

Kaehler geometry of quiver varieties and machine learning

S. Lau, G. Jeffrey. "Kaehler Geometry of Quiver Varieties and Machine Learning."

<https://arxiv.org/abs/2101.11487> and get abstract from the link

<https://hdl.handle.net/2144/44110>

Downloaded from DSpace Repository, DSpace Institution's institutional repository

KÄHLER GEOMETRY OF QUIVER VARIETIES AND MACHINE LEARNING

GEORGE JEFFREYS AND SIU-CHEONG LAU

ABSTRACT. We develop an algebro-geometric formulation for neural networks in machine learning using the moduli space of framed quiver representations. We find natural Hermitian metrics on the universal bundles over the moduli which are compatible with the GIT quotient construction by the general linear group, and show that their Ricci curvatures give a Kähler metric on the moduli. Moreover, we use toric moment maps to construct activation functions, and prove the universal approximation theorem for the multi-variable activation function constructed from the complex projective space.

Keywords: Quiver varieties; Kähler geometry; Representation theory; Moduli spaces; Toric geometry; Neural networks

1. INTRODUCTION

Machine learning by artificial neural networks has made exciting developments and has been applied to many branches of science in recent years. Mathematically, stochastic gradient flow over a matrix space (or called the weight space) is the central tool. The non-convex nature of the cost function has made the problem very interesting. Current research has focused on different types of stochastic gradient flows and finding new types of networks, which have brought great improvements of computational efficiency.

In geometry and physics, the applications of gradient flow and Morse theory have a long history and have brought numerous fundamental breakthroughs. For instance, the gradient flow of the Yang-Mills functional is used to find Hermitian Yang-Mills connections, whose existence in a stable holomorphic vector bundle is proved by Donaldson [Don85] and Uhlenbeck-Yau [UY86]. The celebrated Ricci flow found by Hamilton [Ham82], which is a crucial tool to solve the three-dimensional Poincaré conjecture, is essentially a gradient flow [Per02, Per03]. Its Kähler analog has been an important tool in finding Kähler-Einstein metrics on Fano manifolds [Yau96, Tia97, Don12, CSW18, CDS15a, CDS15b, CDS15c]. In these works, GIT quotients and finite-dimensional models have provided important motivations and guidelines [Don99]. Hamiltonian Floer theory [Flo89], which is essentially Morse theory on the loop space, was invented to solve the Arnold conjecture [FHS95, Ono95, FO99]. Various versions of Floer theory have been crucial ingredients in the study of mirror symmetry.

In this paper, we would like to develop a foundational *algebro-geometric formulation* for neural networks in machine learning. The theory of *quiver representations*, which is a well-developed branch of mathematics motivated from Lie theory and has been an important tool in mathematical physics, will be well suited for this purpose.

A quiver representation assigns to a directed graph Q a bunch of vector spaces for the vertices and a bunch of linear maps for the arrows. Such a construction is in common with neural networks. However, in order to use quiver theory to formulate machine-learning neural networks, there are two main differences between these two subjects that needs to be addressed.

- (1) Compactness of moduli space. A moduli space of quiver representations [Kin94] is defined by identifying *isomorphic* quiver representations using GIT quotients. As a result, the moduli space is *compact* when the quiver has no oriented cycle. On the other hand, the matrix space used in neural networks is non-compact. In machine learning, isomorphic quiver representations may correspond to physically different input or output information and in general cannot be identified.
- (2) Non-linearity. *Activation functions*, which are non-linear maps on the vector spaces over the vertices, serve as a crucial ingredient to achieve machine learning of non-linear functions. Such non-linearity jumps out of the category of quiver representations. This is also related to the first point above. Namely, such non-linear maps are not necessarily equivariant under the group of automorphisms of quiver representations.

For the first point, we shall use framed quiver representations, which were first found by Nakajima [Nak94] in the study of affine Lie algebras. A framed representation assigns to each vertex a vector space together with a choice of ‘framing’ (for instance it is a basis in the simplest situation). In the applications considered here, such a decoration makes sure that isomorphic framed quiver representations correspond to the same physical state. Note that framed quiver moduli \mathcal{M} are also compact when Q has no oriented cycle. *Compactness* is one of the main advantages of our algebro-geometric formulation, which makes sure the convergence of a gradient flow.

In this formulation, the weight matrices are encoded as morphisms between the *universal vector bundles* (over the framed quiver moduli) associated to the vertices. The data flow is encoded by sections of the universal bundles, which are sent from one to another bundles by the morphisms associated to the arrows of Q . The cost function, and hence its gradient flow, is defined on the framed quiver moduli \mathcal{M} .

In particular, the *critical points and the gradient flow are controlled by the topology of \mathcal{M}* (for instance, the Morse inequalities). The topology of a framed quiver moduli is well-understood by the work of Reineke [Rei08] when Q has no oriented cycle. \mathcal{M} is an iterated Grassmann bundle, and its Poincaré polynomial is a product of that of the Grassmannians.

For the purpose of gradient flow, one needs to choose a *Kähler metric* on \mathcal{M} , and also Hermitian metrics on the universal vector bundles. As a result, we have found metrics that are defined by explicit beautiful formulae. These metrics are not just $U_{\bar{d}}$ -equivariant so that they descend to symplectic quotients, but are also $GL_{\bar{d}}$ -equivariant and hence compatible with the GIT construction of \mathcal{M} . Moreover, they are compatible with the iterated Grassmann structure found by Reineke. In application, such metrics would simplify the actual computational algorithm over the quiver moduli. They are summarized as follows.

Theorem 1.1 (Combining Theorem 3.7, 3.15, 3.18). *Let Q be an arbitrary quiver. Fix a vertex $i \in Q_0$. Let ρ be the row vector whose entries are $V_\gamma e^{t(\gamma)}$, where γ is any path whose head $h(\gamma)$ is i (including the trivial path), $t(\gamma)$ denotes its tail, and $V_\gamma \in \text{Hom}(\mathbb{C}^{d_{t(\gamma)}}, \mathbb{C}^{d_{h(\gamma)}})$ is the representing matrix of γ . Then*

$$(\rho_i \rho_i^*)^{-1} = \left(\sum_{h(\gamma)=i} (V_\gamma e^{t(\gamma)}) (V_\gamma e^{t(\gamma)})^* \right)^{-1}$$

is $GL_{\bar{d}}$ -equivariant, and it descends to a metric on the universal bundle \mathcal{V}_i over a certain domain of convergence \mathcal{M}° .

When Q has no oriented cycle, $\mathcal{M}^\circ = \mathcal{M}$. Moreover, the Ricci curvature of the induced metric on $\bigotimes_{i \in Q_0} \mathcal{V}_i$ gives a Kähler metric on \mathcal{M} .

The precise definition of \mathcal{M}° is given in Section 3.3.

In this paper, we focus on the framed quiver moduli defined over complex numbers. In actual applications, we can also restrict to real coefficients. Then the above formula defines a bundle metric over $\mathcal{M}_{\mathbb{R}}$, and the Ricci curvature gives a Riemannian metric on $\mathcal{M}_{\mathbb{R}}$.

Now let us address the second point. Namely, we need to introduce *non-linearity* in addition to the usual theory of quiver representations. By definition, morphisms between universal vector bundles over \mathcal{M} are linear along fibers. They correspond to weight matrices in neural networks. To introduce non-linearity, we shall treat the universal bundles as fiber bundles and construct suitable *fiber-bundle maps* that play the role of activation functions.

One of the commonly-used activation functions is

$$\frac{e^{2x}}{1 + e^{2x}} : \mathbb{R} \rightarrow (0, 1).$$

We observe that this function also appears in the base of the symplectic trivialization of the open dense toric orbit of \mathbb{P}^1 as a toric variety:

$$(\mathbb{C}^\times, \omega_{\mathbb{C}\mathbb{P}^1}) \cong ((0, 1) \times \mathbb{S}^1, \omega_{\text{std}}),$$

or lifted to the universal cover:

$$(\mathbb{C}, \pi^* \omega_{\mathbb{C}\mathbb{P}^1}) \cong ((0, 1) \times \mathbb{R}, \omega_{\text{std}}).$$

Here, $\omega_{\mathbb{C}\mathbb{P}^1}$ denotes the Fubini-Study metric of $\mathbb{C}\mathbb{P}^1$, that is, the standard area form of the unit sphere; $\omega_{\text{std}} = dx \wedge dy$ is the standard symplectic form.

Similarly, another activation function

$$\frac{z}{\sqrt{1 + |z|^2}} : \mathbb{C} \rightarrow \mathbb{C}$$

also arises as a symplectic trivialization: $(\mathbb{C}, \omega_{\mathbb{C}\mathbb{P}^1}) \cong (\{w \in \mathbb{C} : |w| < 1\}, \omega_{\mathbb{C}})$.

Motivated from these observations, we consider

$$\sigma(x) = \left(\frac{e^{2x_i}}{1 + \sum_{j=1}^n e^{2x_j}} \right)_{i=1}^n : \mathbb{R}^n \rightarrow \Delta$$

and

$$\psi(\vec{z}) = \left(\frac{z_i}{\sqrt{1 + \|\vec{z}\|^2}} \right)_{i=1}^n : \mathbb{C}^n \rightarrow \{\vec{w} \in \mathbb{C}^n : \|\vec{w}\| < 1\}$$

as multi-variable activation functions, where Δ denotes the standard simplex with vertices 0 and $\epsilon_1, \dots, \epsilon_n$, the standard basis of \mathbb{R}^n . They arise from symplectomorphisms $(\mathbb{C}^n, \omega_{\mathbb{C}\mathbb{P}^n}) \cong (\{\vec{w} \in \mathbb{C}^n : \|\vec{w}\| < 1\}, \omega_{\mathbb{C}^n})$. More generally, these come from moment maps of *toric varieties* [Gui94, Abr98]. We note that ψ has an advantage of being $U(n)$ -equivariant.

The *universal approximation theorem* (see for instance [Cyb89, Pet99, MM92, Pin99]) provides a theoretical foundation for the success of neural networks. In existing literature, the theorem was proved for single-variable activation functions.

In this paper, we prove the universal approximation theorem for the above multivariable function $\sigma = \sigma_{\mathbb{R}^n}$. Note that σ is the softmax function $\left(\frac{e^{2x_i}}{\sum_{j=0}^n e^{2x_j}} \right)_{i=0}^n$ restricted to the

hyperplane $x_0 = 0$ and composed with the projection along x_0 -direction. We shall restrict to real coefficients in this theorem.

Theorem 1.2 (same as Theorem 5.2). *Let K be a compact set of \mathbb{R}^{d_1} , and $f : K \rightarrow \mathbb{R}^{d_3}$ a continuous function. For any $\epsilon > 0$, there exists $d_2 > 0$ and $W_1 \in \text{Mat}(d_2, d_1)$, $W_2 \in \text{Mat}(d_3, d_2)$, $b \in \mathbb{R}^{d_2}$ such that $\|f_{W_1, W_2, b}^U - f\|_{L^2(K)} < \epsilon$. Here, $f_{W_1, W_2, b}^U(x) = W_2 \cdot \sigma_{\mathbb{R}^{d_1}}(W_1 \cdot x + b)$ is the function coming from the A_3 quiver.*

The A_3 quiver corresponds to the feed-forward network with one input layer, one middle layer and one output layer. See Figure 1.

$$a \xrightarrow{\alpha} b \xrightarrow{\beta} c$$

FIGURE 1. The A_3 quiver.

The above theorem is proved by using the tropical limit of the toric manifold \mathbb{P}^n , and a geometric object that we call a centered polyhedral web, which is an analog of a tropical variety in an integral affine manifold. Since we do not have integral structure in the context here, we need to invent this new notion.

In above, we have focus on explaining non-linearity for a single vector space. We shall *globalize them as non-linear fiber-bundle morphisms* for the universal bundles over \mathcal{M} . This can be achieved with the help of Hermitian metrics on the universal bundles, so that the Fubini-Study metric on \mathbb{P}^n can be globalized as a fiberwise symplectic structure on projective bundles over \mathcal{M} . Actually, the globalization from a single framing vector space V to the universal fiber bundle over \mathcal{M} works for any continuous function $V \rightarrow V$ (and in particular for a symplectomorphism from V to its image). Combining the ingredients explained above together, we can construct a gradient flow over the framed quiver moduli to achieve machine learning. The detail is given in Section 4.

Such an algebro-geometric formulation has several *advantages*. First, the gradient flow under consideration runs in a *compact* manifold. This ensures the existence of absolute extrema, convergence of the flow, and upper bound for the norm of the gradient vector field. Second, because of compactness, the flow is constrained by *topology* of the manifold due to Morse theory. See Section 4.4. Finally, the moduli space has extra *symmetry* coming from framing. If we use activation functions that respect this symmetry (for instance ψ above enjoys $U(n)$ -equivariance), we can perform dimension reduction which improves the effectiveness of the network. (See Proposition 4.16 and 4.17.)

In summary, from this point of view, the success of neural network is resulted from the interplay between algebraic morphisms and (transcendental) symplectomorphisms. Interestingly, such an interplay is also an important feature that occurs in the study of complete integrable systems and mirror symmetry for toric manifolds and flag varieties, see for instance [Gui94, FLTZ12, Abo06, CLL12, NNU10, HKL18].

Some related works. The relation between neural network and quiver representation was investigated in the recent paper [AJ20]. Their work considered the quotient space by $(\mathbb{C}^\times)^N$ of pairs (W, f) , where W is a quiver representation of Q with the dimension vector $\vec{1}$, and f associates each vertex a function $\mathbb{C} \rightarrow \mathbb{C}$ (playing the role of an activation function). Moreover, in dimension $\vec{1}$ (which is a typical case for machine learning), they invented an interesting way of encoding the data flow as a quiver representation. (In our work, the data flow is given as sections of universal bundles over the quiver moduli.)

The approach and the goal of this paper is rather different. We aim at formulating machine learning as a gradient flow over a compact quiver moduli. In [AJ20], quiver representations were used in encoding the data in the network; however, the machine learning process was independent of the quiver moduli. Moreover, ‘double-framing’ was used, and the corresponding moduli space is non-compact.

The map $(W, f) \mapsto W$ gives an infinite-dimensional fiber bundle over the quiver moduli \mathcal{M} , whose fibers are the spaces of choices of activation functions. In a typical program of machine learning, the activation functions are fixed during the optimization process. In order to formulate the program as a gradient flow over the compact moduli \mathcal{M} , we found a non-trivial way by equipping intermediate vertices with additional framings and metrics, so that we can lift f to be a well-defined fiber-bundle map over \mathcal{M} . Note that f is not equivariant under the group action of $(\mathbb{C}^\times)^N$ ($GL_{\vec{d}}$ in the higher rank case). Such a lifting is an important non-trivial step.

Furthermore, we have dealt with representations of general rank \vec{d} , and a class of activation functions coming from toric symplectomorphisms. Different functions (on the same domain and target) are obtained if we deform the toric Kähler metric. To also optimize the activation functions during the learning process (see also [GGL19]), we may consider a gradient flow on $\mathcal{M} \times \mathcal{K}$ where \mathcal{K} denotes the moduli of toric Kähler metrics in the same class. By the celebrated works of [Don99, Sem92], \mathcal{K} is an infinite-dimensional negatively curved symmetric space.

Recently, there is a rising interest of applying geometric techniques to the study of neural networks. For instance, in the works [GBH18, CYRL19], hyperbolic spaces are applied to machine learning in graphs and achieved great performance.

Moreover, the applications of symmetry and group equivariance in neural networks were studied and developed in [CW16, CGW19, CGKW18, CWKW19, CAW⁺19, dHCW20]. Overall, these works aim at capturing symmetry of the input data and designing networks that are adapted to such symmetry. Moreover, homogeneous spaces (in place of vector spaces) have been employed in layers of convolutional neural networks.

In comparison, our paper aims at revealing the geometric nature of neural networks and build a connection with algebraic geometry. The resulting framed quiver moduli, which has interesting topology and metrics, is the main geometric object of interest. Furthermore, we study activation functions that respects the ‘intrinsic symmetry’ over the quiver moduli, which can provide a more effective algorithm by dimension reduction.

In the reverse direction, there are interesting applications of machine learning in frontier geometry and physics. For instance, [HY20] used machine learning to solve problems in computing graph Laplacians, such as recognizing graph Ricci-flatness and predicting the spectral gap. In physics, [HSTT18a, HSTT18b] used deep learning to study AdS/CFT correspondence by discretizing the equation of motion. Since we have formulated machine learning using quiver representations, it will be interesting to find direct relations between quiver gauge theory and these problems that can be attacked via machine learning.

Organization of this paper. In Section 2, we will take a quick review on quiver representations and their moduli spaces. In Section 3, we will construct nice Hermitian metrics on universal bundles over the moduli. For readers who are mainly interested in machine learning, Section 3 can be skipped for the first reading. Then we give an algebro-geometric formulation of neural network using quiver representations in Section 4. In Section 5, we prove the universal approximation theorem for the multivariable activation function σ .

ACKNOWLEDGMENT

We are grateful to Marco Antonio Armenta for informing us about the work [AJ20] and the further useful discussions. We express our gratitude to Shing-Tung Yau for his generous encouragement. The work of S.C. Lau in this paper is partially supported by the Simons collaboration grant.

2. REVIEW OF FRAMED QUIVER MODULI

Let Q be a directed graph. Denote by Q_0, Q_1 the set of vertices and arrows respectively. A quiver representation V with dimension vector $\vec{d} \in \mathbb{Z}_{\geq 0}^{Q_0}$ associates each arrow a with a matrix $V(a)$ of size $d_{h(a)} \times d_{t(a)}$ (where $h(a), t(a)$ denote the head and tail vertices of a respectively). The set of complex quiver representations with dimension \vec{d} form a vector space denoted by $R_{\vec{d}}(Q)$. The moduli space of quiver representations is a GIT quotient of $R_{\vec{d}}(Q)$ by the group of isomorphisms $\mathrm{GL}(\vec{d}) = \prod_{i \in Q_0} \mathrm{GL}(d_i, \mathbb{C})$ [Kin94], where $\mathrm{GL}(\vec{d})$ acts on $R_{\vec{d}}(Q)$ via

$$g \cdot (V(a) : a \in Q_1) = (g_{h(a)} \cdot V(a) \cdot g_{t(a)}^{-1} : a \in Q_1). \quad (2.1)$$

In the applications we consider in this paper, since the vector space over the input and output vertices are equipped with fixed basis with physical meanings, we need to use framed quiver representations [Nak94, Nak01, CB03, Rei08].

Let $\vec{d}, \vec{n} \in \mathbb{Z}_{\geq 0}^{Q_0}$. \vec{n} will be the dimension vector for the framing, which is a linear map $e^{(i)} : \mathbb{C}^{n_i} \rightarrow V_i$ at each $i \in Q_0$ (where $V_i = \mathbb{C}^{d_i}$). Since we will take a quotient by $\mathrm{GL}(\vec{d})$, we shall think of V_i as a vector space without a preferred basis, while \mathbb{C}^{n_i} is equipped with the standard basis.

Definition 2.1. *The vector space of framed representations is given by*

$$R_{\vec{n}, \vec{d}} = R_{\vec{d}} \times \bigoplus_{i \in Q_0} \mathrm{Hom}(\mathbb{C}^{n_i}, \mathbb{C}^{d_i}).$$

It carries a natural action of $\mathrm{GL}(\vec{d})$ given by $g \cdot (V, e) = (g \cdot V, (ge^{(i)} : i \in Q_0))$, where $g \cdot V$ is given by Equation (2.1).

We need to remove unstable framed representations from $R_{\vec{n}, \vec{d}}$ in order to get a nice quotient by $\mathrm{GL}(\vec{d})$.

Theorem 2.2 ([Nak96]). *$(V, e) \in R_{\vec{n}, \vec{d}}$ is called stable if there is no proper subrepresentation U of V which contains $\mathrm{Im} e$. The set of all stable points of $R_{\vec{n}, \vec{d}}$ is denoted by $R_{\vec{n}, \vec{d}}^s$. Then the quotient $\mathcal{M}_{\vec{n}, \vec{d}} := R_{\vec{n}, \vec{d}}^s / \mathrm{GL}(\vec{d})$ is a smooth variety, which is called to be a framed quiver moduli.*

Actually $\mathcal{M}_{\vec{n}, \vec{d}}$ can be formulated as a GIT quotient [CB03, Rei08]. Namely, by adding an extra vertex labeled as ∞ to the quiver and n_i arrows from the vertex ∞ to the vertex i , (V, e) can be identified as a usual representation of this bigger quiver with the dimension vector $(\vec{d}, 1)$. The above stability condition can be rewritten as slope stability, and hence it is a GIT quotient [Kin94].

Since $(d, 1)$ is a primitive vector, $\mathcal{M}_{\vec{n}, \vec{d}}$ is a smooth fine moduli. There are universal vector bundles \mathcal{V}_i over $\mathcal{M}_{\vec{n}, \vec{d}}$ corresponding to each vertex i , with fibers $\mathcal{V}_i|_{[V, e]} = V_i$.

Example 2.3. For the quiver with a single vertex and no arrow, and $n > d$,

$$\mathcal{M}_{n,d} = \text{Gr}(n,d) = \{e \in \text{Hom}(\mathbb{C}^n, \mathbb{C}^d) : e \text{ is surjective}\} / \text{GL}_{\vec{d}}$$

is the (dual) Grassmannian. We have the tautological bundle \mathcal{V} over $\text{Gr}(n,d)$. (Note that this tautological bundle is dual to the one on $\text{Gr}(d,n) \cong \text{Gr}(n,d)$.)

The topology of $\mathcal{M}_{\vec{n},\vec{d}}$ is well-understood. Let's make an ordering of the vertices. Namely the vertices are labeled by $\{1, \dots, N\}$, such that $i < j$ implies there is no arrow going from j to i . Such a labeling exists if Q has no oriented cycle.

Theorem 2.4 (Reineke [Rei08]). Assume Q has no oriented cycle. Consider the chain of iterated Grassmann bundles $M^{(N)} \xrightarrow{p_N} M^{(N-1)} \xrightarrow{p_{N-1}} \dots \xrightarrow{p_2} M^{(1)} \xrightarrow{p_1} \text{pt}$ (where pt denotes a singleton) defined by induction:

$$M^{(i)} = \text{Gr}_{M^{(i-1)}} \left(\frac{\mathbb{C}^{n_i} \oplus \bigoplus_{j \rightarrow i} p_{i-1}^* \dots p_{j+1}^*(S_j), d_i}{d_i} \right) \rightarrow M^{(i-1)},$$

where S_i denotes the tautological bundle on M_i (as a Grassmann bundle over M_{i-1}). (The direct sum is over each arrow $j \rightarrow i$.) Then $\mathcal{M}_{\vec{n},\vec{d}} \cong M^{(N)}$, with universal bundles $\mathcal{V}_i \cong p_N^* \dots p_{i+1}^* S_i$ for all $i \in Q_0$.

Corollary 2.5 (Reineke [Rei08]). The Poincare polynomial of nonempty $\mathcal{M}_{\vec{n},\vec{d}}$ is given by

$$\prod_{i \in Q_0} \binom{n_i + \sum_{j \rightarrow i} d_j}{d_i}_{q^2}$$

where

$$\binom{n}{d}_q = \prod_{k=1}^d \frac{q^{n-d+k} - 1}{q^k - 1}.$$

Remark 2.6. In [Rei08], the framing e goes in the other direction (from \mathbb{C}^{d_i} to \mathbb{C}^{n_i}). The above theorem is stated in the dual way, which is the convention we take for the rest of this paper.

Example 2.7. Consider the A_3 -quiver which has three vertices $i = 1, 2, 3$ and two arrows $a_1 : 1 \rightarrow 2, a_2 : 2 \rightarrow 3$. Suppose $n_1 = d_1, n_2 = d_2 + 1$ and $n_3 = d_3$. Then the iterated Grassmann bundle is $M^{(3)} \rightarrow M^{(2)} \rightarrow M^{(1)}$, where $M^{(1)} = \text{Gr}(d_1, d_1) = \text{pt}$ (and its tautological bundle is the vector space \mathbb{C}^{d_1}); $M^{(2)} = \text{Gr}(d_2 + 1 + d_1, d_2)$ is equipped with the tautological bundle S_2 of rank d_2 ; $M^{(3)} = \text{Gr}_{\text{Gr}(d_2+1+d_1, d_2)}(\mathbb{C}^{d_3} \oplus S_2, d_3)$ is a Grassmannian bundle over $\text{Gr}(d_2 + 1 + d_1, d_2)$ with fibers $\text{Gr}(d_2 + d_3, d_3)$. The corresponding Poincare Polynomial will be

$$\left(\prod_{k=1}^{d_3} \frac{q^{2(d_2+k)} - 1}{q^{2k} - 1} \right) \left(\prod_{k=1}^{d_2} \frac{q^{2(d_1+1+k)} - 1}{q^{2k} - 1} \right).$$

See Figure 2.

$$\begin{array}{ccc}
 1 & \longleftarrow & Gr(d_1, d_1) = \text{pt} \\
 \downarrow a_1 & & \uparrow \\
 2 & \longleftarrow & Gr(d_2 + 1 + d_1, d_2) \\
 \downarrow a_2 & & \uparrow \\
 3 & \longleftarrow & Gr_{Gr(d_2+1+d_1, d_2)}(\underline{\mathbb{C}}^{d_3} \oplus S_2, d_3)
 \end{array}$$

 FIGURE 2. The iterated Grassmann bundles associated to the A_3 quiver.

