

A UNIFIED THEORY OF QUANTUM NEURAL NETWORK LOSS LANDSCAPES

Anonymous authors

Paper under double-blind review

ABSTRACT

Classical neural networks with random initialization famously behave as Gaussian processes in the limit of many neurons, which allows one to completely characterize their training and generalization behavior. No such general understanding exists for quantum neural networks (QNNs), which—outside of certain special cases—are known to not behave as Gaussian processes when randomly initialized. We here prove that QNNs and their first two derivatives instead generally form what we call *Wishart processes*, where certain algebraic properties of the network determine the hyperparameters of the process. This Wishart process description allows us to, for the first time: give necessary and sufficient conditions for a QNN architecture to have a Gaussian process limit; calculate the full gradient distribution, generalizing previously known barren plateau results; and calculate the local minima distribution of algebraically constrained QNNs. Our unified framework suggests a certain simple operational definition for the “trainability” of a given QNN model using a newly introduced, experimentally accessible quantity we call the *degrees of freedom* of the network architecture.

1 INTRODUCTION

1.1 MOTIVATION

One of the miracles of machine learning on classical computers is that simple, gradient-based optimizers can efficiently find the minimum of extremely high-dimensional, nonconvex loss landscapes, allowing for the efficient training of deep neural networks. Over the past decade this has been understood in more and more detail via random matrix theory. In particular, it is now known that the loss landscapes of randomly initialized, wide neural networks are distributed as Gaussian processes with covariance given by the so-called *neural tangent kernel* (NTK) (Neal, 1996; Choromanska et al., 2015; Chaudhari, 2018; Lee et al., 2018). The NTK is completely determined by the neural structure of the network, linking the asymptotic behavior of the network to architectural choices made in its construction. This understanding of classical neural networks has been used to show that wide neural networks train efficiently via gradient descent (Choromanska et al., 2015; Chaudhari, 2018; Jacot et al., 2018; Allen-Zhu et al., 2019). It also explains other emergent phenomena in deep learning, such as the remarkably good generalization performance of neural networks beyond what learning theory predicts (Jacot et al., 2018; Wei et al., 2022).

One might hope that a similar, universal story would exist for *quantum neural networks* (QNNs). These are classes of neural networks where the associated loss function ℓ takes as input a *quantum state* $\rho \succeq \mathbf{0} \in \mathbb{C}^{N \times N}$ with trace 1, conjugates ρ by a unitary operation $U_\theta \in U(N)$ parameterized by $\theta \in \mathbb{R}^p$, and then takes the Frobenius inner product with the Hermitian $O \in \mathbb{C}^{N \times N}$. That is,

$$\ell(\rho; \theta) = \text{Tr} \left(\rho U_\theta^\dagger O U_\theta \right). \quad (1)$$

Such networks are defined by the choice of parameterization for U_θ (typically called the *ansatz*) and the Hermitian O (here called the *objective observable*) (Peruzzo et al., 2014), and are known to efficiently (on a quantum computer) perform learning tasks provably difficult using traditional machine learning methods (Liu et al., 2021; Hastings & O’Donnell, 2022; Anschuetz et al., 2023b; Huang et al., 2024; Anschuetz & Gao, 2024). Unfortunately, it is known that QNNs generally do not have a Gaussian process (or *quantum neural tangent kernel* (QNTK)) asymptotic limit (García-Martín et al., 2023; Girardi & Palma, 2024). Indeed, it is not even obvious whether an equivalently

054 simple universal description exists due to QNN training dynamics differing completely in various
 055 parameter regimes. For instance, the loss landscapes of generic shallow QNNs are known to be
 056 described by so-called *Wishart hypertoroidal random fields* (WHRFs) and dominated by poor local
 057 minima (Anschuetz, 2022; Anschuetz & Kiani, 2022); this is a far cry from the effectively convex
 058 training landscapes of deep quantum networks which, in certain circumstances, *can* be shown to
 059 have a Gaussian process limit (Liu et al., 2022; 2023; You et al., 2022; García-Martín et al., 2023;
 060 Girardi & Palma, 2024; García-Martín et al., 2024). Unfortunately, these deep networks are still
 061 untrainable in practice due to an exponential decay in their gradients known as the *barren plateau*
 062 phenomenon (McClellan et al., 2018), making estimating gradients on a quantum computer asymp-
 063 totically intractable.¹

