\bigoplus_{\omega:i\rightsquigarrow j} \mathrm{Hom}_k(U_i, W_j)$ via

$$\pi : ((V_\alpha), (f_i), (h_i)) \mapsto (h_j V_\omega f_i)_{\omega:i\rightsquigarrow j}.$$

Observation:

$$\Big(\sum_{\omega:i\rightsquigarrow j} h_j V_\omega f_i\Big)_{i,j} : \bigoplus_i U_i \to \bigoplus_j W_j$$

is the network function computed by (V, f, h).

Results II: Quantitative

Theorem

- *Via the embedding* $\mathcal{M} \subset \bigoplus_{\omega:i\rightsquigarrow j} \mathrm{Hom}(U_i, W_j)$, *the network function*
$$\mathcal{M} \to \mathrm{Hom}(\bigoplus_i U_i, \bigoplus_j W_j)$$
is $(q_\omega)_\omega \mapsto (\sum_{\omega:i\rightsquigarrow j} q_\omega)_{i,j}$.
- $\mathcal{M}$ *consists of all* $(q_\omega)_\omega$ *such that* $\mathrm{rk}(q^{(i)}) \leq d_i$ *for all* $i \in Q_0$.
- *Equality of ranks describes* $\mathcal{M}^{\mathrm{simp}} \subset \mathcal{M}$.

Here
$$q^{(i)} = \bigoplus_{\omega:j\rightsquigarrow i} \bigoplus_{\omega':i\rightsquigarrow k} q_{\omega'\omega} : \bigoplus_{\omega:j\rightsquigarrow i} U_j \to \bigoplus_{\omega':i\rightsquigarrow k} W_k.$$

Results II: Quantitative

$\prod_{i\in Q_0} \mathrm{Gr}^{d_i}(\bigoplus_{k\rightsquigarrow i} U_k)$ product of Grassmannians of d_i-codimensional subspaces. Contains $\mathrm{Gr}^{\mathbf{d}}(\mathbf{p})$: tuples of subspaces $(X_i)_i$ such that for all $\alpha : i \to j$:

$$X_i \subset X_j \text{ under inclusion } \bigoplus_{k\rightsquigarrow i} U_k \subset \bigoplus_{k\rightsquigarrow j} U_k.$$

Theorem

$\mathrm{Gr}^{\mathbf{d}}(\mathbf{p})$ isomorphic to tower of Grassmann bundles (induction from source to sink). Carries tautological bundles $\mathcal{X}_i$.

$$\mathcal{M}^{\Theta^+} \simeq \bigoplus_{i\in Q_0} \mathrm{Hom}(\mathcal{X}_i, W_i).$$

Consequently $\sum_i \dim H^i(\mathcal{M}^{\Theta^+})q^{i/2} = \prod_{i\in Q_0} \begin{bmatrix} p_i+\sum_{j\to i} d_j \\ d_i \end{bmatrix}_q$.

Similarly for $\mathcal{M}^{\Theta^+}$.

Methods of proof: representation-theoretic interpretation

$\mathbf{p}$ defines projective representation $P = \bigoplus_i P_i \otimes U_i$, and $\mathbf{q}$ defines injective representation $I = \bigoplus_i I_i \otimes W_i$.

Definition

Define category $\mathcal{C}$ with objects $P \to V \to I$, morphisms are compatible maps of quiver reps.

Observation: quotient stack $[R_{\mathbf{d},\mathbf{p},\mathbf{q}}(Q)/G_{\mathbf{d}}]$ isomorphic to stack of isoclasses $\mathrm{Iso}_{\mathbf{d}}(\mathcal{C})$.