3. HERMITIAN METRIC OVER FRAMED QUIVER MODULI

In constructing fiber-bundle endomorphisms, it will be crucial to consider Kähler metrics on universal bundles. In this section, we find a beautiful formula for the canonical metric on the universal bundle \mathcal{V}_i over $\mathcal{M}_{n,d}$ written in homogeneous coordinates. Using this formula, we then show that the sum of Ricci curvatures over the vertices i give a Kähler metric on $\mathcal{M}_{\vec{n}, \vec{d}}$.

First, let us begin by recalling the typical example $Gr(n, k)$.

3.1 The Grassmannian. Consider

$$Gr(n, k) = \text{Mat}_{n,k}^{\mathbb{C}} //_{\chi=1} U(k) = \{e \in \text{Mat}_{n,k}^{\mathbb{C}} : ee^* = I_k\} / U(k)$$

for $n \geq k$. Here we have used the dual description which better matches the frame convention used in this paper. Namely, $Gr(n, k)$ parametrizes k -dimensional quotient vector spaces of a fixed n -dimensional vector space as opposed to k -dimensional subspaces. The moment map for the standard $U(k)$ -action on $\text{Mat}_{n,k}^{\mathbb{C}}$ is $ee^* : \text{Mat}_{n,k}^{\mathbb{C}} \rightarrow \mathfrak{iu}_k$. We have taken the moment map level $\chi = 1$ in the above symplectic reduction. Note that $U(k)$ is acting on the left, although in the above expression $U(k)$ appears on the right.

Writing $e = (b, p)$ where $b \in \text{Mat}_{k,k}^{\mathbb{C}}$ and $p \in \text{Mat}_{n-k,k}^{\mathbb{C}}$ the moment-map equation $ee^* = I_k$ becomes

$$bb^* + pp^* = I_k.$$

We shall consider the chart defined by

$$U = \{[b, p] \in Gr(n, k) : \det b \neq 0\} \cong \text{Mat}_{n-k,k}^{\mathbb{C}}$$

where the identification is given by the holomorphic coordinates

$$\zeta^h = b^{-1}p \in \text{Mat}_{n-k,k}^{\mathbb{C}}.$$

We also have the symplectic coordinates

$$\zeta^u = (b^*b)^{\frac{1}{2}} b^{-1}p \in \text{Mat}_{n-k,k}^{\mathbb{C}}.$$

The entries of ζ^u are not meromorphic functions. On the other hand, ζ^u has the advantage that it satisfies the moment-map equation

$$b^*b + \zeta^u(\zeta^u)^* = I_k. \quad (3.1)$$

(Note that the first term is b^*b instead of bb^* .)

The construction of ζ^u uses the polar decomposition

$$b = \left(b (b^*b)^{-\frac{1}{2}} \right) (b^*b)^{\frac{1}{2}}$$

where $(b(b^*b)^{-\frac{1}{2}}) \in U(k)$ and $(b^*b)^{\frac{1}{2}} \in \mathfrak{iu}(k)$ is positive definite. We obtain the coordinates ζ^u by observing

$$[b, p] = \left[(b^*b)^{\frac{1}{2}}, \zeta^u \right]$$

using the left- $U(k)$ -action, such that the first component $(b^*b)^{\frac{1}{2}}$ is Hermitian, and is determined ζ^u due to the moment-map equation (3.1).

The two coordinate systems are related by

$$\zeta^u = (b^*b)^{\frac{1}{2}} \cdot \zeta^h. \quad (3.2)$$

Let S be the tautological vector bundle whose fibers are the quotient vector spaces. (This is dual to the tautological bundle of $\text{Gr}(k, n) \cong \text{Gr}(n, k)$.) It can be written as the quotient of the trivial bundle:

$$S = \left(\{ee^* = I_k\} \times \mathbb{C}^k \right) / U(k)$$

where the left action of $U(k)$ on \mathbb{C}^k is the standard one.

We now take the standard metric on \mathbb{C}^k , which is preserved by $U(k)$ and hence descends to a metric H of S .

Denote the standard basis of \mathbb{C}^k by ϵ_j for $j = 1, \dots, k$. Under this metric, we have the lifting of a local Hermitian frame over the chart $U = \{\det b \neq 0\}$ being

$$u_i = b(b^*b)^{-\frac{1}{2}} \cdot \epsilon_i$$

since $[b, p, b(b^*b)^{-\frac{1}{2}} \epsilon_i] \sim [(b^*b)^{\frac{1}{2}}, \zeta^u, \epsilon_i]$.

We also have the lifting of a local holomorphic frame

$$h_i = b\epsilon_i = b(b^*b)^{\frac{1}{2}} b^{-1}u_i$$

since $[b, p, b\epsilon_i] \sim (I_k, \zeta^h, \epsilon_i)$. The two frames are related as follows.

Lemma 3.1. $h_i = u_j a_i^j$ where $(a_i^j) = (b^*b)^{\frac{1}{2}}$, i is indexing the columns and j is indexing the rows.

Proof. Consider

$$(h_1 \dots h_k) = b(b^*b)^{\frac{1}{2}} b^{-1} (u_1 \dots u_k) = (u_1 \dots u_k) (a_i^j).$$

Thus $(a_i^j) = (u_1 \dots u_k)^{-1} b(b^*b)^{\frac{1}{2}} b^{-1} (u_1 \dots u_k)$.

$$(u_1 \dots u_k) = b(b^*b)^{-\frac{1}{2}}$$

since $u_i = b(b^*b)^{-\frac{1}{2}} \epsilon_i$. Result follows. \square

Proposition 3.2. The metric H defined above on the tautological bundle S is represented by the matrix $(I_k + \zeta^h (\zeta^h)^*)^{-1}$ in the local holomorphic frame h_i and the local coordinates ζ^h .

Proof. Using Lemma 3.1,

$$(H(h_i, h_p)) = \sum_{j,l} \left(H(a_i^j u_j, a_p^l u_l) \right) = \sum_{j,l} \left(\bar{a}_i^j a_p^l H(u_j, u_l) \right) = \sum_j \left(\bar{a}_i^j a_p^j \right) = b^*b$$

where i is indexing the rows and p is indexing the columns. By the moment map equation (3.1) and the relation (3.2),

$$b^*b + (b^*b)^{\frac{1}{2}} \cdot \zeta^h \left(\zeta^h \right)^* (b^*b)^{\frac{1}{2}} = I_k.$$

Then

$$I_k + \zeta^h \left(\zeta^h \right)^* = (b^*b)^{-1}.$$

Hence

$$H = \left(I_k + \zeta^h \left(\zeta^h \right)^* \right)^{-1}.$$

□

Example 3.3. Let's consider the simplest example:

$$\mathbb{P}^1 = \text{Gr}(2,1) = (\mathbb{C}^2 - \{0\})/\mathbb{C}^\times = \mathbb{S}^3/U(1).$$

The tautological bundle for $\text{Gr}(2,1)$ is

$$S = (\mathbb{S}^3 \times \mathbb{C})/U(1)$$

where $U(1)$ acts on \mathbb{C} in the standard way, and it acts on both factors on the left. (Note that this is dual to the usual notion of the tautological bundle of $\text{Gr}(1,2) = \mathbb{P}^1$, since we are now considering the family of quotient lines of \mathbb{C}^2 , which are dual to subspaces of \mathbb{C}^2 .)

Let's take the standard metric on \mathbb{C} . We have the local Hermitian frame (over $z_1 \neq 0$) u given by

$$(z_1, z_2, z_1/|z_1|) \stackrel{U(1)}{\sim} (|z_1|, \zeta^u, 1)$$

where ζ^u is the coordinate of \mathbb{P}^1 which belongs to the open unit disc, and $|z_1|$ is determined by the moment-map equation

$$|z_1|^2 + |z_2|^2 = |z_1|^2 + |\zeta^u|^2 = 1.$$

We also have the local holomorphic frame h defined by

$$(z_1, z_2, z_1) \stackrel{\mathbb{C}^\times}{\sim} (1, \zeta^h, 1).$$

The unitary and holomorphic coordinates are related by $\zeta^u = |z_1| \cdot \zeta^h = (1 - |\zeta^u|^2) \cdot \zeta^h$ for $\zeta^h \in \mathbb{C}$. The frames are related by

$$h = |z_1| \cdot u.$$

The Hermitian frame u always have length one. Writing the metric in the holomorphic frame h :

$$\begin{aligned} |h|^2 &= |z_1|^2 = 1 - |\zeta^u|^2 = 1 - |z_1|^2 |\zeta^h|^2 \\ &= 1 - \left(1 - |z_1|^2 |\zeta^h|^2 \right) |\zeta^h|^2 = \dots \\ &= \frac{1}{1 + |\zeta^h|^2}. \end{aligned}$$

This is the standard metric on $\mathcal{O}_{\mathbb{P}^1}(1)$, whose curvature gives the Fubini-Study metric on \mathbb{P}^1 .

3.2 Metric on framed quiver moduli. We have seen that the standard metric on the trivial bundle over $\text{Mat}_{n,k}^{\mathbb{C}}$ descends to give the standard metric on $\text{Gr}(n,k)$. However, it turns out that for the framed quiver moduli, the standard metric on the trivial bundle over $R_{\vec{n},\vec{d}}$ is not good from the GIT quotient point of view, namely it is not equivariant under $\text{GL}_{\vec{d}}$. In this section, we find a nice metric over $R_{\vec{n},\vec{d}}$ which is equivariant under $\text{GL}_{\vec{d}}$.

Recall from the last section that $\mathcal{M}_{\vec{n},\vec{d}} = R_{\vec{n},\vec{d}}^s / \text{GL}_{\vec{d}}$. The universal bundle over the vertex i is given by

$$\mathcal{V}_i = \left(R_{\vec{n},\vec{d}}^s \times \mathbb{C}^{d_i} \right) / \text{GL}_{\vec{d}}$$

where $\text{GL}_{\vec{d}}$ acts diagonally on the left, the factor $\text{GL}(d_i, \mathbb{C})$ of $\text{GL}_{\vec{d}}$ acts on \mathbb{C}^{d_i} in the standard way, and other factors of $\text{GL}_{\vec{d}}$ act trivially on \mathbb{C}^{d_i} .

There is an equivalent description of $\mathcal{M} = \mathcal{M}_{\vec{n},\vec{d}}$ and the universal bundle \mathcal{V}_i in terms of symplectic quotient. Namely, let $\mu : R_{\vec{n},\vec{d}} \rightarrow \mathfrak{iu}_{\vec{d}}$ be the moment map. Explicitly, $\mu = (\mu_i)_{i \in Q_0}$ where

$$\mu_i = e^{(i)} (e^{(i)})^* - \sum_{t(a)=i} V_a^* V_a + \sum_{h(a')=i} V_{a'} V_{a'}^*.$$

Then define

$$\mathcal{M}_{\vec{n},\vec{d}} = \mu^{-1} \{-c\} / \text{U}_{\vec{d}}$$

for the following level c .

Lemma 3.4. *The slope stability condition $(1, \vec{0}) \in \mathbb{C}^{\hat{Q}_0}$ corresponds to the moment-map level $c = (-I_{d_i})_{i \in Q_0} \in \mathfrak{iu}_{\vec{d}}$, where I_k denotes the identity matrix of rank k .*

Proof. The character taken in King's stability [Kin94] corresponding to $(1, \vec{0})$ is

$$(1 + \Sigma \vec{d}) \left((1, \vec{0}) - \frac{1}{1 + \Sigma \vec{d}} (1, \vec{1}) \right) = (\Sigma \vec{d}, -1, \dots, -1)$$

where $\Sigma \vec{d} = \sum_{i \in Q_0} d_i$, and the first entry is over the root vertex. (Note that there is no group action over the root vertex.) Thus we should take c to be $-I_{d_i}$ over each vertex i . \square

The universal bundle over the vertex i is then given by

$$\mathcal{V}_i = \left(\mu^{-1} \{I_{\vec{d}}\} \times \mathbb{C}^{d_i} \right) / \text{U}_{\vec{d}}.$$

Let's review some very basic definitions about group actions.

Definition 3.5. *Suppose a Lie group G acts on a vector bundle $V \xrightarrow{\pi} M$ equivariantly, namely, $g \circ \pi = \pi \circ g$ for all $g \in G$, and the action is fiberwise linear. A metric H on V is said to be G -equivariant if*

$$H_x(v, w) = H_{g \cdot x}(g \cdot v, g \cdot w).$$

Writing in matrix form when $G = \text{GL}(n, \mathbb{C})$, the above equation is $v^* \cdot H_x \cdot w = v^* \cdot (g^* \cdot H_{g \cdot x} \cdot g) \cdot w$, that is,

$$(g^*)^{-1} \cdot H_x \cdot g^{-1} = H_{g \cdot x}. \quad (3.3)$$

The following easily follows from the definition.

Lemma 3.6. *Suppose a Lie group G acts on a vector bundle $V \xrightarrow{\pi} M$ equivariantly and fiberwise linearly, and the action of G on M is free and proper. A Hermitian form H on V descends to the corresponding bundle over the quotient M/G if and only if H is G -equivariant.*

For framed quiver varieties, we have the framing map $e^{(j)}: R_{\vec{n}, \vec{d}}^s \rightarrow \text{Hom}(\mathbb{C}^{n_j}, \mathbb{C}^{d_j})$ for each vertex j . Using this, we cook up a $\text{GL}_{\vec{d}}$ -invariant Hermitian form on the trivial bundle $\underline{\mathbb{C}}^{d_i} \rightarrow R_{\vec{n}, \vec{d}}$, which descends to a metric on $\mathcal{V}_i \rightarrow \mathcal{M}_{\vec{n}, \vec{d}}$.

Theorem 3.7. *Suppose Q has no oriented cycle. Fix $i \in Q_0$. Let ρ be the row vector whose entries are $V_\gamma e^{(t(\gamma))}$, where γ is any path whose head $h(\gamma)$ is i (including the trivial path), $t(\gamma)$ denotes its tail, and $V_\gamma \in \text{Hom}(\mathbb{C}^{d_{t(\gamma)}}, \mathbb{C}^{d_{h(\gamma)}})$ is the representing matrix of γ . This defines a map*

$$V_\gamma e^{(t(\gamma))} : R_{\vec{n}, \vec{d}} \rightarrow \text{Hom}(\mathbb{C}^{n_{t(\gamma)}}, \mathbb{C}^{d_i}).$$

Take

$$\rho\rho^* = \sum_{h(\gamma)=i} \left(V_\gamma e^{(t(\gamma))} \right) \left(V_\gamma e^{(t(\gamma))} \right)^* : R_{\vec{n}, \vec{d}} \rightarrow \text{End}(\mathbb{C}^{d_i}).$$

Then $H = (\rho\rho^*)^{-1}$ is $\text{GL}_{\vec{d}}$ -equivariant, and it descends to a metric on \mathcal{V}_i over $\mathcal{M}_{\vec{n}, \vec{d}}$.

Proof. $(\rho\rho^*)^{-1}$ is $\text{GL}_{\vec{d}}$ -equivariant:

$$\begin{aligned} H_{\left(g_{h(a)} V_a g_{t(a)}^{-1} g_j e^{(j)} \right)_{a \in Q_1, j \in Q_0}} &= \left(\sum_{\gamma} \left(g_{h(\gamma)} V_\gamma e^{(t(\gamma))} \right) \left(g_{h(\gamma)} V_\gamma e^{(t(\gamma))} \right)^* \right)^{-1} \\ &= (g_i^*)^{-1} \left(\sum_{\gamma} \left(V_\gamma e^{(t(\gamma))} \right) \left(V_\gamma e^{(t(\gamma))} \right)^* \right)^{-1} g_i^{-1} \end{aligned}$$

for all $g \in \text{GL}_{\vec{d}}$. By Lemma 3.6, it descends to the bundle \mathcal{V}_i of the quotient.

Then we prove that the matrix-valued function $H_{(V_a, e^{(j)})_{a \in Q_1, j \in Q_0}} = (\rho\rho^*)^{-1}$ defines a positive-definite metric on the trivial bundle $\underline{\mathbb{C}}^{d_i}$ over the moment map level $\mu^{-1}(I_{\vec{d}})$ (rather than the whole $R_{\vec{n}, \vec{d}}$). We prove by induction on the vertices that $\rho\rho^* = I_{d_i} + B$ where B is a semi-positive-definite Hermitian matrix, and hence $\rho\rho^*$ is positive definite (and so does $(\rho\rho^*)^{-1}$).

Since the quiver does not have oriented cycle, Q_0 can be ordered such that $i < j$ whenever there is an arrow $i \rightarrow j$. Let i_0 be the minimal vertex.

At i_0 , there is no incoming arrow (other than the framing), and the moment-map equation reads

$$e^{(i_0)}(e^{(i_0)})^* = I_{d_{i_0}} + \sum_{t(a)=i_0} V_a^* V_a.$$

$\sum_{t(a)=i_0} V_a^* V_a$ is semi-positive definite: $v^* \cdot V_a^* V_a \cdot v = \|V_a \cdot v\|^2 \geq 0$ for any column vector v . Thus the statement is true for $\rho\rho^* = e_{i_0} e_{i_0}^*$.

Suppose the statement is true for all vertices less than $i \in Q_0$. At i , the moment-map equation is

$$e^{(i)}(e^{(i)})^* = I_{d_i} + \sum_{t(a)=i} V_a^* V_a - \sum_{h(a')=i} V_{a'} V_{a'}^*.$$

Then

$$\begin{aligned}
 \rho\rho^* &= e^{(i)}(e^{(i)})^* + \sum_{h(a)=i} V_a \rho_{(t(a))} \rho_{(t(a))}^* V_a^* \\
 &= I_{d_i} + \sum_{t(a)=i} V_a^* V_a - \sum_{h(a')=i} V_{a'} V_{a'}^* + \sum_{h(a')=i} V_{a'} (I_{d_{t(a')}} + B_{t(a')}) V_{a'}^* \\
 &= I_{d_i} + \sum_{t(a)=i} V_a^* V_a + \sum_{h(a')=i} V_{a'} B_{t(a')} V_{a'}^*
 \end{aligned}$$

where $\rho_{(t(a))} = \left(V_\gamma e^{t(\gamma)} \right)_{h(\gamma)=t(a)}$, which by inductive assumption can be written as $I_{d_{t(a')}} + B_{t(a')}$ where $B_{t(a')}$ is semi-positive definite. The matrices $V_a^* V_a$ and $V_{a'} B_{t(a')} V_{a'}^*$ are semi-positive definite:

$$v^* V_{a'} B_{t(a')} V_{a'}^* v = (V_{a'}^* v)^* B_{t(a')} (V_{a'}^* v) \geq 0$$

for all v . This proves the statement for the vertex i . \square

The expression $(\rho\rho^*)^{-1}$ can be understood as follows. ρ^* embeds the dual V_i^* into the dual frame which is a trivial bundle equipped with the standard metric. This gives an induced metric on V_i^* , which is $\rho\rho^*$ written in matrix form. Taking the dual, we get the metric $H_i = (\rho\rho^*)^{-1}$ on V_i .

By construction, the metrics on the dual \mathcal{V}_i^* (still denoted as H_i) have the following nice property. Inductively, it gives nice expressions of H_i in terms of *holomorphic coordinates*.

Proposition 3.8. *Suppose Q has no oriented cycle. For $v, w \in (\mathcal{V}_i)^*$,*

$$H_i(v, w) = H_0((e^{(i)})^*(v), (e^{(i)})^*(w)) + \sum_{h(a)=i} H_{t(a)}(a^*(v), a^*(w))$$

where H_0 denotes the trivial metric on the trivial bundle, and $e^{(i)}, a$ are denoting the holomorphic bundle maps corresponding to the framing and arrow maps respectively.

Proof. The metric on \mathcal{V}_i^* is given by the matrix $\rho\rho^*$. Then the above equation follows from

$$\rho\rho^* = e^{(i)}(e^{(i)})^* + \sum_{h(a)=i} V_a \rho_{(t(a))} \rho_{(t(a))}^* V_a^*.$$

\square

Remark 3.9. *As we have seen, the Grassmannian $\text{Gr}(n, k)$ can be understood as the framed moduli for the quiver which has one vertex and no arrow. The matrix $e \in \text{Hom}(\mathbb{C}^n, \mathbb{C}^k)$ is the framing map. Then the moment map equation implies*

$$\rho\rho^* = ee^* = I_k$$

in the above proposition. This is the standard metric on the trivial bundle \mathbb{C}^k that we have used in the last subsection. In particular $\rho\rho^ = (\rho\rho^*)^{-1}$ in this case. But this is not true for other quivers.*

Remark 3.10. *Note that the above becomes an infinite sum if the quiver has oriented cycles. The $\text{GL}_{\vec{a}}$ -equivariance still holds. We should restrict to the open subset of $R_{n, \vec{d}}^s$ that $(\rho\rho^*)^{-1}$ is convergent. In the next subsection, we will prove that the same expression defines a metric for any given quiver.*

The GIT description will be important to the proof of Theorem 3.14.

There is a residual symmetry $U_{\vec{n}} = \prod_{i \in Q_0} U(n_i)$ acting on $\mathcal{M}_{\vec{n}, \vec{d}}$. Actually, there is a bigger symmetry by the non-compact group $\mathrm{GL}(\hat{W})$ [Rei08]. $U_{\vec{n}}$ is considered here since this is the symmetry of the metric H on \mathcal{V}_i as we shall see.

Definition 3.11. *The right residual action of $U_{\vec{n}}$ on \mathcal{M} is defined as*

$$\left[(V_a, e^{(j)})_{a \in Q_1, j \in Q_0} \right] \cdot g = \left[(V_a, e^{(j)} \circ g_j)_{a \in Q_1, j \in Q_0} \right]$$

for $g = (g_j \in U(d_i))_{j \in Q_0} \in U_{\vec{n}}$.

Since the above commutes with the left action of $\mathrm{GL}_{\vec{d}}$, the action is well-defined on \mathcal{M} .

Lemma 3.12. *There is a canonical lift of the action of $U_{\vec{n}}$ on $\mathcal{M}_{\vec{n}, \vec{d}}$ to the universal bundle \mathcal{V}_i , so that the bundle map $\mathcal{V}_i \rightarrow \mathcal{M}_{\vec{n}, \vec{d}}$ is equivariant.*

Proof. \mathcal{V}_i is the $\mathrm{GL}_{\vec{d}}$ -quotient of the trivial bundle $R_{n,d}^s \times V_i$. $U_{\vec{n}}$ acts on this by acting on the component V_i trivially. This action commutes with the left action of $\mathrm{GL}_{\vec{d}}$ on $R_{n,d}^s \times V_i$, and hence descends to act on \mathcal{V}_i . \square

Lemma 3.13. *The metric defined in Theorem 3.7 are $U_{\vec{n}}$ -invariant.*

Proof. For any $g \in U_{\vec{n}}$, since $g_j g_j^* = I_{n_j}$ for any $j \in Q_0$,

$$\begin{aligned} H_{(V_a, e^{(j)})_{a \in Q_1, j \in Q_0}} \cdot g &= \left(\sum_{\gamma} \left(V_{\gamma} e^{(t(\gamma))} \cdot g_{t(\gamma)} \right) \left(V_{\gamma} e^{(t(\gamma))} \cdot g_{t(\gamma)} \right)^* \right)^{-1} \\ &= \left(\sum_{\gamma} \left(V_{\gamma} e^{(t(\gamma))} \right) \left(V_{\gamma} e^{(t(\gamma))} \right)^* \right)^{-1} = H_{(V_a, e^{(j)})_{a \in Q_1, j \in Q_0}}. \end{aligned}$$

\square

Recall from Theorem 2.4 that $\mathcal{M}_{\vec{n}, \vec{d}}$ is the total space of an iterated Grassmann bundle $M^{(N)} \xrightarrow{p_N} M^{(N-1)} \xrightarrow{p_{N-1}} \dots \xrightarrow{p_2} M^{(1)} \xrightarrow{p_1} \text{pt}$. Moreover, \mathcal{V}_i is the pull-back of the tautological bundle S_i of the Grassmann bundle $M^{(i)} = \mathrm{Gr}_{M^{(i)}}(\underline{\mathbb{C}}^{n_i} \oplus \bigoplus_{j \rightarrow i} p_i^* \dots p_{j+1}^*(S_j), d_i) \rightarrow M^{(i-1)}$. The tautological bundle of the Grassmannian is equipped with a standard metric as illustrated in Section 3.1. Inductively, \mathcal{V}_i is also equipped with a pull-back metric. We show that this equals to the metric H we defined by an explicit formula.

Theorem 3.14. *For all $i \in Q_0$, the metric $H = \left(\sum_{h(\gamma)=i} \left(V_{\gamma} e^{(t(\gamma))} \right) \left(V_{\gamma} e^{(t(\gamma))} \right)^* \right)^{-1}$ equals to the metric on \mathcal{V}_i constructed from the iterated Grassmann bundle.*

Proof. As the proof of Theorem 3.7, we do induction on the vertices, which are totally ordered such that $i < j$ whenever there is an arrow $i \rightarrow j$.

First, we have the GIT fiber-bundle map $\mathcal{M} \rightarrow M^{(j)}$, where $M^{(j)}$ is the framed moduli of the quiver $Q^{(j)}$ (which is obtained by removing all vertices $k > j$ and the corresponding arrows from Q). This map $(V, e) \mapsto (V', e')$ is simply forgetting all the irrelevant arrow maps and frame maps that are not supported on the subquiver $Q^{(j)}$. Note that the stability condition is preserved: any subrepresentation $R' \subset V'$ can be extended to a subrepresentation $R \subset V$ by assigning the whole V_k to the additional vertices $k > j$ (and the arrow maps just come from restriction). If $\mathrm{Im}(e') \subset R'$, then $\mathrm{Im}(e) \subset R$.

Note that here we use the GIT description instead of symplectic reduction since (V', e') no longer satisfies the moment-map equation in defining $M^{(j)}$ (even when (V, e) satisfies the moment-map equation for \mathcal{M}).

We start with the minimal vertex i_0 . The quiver $Q^{(i_0)}$ is simply a single vertex, and the corresponding framed moduli is $M^{(i_0)} = \text{Gr}(n_{i_0}, d_{i_0})$. The universal bundle $\mathcal{V}_{i_0}^{Q^{(i_0)}}$ is the tautological bundle of $M^{(i_0)} = \text{Gr}(n_{i_0}, d_{i_0})$. From the last subsection, the standard metric of $\mathcal{V}_{i_0}^{Q^{(i_0)}}$ is descended from $I_{d_{i_0}}$, which is exactly $(e^{Q^{(i_0)}}(e^{Q^{(i_0)}})^*)^{-1}$ by the moment map equation for $Q^{(i_0)}$. The statement is trivial in this case.

Now consider the vertex i . Denote the vertex right before i by $i-1$. For the quiver $Q^{(i-1)}$, assume that the two metrics on the universal bundle $\mathcal{V}_j^{Q^{(i-1)}}$ agree for every $j \in Q_0^{(i-1)}$. We have the bundle map $\pi : M^{(i)} \rightarrow M^{(i-1)}$, and $\mathcal{V}_j^{Q^{(i)}} = \pi^* \mathcal{V}_j^{Q^{(i-1)}}$ for all $j < i$. Moreover, the metric on $\mathcal{V}_j^{Q^{(i)}}$ is pull-back from $\mathcal{V}_j^{Q^{(i-1)}}$, which equals to $\left(\sum_{h(\gamma)=j} \left(V_\gamma e^{(t(\gamma))} \right) \left(V_\gamma e^{(t(\gamma))} \right)^* \right)^{-1}$ by inductive assumption. The pull-back map does not change the arrow and framing maps of $Q^{(i-1)}$. This proves the statement for $S_j^{Q^{(i)}}$ for $j < i$.

Consider $\mathcal{V}_i^{Q^{(i)}}$, which is the tautological bundle associated to the Grassmannian bundle over $M^{(i-1)}$ parametrizing quotients of

$$(e^{(i)}, (V_a)_{h(a)=i}) : \left(\underline{\mathbb{C}}^{d_i} \oplus \left(\bigoplus_{h(a)=i} \mathcal{V}_{t(a)}^{Q^{(i-1)}} \right) \right) \Big|_{x \in M^{(i-1)}} \rightarrow \underline{\mathbb{C}}^{d_i}.$$

The metric on $(\mathcal{V}_i^{Q^{(i)}})^*$ is induced by the embedding $(e^{(i)}, (V_a)_{h(a)=i})^*$ to $\left(\underline{\mathbb{C}}^{d_i} \oplus \left(\bigoplus_{h(a)=i} \mathcal{V}_{t(a)} \right) \right)^*$, whose metric is given by

$$I_{d_i} \oplus \bigoplus_{h(a)=i} \left(\sum_{h(\gamma)=t(a)} \left(V_\gamma e^{(t(\gamma))} \right) \left(V_\gamma e^{(t(\gamma))} \right)^* \right)$$

by inductive assumption. Thus the induced metric on $(\mathcal{V}_i^{Q^{(i)}})^*$ is

$$\begin{aligned} & (e^{(i)}, (V_a)_{h(a)=i}) \cdot \left(I_{d_i} \oplus \bigoplus_{h(a)=i} \left(\sum_{h(\gamma)=t(a)} \left(V_\gamma e^{(t(\gamma))} \right) \left(V_\gamma e^{(t(\gamma))} \right)^* \right) \right) \cdot (e^{(i)}, (V_a)_{h(a)=i})^* \\ &= e^{(i)}(e^{(i)})^* + \sum_{\substack{h(a)=i \\ h(\gamma)=t(a)}} V_a \cdot \left(V_\gamma e^{(t(\gamma))} \right) \left(V_\gamma e^{(t(\gamma))} \right)^* \cdot V_a^* \\ &= \sum_{h(\gamma)=i} \left(V_\gamma e^{(t(\gamma))} \right) \left(V_\gamma e^{(t(\gamma))} \right)^*. \end{aligned}$$

Taking reciprocal gives the metric on $\mathcal{V}_i^{Q^{(i)}}$. This proves the metric has the given expression. \square

Theorem 3.15. *The Ricci curvature of the metric on $\bigotimes_{i \in Q_0} \mathcal{V}_i$ given in Theorem 3.7 defines a Kähler metric on \mathcal{M} .*

Proof. As in Theorem 3.7, denote $\rho = \rho^{(i)} = \left(V_\gamma e^{t(\gamma)} \right)_{\gamma: h(\gamma)=i}$ which is a matrix-valued function on the vector space $R_{\vec{n}, \vec{d}}$. At each point of $R_{\vec{n}, \vec{d}}$, ρ is a linear map from

$$\widehat{W}_i := \bigoplus_{\gamma: h(\gamma)=i} \mathbb{C}^{n_{t(\gamma)}} \quad (3.4)$$

to V_i . The Ricci curvature of the metric $(\rho\rho^*)^{-1}$ is given by $i\partial\bar{\partial} \log \det \rho\rho^*$. We have

$$\begin{aligned} \partial\bar{\partial} \log \det \rho\rho^* &= \partial \left(\text{tr} \left((\rho\rho^*)^{-1} \bar{\partial} (\rho\rho^*) \right) \right) \\ &= \text{tr} \left(\partial \left((\rho\rho^*)^{-1} \rho (\partial\rho)^* \right) \right) \\ &= \text{tr} \left((\rho\rho^*)^{-1} \partial\rho (\partial\rho)^* + \left(\partial (\rho\rho^*)^{-1} \right) \rho (\partial\rho)^* \right) \\ &= \text{tr} \left((\partial\rho)^* (\rho\rho^*)^{-1} \partial\rho \right) - \text{tr} \left((\rho\rho^*)^{-1} (\partial (\rho\rho^*)) (\rho\rho^*)^{-1} \rho (\partial\rho)^* \right) \\ &= \text{tr} \left((\partial\rho)^* (\rho\rho^*)^{-1} \partial\rho \right) - \text{tr} \left((\rho\rho^*)^{-1} \rho (\partial\rho)^* (\rho\rho^*)^{-1} ((\partial\rho) \rho^*) \right) \\ &= \text{tr} \left((\partial\rho)^* (\rho\rho^*)^{-1} \partial\rho \right) - \text{tr} \left(\left(\partial\rho \cdot \left(\rho^* (\rho\rho^*)^{-\frac{1}{2}} \right) \right)^* (\rho\rho^*)^{-1} \left(\partial\rho \cdot \left(\rho^* (\rho\rho^*)^{-\frac{1}{2}} \right) \right) \right) \end{aligned}$$

where $\bar{\partial}\rho = 0$ since the matrix ρ has polynomial entries in holomorphic coordinates.

We can take the singular value decomposition

$$\rho = U \cdot (\text{diag}(\lambda_1, \dots, \lambda_{d_i}) \ 0) \cdot V^*$$

where $U \in U(d_i)$, $V \in U(\dim \widehat{W}_i)$, and $\lambda_i > 0$. ($\lambda_i \neq 0$ since ρ is surjective.) Then

$$\begin{aligned} \rho\rho^* &= U \left(\text{diag}(\lambda_1^2, \dots, \lambda_{d_i}^2) \right) U^* \\ \rho^* (\rho\rho^*)^{-\frac{1}{2}} &= V \left(\begin{array}{c} \text{diag}(\lambda_1, \dots, \lambda_{\alpha(i)}) \\ 0 \end{array} \right) \left(\text{diag}(\lambda_1^{-1}, \dots, \lambda_{\alpha(i)}^{-1}) \right) U^* = V \left(\begin{array}{c} I_{\alpha(i)} \\ 0 \end{array} \right) U^*. \end{aligned}$$

In other words, $\rho^* = \left(\rho^* (\rho\rho^*)^{-\frac{1}{2}} \right) (\rho\rho^*)^{\frac{1}{2}}$ is decomposed into the rescaling $\rho\rho^*$ and the orthogonal embedding $\left(\rho^* (\rho\rho^*)^{-\frac{1}{2}} \right)$ to $\text{Im} \rho^* \subset \widehat{W}_i$.

Now take a vector $v \in T^{1,0}R_{\alpha, d} \cong TR_{\alpha, d}$, and evaluate the above two-form by (v, \bar{v}) . The first term $\text{tr} \left((\partial_v \rho)^* (\rho\rho^*)^{-1} \partial_v \rho \right)$ is the square norm of the linear map

$$\partial_v \rho: \left(\widehat{W}_i, h_{\text{std}} \right) \rightarrow \left(V_i, h_{(\rho\rho^*)^{-1}} \right).$$

Namely we take the standard basis in \widehat{W}_i (which is orthonormal under the standard metric h_{std}), map it to V_i by $\partial_v \rho$, and take the sum of their square norms with respect to the metric $h_{(\rho\rho^*)^{-1}}$.

The second term

$$\text{tr} \left(\left(\partial_v \rho \cdot \left(\rho^* (\rho\rho^*)^{-\frac{1}{2}} \right) \right)^* (\rho\rho^*)^{-1} \left(\partial_v \rho \cdot \left(\rho^* (\rho\rho^*)^{-\frac{1}{2}} \right) \right) \right)$$

is the square norm of the following component $(\partial_v \rho)_1$ of $\partial_v \rho$. Namely, we decompose

$$\widehat{W}_i = (\text{Im} \rho^*) \oplus (\text{Im} \rho^*)^\perp$$

and write $\partial_v \rho = ((\partial_v \rho)_1, (\partial_v \rho)_2)$ where $(\partial_v \rho)_1: \text{Im} \rho^* \rightarrow V_i$ and $(\partial_v \rho)_2: (\text{Im} \rho^*)^\perp \rightarrow V_i$.

We have

$$\partial_v \bar{\partial}_v \log \det \rho \rho^* = \|\partial_v \rho\|_H^2 - \|(\partial_v \rho)_1\|_H^2 = \|(\partial_v \rho)_2\|_H^2 \geq 0 \quad (3.5)$$

for all $v \in T^{1,0}$. (H stands for the metric $(\rho \rho^*)^{-1}$.) This proves that the Ricci curvature of the metric for each i is semi-positive definite.

Now consider

$$H_T = \sum_{i \in Q_0} \partial \bar{\partial} \log \det \rho^{(i)} \left(\rho^{(i)} \right)^*.$$

Suppose it is zero when evaluated at (v, \bar{v}) . Then each individual term equals to zero. This forces $\left(\partial_v \rho^{(i)} \right)_2 = 0$ for all i , that is, image of $\left(\partial_v \rho^{(i)} \right)^* = \partial_v (\rho^{(i)})^*$ sits in the image of $(\rho^{(i)})^*$. This exactly means v descends to the zero tangent vector in the quotient \mathcal{M} : v does not alter the subspaces given by $(\rho^{(i)})^* : V_i \rightarrow \hat{W}_i$ for all i . By the identification of \mathcal{M} as a quiver Grassmannian [Rei08], it means v does not change the position of the point $((\rho^{(i)})^* : i \in Q_0)$ in the quiver Grassmannian, and hence must be the zero tangent vector. This proves the above expression is positive definite. \square

By Equation (3.5), the metric on $T\mathcal{M}$ produced from the Ricci curvature of \mathcal{V}_i is

$$H_T(v, v) = \sum_{i \in Q_0} \|(\partial_v \rho^{(i)})_2\|_{H_i}^2.$$

We have the tautological exact sequence of vector bundles over \mathcal{M} :

$$0 \rightarrow \bigoplus_{i \in Q_0} \text{End}(\mathcal{V}_i) \rightarrow \bigoplus_{a \in Q_1} \text{Hom}(\mathcal{V}_{t_a}, \mathcal{V}_{h_a}) \oplus \bigoplus_{i \in Q_0} \mathcal{V}_i^{n_i} \rightarrow T\mathcal{M} \rightarrow 0$$

where the second arrow is given by sending $X|_{[\phi, e] \in \mathcal{M}} \in \bigoplus_{i \in Q_0} \text{End}(\mathcal{V}_i)$ to

$$\left((X_{h_a} \phi_a - \phi_a X_{t_a})_{a \in Q_1}, (X_i e^{(i)})_{i \in Q_0} \right)$$

(which is the derivative of the action of $\text{GL}_{\vec{d}}$), and $T\mathcal{M}$ is obtained as the quotient bundle (of the middle one by the first one). $H(v, v)$ can be defined for $v \in \bigoplus_{a \in Q_1} \text{Hom}(\mathcal{V}_{t_a}, \mathcal{V}_{h_a}) \oplus \bigoplus_{i \in Q_0} \mathcal{V}_i^{n_i}$. H_{TM} is zero on $\bigoplus_{i \in Q_0} \text{End}(\mathcal{V}_i)$: the action of $\text{GL}(d_j, \mathbb{C})$ does not change $\rho^{(i)}$ for $j \neq i$; for $X \in \text{End}(\mathcal{V}_i)$, $(X \cdot \rho^{(i)} \cdot w)^* v = w^* (\rho^{(i)})^* (X^* \cdot v) = 0$ for all $v \in \mathcal{V}_i, w \in (\text{Im}(\rho^{(i)})^*)^\perp$, and so $(\partial_X \rho^{(i)})_2 = 0$.

3.3 Quiver with oriented cycles. When the quiver Q has an oriented cycle, the framed moduli $\mathcal{M}_{\vec{n}, \vec{d}}$ is no longer projective. Examples of such quivers were studied algebraically by [Fed13, ER09] along the line of Reineke.

On the other hand, the metric given in Theorem 3.7 still makes sense for quiver with oriented cycles, as long as we stay in the domain of convergence and prove that it is positive-definite. Below we will prove this for any given quiver.

Denote the moment-map level by $R_{n, \vec{d}}^{\mu=1} = \{\mu = 1\} \subset R_{\vec{n}, \vec{d}}$. Let $\|A\| = \sup_{\|v\|=1} \|A \cdot v\|$ be the operator norm of a matrix A . We take the following open subset of $R_{n, \vec{d}}^{\mu=1}$.

Definition 3.16. *Define*

$$R_{\vec{n}, \vec{d}}^{\mu=1, \circ} := \{(V, e) \in R_{n, \vec{d}}^{\mu=1} : \|V(\gamma)\| < 1 \text{ for every oriented cycle } \gamma\}$$

where $V(\gamma) = V(a_k) \dots V(a_1)$ for $\gamma = a_k \dots a_1$.

$$R_{\vec{n}, \vec{d}}^\circ := \text{GL}_{\vec{d}} \cdot R_{\vec{n}, \vec{d}}^{\mu=1, \circ}.$$

and

$$\mathcal{M}_{\vec{n}, \vec{d}}^\circ := R_{\vec{n}, \vec{d}}^{\mu=I, \circ} / U_{\vec{d}} = R_{\vec{n}, \vec{d}}^\circ / \mathrm{GL}_{\vec{d}}.$$

The above definition of $\mathcal{M}_{\vec{n}, \vec{d}}^\circ$ makes sense because of the following.

Lemma 3.17. $R_{\vec{n}, \vec{d}}^{\mu=I, \circ}$ is invariant under $U_{\vec{d}}$.

Proof. For every oriented cycle γ at $i \in Q_0$, $(g \cdot V)(\gamma) = g_i V(\gamma) g_i^{-1}$, and hence the condition $\|V(\gamma)\| < 1$ is respected for $g \in U_{\vec{d}}$. \square

The main theorem in this section is the following.

Theorem 3.18. Let Q be an arbitrary quiver. As in Theorem 3.7, for each $i \in Q_0$, set

$$H_i = (\rho_i \rho_i^*)^{-1} = \left(\sum_{h(\gamma)=i} \left(V_\gamma e^{t(\gamma)} \right) \left(V_\gamma e^{t(\gamma)} \right)^* \right)^{-1}$$

which is an infinite sum, whose terms are ordered by the length of the path γ . (There are just finitely many paths under each fixed length.) This gives a convergent function $H_i : R_{\vec{n}, \vec{d}}^\circ \rightarrow \mathrm{End}(\mathbb{C}^{d_i})$. H_i is $\mathrm{GL}_{\vec{d}}$ -equivariant, and it descends to a metric on \mathcal{V}_i over $\mathcal{M}_{\vec{n}, \vec{d}}^\circ$.

We break into several steps to prove the above theorem. First, consider the convergence.

Lemma 3.19. $\rho \rho^*$ is absolutely convergent over $R_{\vec{n}, \vec{d}}^{\mu=I, \circ}$. Hence $(\rho \rho^*)^{-1}$ is well-defined and $\mathrm{GL}_{\vec{d}}$ -equivariant on $R_{\vec{n}, \vec{d}}^\circ$.

Proof. For $(V, e) \in R_{\vec{n}, \vec{d}}^{\mu=I, \circ}$, we consider the expression

$$\sum_{h(\gamma)=i} \|V_\gamma e^{t(\gamma)}\| \left\| \left(V_\gamma e^{t(\gamma)} \right)^* \right\| \leq \sum_{h(\gamma)=i} \|V_\gamma e^{t(\gamma)}\|^2.$$

There are only finitely many paths $\gamma_1, \dots, \gamma_k$ with $h(\gamma_l) = i$ which do not contain any oriented cycle. Any other path (with $h(\gamma) = i$) can be written as concatenation of one of these γ_l and some oriented cycles at some vertices. Thus

$$\sum_{h(\gamma)=i} \|V_\gamma e^{t(\gamma)}\|^2 \leq \sum_{l=1}^k \|\gamma_l e^{t(\gamma_l)}\|^2 \sum_{p=0}^{\infty} (1-\epsilon)^p = \sum_{l=1}^k \|\gamma_l e^{t(\gamma_l)}\|^2 \sum_{p=0}^{\infty} (1-\epsilon)^p = \sum_{l=1}^k \frac{\|\gamma_l e^{t(\gamma_l)}\|^2}{\epsilon} < \infty$$

where given V , there is a fixed $\epsilon \in (0, 1)$ such that $\|V_\gamma\|^2 < 1 - \epsilon$ for all oriented cycles γ . Hence $\rho \rho^*$ is absolutely convergent for every $(V, e) \in R_{\vec{n}, \vec{d}}^{\mu=I, \circ}$.

Every element in $R_{\vec{n}, \vec{d}}^\circ$ can be written as $g \cdot (V, e)$ for $g \in \mathrm{GL}_{\vec{d}}$ and $(V, e) \in R_{\vec{n}, \vec{d}}^{\mu=I, \circ}$. $(\rho \rho^*)^{-1}|_{g \cdot (V, e)} = (g_i^*)^{-1} (\rho \rho^*)^{-1}|_{(V, e)} g_i^{-1}$ where $(\rho \rho^*)^{-1}|_{(V, e)}$ is convergent. \square

It remains to prove positive definiteness of H_i . First we consider the following specific quiver which is simply a single oriented cycle.

Lemma 3.20. If Q is a single oriented cycle with N vertices, then for each vertex i , $\rho_i \rho_i^*$ is positive definite. (In particular, when $N = 1$, Q consists of one vertex and a self loop.)

Proof. By symmetry, we just need to prove for $i = N$. Let $l = a_N \dots a_1$, where a_k is the arrow $(k-1) \rightarrow k$.

$$\rho_N \rho_N^* = \sum_{k=1}^N \left(a_N \dots a_{k+1} e_k e_k^* a_{k+1}^* \dots a_N^* + \sum_{p>0} l^p a_N \dots a_{k+1} e_k e_k^* a_{k+1}^* \dots a_N^* (l^p)^* \right).$$

The moment map equation at the vertex k is $e_k e_k^* = I - a_k a_k^* + a_{k+1}^* a_{k+1}$. Then the first term gives

$$\begin{aligned} & \sum_{k=1}^N a_N \dots a_{k+1} (I - a_k a_k^* + a_{k+1}^* a_{k+1}) a_{k+1}^* \dots a_N^* \\ &= \sum_{k=1}^N a_N \dots a_{k+1} a_{k+1}^* \dots a_N^* - \sum_{k=1}^N a_N \dots a_{k+1} a_k a_k^* a_{k+1}^* \dots a_N^* + \sum_{k=1}^N a_N \dots a_{k+1} a_{k+1}^* a_{k+1} a_{k+1}^* \dots a_N^* \\ &= I - ll^* + \sum_{k=1}^N a_N \dots a_{k+1} a_{k+1}^* a_{k+1} a_{k+1}^* \dots a_N^*. \end{aligned}$$

Similarly, the second term gives

$$\begin{aligned} & \sum_{k=1}^N \sum_{p>0} l^p a_N \dots a_{k+1} (I - a_k a_k^* + a_{k+1}^* a_{k+1}) a_{k+1}^* \dots a_N^* (l^p)^* \\ &= \sum_{p>0} l^p (l^p)^* - \sum_{p>0} l^p ll^* (l^p)^* + \sum_{p>0} \sum_{k=1}^N l^p a_N \dots a_{k+1} a_{k+1}^* a_{k+1} a_{k+1}^* \dots a_N^* (l^p)^* \\ &= ll^* + \sum_{p>0} \sum_{k=1}^N l^p a_N \dots a_{k+1} a_{k+1}^* a_{k+1} a_{k+1}^* \dots a_N^* (l^p)^*. \end{aligned}$$

Combining the two terms,

$$\rho_N \rho_N^* = I + \sum_{p \geq 0} \sum_{k=1}^N l^p a_N \dots a_{k+1} a_{k+1}^* a_{k+1} a_{k+1}^* \dots a_N^* (l^p)^*$$

and the second term is semi-positive-definite. Hence $\rho_N \rho_N^*$ is positive definite. \square

The following is the key lemma to prove positive-definiteness for a general quiver.