- Simples: $P \twoheadrightarrow V \hookrightarrow I$
- Θ^+-stables: $P \twoheadrightarrow V \to I$
- Θ^--stables: $P \to V \hookrightarrow I$

Well-defined map $\mathrm{Iso}_{\mathbf{d}}(\mathcal{C}) \to \mathrm{Hom}_Q(P, I)$:

$$[P \xrightarrow{f} V \xrightarrow{h} I] \mapsto h \circ f.$$

$$\mathcal{M} \simeq \{\varphi \in \mathrm{Hom}_Q(P, I) \,|\, \mathbf{rk}(\varphi) \leq \mathbf{d}\},$$

$$\mathcal{M}^{\mathrm{simp}} \simeq \{\varphi \in \mathrm{Hom}_Q(P, I) \,|\, \mathbf{rk}(\varphi) = \mathbf{d}\},$$

$$\mathcal{M}^{\Theta^+} \simeq \{(U, \varphi) \in \mathrm{Gr}^{\mathbf{d}}(P) \times \mathrm{Hom}_Q(P, I) \,|\, U \subset \mathrm{Ker}(\varphi)\},$$

$$\mathcal{M}^{\Theta^-} \simeq \{(\varphi, U') \in \mathrm{Hom}_Q(P, I) \times \mathrm{Gr}_{\mathbf{d}}(I) \,|\, \mathrm{Im}(\varphi) \subset U'\}.$$

Towards more realistic neural networks

In "reality", base field $= \mathbb{R}$, cector spaces $V_i = \mathbb{R}$ one-dimensional.
Neurons carry (non-linear!) *activation function* $a_j : \mathbb{R} \to \mathbb{R}$:

$$\begin{array}{ccccc} V_{i_1} & & \ldots & & V_{i_m} \\ & f_1 \searrow & & f_m \swarrow & \\ & & {\color{red} V_j} & & \\ & h_1 \swarrow & & h_n \searrow & \\ V_{k_1} & & \ldots & & W_{i_n} \end{array}$$

computing

$$(v_{i_1}, \ldots, v_{i_m}) \mapsto (h_l({\color{red} a_j(} \sum_k f_k(v_{i_k}){\color{red})}))_l.$$

Typically $a_j(x) = \mathrm{ReLU}(x) := \max(0, x)$.

$\rightsquigarrow$ symmetry broken from $G_{\mathbf{d}} = (\mathbb{R}^*)^{Q_0}$ to $(\mathbb{R}^+)^{Q_0}$.

Towards more realistic neural networks

In this case, use description of quiver moduli via symplectic reduction (A King). It works over $\mathbb{R}$!
Momentum map

$$\mu : R_{\mathbf{1},\mathbf{p},\mathbf{q}}(Q)_{\mathbb{R}} \to \mathbb{R}^{Q_0},$$

$$\mu((V_\alpha),(f_i),(h_i)) \mapsto (\sum_{\alpha:\to i} V_\alpha^2 - \sum_{\alpha:i\to} V_\alpha^2 + ||f_i||^2 - ||h_i||^2)_i.$$

Fact

$$\mathcal{M}_{\mathbb{R}} \simeq \mu^{-1}(0)/(\underbrace{O_1(\mathbb{R})}_{=\pm 1})^{Q_0}.$$

Thus the non-linear network functions are well-defined on

$$R_{\mathbf{1},\mathbf{p},\mathbf{q}}//(\mathbb{R}^+)^{Q_0} := \mu^{-1}(0).$$

Towards more realistic neural networks

$$R_{\mathbf{1},\mathbf{p},\mathbf{q}}//(\mathbb{R}^+)^{Q_0} := \mu^{-1}(0).$$

Theorem

$R_{\mathbf{1},\mathbf{p},\mathbf{q}}//(\mathbb{R}^+)^{Q_0}$ admits explicit description as set of tuples

$$((V_\alpha \in \mathbb{R})_\alpha, (f_i \in U_i^*)_i, (h_i \in W_i)_i)$$

such that for all $i \in Q_0$:

$$\sum_{\alpha:\to i} V_\alpha^2 + |f_i|^2 = \sum_{\alpha:i\to} V_\alpha^2 + |h_i|^2.$$

Final frame

Thank you!