Lemma 3.21. *Suppose Q has the property that for every $i \in Q_0$, restricted to the intersection of the moment map locus and $R_{n,d}^{Q,\circ}$, $\rho_i^Q (\rho_i^Q)^* = I + B_i$ for some semi-positive definite matrix B_i . Let*

Q' be obtained by concatenating to Q a chain $x_1 \xrightarrow{a'_1} \dots \xrightarrow{a'_{k-1}} x_k$ where x_1 and x_k are certain vertices in Q . (x_1 can be equal to x_k , meaning what we have added is an oriented cycle. When $k = 1$, we have added a loop; $a'_0 = a'_1$. When $k = 2$, there is no intermediate vertex in the chain.) Then Q' has the same property. Namely, for every $i \in Q'_0$, restricted to the intersection of the moment map locus and $R_{n,d}^{Q',\circ}$, $\rho_i^{Q'} (\rho_i^{Q'})^ = I + B'_i$ for some semi-positive-definite matrix B'_i .*

Proof. First, consider the case that the given vertex $i \in Q'_0$ belongs to Q_0 . For the original quiver Q , $\rho_i^Q (\rho_i^Q)^* = I + B_i$ where B_i is semi-positive definite. After concatenating the chain, the terms in ρ_i^Q ,

$$\gamma_Q^{x_1 \rightarrow i} e^{(1)} (e^{(1)})^* (\gamma_Q^{x_1 \rightarrow i})^* \text{ and } \gamma_Q^{x_k \rightarrow i} e^{(k)} (e^{(k)})^* (\gamma_Q^{x_k \rightarrow i})^*$$

where $\gamma_Q^{x_1 \rightarrow i}$ ($\gamma_Q^{x_k \rightarrow i}$ resp.) is a path from x_1 (from x_k resp.) to i , get affected. Namely, the moment map equation for $e^{(1)}(e^{(1)})^*$ (or $e^{(k)}(e^{(k)})^*$) gets an extra term $(a'_1)^* a'_1 (-a'_{k-1}(a'_{k-1})^*$ resp.). (If $k = 1$, $x_1 = x_k$ and $a'_1 = a'_0$, and the moment map equation for $e^{(1)}(e^{(1)})^*$ gets both the extra terms $(a'_0)^* a'_0$ and $-a'_0(a'_0)^*$.) As a result, we have an extra negative term

$$-\gamma_Q^{x_k \rightarrow i} a'_{k-1} (a'_{k-1})^* (\gamma_Q^{x_k \rightarrow i})^* \quad (3.6)$$

for each path $\gamma_Q^{x_k \rightarrow i}$. We shall show that these negative terms can be canceled.

We also have additional paths in Q' heading to i , which can be divided into the following types:

- (1) $\gamma_Q^{x_k \rightarrow i} a'_{k-1} \dots a'_l e^{(x_l)}$ for $l = 2, \dots, k-1$. (This is an empty case when $k = 1, 2$.)
- (2) $\gamma_Q^{x_k \rightarrow i} \gamma_{\text{new}} \gamma_Q^{j \rightarrow x_1} e_j$ where $\gamma_{\text{new}} := a'_{k-1} \dots a'_1$ and $\gamma_Q^{j \rightarrow x_1}$ is any path in Q from $j \in Q_0$ to x_1 . ($\gamma_Q^{x_1 \rightarrow x_1}$ can be the trivial path. $\gamma_{\text{new}} = a'_0$ when $k = 1$.)
- (3) $\gamma_Q^{x_k \rightarrow i} \left(\prod_{r=1}^p \gamma_{\text{new}} \gamma_{Q,r}^{x_k \rightarrow x_1} \right) a'_{k-1} \dots a'_l e^{(x_l)}$ for some $p > 0$ and $l = 2, \dots, k-1$.
- (4) $\gamma_Q^{x_k \rightarrow i} \left(\prod_{r=1}^p \gamma_{\text{new}} \gamma_{Q,r}^{x_k \rightarrow x_1} \right) \gamma_{\text{new}} \gamma_Q^{j \rightarrow x_1} e_j$ for some $p > 0$ and $j \in Q_0$.

For (1), by the moment-map equation $e^{(x_l)}(e^{(x_l)})^* = I + (a'_l)^* a'_l - a'_{l-1}(a'_{l-1})^*$, we have

$$\begin{aligned} & \gamma_Q^{x_k \rightarrow i} a'_{k-1} \dots a'_l \cdot e^{(x_l)}(e^{(x_l)})^* \cdot (\gamma_Q^{x_k \rightarrow i} a'_{k-1} \dots a'_l)^* \\ &= \gamma_Q^{x_k \rightarrow i} a'_{k-1} \dots a'_l (\gamma_Q^{x_k \rightarrow i} a'_{k-1} \dots a'_l)^* + \gamma_Q^{x_k \rightarrow i} a'_{k-1} \dots a'_l (a'_l)^* a'_l (\gamma_Q^{x_k \rightarrow i} a'_{k-1} \dots a'_l)^* \\ & \quad - \gamma_Q^{x_k \rightarrow i} a'_{k-1} \dots a'_l a'_{l-1} (a'_{l-1})^* (\gamma_Q^{x_k \rightarrow i} a'_{k-1} \dots a'_l)^*. \end{aligned}$$

The first term above for $l = k-1$ cancel with the extra negative term (3.6) for ρ_i^Q . The third term (which is negative) for $l \in \{3, \dots, k-1\}$ cancel with the first term of $l-1$. As a result, after combining (1) with the modified ρ_i^Q , the remaining negative terms are

$$-\gamma_Q^{x_k \rightarrow i} a'_{k-1} \dots a'_1 (\gamma_Q^{x_k \rightarrow i} a'_{k-1} \dots a'_1)^* = -\gamma_Q^{x_k \rightarrow i} \gamma_{\text{new}} (\gamma_Q^{x_k \rightarrow i} \gamma_{\text{new}})^*. \quad (3.7)$$

(For the case $k = 1$, this trivially holds since $\gamma_{\text{new}} = a'_0$, and the above equals to (3.6).)

Now consider (2): $\gamma_Q^{x_k \rightarrow i} \gamma_{\text{new}} \gamma_Q^{j \rightarrow x_1} e_j e_j^* (\gamma_Q^{x_k \rightarrow i} \gamma_{\text{new}} \gamma_Q^{j \rightarrow x_1})^*$. The moment map equation for $e_j e_j^*$ when $j \neq x_1$, x_k are the same for Q and Q' . Summing $\gamma_Q^{x_k \rightarrow i} \gamma_{\text{new}} \gamma_Q^{j \rightarrow x_1} e_j e_j^* (\gamma_Q^{x_k \rightarrow i} \gamma_{\text{new}} \gamma_Q^{j \rightarrow x_1})^*$ over arbitrary $\gamma_Q^{j \rightarrow x_1}$ and $j \in Q_0$, we obtain

$$(\gamma_Q^{x_k \rightarrow i} \gamma_{\text{new}}) \cdot \rho_{x_1}^Q \cdot (\gamma_Q^{x_k \rightarrow i} \gamma_{\text{new}})^* = (\gamma_Q^{x_k \rightarrow i} \gamma_{\text{new}}) \cdot (I + B_1) \cdot (\gamma_Q^{x_k \rightarrow i} \gamma_{\text{new}})^*$$

plus

$$\sum_{\gamma_Q^{x_1 \rightarrow x_1}} \gamma_Q^{x_k \rightarrow i} \gamma_{\text{new}} \gamma_Q^{x_1 \rightarrow x_1} (a'_1)^* a'_1 (\gamma_Q^{x_k \rightarrow i} \gamma_{\text{new}} \gamma_Q^{x_1 \rightarrow x_1})^* - \sum_{\gamma_Q^{x_k \rightarrow x_1}} \gamma_Q^{x_k \rightarrow i} \gamma_{\text{new}} \gamma_Q^{x_k \rightarrow x_1} a'_{k-1} (a'_{k-1})^* (\gamma_Q^{x_k \rightarrow i} \gamma_{\text{new}} \gamma_Q^{x_k \rightarrow x_1})^*$$

which is due to the additional terms in the moment map equations for $e^{(x_1)}(e^{(x_1)})^*$ and $e^{(x_k)}(e^{(x_k)})^*$. (When $k = 1$, $\gamma_Q^{x_0 \rightarrow x_0}$ being the trivial path at x_0 is one of the possibilities.) In above, B_1 is semi-positive-definite by the assumption on Q . Then the negative terms (3.7) cancel with the first term. After combining the paths in Q and (1) and (2), the remaining negative terms are

$$-\sum_{\gamma_Q^{x_k \rightarrow x_1}} \gamma_Q^{x_k \rightarrow i} \gamma_{\text{new}} \gamma_Q^{x_k \rightarrow x_1} a'_{k-1} (a'_{k-1})^* (\gamma_Q^{x_k \rightarrow i} \gamma_{\text{new}} \gamma_Q^{x_k \rightarrow x_1})^*. \quad (3.8)$$

(If there is no path $\gamma_Q^{x_k \rightarrow x_1}$ in Q from x_k to x_1 , then this is zero, and we do not have (3) nor (4). We stop here and get that $\rho_i^{Q'}(\rho_i^{Q'})^* = I + B'_i$ for a semi-positive-definite matrix B'_i .)

The terms in (3) for $p = 1$ and $l = k - 1$ cancel with the above (3.8):

$$\begin{aligned} & \gamma_Q^{x_k \rightarrow i} \left(\prod_{r=1}^p \gamma_{\text{new}} \gamma_{Q,r}^{x_k \rightarrow x_1} \right) a'_{k-1} \dots a'_l \cdot e^{(x_l)}(e^{(x_l)})^* \cdot \left(\gamma_Q^{x_k \rightarrow i} \left(\prod_{r=1}^p \gamma_{\text{new}} \gamma_{Q,r}^{x_k \rightarrow x_1} \right) a'_{k-1} \dots a'_l \right)^* \\ &= \gamma_Q^{x_k \rightarrow i} \left(\prod_{r=1}^p \gamma_{\text{new}} \gamma_{Q,r}^{x_k \rightarrow x_1} \right) a'_{k-1} \dots a'_l \left(\gamma_Q^{x_k \rightarrow i} \left(\prod_{r=1}^p \gamma_{\text{new}} \gamma_{Q,r}^{x_k \rightarrow x_1} \right) a'_{k-1} \dots a'_l \right)^* \\ & \quad + \gamma_Q^{x_k \rightarrow i} \left(\prod_{r=1}^p \gamma_{\text{new}} \gamma_{Q,r}^{x_k \rightarrow x_1} \right) a'_{k-1} \dots a'_l (a'_l)^* a'_l \left(\gamma_Q^{x_k \rightarrow i} \left(\prod_{r=1}^p \gamma_{\text{new}} \gamma_{Q,r}^{x_k \rightarrow x_1} \right) a'_{k-1} \dots a'_l \right)^* \\ & \quad - \gamma_Q^{x_k \rightarrow i} \left(\prod_{r=1}^p \gamma_{\text{new}} \gamma_{Q,r}^{x_k \rightarrow x_1} \right) a'_{k-1} \dots a'_l a'_{l-1} (a'_{l-1})^* \left(\gamma_Q^{x_k \rightarrow i} \left(\prod_{r=1}^p \gamma_{\text{new}} \gamma_{Q,r}^{x_k \rightarrow x_1} \right) a'_{k-1} \dots a'_l \right)^* . \end{aligned}$$

Like in (1), for each p , the third term (which is negative) for $l \in \{3, \dots, k-1\}$ cancel with the first term for $l-1$. (When $p = 0$ and $l = k-1$, the third term is exactly (3.8).) Then the remaining negative terms are (3.7) modified by inserting the loops $\left(\prod_{r=1}^p \gamma_{\text{new}} \gamma_{Q,r}^{x_k \rightarrow x_1} \right)$ for $p > 0$, that is, $-\gamma_Q^{x_k \rightarrow i} \left(\prod_{r=1}^p \gamma_{\text{new}} \gamma_{Q,r}^{x_k \rightarrow x_1} \right) \gamma_{\text{new}}(\gamma_Q^{x_k \rightarrow i} \left(\prod_{r=1}^p \gamma_{\text{new}} \gamma_{Q,r}^{x_k \rightarrow x_1} \right) \gamma_{\text{new}})^*$. Then like in (2), these negative terms cancel with terms in (4). Summing up to finite p , the only negative terms left are (3.8) modified by inserting the loops $\left(\prod_{r=1}^p \gamma_{\text{new}} \gamma_{Q,r}^{x_k \rightarrow x_1} \right)$:

$$- \sum_{\gamma_Q^{x_k \rightarrow x_1}} \gamma_Q^{x_k \rightarrow i} \left(\prod_{r=1}^p \gamma_{\text{new}} \gamma_{Q,r}^{x_k \rightarrow x_1} \right) \gamma_{\text{new}} \gamma_Q^{x_k \rightarrow x_1} a'_{k-1} (a'_{k-1})^* \left(\gamma_Q^{x_k \rightarrow i} \left(\prod_{r=1}^p \gamma_{\text{new}} \gamma_{Q,r}^{x_k \rightarrow x_1} \right) \gamma_{\text{new}} \gamma_Q^{x_k \rightarrow x_1} \right)^*$$

which cancel with terms in (3) for $(p+1)$. As $p \rightarrow \infty$, $\left\| \left(\prod_{r=1}^p \gamma_{\text{new}} \gamma_{Q,r}^{x_k \rightarrow x_1} \right) \right\| \rightarrow 0$. This finishes the proof that $\rho_i^{Q'}(\rho_i^{Q'})^* = I + B_i$, for $i \in Q_0$.

For the case that $i = x_j$ for $j = 2, \dots, k-1$, the proof is similar. (We do not need to consider this case when $k = 1, 2$.) The paths in Q' heading to x_j are divided into the following types:

- (1) $a'_{j-1} \dots a'_l e^{(x_l)}$ for $l = 2, \dots, j$.
- (2) $a'_{j-1} \dots a'_1 \left(\prod_{r=1}^p \gamma_{Q,r}^{x_k \rightarrow x_1} \gamma_{\text{new}} \right) \gamma_Q^{j \rightarrow x_1} e_j$ for some $p \geq 0, j \in Q_0$.
- (3) $a'_{j-1} \dots a'_1 \left(\prod_{r=1}^p \gamma_{Q,r}^{x_k \rightarrow x_1} \gamma_{\text{new}} \right) \gamma_Q^{x_k \rightarrow x_1} a'_{k-1} \dots a'_l e^{(x_l)}$ for some $p \geq 0, l = 2, \dots, k-1$.

The cancellation is similar and we do not repeat here. \square

Similarly, adding a chain at a single vertex of Q preserves the positive-definiteness property.

Lemma 3.22. *Suppose Q as in Lemma 3.21. Let Q' be obtained by concatenating to Q a chain $x_1 \xrightarrow{a'_1} \dots \xrightarrow{a'_{k-1}} x_k$ at either x_1 or x_k in Q_0 . Then for every $i \in Q'_0$, $\rho_i^{Q'}(\rho_i^{Q'})^* = I + B'_i$ for some semi-positive-definite matrix B'_i .*

Proof. The proof in this case is simpler than that of Lemma 3.21, since there is no new oriented cycle.

Consider the case that $x_k \in Q_0$. If i belongs to the chain, then the only paths that head to i are contained in the chain. Since no oriented cycle is involved in all such paths, Theorem 3.7 already gives the result.

If $i \in Q_0$, then

$$\begin{aligned}
\rho_i^{Q'}(\rho_i^{Q'})^* &= \sum_{j \in Q_0 - \{x_k\}} \sum_{\gamma_Q^{j \rightarrow i}} (\gamma_Q^{j \rightarrow i} e^{(j)}) (\gamma_Q^{j \rightarrow i} e^{(j)})^* + \sum_{\gamma_Q^{x_k \rightarrow i}} \gamma_Q^{x_k \rightarrow i} \cdot e^{(x_k)} (e^{(x_k)})^* \cdot (\gamma_Q^{x_k \rightarrow i})^* \\
&\quad + \sum_{\gamma_Q^{x_k \rightarrow i}} \sum_{r=1}^{k-1} (\gamma_Q^{x_k \rightarrow i} a_{k-1} \dots a_r) e_{x_r} e_{x_r}^* \cdot (\gamma_Q^{x_k \rightarrow i} a'_{k-1} \dots a'_r)^* \\
&= \rho_i^Q(\rho_i^Q)^* - \sum_{\gamma_Q^{x_k \rightarrow i}} \gamma_Q^{x_k \rightarrow i} (a'_{k-1} (a'_{k-1})^*) (\gamma_Q^{x_k \rightarrow i})^* \\
&\quad + \sum_{\gamma_Q^{x_k \rightarrow i}} \sum_{r=1}^{k-1} (\gamma_Q^{x_k \rightarrow i} a'_{k-1} \dots a'_r) (I - a'_{r-1} (a'_{r-1})^* + (a'_r)^* a'_r) \cdot (\gamma_Q^{x_k \rightarrow i} a'_{k-1} \dots a'_r)^* \\
&= \rho_i^Q(\rho_i^Q)^* + \sum_{\gamma_Q^{x_k \rightarrow i}} \sum_{r=1}^{k-1} (\gamma_Q^{x_k \rightarrow i} a'_{k-1} \dots a'_r) (a'_r)^* a'_r \cdot (\gamma_Q^{x_k \rightarrow i} a'_{k-1} \dots a'_r)^*.
\end{aligned}$$

Since $\rho_i^Q(\rho_i^Q)^* = I + B_i$ for some semi-positive-definite matrix B_i , and the second term is semi-positive-definite, $\rho_i^{Q'}(\rho_i^{Q'})^*$ satisfies the requirement.

The case that $x_1 \in Q_0$ is similar and the proof is omitted. \square

Proof of Theorem 3.18. Without loss of generality, suppose Q is connected. (Otherwise $\mathcal{M}_{\vec{n}, \vec{d}}$ and \mathcal{V}_i decompose into products coming from the connected components, and we just need to study each component.) The case without oriented cycle is given in Theorem 3.7. Suppose Q has at least one oriented cycle. By Lemma 3.20, the statement is true for this oriented cycle as a quiver. There must be additional arrows if this single oriented cycle is not yet the whole Q . Then we can either add a chain as in Lemma 3.21 or 3.22, and the statement still holds. (Both the cases of loop at a vertex or multiple edge are covered by Lemma 3.21.) Inductively the statement holds for Q . \square

Example 3.23. For the A_2 -quiver, $\rho_i \rho_i^*$ gives a metric for all i . By Lemma 3.21, this is still true if we add an oriented cycle with arrows l_1, \dots, l_p . See Figure 3.

Note that the following equality still holds over $R_{n,d}^\circ$:


$$\rho \rho^* = e^{(i)} (e^{(i)})^* + \sum_{h(a)=i} V_a \rho_{(t(a))} \rho_{(t(a))}^* V_a^*.$$

Thus Proposition 3.8 still holds for quivers with oriented cycles.

Proposition 3.24. For any quiver Q and every $v, w \in (\mathcal{V}_i)^*$,

$$H_i(v, w) = H_0((e^{(i)})^*(v), (e^{(i)})^*(w)) + \sum_{h(a)=i} H_{t(a)}(a^*(v), a^*(w)).$$

Now we consider a version of Theorem 3.15 in this case. We define \hat{W}_i by Equation (3.4). But this time, it is an infinite direct sum of Hilbert spaces (meaning that it consists of infinite sequence $w = (w_\gamma : h(\gamma) = i)$ with $\|w\|^2 = \sum_\gamma \|w_\gamma\|^2 < \infty$).


 FIGURE 3. A_2 modified by adding an oriented cycle.

Lemma 3.25. For each $(V, e) \in R_{n,d}^\circ$, $\rho_i(V, e)$ defines a bounded linear map $\hat{W}_i \rightarrow V_i$. Its adjoint $\rho_i(V, e)^* : V_i \rightarrow \hat{W}_i$ has a singular-value decomposition.

Proof. For $w = (w_\gamma : h(\gamma) = i) \in \hat{W}_i$, $\rho_i(V, e)$ maps it to

$$\sum_{\gamma: h(\gamma)=i} V_\gamma e^{t(\gamma)} w_\gamma.$$

Like in the proof of Lemma 3.19, consider $\sum_{\gamma: h(\gamma)=i} \|V_\gamma e^{t(\gamma)}\| \|w_\gamma\| \leq \|w\| \sum_{\gamma: h(\gamma)=i} \|V_\gamma e^{t(\gamma)}\| < +\infty$. This also shows that if $\|w\| = 1$, then the image of w is also bounded. $\rho_i(V, e)^*$ has image being finite-dimensional (since V_i is finite-dimensional), and hence is a compact operator. Thus it has a singular-value decomposition. \square

Proposition 3.26. The Ricci curvature of the metric given by $(\rho_i \rho_i^*)^{-1}$ is semi-positive definite on \mathcal{M}° .

Proof. By the previous lemma, the proof of Theorem 3.15 on semi-positive definiteness still works. Namely,

$$\partial \bar{\partial} \log \det \rho \rho^* = \text{tr} \left((\partial \rho)^* (\rho \rho^*)^{-1} \partial \rho \right) - \text{tr} \left(\left(\partial \rho \cdot \left(\rho^* (\rho \rho^*)^{-\frac{1}{2}} \right) \right)^* (\rho \rho^*)^{-1} \left(\partial \rho \cdot \left(\rho^* (\rho \rho^*)^{-\frac{1}{2}} \right) \right) \right).$$

Note that the two terms on the RHS are finite: $\text{tr} \left((\partial \rho)^* (\rho \rho^*)^{-1} \partial \rho \right) = \text{tr} \left(\rho \rho^* \right)^{-1} \partial \rho \left((\partial \rho)^* \right) = \sum_{j=1}^{d_i} \langle \partial \rho \left((\partial \rho)^* \epsilon_j \right), (\rho \rho^*)^{-1} \epsilon_j \rangle_{V_i}$ which is a finite sum, and similar for the second term. $\rho^* = \left(\rho^* (\rho \rho^*)^{-\frac{1}{2}} \right) (\rho \rho^*)^{\frac{1}{2}}$ is decomposed into the rescaling $(\rho \rho^*)^{\frac{1}{2}}$ and the orthogonal embedding $\left(\rho^* (\rho \rho^*)^{-\frac{1}{2}} \right)$ to $\text{Im} \rho^* \subset \hat{W}_i$. Then the above equals to $\|(\partial \rho)_2\|_H^2 \geq 0$ as in Theorem 3.15. \square

4. FIBERWISE NONLINEARITY

In the mathematical study of quivers, we mostly focused on linear representations. In particular, the morphisms between universal vector bundles are linear along fibers. On

the other hand, nonlinear ‘activation functions’ play a key role in machine learning. In this section, we construct some natural non-linear fiber-bundle endomorphisms of the universal bundles \mathcal{V}_i over $\mathcal{M}_{\vec{n}, \vec{d}}$ by using fiberwise symplectomorphisms.

For simplicity, we shall take $n_i = d_i + 1$ for all $i \in Q_0$ in this section.

4.1 Activation functions arising from toric moment maps and symplectomorphisms.

In this section, we make the observation that several activation functions commonly used in machine learning actually belong to a much bigger class, namely the T -equivariant symplectomorphisms on open subsets of a symplectic toric variety.

First, let’s recall the basic setup for toric varieties. Let’s equip \mathbb{C}^m with the standard Kähler structure. We obtain a symplectic toric variety (X, ω_X) as a symplectic quotient by the real torus T^{m-d} . We assume X is smooth. The T^{m-d} -action can be specified by an injective homomorphism $\mathbb{Z}^{m-d} \rightarrow \mathbb{Z}^m$, which induces a map $T^{m-d} \rightarrow T^m$, and T^m acts on \mathbb{C}^m by coordinate-wise multiplication. We assume that the quotient of \mathbb{Z}^m by the image of \mathbb{Z}^{m-d} is again a lattice, which we identify as \mathbb{Z}^d . We denote by $v_i \in \mathbb{Z}^d$ the images of the standard basic vectors of \mathbb{Z}^m under the quotient map $\mathbb{Z}^m \rightarrow \mathbb{Z}^d$.

The residual action of T^d on X gives a moment-map fibration over a polytope P , which is given by the intersection of m half-spaces in \mathbb{R}^d :

$$\{x \in \mathbb{R}^d : \ell_j(x) := v_j \cdot x - c_j \geq 0\}$$

where the constants $c_j \in \mathbb{R}^d$ are determined by the level taken in the symplectic quotient. We assume that the level is chosen such that for all j , $\{\ell_j(x) = 0\} \cap P$ is a (non-empty) codimension-one boundary of the polytope P .

Consider the open toric orbit of X , which can be identified as $(\mathbb{C}^\times)^d$ by fixing a basis of \mathbb{Z}^d . Denote by ω_X the Kähler form induced on the symplectic quotient. Let $\omega_{\text{std}} = \sum_{i=1}^d dx_i \wedge d\theta_i$ be the standard symplectic form on $\mathbb{R}^d \times T^d$ (where T^d denotes the real d -torus). The symplectic form $\omega_X|_{(\mathbb{C}^\times)^d}$ has an explicit description by the following beautiful formula.

Theorem 4.1 ([Gui94, Abr98]).

$$\left(\frac{1}{2} \left(\sum_{i=1}^m v_i \log \ell_i(x) \right), \text{Id} \right) : (P^\circ \times T^d, \omega_{\text{std}}|_{P^\circ \times T^d}) \rightarrow \mathbb{R}^d \times T^d \stackrel{\text{exp}}{\cong} ((\mathbb{C}^\times)^d, \omega_X|_{(\mathbb{C}^\times)^d})$$

is a symplectomorphism.

Taking the universal cover $\mathbb{R}^d \rightarrow T^d$ and lifting the above, one obtains the following.

Corollary 4.2. *The inverse of $(\frac{1}{2} (\sum_{i=1}^m v_i \log \ell_i(x)), \text{Id})$ gives a symplectomorphism*


$$\sigma_{\mathbb{C}} = (\sigma(\text{Re}(\vec{z})), \text{Im}(\vec{z})) : (\mathbb{C}^d, \exp^* \omega_X) \rightarrow (P^\circ \times \mathbb{R}^d, \omega_{\text{std}})$$

where $\exp : \mathbb{C}^d \rightarrow (\mathbb{C}^\times)^d \subset X$, and $\sigma : \mathbb{R}^d \rightarrow P^\circ$ is the inverse of $(\frac{1}{2} (\sum_{i=1}^m v_i \log \ell_i(x)), \text{Id})$.

Example 4.3. For \mathbb{C}^d , the moment polytope P is $\mathbb{R}_{\geq 0}^d = \{x_i \geq 0 : i = 1, \dots, d\}$. $v = (1, \dots, 1)$.

The above map is simply $\left(\frac{\log x_i}{2} \right)_{i=1}^d : \mathbb{R}_{> 0}^d \rightarrow \mathbb{R}^d$. The symplectomorphism $((\mathbb{C}^\times)^d, \omega_{\mathbb{C}^d}|_{(\mathbb{C}^\times)^d}) \xrightarrow{\cong} (\mathbb{R}_{> 0}^d \times T^d, \omega_{\text{std}}|_{P^\circ \times T^d})$ is $\left(|z_i|^2, \frac{\log z_i - \log \bar{z}_i}{2i} \right)_{i=1}^d$.

Example 4.4. For the complex projective space \mathbb{P}^d , the corresponding moment polytope P (for a chosen level) is the d -simplex given by $\ell_i \geq 0$ where $\ell_i(x) = x_i$ for $i = 1, \dots, d$, and $\ell_{d+1}(x) =$


 FIGURE 4. \mathbb{C} as a Covering Space Mapped to an Open Strip

$1 - x_1 - x_2 - \dots - x_d$. The generators are $v_i = \epsilon_i$ for $i = 1, \dots, d$ (the standard basis) and $v_{d+1} = -\sum_{i=1}^d v_i$. We have

$$d \left(\frac{1}{2} \sum_{i=1}^{d+1} \ell_i(x) \log \ell_i(x) \right) = \frac{1}{2} \sum_{i=1}^d \log \left(\frac{x_i}{1 - \sum_{j=1}^d x_j} \right) dx_i$$

as a map $P^\circ \rightarrow \mathbb{R}^d$. By direct computation, the inverse of this map equals to

$$\sigma(\vec{r}) = \left(\frac{e^{2r_i}}{1 + \sum_{j=1}^d e^{2r_j}} \right)_{i=1}^d : \mathbb{R}^d \rightarrow P^\circ. \quad (4.1)$$

Written in terms of the complex coordinates $\vec{z} \in (\mathbb{C}^\times)^d$, the symplectomorphism $((\mathbb{C}^\times)^d, \omega_{\mathbb{P}^n}|_{(\mathbb{C}^\times)^d}) \rightarrow (P^\circ \times T^d, \omega_{\text{std}}|_{P^\circ \times T^d})$ is given by

$$\left(\frac{|z_i|^2}{1 + \sum_{j=1}^d |z_j|^2}, \frac{\log z_i - \log \bar{z}_i}{2i} \right)_{i=1}^d.$$

Pulling back by $\mathbb{C}^d \rightarrow (\mathbb{C}^\times)^d$, we have the symplectomorphism $\sigma_{\mathbb{C}} = \left(\left(\frac{e^{2r_i}}{1 + \sum_{j=1}^d e^{2r_j}} \right)_{i=1}^d, \text{Id} \right) :$

$(\mathbb{C}^d, \exp^* \omega_X) \rightarrow (P^\circ \times \mathbb{R}^d, \omega_{\text{std}})$.

When $d = 1$, $\frac{e^{2r}}{1+e^{2r}}$ is a commonly-used activation function. See Figure 4. By taking a direct product, $\left(\frac{e^{2r_i}}{1+e^{2r_i}} \right)_{i=1}^d : \mathbb{R}^d \rightarrow [0, 1]^d$ corresponds to $(\mathbb{P}^1)^d$.

Remark 4.5. In above, we have taken the quotient Kähler structure from \mathbb{C}^m . For a general toric Kähler structure, the symplectomorphism in Theorem 4.1 is given by

$$\left(d \left(\frac{1}{2} \left(\sum_{i=1}^m \ell_i \log \ell_i(x) \right) + h \right), \text{Id} \right)$$

where h is a smooth function on the closed polytope P such that the Hessian of $(\frac{1}{2}(\sum_{i=1}^m \ell_i \log \ell_i(x)) + h)$ is positive definite in P° [Abr98].

In particular, for a general projective toric variety X , we can take an embedding of X to \mathbb{P}^N by toric holomorphic sections of a very ample line bundle L , and use the induced toric Kähler structure from \mathbb{P}^N . Then the symplectomorphism $\sigma_{\mathbb{C}}$ is given by (σ, Id) where

$$\sigma(\vec{r}) = \frac{\sum_{i=1}^d e^{2(\vec{u}_i, \vec{r})} \vec{u}_i}{\sum_{j=1}^d e^{2(\vec{u}_j, \vec{r})}} : \mathbb{R}^d \rightarrow P^\circ,$$

\vec{u}_i are points such that their convex hull equal to P , and (\vec{u}, \vec{r}) is the standard dot product on \mathbb{R}^d . See [Ful93, Section 4.2].

Now we have the symplectomorphisms $\phi_X : (P^\circ \times T^d, \omega_{\text{std}}|_{P^\circ \times T^d}) \xrightarrow{\cong} ((\mathbb{C}^\times)^d, \omega_X|_{(\mathbb{C}^\times)^d})$ and $\phi_{\mathbb{C}^d} : \mathbb{R}_{>0}^d \times T^d \xrightarrow{\cong} ((\mathbb{C}^\times)^d, \omega_{\mathbb{C}^d}|_{(\mathbb{C}^\times)^d})$ (Example 4.3). For the toric structure of X , let's arrange the order of the indices such that the first d vectors v_i for $i = 1, \dots, d$ form a basis of \mathbb{Z}^d . (We assume $m \geq d$.) Moreover, we take the first d constants $c_j = 0$ for $j = 1, \dots, d$. Then $P^\circ \subset \mathbb{R}_{>0}^d$.

Consider the composition $\phi_{\mathbb{C}^d} \circ \phi_X^{-1} : ((\mathbb{C}^\times)^d, \omega_X|_{(\mathbb{C}^\times)^d}) \rightarrow ((\mathbb{C}^\times)^d, \omega_{\mathbb{C}^d}|_{(\mathbb{C}^\times)^d})$. It is a symplectomorphism onto the image $\phi_{\mathbb{C}^d}(P^\circ \times T^d)$.

Proposition 4.6. $\phi_{\mathbb{C}^d} \circ \phi_X^{-1}$ extends to a T -equivariant symplectomorphism

$$\psi : (\mathbb{C}^d, \omega_X|_{\mathbb{C}^d}) \xrightarrow{\cong} (\pi_{\mathbb{C}^d}^{-1}(P - B), \omega_{\mathbb{C}^d}|_{\pi_{\mathbb{C}^d}^{-1}(P - B)})$$

where $B = \bigcup_{i=d+1}^m \{\ell_i(x) = 0\}$, and $\pi_{\mathbb{C}^d} = (|z_i|^2)_{i=1}^d : \mathbb{C}^d \rightarrow \mathbb{R}_{\geq 0}^d$ is the moment map for \mathbb{C}^d .

Proof. ϕ_X is given by $2r_i^X = \log x_i + \sum_{j=d+1}^m (v_j^{(i)} \log \ell_j(x))$, where $v_j = (v_j^{(1)}, \dots, v_j^{(d)})$. $\phi_{\mathbb{C}^d}$ is given by $2r_i^{\mathbb{C}^d} = \log x_i$. Hence $e^{2r_i^X} = e^{2r_i^{\mathbb{C}^d}} \prod_{j=d+1}^m \ell_j^{v_j^{(i)}}(e^{2r_1^{\mathbb{C}^d}}, \dots, e^{2r_d^{\mathbb{C}^d}})$. In terms of the complex coordinates, this gives

$$z_i^X = z_i^{\mathbb{C}^d} \left(\prod_{j=d+1}^m \ell_j^{v_j^{(i)}}(|z_1^{\mathbb{C}^d}|^2, \dots, |z_d^{\mathbb{C}^d}|^2) \right)^{1/2}.$$

It is obviously well-defined over $\pi_{\mathbb{C}^d}^{-1}(P - B)$. We need to show that it has inverse, which gives the required extension of $\phi_{\mathbb{C}^d} \circ \phi_X^{-1}$. Since $(\phi_{\mathbb{C}^d} \circ \phi_X^{-1})^*(\omega_{\mathbb{C}^d}) = \omega_X$ on $(\mathbb{C}^\times)^d$, this still holds over \mathbb{C}^d as the equality is a closed condition.

Consider the Jacobian of $\bar{z}^X(\bar{z}^{\mathbb{C}^d}, \overline{\bar{z}^{\mathbb{C}^d}})$. We shall show it is positive definite, and hence invertible. To simplify, we write $z = z^{\mathbb{C}^d}$. Denote $G = \frac{1}{2}(\sum_{i=1}^m \ell_i(x) \log \ell_i(x) - v \cdot x)$. For

any non-zero vector (a_1, \dots, a_d) ,

$$\begin{aligned}
& \sum_{i,j} \bar{a}_j a_i \partial_{z_i} z_j^X \\
&= \sum_{i,j} \bar{a}_j a_i \partial_{z_i} \left(\exp(\partial_{x_j}|_{x_p=z_p \bar{z}_p} G) \cdot \frac{z_j}{|z_j|} \right) \\
&= \sum_{i,j} \bar{a}_j a_i \exp(\partial_{x_j}|_{x_p=z_p \bar{z}_p} G) \cdot \partial_{z_i} \left(\frac{z_j}{|z_j|} \right) + \sum_{i,j} \bar{a}_j a_i \frac{z_j}{|z_j|} \cdot \exp(\partial_{x_j}|_{x_p=z_p \bar{z}_p} G) \cdot \frac{\partial(z_i \bar{z}_i)}{\partial z_i} \cdot \frac{\partial^2 G}{\partial x_i \partial x_j} \\
&= \sum_i |a_i|^2 |z_i^X| \cdot \left(\frac{1}{2|z_i|} \right) + \sum_{i,j} |z_j|^{-1} \cdot |z_j^X| \cdot \bar{a}_j z_j \cdot a_i \bar{z}_i \cdot \frac{\partial^2 G}{\partial x_i \partial x_j}.
\end{aligned}$$

Note that $|z_j|^{-1} |z_j^X| = \left(\prod_{k=d+1}^m \ell_k^{v_k(j)} \left(|z_1^{\mathbb{C}^d}|^2, \dots, |z_d^{\mathbb{C}^d}|^2 \right) \right)^{1/2}$ which is positive. Let $c = \min\{|z_j|^{-1} |z_j^X| : j = 1, \dots, d\}$. Then the second term is no less than $ca_i \bar{z}_i \cdot \frac{\partial^2 G}{\partial x_i \partial x_j}$. Since $\frac{\partial^2 G}{\partial x_i \partial x_j}$ is positive definite on P° , it is semi-positive definite on $P - B$. Thus this term is non-negative. The first term is positive. Thus $\sum_{i,j} \bar{a}_j a_i \partial_{z_i} z_j^X > 0$. Similarly $\sum_{i,j} a_j \bar{a}_i \partial_{z_i} z_j^X$. Hence the Jacobian is positive-definite and hence invertible. \square

Example 4.7. We continue to consider \mathbb{P}^d . From Example 4.4, $x_i = \frac{|z_i^{\mathbb{P}^d}|^2}{1 + \sum_{j=1}^d |z_j^{\mathbb{P}^d}|^2}$. $\pi_{\mathbb{C}^d}^{-1}(P - B) = \{\|\bar{z}^{\mathbb{C}^d}\| < 1\}$. From Example 4.3, $x_i = |z_i^{\mathbb{C}^d}|^2$. Hence the symplectomorphism $(\mathbb{C}^d, \omega_{\mathbb{P}^d}|_{\mathbb{C}^d}) \xrightarrow{\cong} (\{\|\bar{z}^{\mathbb{C}^d}\| < 1\}, \omega_{\mathbb{C}^d}|_{\pi_{\mathbb{C}^d}^{-1}(P-B)})$ is

$$z_i^{\mathbb{C}^d} = \frac{z_i^{\mathbb{P}^d}}{\sqrt{1 + \|\bar{z}^{\mathbb{P}^d}\|^2}}. \quad (4.2)$$

When $d = 1$, this gives $\frac{z}{\sqrt{1+|z|^2}}$ which is another activation function used in machine learning. (z is restricted in \mathbb{R} in most algorithms.) The symplectomorphism can be easily checked in this case: ($z = z^{\mathbb{P}^1}$ for simplicity)


$$\begin{aligned}
dz^{\mathbb{C}} \wedge \overline{dz^{\mathbb{C}}} &= d \frac{z}{\sqrt{1+|z|^2}} \wedge d \frac{\bar{z}}{\sqrt{1+|z|^2}} \\
&= \frac{(1+z\bar{z})dz - (|z|^2 dz + z^2 \overline{dz})/2}{(1+z\bar{z})^{3/2}} \wedge \frac{(1+z\bar{z})\overline{dz} - (|z|^2 \overline{dz} + \bar{z}^2 dz)/2}{(1+z\bar{z})^{3/2}} = \frac{dz \wedge \overline{dz}}{(1+z\bar{z})^2}
\end{aligned}$$

giving the Fubini-Study metric. See Figure 5.

By taking the direct product, the symplectomorphism for the case $X = (\mathbb{P}^1)^d$ is $\left(\frac{z_j}{\sqrt{1+|z_j|^2}} \right)_{i=1}^d$.

Due to the nice fact that the Fubini-Study metric on \mathbb{P}^d is $U(d)$ -invariant (and so does the standard metric on \mathbb{C}^d), we have the following (which is not true for $(\mathbb{P}^1)^d$ nor general toric manifolds).

Lemma 4.8. For $X = \mathbb{P}^d$, the symplectomorphism $\psi : (\mathbb{C}^d, \omega_{\mathbb{P}^d}|_{\mathbb{C}^d}) \xrightarrow{\cong} \left(\pi_{\mathbb{C}^d}^{-1}(P - B), \omega_{\mathbb{C}^d}|_{\pi_{\mathbb{C}^d}^{-1}(P-B)} \right)$ in Proposition 4.6 is $U(d)$ -equivariant.

FIGURE 5. \mathbb{C} as a Chart in \mathbb{P}^1 Mapped to an Open Disk

Proof. Using the explicit expression (4.2),

$$\psi(U \cdot \bar{z}^{\mathbb{P}^d}) = \frac{U \cdot \bar{z}^{\mathbb{P}^d}}{\sqrt{1 + \|U \cdot \bar{z}^{\mathbb{P}^d}\|^2}} = \frac{U \cdot \bar{z}^{\mathbb{P}^d}}{\sqrt{1 + \|\bar{z}^{\mathbb{P}^d}\|^2}} = U \cdot \psi(\bar{z}^{\mathbb{P}^d})$$

for all $U \in U(d)$. □

As explained in Remark 4.5, we can also equip X with another T -invariant Kähler form (that do not come from the standard Kähler structure on \mathbb{C}^m). The symplectomorphism is given by

$$\left(\frac{1}{2}d \left(\sum_{i=1}^m \ell_i(x) \log \ell_i(x) + h(x) \right), \text{Id} \right) : (P^o \times T^d, \omega_{\text{std}}|_{P^o \times T^d}) \rightarrow \mathbb{R}^d \times T^d \stackrel{\text{exp}}{\cong} ((\mathbb{C}^\times)^d, \omega^h|_{(\mathbb{C}^\times)^d}).$$

Thus the toric construction is rather flexible.

Example 4.9. The ‘softplus’ function $x = \log(1 + e^{2y}) : \mathbb{R} \rightarrow \mathbb{R}_{>0}$ gives an example of such a Kähler structure on \mathbb{C} (by identifying it with $\mathbb{R}_{>0} \times \mathbb{R}$ with the standard symplectic structure). The inverse is $y = \frac{1}{2} \log(e^x - 1)$, whose difference with $\frac{1}{2} \log x$ is $h' = \frac{1}{2} \log \frac{e^x - 1}{x} = \frac{1}{2} \log \left(1 + \sum_{k=1}^{\infty} \frac{x^k}{(k+1)!} \right)$ which is indeed a smooth function on $\mathbb{R}_{\geq 0}$. Moreover, $y' = \frac{e^x}{2(e^x - 1)} > 0$ on $\mathbb{R}_{>0}$.

4.2 Symplectomorphisms of fiber bundles over the moduli. In the last subsection, we have exhibited various symplectic embeddings for \mathbb{C}^n . Now we want to make a family version of these maps over the framed quiver moduli $\mathcal{M}_{\vec{n}, \vec{d}}$. The last subsection can be understood as constructing self-maps on a fiber of a vector bundle over $\mathcal{M}_{\vec{n}, \vec{d}}$.

To globalize (4.2), we consider the universal bundle \mathcal{V}_i equipped with a Hermitian metric H_i . (We have constructed a nice Hermitian metric on \mathcal{V}_i in Section 3.1.) We have a fiberwise symplectic structure $\omega_{\mathcal{V}_i}$ induced from the Hermitian metric. Moreover, we have the projective bundle $\mathbb{P}(\mathcal{V}_i \oplus \mathcal{O}_{\mathcal{M}})$ which is a fiberwise compactification of \mathcal{V}_i . Then the fiber bundle $\mathbb{P}(\mathcal{V}_i \oplus \mathcal{O}_{\mathcal{M}})$ is equipped with a fiberwise Kähler metric $\omega_{\mathbb{P}(\mathcal{V}_i \oplus \mathcal{O}_{\mathcal{M}})}$ induced from H_i (namely, $\frac{i}{2} \partial \bar{\partial} \log(H_i \oplus H_0)$ where H_0 is the trivial metric on $\mathcal{O}_{\mathcal{M}}$).

Proposition 4.10. *There is a fiberwise symplectomorphism*

$$\psi_{\mathcal{V}_i} : (\mathcal{V}_i, \omega_{\mathbb{P}(\mathcal{V}_i \oplus \mathcal{O}_{\mathcal{M}})}|_{\mathcal{V}_i}) \xrightarrow{\cong} \left(\{v \in \mathcal{V}_i : H_i(v, v) < 1\}, \omega_{\mathcal{V}_i}|_{\{H_i(v, v) < 1\}} \right).$$

Proof. For each $p \in \mathcal{M}$, we have computed the symplectomorphism

$$(\mathcal{V}_i|_p, \omega_{\mathbb{P}(\mathcal{V}_i \oplus \mathcal{O}_{\mathcal{M}})}|_{\mathcal{V}_i|_p}) \xrightarrow{\cong} \left(\{v \in \mathcal{V}_i|_p : H_i(v, v) < 1\}, \omega_{\mathcal{V}_i}|_{\{H_i(v, v) < 1\}} \right)$$

in (4.2), with the metric given by $H_i|_p$ here. Thus

$$\psi_{\mathcal{V}_i}(v) = \frac{v}{\sqrt{1 + H_i(v, v)}} \tag{4.3}$$

gives a fiberwise symplectomorphism whose image is $\{H_i(v, v) < 1\}$. \square

Recall that the universal bundle $\mathcal{V}_i \rightarrow \mathcal{M}$ admits an action of $U_{\bar{n}}$ coming from framing (Definition 3.11). One advantage of $\psi_{\mathcal{V}_i}$ is that it is equivariant under this action.

Lemma 4.11. *For $g \in U_{\bar{n}}$,*

$$\psi_{\mathcal{V}_i} \circ g = g \circ \psi_{\mathcal{V}_i}$$

if we used the metric H_i given in Theorem 3.7.

Proof. This follows from (4.3) and $H_i(g \cdot v, g \cdot v) = H_i(v, v)$ by Lemma 3.13. \square

Remark 4.12. *Equation (4.3) has an alternative derivation using the framing. Namely, we have the surjective morphism $\rho : \underline{\hat{W}}_i \rightarrow \mathcal{V}_i$ (see Equation (3.4)), whose dual give a fiberwise-linear embedding $\rho^* : \mathcal{V}_i \xrightarrow{H_i} \mathcal{V}_i^* \rightarrow \underline{\hat{W}}_i^*$. (The underline means the trivial bundle over \mathcal{M} associated with the vector space.) Then $\mathbb{P}(\underline{\hat{W}}_i^* \oplus \mathcal{O})$ (with the standard metric) induces a fiberwise Kähler form on \mathcal{V}_i^* , and we have a fiberwise symplectic embedding $(\mathcal{V}_i, \omega_{\mathbb{P}(\underline{\hat{W}}_i^* \oplus \mathcal{O})}) \hookrightarrow (\mathcal{V}_i, \omega_{\mathcal{V}_i})$. This gives*

$$\frac{\rho\rho^* \cdot H_i \cdot v}{\sqrt{1 + H_0(\rho^* \cdot H_i \cdot v, \rho^* \cdot H_i \cdot v)}} = \frac{\rho\rho^* \cdot H_i \cdot v}{\sqrt{1 + v^* \cdot H_i^* \cdot \rho\rho^* \cdot H_i \cdot v}}.$$

Now if we use the metric $H_i = (\rho\rho^)^{-1}$ given by Theorem 3.7, then the above equals to the expression in (4.3).*

We also have a more flexible construction using the framing, which globalize *any given non-linear continuous map* $\sigma_{\mathbb{C}} : \mathbb{C}^{n_i} \rightarrow \mathbb{C}^{n_i}$. Namely, $\sigma_{\mathbb{C}}$ can be regarded as a fiberwise non-linear self-map on the trivial bundle $\underline{\mathbb{C}}^{n_i} \rightarrow \underline{\mathbb{C}}^{n_i}$ (still denoted by $\sigma_{\mathbb{C}}$). Then we take the composition

$$\mathcal{V}_i \xrightarrow{H_i} \mathcal{V}_i^* \xrightarrow{(e^{(i)})^*} \underline{\mathbb{C}}^{n_i} \xrightarrow{\sigma_{\mathbb{C}}} \underline{\mathbb{C}}^{n_i} \xrightarrow{e^{(i)}} \mathcal{V}_i$$

and denote it by $\sigma_{\mathcal{V}_i}$. See Figure 6.

$$\begin{array}{ccc} \sigma_{\mathbb{C}} \hookrightarrow \underline{\mathbb{C}}^n & \xrightarrow{\cong} & \underline{\mathbb{C}}^{n^*} \\ & \downarrow e^{(i)} & \uparrow (e^{(i)})^* \\ \mathcal{V}_i & \xrightarrow{\cong} & \mathcal{V}_i^* \end{array}$$

FIGURE 6. Using the framing to globalize an activation function $\sigma_{\mathbb{C}} : \mathbb{C}^n \rightarrow \mathbb{C}^n$ over \mathcal{M} .

It is easy to get the following explicit expression in terms of $\sigma_{\mathbb{C}}$.

Lemma 4.13. *The above fiber-bundle map $\sigma_{\mathcal{V}_i} : \mathcal{V}_i \rightarrow \mathcal{V}_i$ equals to*

$$\sigma_{\mathcal{V}_i}(v) = \sum_{k=1}^{n_i} (\sigma_{\mathbb{C}})_k \left(H_i(e_1^{(i)}, v), \dots, H_i(e_{n_i}^{(i)}, v) \right) \cdot e_k^{(i)} \quad (4.4)$$

where we write $\sigma_{\mathbb{C}} = ((\sigma_{\mathbb{C}})_1, \dots, (\sigma_{\mathbb{C}})_{n_i})$ and $e^{(i)} = (e_1^{(i)} \dots e_{n_i}^{(i)})$.

For instance, we can take σ to be the one coming from the symplectomorphism in Corollary 4.2. We shall prove the universal approximation theorem for such σ in Section 5.

4.3 A machine learning program using the framed quiver moduli. Let Q be a digraph and denote by $\mathbb{C} \cdot Q$ its path algebra over \mathbb{C} . Let $\vec{d} \in \mathbb{Z}_{\geq 0}^{Q_0}$ be a dimension vector. We take the framing dimension vector to be $\vec{n} = \vec{d} + \vec{1}$, where $\vec{1}_i := 1$ for all $i \in Q_0$. The one additional framing vector is used for translation (called a ‘bias’ vector).

We fix a collection of input vertices and a collection of output vertices $I_{\text{in}}, I_{\text{out}} \subset Q_0$, and

$$\gamma \in I_{\text{out}} \cdot (\mathbb{C} \cdot Q) \cdot I_{\text{in}} = \bigoplus_{i \in I_{\text{in}}, j \in I_{\text{out}}} j \cdot (\mathbb{C} \cdot Q) \cdot i.$$

(The trivial path at a vertex i is again denoted by i .)

4.3.1 Machine learning using a flat space. Let’s first formulate a typical machine learning program in the quiver setup. The following flat space

$$U \cong \prod_{a \in Q_1} \text{Hom}(V_{t(a)}, V_{h(a)}) \times \prod_{i \in Q_0} V_i$$

is used frequently in the subject.

Lemma 4.14. *The open subset*

$$U = \{[V, e] \in \mathcal{M}_{\vec{d} + \vec{1}, \vec{d}} : (e_1^{(i)}, \dots, e_{d_i}^{(i)}) = I_{d_i} \text{ for all } i \in Q_0\} \subset \mathcal{M}$$

gives a coordinate chart of \mathcal{M} . (I_{d_i} denotes the identity matrix of rank d_i .)

Proof. First, such (V, e) are stable: $\text{Im}(e)$ is the whole V . Second, for distinct $(V, e), (V', e')$ satisfying the above condition, $[V, e] \neq [V', e']$: since they are stable, their orbits are closed. Suppose $g \cdot (V, e) = (V', e')$ for some $g \in \text{GL}_{\vec{d}}$. Then $g_i \cdot e^{(i)} = (e')^{(i)}$. But since $(e_1^{(i)}, \dots, e_{d_i}^{(i)}) = I_{d_i} = ((e')_1^{(i)}, \dots, (e')_{d_i}^{(i)})$, this forces $g = \text{Id}$ and so $(V, e) = (V', e')$, contradicting that they are distinct.

Then we have the chart map $U \xrightarrow{\cong} \prod_{a \in Q_1} \text{Hom}(V_{t(a)}, V_{h(a)}) \times \prod_{i \in Q_0} V_i$ defined by

$$W_a := (e_1^{(i)}, \dots, e_{d_i}^{(i)})^{-1} V(a), \quad b_i := (e_1^{(i)}, \dots, e_{d_i}^{(i)})^{-1} e_{d_i+1}^{(i)}.$$

□

Now we fix a path γ from the input vertices I_{in} to the output vertices I_{out} . For each element $[V, e] \in U$, by composing the affine linear maps $V_a(\cdot) + e_{d_{t(a)}+1}^{(t(a))}$ attached to arrows a in the path γ , together with some non-linear functions $\sigma_i : V_i \rightarrow V_i$ that are called ‘activation functions’, one obtains a non-linear function

$$f_{[V, e]}^U : \bigoplus_{i \in I_{\text{in}}} \mathbb{C}^{d_i} \xrightarrow{(e_1^{(i)}, \dots, e_{d_i}^{(i)})} \bigoplus_{i \in I_{\text{in}}} V_i \rightarrow \bigoplus_{j \in I_{\text{out}}} V_j \xrightarrow{(e_1^{(j)}, \dots, e_{d_j}^{(j)})} \bigoplus_{j \in I_{\text{out}}} \mathbb{C}^{d_j}$$

which is used to approximate a non-explicitly given function f . A stochastic gradient flow of the error function on U is employed to find the optimal point in U .

Let’s write down the symmetry in Lemma 4.11 in the chart U .

Lemma 4.15. *The chart U is invariant under the right action by*

$$\prod_{i \in Q_0} U(d_i) \subset \prod_{i \in Q_0} U(d_i + 1) = U_{\vec{n}}$$

where $U(d_i) \subset U(d_i + 1)$ is embedded as $\begin{pmatrix} U(d_i) & 0 \\ 0 & 1 \end{pmatrix}$.

Consider the trivialization $\mathcal{V}_i|_U \cong \left(\prod_{a \in Q_1} \text{Hom}(V_{t(a)}, V_{h(a)}) \times \prod_{j \in Q_0} V_j \right) \times V_i$. The right action of $g \in U(d_i)$ on $\mathcal{V}_i|_U$ is given by

$$(W_a, b_j, v) \begin{matrix} a \in Q_1 \\ j \in Q_0, b_j \in V_j \\ v \in V_i \end{matrix} \cdot g = (g^{-1} \cdot W_a, g^{-1} \cdot b_j, g^{-1} v) \begin{matrix} a \in Q_1 \\ j \in Q_0, b_j \in V_j \\ v \in V_i \end{matrix}$$

where $g^{-1} \cdot W_a$ equals to $g^{-1} W_a$ if $h(a) = i$, $W_a g$ if $t(a) = i$, and W_a otherwise; $g^{-1} \cdot b_j$ equals to $g^{-1} b_j$ if $j = i$, and b_j if $j \neq i$.

Proof. U consists of points $[V, e]$ where $(e_1^{(j)}, \dots, e_{d_i}^{(j)})$ are invertible for all $j \in Q_0$. This property is invariant under the action of $\prod_{i \in Q_0} U(d_i)$. Hence U is an invariant subset.

$(W_a, b_j, v) \begin{matrix} a \in Q_1 \\ j \in Q_0, b_j \in V_j \\ v \in V_i \end{matrix}$ corresponds to the point $[W, e, v] \in \mathcal{V}_i = (R_{n,d}^s \times V_i) / \text{GL}_{\vec{d}}$ where $e^{(j)} = (I_{d_j}, b_j)$.

$$\begin{aligned} [W, e, v] \cdot g &= [W, e \cdot g, v] = [W, (g b_j)_{j \in Q_1}, v] \\ &= [g^{-1} \cdot W, (I_{d_j}, g^{-1} \cdot b_j)_{j \in Q_1}, g^{-1} \cdot v] \end{aligned}$$

where the left action by g^{-1} (as an element in $\text{GL}_{\vec{d}}$) is as specified by definition. \square

Now we prove an important symmetric property that the activation function (4.2) enjoys, which can be used to reduce the dimensions.

Proposition 4.16. *Suppose the activation functions $\sigma_i : V_i \rightarrow V_i$ are taken to be the one given in Equation (4.2). Then f^U is invariant under $\prod_{j \notin I_{\text{in}} \cup I_{\text{out}}} U(d_j) \subset U_{\vec{n}}$. (See the embedding in Lemma 4.15.)*

Proof. The terms of f^U are of the form $\sigma_{h(a_k)} \left(W_{a_k} \dots \left(\sigma_{h(a_1)} \left(W_{a_1}(v) + b_{h(a_1)} \right) \dots \right) + b_{h(a_k)} \right)$. By the above lemma, for $g \in U(d_{h(a_i)})$ where $h(a_i) \notin I_{\text{in}} \cup I_{\text{out}}$, the action of g results in $\sigma_{h(a_i)} \mapsto g \cdot \sigma_{h(a_i)} g^{-1}$ in the above expression and does not affect any other part. By Lemma 4.8, σ_i is $U(d_i)$ -equivariant, and hence f^U remains invariant. \square

By the above proposition, we can descend f^U to the orbit space $U / \prod_{j \notin I_{\text{in}} \cup I_{\text{out}}} U(d_j)$ to reduce the dimensions. However, the quotient space will be highly singular. Instead, we can realize the dimension reduction by restricting f^U to a submanifold U' of U whose orbit occupies the whole U , and do stochastic gradient flow on U' instead of U . The following gives one simple possibility.

Proposition 4.17. *Consider a subset $\{a_1, \dots, a_p\}$ of arrows whose heads and tails do not belong to $I_{\text{in}} \cup I_{\text{out}}$, and for any two distinct arrows a_1, a_2 in the subset, $t(a_1) \neq h(a_2)$. Then the vector subspace*

$$U' := \left\{ (W_a, b_j) \begin{matrix} a \in Q_1 \\ j \in Q_0, b_j \in V_j \end{matrix} \in U : W_{a_i} \text{ is of the form } (D W'_{a_i}) \ \forall i = 1, \dots, p \right\}$$

has its orbit being the whole U , that is, $\left(\prod_{j \notin I_{\text{in}} \cup I_{\text{out}}} U(d_j) \right) \cdot U' = U$. In above, D is a diagonal matrix (of the maximum possible size) and W'_{a_i} is any matrix occupying the rest.

Proof. This follows from the singular-value decomposition of a matrix W as $A \cdot (D W') \cdot B$ where A and B are unitary matrices of appropriate sizes. \square

4.3.2 *Machine learning using the quiver moduli.* The quiver moduli $\mathcal{M} = \mathcal{M}_{\vec{n}, \vec{d}}$ gives a compactification of U . Compactness is important for the formulation of Morse theory and convergence of a gradient flow. We would like to use the whole \mathcal{M} in application of machine learning. Non-trivial metrics over the moduli will play a crucial role.

First, consider the situation before adding in activation functions. Each arrow $a \in Q_1$ is associated with a vector-bundle morphism $a_{\mathcal{M}} : \mathcal{V}_{t(a)} \rightarrow \mathcal{V}_{h(a)}$. For each path $a_k \dots a_1$, we take the map

$$(a_k)_{\mathcal{M}} \left(\dots \left((a_2)_{\mathcal{M}} \left((a_1)_{\mathcal{M}}(v) + e_{d_{h(a_1)}+1}^{(h(a_1))} \right) + e_{d_{h(a_2)}+1}^{(h(a_2))} \right) \dots \right) + e_{d_{h(a_k)}+1}^{(h(a_k))} \quad (4.5)$$

which is fiberwise affine linear. Thus a path γ gives an affine bundle morphism

$$\gamma_{\mathcal{M}} : \mathcal{V}_{I_{\text{in}}} := \bigoplus_{i \in I_{\text{in}}} \mathcal{V}_i \rightarrow \mathcal{V}_{I_{\text{out}}} := \bigoplus_{j \in I_{\text{out}}} \mathcal{V}_j.$$

Then we have

$$L_{\gamma} \left(\left(s_k^{(i)} \right)_{\substack{i \in I_{\text{in}} \\ k \in \{1, \dots, d_i\}}} \right) = \left(H_j \left(e_p^{(j)}, \sum_{i \in I_{\text{in}}} \gamma_{\mathcal{M}} \cdot \sum_{k=1}^{d_i} s_k^{(i)} e_k^{(i)} \right) \right)_{\substack{j \in I_{\text{out}}, \\ p \in \{1, \dots, d_j\}}} : \bigoplus_{i \in I_{\text{in}}} \mathbb{C}^{d_i} \rightarrow \bigoplus_{j \in I_{\text{out}}} \mathbb{C}^{d_j}. \quad (4.6)$$

Suppose a continuous function $f = K \rightarrow \bigoplus_{j \in I_{\text{out}}} \mathbb{C}^{d_j}$ is given, where K is a compact subset of $\bigoplus_{i \in I_{\text{in}}} \mathbb{C}^{d_i}$. Then the fiberwise integral

$$\mathcal{E} := \int_K \|f - L_{\gamma}\|_{\mathcal{V}_{I_{\text{out}}}}^2 d\mu_K \quad (4.7)$$

gives a smooth function on \mathcal{M} .

Remark 4.18. *Alternatively, we can define the fiber-bundle morphism*

$$f_{\mathcal{M}} : \sum_{j \in I_{\text{out}}} \sum_{l=1}^{d_j} f_l^{(j)} \cdot e_l^{(j)} : \underline{K} \rightarrow \mathcal{V}_{I_{\text{out}}}$$

and take

$$\int_K \left\| f_{\mathcal{M}} \left(\left(s_k^{(i)} \right)_{\substack{i \in I_{\text{in}} \\ k \in \{1, \dots, d_i\}}} \right) - \sum_{i \in I_{\text{in}}} \gamma_{\mathcal{M}} \cdot \sum_{k=1}^{d_i} s_k^{(i)} e_k^{(i)} \right\|_{\mathcal{V}_{I_{\text{out}}}}^2 d\mu_K.$$

However, with such a definition, we need to worry that $(e_1^{(j)}, \dots, e_{d_j}^{(j)})$ for some output j degenerates (as a frame), in which case approximating f and approximating $f_{\mathcal{M}}$ are different.

We have a gradient flow $r : \mathbb{R} \rightarrow \mathcal{M}$ which can be used to minimize \mathcal{E} :

$$\frac{dr}{dt} = -(\nabla \mathcal{E})(r(t))$$

where $\nabla \mathcal{E} = (d\mathcal{E})^{\#_g}$ where $(\cdot)^{\#_g} : T^* \mathcal{M} \xrightarrow{\cong} T\mathcal{M}$ is the identification by a metric g on \mathcal{M} . ($(\nabla \mathcal{E})^p = g^{pq} \partial_q \mathcal{E}$ in local coordinates.) g can be taken to be the induced metric from the trivial metric on the vector space $R_{\vec{n}, \vec{d}}$ via symplectic reduction. Alternatively, g can be taken to be the metric given by the Ricci curvature in Theorem 3.15 (when Q has no oriented cycle), which has a better expression in homogeneous coordinates.

Note that $(L_{\gamma})|_{[V, e]}$ is affine linear on $\bigoplus_{i \in I_{\text{in}}} \mathbb{C}^{d_i}$ for every $[V, e] \in \mathcal{M}$, which is not good enough for the purpose of approximating f . We introduce fiberwise non-linearity below.

Definition 4.19. Let A be a finite set whose every element is associated with two vertices (head h and tail t) in Q_0 . Elements in A are called activation arrows. (These are not arrows in Q_1 .) The semiring generated by Q_1 and A , denoted by $\Gamma(Q, A)$, has the underlying vector space spanned by the independent set $\coprod_{p=0}^{\infty} S_p$ where S_p is defined inductively as follows.

- (1) S_0 consists of all paths of Q .
- (2) Suppose S_p has been defined, and each element in S_p has a head h and a tail t . S_{p+1} consists of $\gamma \cdot \alpha \cdot \tilde{\gamma}$, where $\alpha \in A$, γ is any path of Q with $h(\alpha) = t(\gamma)$, and $\tilde{\gamma} \in t(\alpha) \cdot (\mathbb{C} \cdot S_p)$. ($\mathbb{C} \cdot S_p$ denotes the vector space generated by S_p ; $t(\alpha) \cdot (\mathbb{C} \cdot S_p)$ is the subspace generated by elements of S_p with head being $t(\alpha)$.) The above element has the head $h(\gamma)$ and the tail $t(\tilde{\gamma})$.

The above vector space $\Gamma(Q, A)$ has an obvious product by concatenation. ($s_1 \cdot s_2 = 0$ if $h(s_2) \neq t(s_1)$.) Note that for $\alpha \in A$, $(s_1 + cs_2) \cdot \alpha = s_1 \cdot \alpha + cs_2 \cdot \alpha$, but $\alpha \cdot (s_1 + cs_2) \neq \alpha \cdot s_1 + c\alpha \cdot s_2$.

Now, suppose each element $\alpha \in A$ is associated with a fiber-bundle morphism

$$\alpha_{\mathcal{M}} : \mathcal{V}_{t(\alpha)} \rightarrow \mathcal{V}_{h(\alpha)}.$$

Then by composing the corresponding affine linear morphisms (as in Equation (4.5)) and $\alpha_{\mathcal{M}}$, an element $\tilde{\gamma} \in \Gamma(Q, A)$ induces a fiber-bundle morphism

$$\tilde{\gamma}_{\mathcal{M}} : \mathcal{V}_{t(\tilde{\gamma})} \rightarrow \mathcal{V}_{h(\tilde{\gamma})}.$$

In particular, if we fix $\tilde{\gamma} \in I_{\text{out}} \cdot \Gamma(Q, A) \cdot I_{\text{in}}$, then we define

$$f_{\tilde{\gamma}} : \bigoplus_{i \in I_{\text{in}}} \mathbb{C}^{d_i} \rightarrow \bigoplus_{j \in I_{\text{out}}} \mathbb{C}^{d_j}$$

like in the definition of L_{γ} in (4.6) by replacing $\gamma_{\mathcal{M}}$ by $\tilde{\gamma}_{\mathcal{M}}$. Then a stochastic gradient flow for the corresponding error function (by replacing L_{γ} by $f_{\tilde{\gamma}}$ in (4.7)) can be carried out.

For now, we set A to be Q_0 as a set, and each element \mathfrak{o}_i has the head and tail being $i \in Q_0$. We can associate \mathfrak{o}_i with the fiber-bundle morphism $\psi_{\mathcal{V}_i}$ given in (4.3), or $\sigma_{\mathcal{V}_i}$ in (4.4). Then we obtain $f_{\tilde{\gamma}}$ above which is non-linear along fibers for the purpose of machine learning.

If we associate \mathfrak{o}_i with $\psi_{\mathcal{V}_i}$ given in (4.3), then the symmetry of $\prod_{j \notin I_{\text{in}} \cup I_{\text{out}}} U(d_j) \subset U_{\vec{n}}$ is respected, by Lemma 4.11. The proof is similar to that for Proposition 4.16 and is omitted here.

Proposition 4.20. Suppose \mathfrak{o}_i is assigned as $\psi_{\mathcal{V}_i}$ given in (4.3), and H_i is taken to be the metric in Theorem 3.7. Then $f_{\tilde{\gamma}}$ is invariant under $\prod_{j \notin I_{\text{in}} \cup I_{\text{out}}} U(d_j) \subset U_{\vec{n}}$. (See the embedding in Lemma 4.15.)

Remark 4.21. Since we are taking affine linear morphisms $(a_i)_{\mathcal{M}}(v) + e_{d_{h(a_i)}+1}^{(h(a_i))}$ for the arrows which involves the term $e_{d_{h(a_i)}+1}^{(h(a_i))}$, only the symmetry $U(d_j)$ rather than $U(d_j + 1) = U(n_j)$ is respected. If the bias vector $e_{d_{h(a_i)}+1}^{(h(a_i))}$ is not used in the program, then $f_{\tilde{\gamma}}$ will be invariant under the bigger group $\prod_{j \notin I_{\text{in}} \cup I_{\text{out}}} U(n_j)$.

4.3.3 A simple example. Recall the quiver in Example 2.7. Let's take $n_1 = d_1, n_2 = d_2 + 1, n_3 = d_3$ instead of $n_i = d_i + 1 \forall i$, since we do not need to use bias vectors at the input and output vertices. The path γ is simply $a_2 a_1$, and $\tilde{\gamma} = a_2 \cdot \mathfrak{o}_1 \cdot a_1$.

We have the universal bundles $\mathcal{V}_1 \cong \mathbb{C}^{d_1}$, \mathcal{V}_2 and \mathcal{V}_3 . We have $\rho^{(1)} = e^{(1)}$, $\rho^{(2)} = (e^{(2)} a_1 e^{(1)})$, $\rho^{(3)} = (e^{(3)} a_2 e^{(2)} a_2 a_1 e^{(1)})$. In terms of homogeneous coordinates (namely

the coordinates $(e^{(1)}, e^{(2)}, e^{(3)}, a_1, a_2)$ on the vector space $R_{\vec{n}, \vec{d}}$ (where each entry is a matrix of suitable size), the metrics given in Theorem 3.7 are

$$H_1 = \left(e^{(1)}(e^{(1)})^* \right)^{-1}, \quad H_2 = \left(e^{(2)}(e^{(2)})^* + a_1 e^{(1)}(e^{(1)})^* a_1^* \right)^{-1},$$

$$H_3 = \left(e^{(3)}(e^{(3)})^* + a_2 e^{(2)}(e^{(2)})^* a_2^* + a_2 a_1 e^{(1)}(e^{(1)})^* a_1^* a_2^* \right)^{-1}.$$

on $\mathcal{V}_i, i = 1, 2, 3$ respectively. The activation functions we constructed in the previous subsection are

$$\psi_i(v) = \frac{v}{\sqrt{1 + v^* H_i v}} \text{ and } \sigma_i(v) = \sum_{k=1}^{d_i+1} (\sigma_{\mathbb{C}})_k \left((e_1^{(i)})^* H_i v, \dots, (e_{d_i+1}^{(i)})^* H_i v \right) \cdot e_k^{(i)},$$

where we can take $(\sigma_{\mathbb{C}})_k(\vec{z}) = \frac{e^{2\operatorname{Re}(z_k)}}{1 + \sum_{j=1}^{d_i+1} e^{2\operatorname{Re}(z_j)}} + \mathbf{i} \operatorname{Im}(z_k)$ for instance. Both have the GL-equivariance property $\psi_i(g \cdot v) = g \cdot \psi_i(v)$ and $\sigma_i(g \cdot v) = g \cdot \sigma_i(v)$.

The function (over \mathcal{M}) cooked up from this quiver is $f_{\vec{\gamma}} : \mathbb{C}^{d_1} \rightarrow \mathbb{C}^{d_3}$,

$$f_{\vec{\gamma}}(s_1, \dots, s_{d_1}) = \left(H_3 \left(e_p^{(3)}, a_2 \cdot \sigma_2 \left(a_1 \cdot \sum_{k=1}^{d_1} s_k e_k^{(1)} + e_{d_2+1}^{(2)} \right) \right) \right)_{p=1}^{d_3}$$

if we use σ_2 as the activation function, or the same expression with σ_2 replaced by ψ_2 . Then we run a stochastic gradient flow on \mathcal{M} (or on the vector space $R_{\vec{n}, \vec{d}}$ upstairs) to minimize the distance of $f_{\vec{\gamma}}$ and a function f coming from reality.

To run the gradient flow, we need to take a metric on \mathcal{M} . Recall that we have the metric on the tangent bundle of \mathcal{M} coming from the Ricci curvatures of H_i :

$$H_T(v, w) = \sum_{i=1}^3 \operatorname{tr} \left(\left(\partial_v \rho^{(i)} \right)^* \cdot H_i \cdot \partial_w \rho^{(i)} \right) - \sum_{i=1}^3 \operatorname{tr} \left(\left(\partial_v \rho^{(i)} \cdot (\rho^{(i)})^* \cdot H_i^{\frac{1}{2}} \right)^* \cdot H_i \cdot \left(\partial_w \rho^{(i)} \cdot (\rho^{(i)})^* \cdot H_i^{\frac{1}{2}} \right) \right). \quad (4.8)$$

Consider the open subset of the vector space $R_{\vec{n}, \vec{d}}$ in which $(e_1^{(i)}, \dots, e_{d_i}^{(i)})$ is invertible. (This is the preimage of the chart $U \subset \mathcal{M}$.) A tangent vector v of \mathcal{M} is lifted as $(\delta a_1, \delta a_2, \delta e_{d_2+1}^{(2)})$ (and all other components are set to be zero). Since $\rho^{(1)} = e^{(1)}$, $\rho^{(2)} = (e^{(2)}, a_1 e^{(1)})$, $\rho^{(3)} = (e^{(3)}, a_2 e^{(2)}, a_2 a_1 e^{(1)})$, we have $\partial_v \rho^{(1)} = 0$, $\partial_v \rho^{(2)} = \left((0 \ \delta e_{d_2+1}^{(2)}), (\delta a_1) e^{(1)} \right)$, and

$$\partial_v \rho^{(3)} = \left(0, (\delta a_2) e^{(2)} + (0 \ a_2 \delta e_{d_2+1}^{(2)}), (\delta a_2) a_1 e^{(1)} + a_2 (\delta a_1) e^{(1)} \right).$$

Then the above metric H_T can be computed explicitly in terms of the homogeneous coordinates.

Remark 4.22. *In above, if we use trivial metrics over the vector space $R_{\vec{n}, \vec{d}}$ instead of H_i and H_T , the expressions will get simpler; however they will only be $U_{\vec{d}}$ -equivariant rather than $\operatorname{GL}_{\vec{d}}$ -equivariant. Then we need to restrict to the moment-map level $\mu^{-1}(I) \subset R_{\vec{n}, \vec{d}}$ and its tangent bundle, in order to stay in the same moduli \mathcal{M} downstairs. This would increase the computational complexity.*

We can also write in inhomogeneous coordinates (W_1, W_2, b) in the chart $U \subset \mathcal{M}$, where $(e_1^{(i)}, \dots, e_{d_i}^{(i)}) = I_{d_i}$ for all $i = 1, 2, 3$. W_1, W_2 are the matrices of the arrows a_1, a_2 and $e_{d_2+1}^{(2)} = b$ is the bias vector. Then $H_1 = I$, $H_2 = (I + b b^* + W_1 W_1^*)^{-1}$, and $H_3 =$

$(I + W_2W_2^* + W_2bb^*W_2^* + W_2W_1W_1^*W_2^*)^{-1}$. Then we can run the gradient flow in $U \subset \mathcal{M}$ (which has lower dimensions than $R_{\vec{n}, \vec{d}}$).

Note that when W_1, W_2, b are close to zero, H_i are close to the identity matrix, and σ_i is close to $\sigma_{\mathbb{C}}$. Moreover, the second term of (4.8) is close to zero, and the first term is close to the standard metric. Thus when W_1, W_2, b are small, the function $f_{\vec{\gamma}}$ is close to the commonly used one

$$f_{W_1, W_2, b}^U(v) = W_2 \cdot (\sigma_{\mathbb{C}}(W_1 \cdot v + b)), \quad (4.9)$$

and the gradient flow is close to the usual one on the flat space U (see Section 4.3.1). The additional terms can be understood as modifications to ensure the flow converges in \mathcal{M} .

4.4 A discussion on Morse inequalities. By the work of Reineke [Rei08], the framed quiver moduli \mathcal{M} is a tower of Grassmannians (Theorem 2.4), and hence its Poincaré polynomial is a product of that of Grassmannians (Corollary 2.5). Such topological invariants give important information about a gradient flow on \mathcal{M} .

In particular the Morse inequalities for a Morse function \mathcal{E} on a compact manifold \mathcal{M} state as follows. Let $c^j(\mathcal{E})$ be the number of critical points of index j for \mathcal{E} . Then for every j ,

$$c^j(\mathcal{E}) \geq h^j(\mathcal{M})$$

where $h^j(\mathcal{M})$ denotes the cohomological numbers (which are coefficients of the Poincaré polynomial).

Given a gradient flow, which is a path $\gamma : \mathbb{R} \rightarrow \mathcal{M}$ satisfying the gradient flow equation, $\lim_{t \rightarrow \pm\infty} \gamma(t)$ are critical points. Moreover, critical points carry important effect to the rate of the gradient flow. Namely, when the flow γ gets close to a critical point with index being $1, \dots, \dim \mathcal{M} - 1$, $\|\gamma'(t)\| = \|\text{grad } \mathcal{E}(\gamma(t))\|$ becomes small. In other words the flow slows down when it passes through a neighborhood of a critical point. Such a slowing-down effect of saddle points was studied in machine learning in [PDGB14, DPG⁺14].

The cohomological numbers $h^j(\mathcal{M})$ give the minimum number of critical points and hence are important invariants of a neural network (which simply means a directed graph Q together with a dimension vector \vec{d} here). Over \mathbb{C} , $h^j(\mathcal{M}) = 0$ when j is odd. Thus the Euler characteristic $\chi_{Q, \vec{d}}$ equals to $\sum_j h^j(\mathcal{M})$, which is the minimal total number of critical points. It is computed by simply setting $q = 1$ in Corollary 2.5. Another important invariant is $\dim \mathcal{M}$ (that is, the number of training parameters of the network), which is simply

$$\mathcal{D}_{Q, \vec{d}} = \sum_{i \in Q_0} d_i \binom{n_i + \sum_{\substack{j \rightarrow i \\ j \neq i}} d_j}{d_i}$$

using the notation of Corollary 2.5. (We take $n_i = d_i + 1$ in this section.) These are illustrated in the two practical examples below.

Example 4.23. Consider Q being the A_{k+2} quiver, which has $k+2$ vertices labeled by $0, \dots, k+1$, and there is exactly one arrow from i to $i+1$ for $i = 0, \dots, k$, and no arrow otherwise. Set $d_{-1} := 0$. Then the minimal total number of critical points is

$$\chi_{Q, \vec{d}} = \prod_{i=-1}^k \binom{d_i + d_{i+1} + 1}{d_{i+1}}$$

and

$$\mathcal{D}_{Q,\vec{d}} = \sum_{i=-1}^k d_{i+1}(d_i + 1).$$

Example 4.24. Now consider the following quiver A'_{k+2} , which has vertices labeled by $0, \dots, k+1$, and there is one arrow from vertex i to vertex j for every $i < j$. Set $d_{-1} := 0$. Then

$$\chi_{Q,\vec{d}} = \prod_{i=-1}^k \binom{\sum_{j=0}^{i+1} d_j + 1}{d_{i+1}}$$

and

$$\mathcal{D}_{Q,\vec{d}} = \sum_{i=-1}^k d_{i+1} \left(\sum_{j=0}^i d_j + 1 \right).$$

Figure 7 shows the graph of $\log \chi_{Q,\vec{d}}$ versus $\mathcal{D}_{Q,\vec{d}}$ for the two examples, where we set $k = 3, d_1 = 600, d_5 = 10$, and $d_2 = d_3 = d_4$.

Note that the quiver denoted by A'_{k+2} in Example 4.24 is a simple analog of the network known as ResNet, which adds arrows to the A_{k+2} -quiver that skip the middle vertices to get around with the ‘gradient-vanishing problem’. Namely, in the A_{k+2} case, the derivatives of \mathcal{E} with respect to matrix entries for arrows in the early stage are typically very small by chain rule, which is not good for the flow rate. Arrows that skip the middle vertices are added, so that there are short paths which involve the early arrows.

From Figure 7, we see that in the same dimensions, the minimal number of critical points in \mathcal{M} is smaller for A'_5 than that for A_5 . (We have numerically verified this for general k .) This gives a supporting evidence that $\chi_{Q,\vec{d}}$ is an important invariant in applications to machine learning.




FIGURE 7. A plot of $\log \chi_{Q,\vec{d}}$ (y-axis) versus $\mathcal{D}_{Q,\vec{d}}$ (x-axis) for A_5 and A'_5 .

4.5 A remark on Abelianization. In many basic neural networks, each vertex of Q is associated with a vector space of only dimension one. When $\vec{d} = \vec{1}$, that is, all entries of the dimension vector equal to one, $\mathcal{M}_{\vec{n},\vec{1}}$ is a quotient by the Abelian group $(\mathbb{C}^\times)^{\Sigma_{\vec{d}}}$, and hence a toric variety. Indeed, by Theorem 2.4, $\mathcal{M}_{\vec{n},\vec{1}}$ is a tower of projective spaces \mathbb{P}^k for a sequence of k .

Given Q and \vec{d} , we can always construct a bigger quiver $Q^{\text{Ab},\vec{d}}$ as follows. For each vertex $i \in Q_0$, we make d_i copies indexed by (i, p) for $p = 1, \dots, d_i$. For each arrow of Q from i to j , we make a corresponding arrow for $Q^{\text{Ab},\vec{d}}$ from (i, p) to (j, q) for every $p = 1, \dots, d_i$ and $q = 1, \dots, d_j$. See Figure 8 for an example.




FIGURE 8. An example of Abelianization. The LHS shows a quiver Q together with the dimension vector \vec{d} . The RHS shows $Q^{\text{Ab},\vec{d}}$.

Given a dimension vector $\vec{n} \in Q_0$, define $\vec{n}^{\text{Ab}} \in Q_0^{\text{Ab},\vec{d}}$ by $\vec{n}^{\text{Ab}}_{(i,p)} = \vec{n}_i$ for all p . The relation between $\mathcal{M}_{\vec{n},\vec{d}}^Q$ and the toric variety $\mathcal{M}_{\vec{n}^{\text{Ab}},\vec{1}}^{Q^{\text{Ab},\vec{d}}}$ is known as Abelianization and is well-studied in [Mar00]. The basic example is $\text{Gr}(n, d)$ (which is the framed moduli for the quiver with a single vertex), whose Abelianization is $(\mathbb{P}^n)^d$ (the disconnected quiver with d vertices and no arrow).

Namely, the moduli spaces $\mathcal{M}_{\vec{n},\vec{d}}^Q$ and $\mathcal{M}_{\vec{n}^{\text{Ab}},\vec{1}}^{Q^{\text{Ab},\vec{d}}}$ are GIT quotients of the same vector space $R_{\vec{n},\vec{d}}^Q = R_{\vec{n}^{\text{Ab}},\vec{1}}^{Q^{\text{Ab},\vec{d}}}$ by $\text{GL}_{\vec{d}}$ and $(\mathbb{C}^\times)^{\Sigma_{\vec{d}}}$ respectively. More precisely, we have the fiber bundle $\mu_{U_{\vec{d}}}^{-1}(\{I\})/T \rightarrow \mathcal{M}_{\vec{n},\vec{d}}^Q$ with fibers being a product of complete flags $U_{\vec{d}}/T^{\Sigma_{\vec{d}}}$, and the inclusion $\mu_{U_{\vec{d}}}^{-1}(\{I\})/T \subset \mathcal{M}_{\vec{n}^{\text{Ab}},\vec{1}}^{Q^{\text{Ab},\vec{d}}}$. The universal bundles \mathcal{V}_i over $\mathcal{M}_{\vec{n},\vec{d}}^Q$ is descended from the direct sum of universal line bundles $\bigoplus_{p=1}^{d_i} \mathcal{V}_{(i,p)}$ over $\mathcal{M}_{\vec{n}^{\text{Ab}},\vec{1}}^{Q^{\text{Ab},\vec{d}}}$ (restricted to the above subset). The cohomology of $\mathcal{M}_{\vec{n},\vec{d}}^Q$ is generated by the Chern classes $c_k(\mathcal{V}_i)$, which can be written as the k -th elementary symmetric polynomials in $c_1(\mathcal{V}_{(i,p)})$ for $p = 1, \dots, d_i$. On the other side, the cohomology of $\mathcal{M}_{\vec{n}^{\text{Ab}},\vec{1}}^{Q^{\text{Ab},\vec{d}}}$ is generated by $c_1(\mathcal{V}_{(i,p)})$.

Note that the functions $f_{\vec{\gamma}}^Q$ and \mathcal{E}^Q over $\mathcal{M}_{\vec{n},\vec{d}}^Q$ constructed in Section 4.3.2 cannot be lifted to $\mathcal{M}_{\vec{n}^{\text{Ab}},\vec{1}}^{Q^{\text{Ab},\vec{d}}}$. The reason is that, $f_{\vec{\gamma}}^Q$ and \mathcal{E}^Q are $\text{GL}_{\vec{d}}$ -equivariant functions on $R_{n,d}^{Q,S}$, the subset of stable representations, rather than the whole $R_{n,d}^Q$. The definition of $f_{\vec{\gamma}}^Q$ involves the metrics on the universal bundles \mathcal{V}_i , which take the expression $(\rho_i \rho_i^*)^{-1}$, and it is only defined over $R_{n,d}^{Q,S}$ where ρ_i is surjective. Rather, we have the functions $f_{\vec{\gamma}}^{Q^{\text{Ab},\vec{d}}}$ and $\mathcal{E}^{Q^{\text{Ab},\vec{d}}}$ on $\mathcal{M}_{\vec{n}^{\text{Ab}},\vec{1}}^{Q^{\text{Ab},\vec{d}}}$, which uses the metrics $(\rho_{i,p} \rho_{i,p}^*)^{-1}$ on the line bundles $\mathcal{V}_{i,p}$.

Previously we have taken $\vec{n} = \vec{d} + \vec{1}$ for $\mathcal{M}_{n,d}^Q$. After Abelianization, the dimension vectors \vec{n}^{Ab} and $\vec{1}$ for $\mathcal{M}_{\vec{n}^{Ab}, \vec{1}}^{Q^{Ab, \vec{d}}}$ no longer satisfy such equality. This is actually not a problem, since the function $f_{\vec{\gamma}}^{Q^{Ab, \vec{d}}}$ defined on $\mathcal{M}_{\vec{2}, \vec{1}}^{Q^{Ab, \vec{d}}}$ can be lifted to $\mathcal{M}_{\vec{n}^{Ab}, \vec{1}}^{Q^{Ab, \vec{d}}}$. Alternatively, we can set the k -th framing vectors to be zero for all $k = 2, \dots, d_i + 1$ and for all vertices $(i, p) \in Q_0^{Ab, \vec{d}}$. This gives a subvariety of $\mathcal{M}_{\vec{n}^{Ab}, \vec{1}}^{Q^{Ab, \vec{d}}}$ which is isomorphic to $\mathcal{M}_{\vec{2}, \vec{1}}^{Q^{Ab, \vec{d}}}$.

5. UNIVERSAL APPROXIMATION THEOREM

In Section 4.1, we have introduced the multi-variable functions σ coming from moment maps of toric varieties. For instance, $\sigma_k(x) = \frac{e^{2x_k}}{1 + \sum_{j=1}^d e^{2x_j}}$ for $X = \mathbb{P}^d$. In this section, we will give a theoretical basis for using this as an activation function, by proving the universal approximation theorem for this function.

The universal approximation theorem ensures that in theory, any given function on a compact set can be approximated (as close as you want) by the functions produced from directed graphs (denoted by $f_{[V,e]}^U$ in Section 4.3.1). There are several different versions of this theorem [Cyb89, LLPS93, LJ18, LPW+17]. To the authors' knowledge, the past works have focused on proving the theorem for single-variable activation functions.

In the work of Cybenko in proving the theorem below, rescaling on the domain of the activation function plays a key role. The rescaling technique will also be very useful in our situation.

Theorem 5.1 ([Cyb89]). *Let $\phi : \mathbb{R} \rightarrow \mathbb{R}$ be any continuous function with $\lim_{x \rightarrow \infty} \phi(x) = 1$ and $\lim_{x \rightarrow -\infty} \phi(x) = 0$. Let K be a compact set in \mathbb{R}^d . Then the collection of functions $G : K \rightarrow \mathbb{R}$ of the form*

$$G(x) = \sum_{j=1}^N \alpha_j \phi(y_j^T x + \theta_j)$$

where $N \in \mathbb{Z}_{>0}$, $y_j \in \mathbb{R}^d$, and $\theta_j, \alpha_j \in \mathbb{R}$, are dense in the space of continuous functions $C(K)$.

The above function G can be understood as $f_{W_1, W_2, b}^U$ (4.9) produced from the graph A_3 , when the dimension at the output vertex is $d_3 = 1$, and $\sigma : \mathbb{R}^{d_2} \rightarrow \mathbb{R}^{d_2}$ is taken to be $\sigma(\vec{x}) = (\phi(x_1), \dots, \phi(x_{d_2}))$. (Take $d_2 = N$, $W_1 = (y_j^T)_{j=1, \dots, N}$, $b = (\theta_j)_{j=1, \dots, N}$, and $W_2 = (\alpha_j)_{j=1, \dots, N}$.) For general dimension d_3 , we simply have (G_1, \dots, G_{d_3}) , where G_i are of the same form as above (with different $\alpha_{j,i}$) which can be used to approximate any given continuous function $K \rightarrow \mathbb{R}^{d_3}$.

We will prove the following theorem. Consider the quiver with three vertices as in Section 4.3.3, and the function $f_{W_1, W_2, b}^U(v) = W_2(\sigma(W_1 \cdot v + b))$ in (4.9), where σ is the multi-variable activation function on \mathbb{R}^{d_2} made from \mathbb{P}^{d_2} .

Theorem 5.2. *Let K be a compact set of \mathbb{R}^{d_1} , and $f : K \rightarrow \mathbb{R}^{d_3}$ a continuous function. For any $\epsilon > 0$, there exists $d_2 > 0$ and $W_1 \in \text{Mat}(d_2, d_1)$, $W_2 \in \text{Mat}(d_3, d_2)$, $b \in \mathbb{R}^{d_2}$ such that $\|f_{W_1, W_2, b}^U - f\|_{L^2(K)} < \epsilon$.*

The compact set is given as a subset in \mathbb{R}^n . Thus from now on we restrict to the real field, which will suffice for the theorem. This means we take real-valued matrices and the real part σ of $\sigma_{\mathbb{C}}$.

5.1 Tropical limit. A crucial idea in the work of [Cyb89] is to compose ϕ with a rescaling, so that it tends to a step function in the limit. We can apply such a rescaling to the multi-variable function $\sigma : \mathbb{R}^d \rightarrow P^\circ$. This is well-known in toric geometry and is called the tropical limit.

Let Σ be the dual fan of the moment polytope P . Σ is the collection of cones that are dual to the boundary strata of the polytope P . In particular, maximal cones of Σ are one-to-one corresponding to corners of P .

We assume that $|\Sigma| = \mathbb{R}^d$. \mathbb{R}^d is stratified into the relative interiors of cones in Σ . We recall the following interesting fact from toric geometry. It plays an important role in the study of holomorphic discs and Lagrangian Floer theory for toric varieties.

Lemma 5.3. *Let X_Σ be a toric variety (equipped with any toric Kähler form). Let C be a cone in Σ . For any $x \neq 0$ which lies in the relative interior of C , $p_C = \lim_{t \rightarrow -\infty} \sigma(tx)$ exists and equal to a point (which is independent of x) in the boundary stratum of P that is dual to C .*

In other words, the family of functions $\sigma_t(x) = \sigma(tx)$ converges (as $t \rightarrow +\infty$) to the discontinuous function σ_∞ , where $\sigma_\infty(x) = p_C$ if x belongs to the relative interior of C .

Proof. The cone C corresponds to a complex torus orbit of the toric variety X . To be more explicit, consider a maximal cone C^{\max} that contains C . Without loss of generality, let $C^{\max} = \mathbb{R}_{\geq 0} \cdot \{v_1, \dots, v_d\}$, and $C = \mathbb{R}_{\geq 0} \cdot \{v_1, \dots, v_k\}$. C^{\max} gives a local chart \mathbb{C}^d of the toric variety, and the complex torus orbit corresponding to C is given by $z_1 = \dots = z_k = 0$.

We have a special point given by $z_1 = \dots = z_k = 0, z_{k+1} = \dots = z_d = 1$ in the orbit. We assert that $p_C = \lim_{t \rightarrow -\infty} \sigma(tx) \in P$ (for any x in the relative interior of C) is the moment-map image of this point.

To see this, we write $x = \sum_{i=1}^k x_i v_i$ where $x_i \neq 0$ for all $i = 1, \dots, k$. Consider the lifting of $tx = (tx_1, \dots, tx_k, 0, \dots, 0) : (e^{tx_1}, \dots, e^{tx_k}, 1, \dots, 1)$ in the chart $\mathbb{C}^d \subset X$. Then $\sigma(tx)$ is the moment-map image of $(e^{tx_1}, \dots, e^{tx_k}, 1, \dots, 1)$. Taking $t \rightarrow -\infty$, $(e^{tx_1}, \dots, e^{tx_k}, 1, \dots, 1) \rightarrow (z_1 = \dots = z_k = 0, z_{k+1} = \dots = z_d = 1)$. Thus $\sigma(tx)$ converges to the above special point p_C . \square

In terms of solving equations, p_C is the solution of the simultaneous equations $x_1 = \dots = x_k = 0$ and $x_i \prod_{j=d+1}^m \ell_j^{v_{ji}}(0, \dots, 0, x_{k+1}, \dots, x_d) = 1$ for $i = k+1, \dots, m$. (We have used the dual basis of $\{v_1, \dots, v_d\}$ to write the coordinates of P , and $v_j = \sum_{i=1}^d v_{ji} v_i$.) By above, the solution exists and is unique.

Example 5.4. Consider $X = \mathbb{P}^d$. Denote the coordinates of \mathbb{R}^d by (x_1, \dots, x_d) , and set $x_0 \equiv 0$. The $(d+1-l)$ -cones C of Σ are given by $\{x_{i_1} = \dots = x_{i_l} > x_k \text{ for all } k \in \{0, \dots, d\} - \{i_1, \dots, i_l\}\}$,

where $i_1, \dots, i_l \in \{0, \dots, d\}$ are fixed, $l = 1, \dots, d+1$. For $\sigma = \left(\frac{e^{2x_p}}{1 + \sum_{j=1}^d e^{2x_j}} \right)_{p=1}^d$, the point

$p_C = \lim_{t \rightarrow -\infty} \sigma(tx)$ has coordinates $(p_C)_{i_r} = 1/l$ for $r = 1, \dots, l$ and $i_r \neq 0$, and $(p_C)_j = 0$ for all other $j \neq i_1, \dots, i_l, 0$. In particular, for the maximal cones $S_i = \{x_i > x_k \text{ for all } k \in \{0, \dots, d\} - \{i\}\}$, $p_{S_i} = \epsilon_i$ for $i = 0, \dots, d$ where $\epsilon_0 := 0$ and $\{\epsilon_1, \dots, \epsilon_d\}$ denotes the standard basis.

Corollary 5.5. *Let K be any compact set in \mathbb{R}^d . For any $\epsilon > 0$ and an open neighborhood of the union U of codimension-one strata of Σ , there exists $t \gg 0$ such that $|\sigma_t(x) - \sigma_\infty(x)| < \epsilon$ for all $x \in K - U$. (σ_∞ is defined in Lemma 5.3.)*




FIGURE 9. The left shows the fan picture of \mathbb{P}^2 , and the right shows the moment-map polytope. The dots show the limit points p_C for each cone C of the fan.

Proof. Any $x \in K - U$ belongs to one of the maximal cones C . By Lemma 5.3, $\sigma_t(x)$ converges to p_C . Moreover, both σ_t and σ_∞ are continuous on $K - U$. Then the result follows from the compactness of $K - U$. \square

In order to prove Theorem 5.2, we consider a particular type of polyhedral decompositions of \mathbb{R}^n , which we call to be a centered simplicial web.

5.2 Centered polyhedral web.

Definition 5.6. A centered simplicial web with N ordered compact chambers in \mathbb{R}^n is a polyhedral decomposition of \mathbb{R}^n whose vertices are all trivalent, defined inductively on the number of compact chambers as follows.

A centered simplicial web with zero compact chamber is the polyhedral decomposition given by the fan of \mathbb{P}^n , up to an affine linear isomorphism in $GL(n, \mathbb{R}) \times \mathbb{R}^n$.

Now suppose the notion of a centered simplicial web with N ordered compact chambers has been defined, which has exactly $(n + 1)$ non-compact rays (which we call the outer rays), whose corresponding infinite lines intersect at exactly one point called the N -th center that lies in the union of the N compact chambers. Moreover, the web is required to have $(n + 1)$ non-compact chambers; each non-compact chamber is adjacent to n outer rays and opposite to the remaining one outer ray. ('Opposite' here means that the non-compact chamber is disjoint from the corresponding outer ray, whose infinite line intersects the chamber at a half-line.) The outer rays are one-to-one corresponding to their opposite non-compact chambers.

A centered simplicial web with $(N + 1)$ ordered compact chambers is defined as follows. First, take a centered simplicial web with N ordered compact chambers. Second, we choose a non-compact chamber, and denotes the direction of its opposite ray by a non-zero vector v . Third, we take an affine hyperplane which intersects all the relative interior of the n adjacent rays of the non-compact chamber. This bounds a new compact chamber and the intersection points V_i are the new vertices. Finally, we choose the $(N + 1)$ -th center to be $c_{N+1} = c_N - tv$, where c_N is the N -th center, and $t \in \mathbb{R}_{>0}$ is taken such that c_{N+1} lies in the union of the compact chambers (including the new one). Then we have n new rays emanated from the vertices V_i whose infinite lines pass through c_{N+1} . This gives a new web with $(N + 1)$ ordered compact chambers, and it still has $(n + 1)$ non-compact chambers, each of which is adjacent to n outer rays and opposite to one outer ray.

See Figure 10 for some examples of centered simplicial webs in \mathbb{R}^2 .




FIGURE 10. Examples of centered simplicial web in \mathbb{R}^2 . They have zero, one, and two compact chambers respectively.

From the above definition, there is a one-to-one correspondence between the centers and compact chambers. Moreover, the $(k + 1)$ -th compact chamber C_{k+1} (for $k = 1, \dots, N - 1$) is associated with a one-strata of the web, which is a subset of the opposite ray of the non-compact chamber containing C_{k+1} in the $(k + 1)$ -th inductive step. Furthermore, both the k -th and $(k + 1)$ -th centers lie in the infinite line of the associated 1-strata of C_{k+1} .

Remark 5.7. *The above notion is closely related to tropical subvarieties. In the tropical context, there is an integral structure on the ambient space and the balancing condition (whose definition requires the integral structure) is imposed on each vertex of a tropical variety. However, we do not have an integral structure here, since σ_C is defined on the universal cover \mathbb{C}^d rather than $(\mathbb{C}^\times)^d$ (see Corollary 4.2 and Example 4.4). It means affine linear maps are taken over \mathbb{R} rather than over \mathbb{Z} . Instead of the balancing condition, we impose the notion of centers in the above definition.*

Theorem 5.8. *Given a centered simplicial web A with N ordered compact chambers in \mathbb{R}^n , there exists an affine-linear embedding $L : \mathbb{R}^n \rightarrow \mathbb{R}^d$, where $d = n + N$, such that the L -preimage of the fan of \mathbb{P}^d in \mathbb{R}^d equals to the web A .*

Proof. We shall prove the following statement: given a centered simplicial web A with $N \geq 1$ compact chambers in \mathbb{R}^n , there exists a centered simplicial web B with $N - 1$ compact chambers in \mathbb{R}^{n+1} such that the intersection of $\mathbb{R}^n \cong \mathbb{R}^n \times \{0\}$ with B equals to A . Then by applying this statement N times, we obtain a web $B^{(N)}$ with zero compact chamber in \mathbb{R}^{n+N} whose intersection with $\mathbb{R}^n \times \{0\}$ gives A . By an affine linear isomorphism on \mathbb{R}^{n+N} , $B^{(N)}$ is identified with the fan of \mathbb{P}^{n+N} . The required map L is given by the composition of the inclusion $\mathbb{R}^n \times \{0\} \subset \mathbb{R}^{n+N}$ with this linear isomorphism.

The above statement is proved by induction on N . First consider the case $N = 1$. We take a point V away from $\mathbb{R}^n \times \{0\} \subset \mathbb{R}^{n+1}$. Then we take a cone at V over the compact simplicial chamber of A . Moreover, the line joining V with the given center of A intersects with the complement of the cone and produces a ray emanated from V . This gives a simplicial web in \mathbb{R}^{n+1} with no compact chamber, whose intersection with $\mathbb{R}^n \times \{0\}$ is exactly A . (See Figure 11.)

Now suppose it is true for N . Consider a centered simplicial web A in $\mathbb{R}^n \cong \mathbb{R}^n \times \{0\}$ with $N + 1$ ordered compact chambers. We can take away the $(N + 1)$ -th compact chamber $C = C_{N+1}$ (and forget the corresponding center c_{N+1}) and obtain a centered simplicial web A' with N compact chambers. By inductive hypothesis, there exists a

centered simplicial web B' in \mathbb{R}^{n+1} with $N - 1$ compact chambers whose intersection with $\mathbb{R}^n \times \{0\}$ gives A' .

The compact chamber C of A is contained in a non-compact chamber C' of A' , which is the intersection of $\mathbb{R}^n \times \{0\}$ with a non-compact chamber D' of B' . C' is opposite to an outer ray R of A' , which is the intersection of $\mathbb{R}^n \times \{0\}$ with a non-compact 2-plane P of B' . Note that the last two centers c_N^A and c_{N+1}^A of A are contained in the line of R , and hence contained in the infinite 2-plane of P . Consider the two rays of B that are adjacent to P . Denote the one which is opposite to the chamber D' by L . The other one is denoted by L' , which must be adjacent to D' .

Now we construct a web B in \mathbb{R}^{n+1} whose intersection with $\mathbb{R}^n \times \{0\}$ gives A . A point V in the relative interior of the ray L' is taken to be a new vertex. Consider the line passing through V and the last center c_{N+1}^A . This line lies in the infinite 2-plane of P . Thus for a generic choice of V , it must intersect with the infinite line of L at a point, which we shall define as the new center c_N^B for B . V is taken far away enough in the ray L' so that the intersection point c_N^B equals to $c_{N-1}^B - v$ for some vector v in the direction of L .

Consider the n outer rays of A' that are adjacent to the chamber C' . They are the intersections of $\mathbb{R}^n \times \{0\}$ with the corresponding n outer-2-planes of B' that are adjacent to L' . The last chamber C of A is formed by the hyperplane through the vertices taken in the relative interior of the n outer rays of A' . The lines joining V to these vertices in A lie in the outer-2-planes of B' , and hence intersect with the corresponding outer rays of B' at certain points, which we take to be new vertices of B . The hyperplane through V and these new vertices bounds a new chamber. The lines joining c_N^B with the new vertices produce the outer rays of the new web B . This gives B whose intersection with $\mathbb{R}^n \times \{0\}$ equals to A . \square




FIGURE 11. Construction of a tropical web with zero compact chamber whose intersection with a hyperplane equals to a given tropical web with one compact chamber.

The inductive step in the above proof is illustrated by Figure 12.

Next, we consider polytopes rather than simplices and the corresponding webs formed from polytopes. Motivated from the well-known fact below, we define a centered polyhedral web to be the intersection of a centered simplicial web with an affine subspace.

Proposition 5.9. *For a polytope P with m facets in \mathbb{R}^n where $m > n + 1$, there exists a simplex S in \mathbb{R}^{m-1} such that $S \cap (\mathbb{R}^n \times \{0\}) = P$ (where \mathbb{R}^n is identified with $\mathbb{R}^n \times \{0\}$).*




FIGURE 12. Construction of a tropical web with 3 compact chambers whose intersection with a hyperplane equals to a given tropical web with 2 compact chambers.

The simplex S in Proposition 5.9 can be constructed as follows. Without loss of generality suppose $0 \in P$. Consider the dual polytope $P^\vee = \{v \in (\mathbb{R}^n)^* : (v, v) \leq 1 \text{ for all } v \in P\}$, which is the convex hull of its vertices v_i for $i = 1, \dots, m$. Then we have a surjective map from the standard simplex $S^\vee = \{\sum_i a_i \epsilon_i^* \in (\mathbb{R}_{\geq 0}^m)^* : \sum_i a_i = 1\}$ (where $\{\epsilon_i^* : i = 1, \dots, m\}$ is the standard basis) to P^\vee by sending ϵ_i^* to v_i . S^\vee in the affine subspace $\{\sum_i a_i = 1\}$ can be identified as a simplex in $(\mathbb{R}_{\geq 0}^{m-1})^*$ by the projection along the direction $-\sum_i \epsilon_i$. Then the dual linear map gives the desired linear injection $\mathbb{R}^n \rightarrow \mathbb{R}^{m-1}$ which sends P into S . By composing with a linear isomorphism, the image of \mathbb{R}^n can be made to be $\mathbb{R}^n \times \{0\}$.

As a result, a centered polygonal web with one compact chamber (which is constructed by taking a polygon with a chosen center c and outer rays at vertices whose lines pass through c) can be obtained as an intersection with $\mathbb{R}^n \times \{0\}$ of a centered simplicial web with one compact chamber in \mathbb{R}^{m-1} (where m is the number of non-compact chambers). See the left of Figure 13 for an example.




FIGURE 13. The left shows an example of a polyhedral web with one compact chamber, which is given as an intersection of a simplicial web with a subspace. The right two figures show a concentric simplicial and a polygonal web.

The following degenerate configuration will be helpful. In Definition 5.6, suppose we take all the centers to be the same. Moreover, suppose the new hyperplane introduced to bound a new chamber is allowed to intersect the outer rays at the original vertices (rather than their relative interior). Then we can construct the following configuration.

Definition 5.10. *Let's take $(n + 1)$ rays emanated from $0 \in \mathbb{R}^n$, such that any n of them are linearly independent. For each ray, we take a sequence of distinct points $V_{i,k}$ (where $i = 1, \dots, n + 1$ is indexing the ray) such that $V_{i,k+1} - V_{i,k}$ is pointing in the ray direction. Then for each $k > 0$, we take a simplex with vertices at $V_{i,k}$ for $i = 1, \dots, n + 1$. This gives a polyhedral decomposition of \mathbb{R}^n . This is called a concentric simplicial web.*

By taking an intersection of the above degenerate simplicial configuration with a subspace (that passes through the center 0), we get a configuration made from a sequence of polytopes whose vertices lie in a fixed collection of p rays, where p is the number of vertices of each polytope. We call this a concentric polyhedral web. See the right of Figure 13.

5.3 Proof of the approximation theorem. We are now ready to prove Theorem 5.2.

Proof of Theorem 5.2. For any $\delta > 0$, we can take a concentric polyhedral web B in \mathbb{R}^{d_1} , such that for every chamber C of B , $C \cap K$ is contained in a δ -ball. B is constructed as follows. Without loss of generality, let $K \ni 0$. First, we take a polytope P that lies in a δ -ball centered at $0 \in \mathbb{R}^{d_1}$. P induces a subdivision on the unit sphere S^{d_1-1} by projecting its boundary strata onto S^{d_1-1} . P is taken with sufficiently many vertices such that the induced subdivision on the unit sphere S^{d_1-1} lies in a δ' -ball for a chosen δ' . By Proposition 5.9, $P = S \cap (\mathbb{R}^{d_1} \times \{0\})$ for some simplex S in \mathbb{R}^{m-1} where m is the number of facets of P . We take $0 \in S$ to be the center. Then we take rays from 0 through the vertices of S and construct a concentric simplicial web. Since K is compact, by taking δ' sufficiently small, and the sequences of vertices in the rays sufficiently close to each other, the resulting concentric polyhedral web B can be made such that every chamber intersects K in a δ -ball.

Next, we take a centered simplicial web A in \mathbb{R}^{m-1} whose centers are chosen sufficiently close to each other, and the vertices in the inductive steps are taken such that $A \cap K$ is sufficiently close $B \cap K$. Namely, for every chamber C of A , $C \cap K$ lies in the δ -neighborhood of $C' \cap K$ for the corresponding chamber C' of B . In particular, $C \cap K$ is contained in a 2δ -ball.

By Theorem 5.8, there exists $L' : \mathbb{R}^{m-1} \rightarrow \mathbb{R}^{d_2}$ such that A is the L' -preimage of the fan $\Sigma_{\mathbb{P}^{d_2}}$. By composing L' with $\mathbb{R}^{d_1} \times \{0\} \subset \mathbb{R}^{m-1}$, we obtain $L : \mathbb{R}^{d_1} \rightarrow \mathbb{R}^{d_2}$ such that for every maximal cone S_i of $\Sigma_{\mathbb{P}^{d_2}}$, $L^{-1}(S_i) \cap K$ lies in a (2δ) -ball.

Since f is uniformly continuous in K , for every $\epsilon > 0$, δ can be taken such that $|f(x) - f(y)| < \epsilon$ for every $x, y \in K$ lying in a 2δ -ball. In particular, we have a step function $s = \sum_C r_C \delta_C$ supported over A (where $\delta_C(x) = 1$ for $x \in C$ and 0 otherwise, and C are chambers of A) such that $\|f - s\|_{L^2(K)} < \epsilon \sqrt{\text{Vol}(K)}$.

We have the step function σ_∞ which sends the interior of the maximal cones S_i of $\Sigma_{\mathbb{P}^{d_2}}$ to e_i for $i = 0, \dots, d_2$, where $e_0 = 0$. (See Example 5.4.) The cone S_i corresponds to chambers $L^{-1}(S_i)$ of A under L . Since $\{e_i - e_0 : i = 1, \dots, d_2\}$ forms a basis, there exists a unique affine linear map $W_2 : \mathbb{R}^{d_2} \rightarrow \mathbb{R}^{d_3}$ which sends e_i to $r_{L^{-1}(S_i)}$ for all $i = 0, \dots, d_2$. Thus $s = W_2 \circ \sigma_\infty \circ L$.

Finally, by Corollary 5.5, there exists $t \gg 0$ such that $|W_2 \circ \sigma_t - W_2 \circ \sigma_\infty| < \epsilon/2$ on $L(K) - U$, where U is an arbitrary open neighborhood of the codimension-one strata

of $\Sigma_{\mathbb{P}_2^d}$. Moreover, $|W_2 \circ \sigma_t - W_2 \circ \sigma_\infty|$ is bounded. Hence by taking $\text{Vol}(L^{-1}(U) \cap K)$ sufficiently small, we have $\|W_2 \circ \sigma_t \circ L - W_2 \circ \sigma_\infty \circ L\|_{L^2(K)} < \epsilon$. In conclusion, $\|W_2 \circ \sigma \circ (tL) - f\|_{L^2(K)} \leq \|W_2 \circ \sigma \circ (tL) - W_2 \circ \sigma_\infty \circ L\|_{L^2(K)} + \|W_2 \circ \sigma_\infty \circ L - f\|_{L^2(K)}$ can be made arbitrarily small. \square

REFERENCES

- [Abo06] M. Abouzaid, *Homogeneous coordinate rings and mirror symmetry for toric varieties*, *Geom. Topol.* **10** (2006), 1097–1157.
- [Abr98] M. Abreu, *Kähler geometry of toric varieties and extremal metrics*, *Internat. J. Math.* **9** (1998), no. 6, 641–651.
- [AJ20] M. A. Armenta and P.-M. Jodoin, *The representation theory of neural networks*, preprint (2020), [arXiv:2007.12213](https://arxiv.org/abs/2007.12213).
- [CAW⁺19] M.C.N. Cheng, V. Anagiannis, M. Weiler, P. de Haan, T.S. Cohen, and M. Welling, *Covariance in physics and convolutional neural networks*, preprint (2019), [arXiv:1906.02481](https://arxiv.org/abs/1906.02481).
- [CB03] W. Crawley-Boevey, *Normality of Marsden-Weinstein reductions for representations of quivers*, *Math. Ann.* **325** (2003), no. 1, 55–79. MR 1957264
- [CDS15a] X. Chen, S. Donaldson, and S. Sun, *Kähler-Einstein metrics on Fano manifolds. I: Approximation of metrics with cone singularities*, *J. Amer. Math. Soc.* **28** (2015), no. 1, 183–197.
- [CDS15b] ———, *Kähler-Einstein metrics on Fano manifolds. II: Limits with cone angle less than 2π* , *J. Amer. Math. Soc.* **28** (2015), no. 1, 199–234.
- [CDS15c] ———, *Kähler-Einstein metrics on Fano manifolds. III: Limits as cone angle approaches 2π and completion of the main proof*, *J. Amer. Math. Soc.* **28** (2015), no. 1, 235–278.
- [CGKW18] T.S. Cohen, M. Geiger, J. Koehler, and M. Welling, *Spherical cnns*, ICLR (2018).
- [CGW19] T.S. Cohen, M. Geiger, and M. Weiler, *A general theory of equivariant cnns on homogeneous spaces*, NeurIPS (2019), [arXiv:1811.02017](https://arxiv.org/abs/1811.02017).
- [CLL12] K. Chan, S.-C. Lau, and N.C. Leung, *SYZ mirror symmetry for toric Calabi-Yau manifolds*, *J. Differential Geom.* **90** (2012), no. 2, 177–250.
- [CSW18] X. Chen, S. Sun, and B. Wang, *Kähler-Ricci flow, Kähler-Einstein metric, and K-stability*, *Geom. Topol.* **22** (2018), no. 6, 3145–3173.
- [CW16] T.S. Cohen and M. Welling, *Group equivariant convolutional networks*, Proceedings of The 33rd International Conference on Machine Learning, vol. 48, 2016, pp. 2990–2999.
- [CWKW19] T.S. Cohen, M. Weiler, B. Kicanaoglu, and M. Welling, *Gauge equivariant convolutional networks and the icosahedral cnn*, Proceedings of the International Conference on Machine Learning (ICML), 2019.
- [Cyb89] G. Cybenko, *Approximation by superpositions of a sigmoidal function*, *Math. Control Signals Systems* **2** (1989), no. 4, 303–314.
- [CYRL19] I. Chami, Z. Ying, C. Ré, and J. Leskovec, *Hyperbolic graph convolutional neural networks*, Advances in neural information processing systems (2019).
- [dHCW20] P. de Haan, T. Cohen, and M. Welling, *Natural graph networks*, preprint (2020), [arXiv:2007.08349](https://arxiv.org/abs/2007.08349).
- [Don85] S.K. Donaldson, *Anti self-dual Yang-Mills connections over complex algebraic surfaces and stable vector bundles*, *Proc. London Math. Soc.* (3) **50** (1985), no. 1, 1–26.
- [Don99] ———, *Symmetric spaces, Kähler geometry and Hamiltonian dynamics*, Northern California Symplectic Geometry Seminar, Amer. Math. Soc. Transl. Ser. 2, vol. 196, Amer. Math. Soc., Providence, RI, 1999, pp. 13–33.
- [Don12] ———, *Stability, birational transformations and the Kahler-Einstein problem*, Surveys in differential geometry. Vol. XVII, *Surv. Differ. Geom.*, vol. 17, Int. Press, Boston, MA, 2012, pp. 203–228.
- [DPG⁺14] Y. Dauphin, R. Pascanu, C. Gulcehre, K. Cho, S. Ganguli, and Y. Bengio, *Identifying and attacking the saddle point problem in high-dimensional non-convex optimization*, preprint (2014), [arXiv:1406.2572](https://arxiv.org/abs/1406.2572).
- [ER09] J. Engel and M. Reineke, *Smooth models of quiver moduli*, *Math. Z.* **262** (2009), no. 4, 817–848.
- [Fed13] S. Fedotov, *Framed moduli and Grassmannians of submodules*, *Trans. Amer. Math. Soc.* **365** (2013), no. 8, 4153–4179.
- [FHS95] A. Floer, H. Hofer, and D. Salamon, *Transversality in elliptic Morse theory for the symplectic action*, *Duke Math. J.* **80** (1995), no. 1, 251–292.
- [Flo89] A. Floer, *Witten’s complex and infinite-dimensional Morse theory*, *J. Differential Geom.* **30** (1989), no. 1, 207–221.

- [FLTZ12] B. Fang, C.-C. M. Liu, D. Treumann, and E. Zaslow, *T-duality and homological mirror symmetry for toric varieties*, *Adv. Math.* **229** (2012), no. 3, 1875–1911.
- [FO99] K. Fukaya and K. Ono, *Arnold conjecture and Gromov-Witten invariant*, *Topology* **38** (1999), no. 5, 933–1048.
- [Ful93] W. Fulton, *Introduction to toric varieties*, *Annals of Mathematics Studies*, vol. 131, Princeton University Press, Princeton, NJ, 1993, The William H. Roever Lectures in Geometry. MR 1234037 (94g:14028)
- [GBH18] O. Ganea, G. Becigneul, and T. Hofmann, *Hyperbolic neural networks*, *Advances in Neural Information Processing Systems* (S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, eds.), vol. 31, Curran Associates, Inc., 2018, pp. 5345–5355.
- [GGL19] M. Goyal, R. Goyal, and B. Lall, *Learning activation functions: A new paradigm for understanding neural networks*, preprint (2019), [arXiv:1906.09529](https://arxiv.org/abs/1906.09529).
- [Gui94] V. Guillemin, *Kaehler structures on toric varieties*, *J. Differential Geom.* **40** (1994), no. 2, 285–309.
- [Ham82] R.S. Hamilton, *Three-manifolds with positive Ricci curvature*, *J. Differential Geometry* **17** (1982), no. 2, 255–306.
- [HKL18] H. Hansol, Y. Kim, and S.-C. Lau, *Immersed two-spheres and SYZ with application to Grassmannians*, [arXiv:1805.11738](https://arxiv.org/abs/1805.11738) (2018).
- [HSTT18a] K. Hashimoto, S. Sugishita, A. Tanaka, and A. Tomiya, *Deep learning and holographic QCD*, *Phys. Rev. D* **98** (2018), no. 10, 106014, 15.
- [HSTT18b] ———, *Deep learning and the ads/cft correspondence*, *Physical Review D* **98** (2018), no. 4.
- [HY20] Y.-H. He and S.-T. Yau, *Graph laplacians, riemannian manifolds and their machine-learning*, preprint (2020), [arXiv:2006.16619](https://arxiv.org/abs/2006.16619).
- [Kin94] A.D. King, *Moduli of representations of finite-dimensional algebras*, *Quart. J. Math. Oxford Ser. (2)* **45** (1994), no. 180, 515–530.
- [LJ18] H. Lin and S. Jegelka, *Resnet with one-neuron hidden layers is a universal approximator*, *Advances in Neural Information Processing Systems* (S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, eds.), vol. 31, Curran Associates, Inc., 2018, pp. 6169–6178.
- [LLPS93] M. Leshno, V.Ya. Lin, A. Pinkus, and S. Schocken, *Multilayer feedforward networks with a non-polynomial activation function can approximate any function*, *Neural Networks* **6** (1993), no. 6, 861 – 867.
- [LPW+17] Z. Lu, H. Pu, F. Wang, Z. Hu, and L. Wang, *The expressive power of neural networks: A view from the width*, preprint (2017), [arXiv:1709.02540](https://arxiv.org/abs/1709.02540).
- [Mar00] S. Martin, *Symplectic quotients by a nonabelian group and by its maximal torus*, preprint (2000), [arXiv:math/0001002](https://arxiv.org/abs/math/0001002).
- [MM92] H.N. Mhaskar and C.A. Micchelli, *Approximation by superposition of sigmoidal and radial basis functions*, *Adv. in Appl. Math.* **13** (1992), no. 3, 350–373.
- [Nak94] H. Nakajima, *Instantons on ALE spaces, quiver varieties, and Kac-Moody algebras*, *Duke Math. J.* **76** (1994), no. 2, 365–416.
- [Nak96] ———, *Varieties associated with quivers*, *Representation theory of algebras and related topics* (Mexico City, 1994), *CMS Conf. Proc.*, vol. 19, Amer. Math. Soc., Providence, RI, 1996, pp. 139–157.
- [Nak01] ———, *Quiver varieties and finite-dimensional representations of quantum affine algebras*, *J. Amer. Math. Soc.* **14** (2001), no. 1, 145–238.
- [NNU10] T. Nishinou, Y. Nohara, and K. Ueda, *Toric degenerations of Gelfand-Cetlin systems and potential functions*, *Adv. Math.* **224** (2010), no. 2, 648–706.
- [Ono95] K. Ono, *On the Arnold conjecture for weakly monotone symplectic manifolds*, *Invent. Math.* **119** (1995), no. 3, 519–537.
- [PDGB14] R. Pascanu, Y.N. Dauphin, S. Ganguli, and Y. Bengio, *On the saddle point problem for non-convex optimization*, preprint (2014), [arXiv:1405.4604](https://arxiv.org/abs/1405.4604).
- [Per02] G. Perelman, *The entropy formula for the ricci flow and its geometric applications*, preprint (2002), [arXiv:math/0211159](https://arxiv.org/abs/math/0211159).
- [Per03] ———, *Ricci flow with surgery on three-manifolds*, preprint (2003), [arXiv:math/0303109](https://arxiv.org/abs/math/0303109).
- [Pet99] P.P. Petrushev, *Approximation by ridge functions and neural networks*, *SIAM J. Math. Anal.* **30** (1999), no. 1, 155–189.
- [Pin99] A. Pinkus, *Approximation theory of the MLP model in neural networks*, *Acta numerica*, 1999, *Acta Numer.*, vol. 8, Cambridge Univ. Press, Cambridge, 1999, pp. 143–195.

- [Rei08] M. Reineke, *Framed quiver moduli, cohomology, and quantum groups*, J. Algebra **320** (2008), no. 1, 94–115.
- [Sem92] S. Semmes, *Complex Monge-Ampère and symplectic manifolds*, Amer. J. Math. **114** (1992), no. 3, 495–550.
- [Tia97] G. Tian, *Kähler-Einstein metrics with positive scalar curvature*, Invent. Math. **130** (1997), no. 1, 1–37.
- [UY86] K. Uhlenbeck and S.-T. Yau, *On the existence of Hermitian-Yang-Mills connections in stable vector bundles*, Comm. Pure Appl. Math. **39** (1986), no. S, suppl., S257–S293, Frontiers of the mathematical sciences: 1985 (New York, 1985).
- [Yau96] S.-T. Yau, *Review of kähler-einstein metrics in algebraic geometry*, Israel Math. Conference Proceedings, Bar Ilan Univ., 1996, pp. 433–443.

DEPARTMENT OF MATHEMATICS AND STATISTICS, BOSTON UNIVERSITY, 111 CUMMINGTON MALL, BOSTON MA 02215, USA

Email address: georgej@bu.edu

DEPARTMENT OF MATHEMATICS AND STATISTICS, BOSTON UNIVERSITY, 111 CUMMINGTON MALL, BOSTON MA 02215, USA

Email address: lau@math.bu.edu