064 The presence of barren plateaus or poor local minima in QNN loss landscapes paints a pessimistic
 065 picture for the practical utility of generic QNNs. However, these negative results can be circum-
 066 vented by considering *structured* QNNs, further complicating any unifying theory of the asymptotic
 067 training behavior of QNNs. For instance, U_θ may be constrained to belong to some low-dimensional
 068 Lie subgroup G of the full unitary group, and \mathcal{O} in the generating algebra of G . This is known as
 069 the *Lie algebra-supported ansatz* (LASA) setting, and has been shown to be capable of preventing
 070 the barren plateau phenomenon from occurring (Fontana et al., 2024; Ragone et al., 2024). Though
 071 many of the initial proposals for such structured networks have since been “dequantized” (Anschuetz
 072 et al., 2023a; Goh et al., 2023; Cerezo et al., 2023)—i.e., efficient classical algorithms have been
 073 found which simulate such networks—there still do exist unconditionally provable quantum advan-
 074 tages in using such networks for certain learning tasks (Anschuetz et al., 2023b; Anschuetz & Gao,
 075 2024). Though this is perhaps the most promising direction in QNN research, nothing concretely
 076 is known about the loss landscapes of such networks beyond the variance of the loss function over
 077 parameter space.

078 1.2 CONTRIBUTIONS

079 Fully understanding how the various phenomenologies of QNN loss landscapes are related is im-
 080 portant if we ever hope to have as deep an understanding of QNNs as the neural tangent kernel has
 081 enabled for classical neural networks. Motivated by this, we here for the first time prove a concise
 082 asymptotic limit for the loss functions of effectively all QNNs with approximately uniformly ran-
 083 dom initialization, and show that it unifies much of our previous understanding of QNN training
 084 behavior. We then use this new asymptotic description to prove a variety of novel results on the
 085 asymptotic training behavior of QNNs.

086 We achieve this by demonstrating that all QNNs have a natural algebraic structure described by
 087 a *Jordan subalgebra* \mathcal{A} of the complex Hermitian matrices $\mathcal{H}^N(\mathbb{C})$ generated by $\left\{U_\theta^\dagger \mathcal{O} U_\theta\right\}_\theta$
 088 under the anticommutator $\{A, B\} = AB + BA$. The Jordan subalgebras of $\mathcal{H}^N(\mathbb{C})$ have been
 089 completely classified in the sense that they are expressible as a direct sum of subalgebras:
 090

$$091 \mathcal{A} \cong \bigoplus_{\alpha} \mathcal{A}_{\alpha}, \quad (2)$$

092 where for our purposes each \mathcal{A}_{α} is isomorphic to the algebra of $N_{\alpha} \times N_{\alpha}$ Hermitian matrices
 093 $\mathcal{H}^{N_{\alpha}}(\mathbb{F}_{\alpha})$ over a field \mathbb{F}_{α} that is one of the reals \mathbb{R} , complex numbers \mathbb{C} , or quaternions \mathbb{H} . This
 094 allows us to show that the loss functions and first two derivatives of randomly initialized QNNs have
 095 a concise asymptotic description in terms of *Wishart-distributed random matrices* over \mathbb{F}_{α} :
 096

$$097 \mathbf{W}_{\alpha} = \mathbf{X}_{\alpha} \mathbf{X}_{\alpha}^{\dagger}, \quad (3)$$

098 where \mathbf{X}_{α} is $N_{\alpha} \times r_{\alpha}$ with i.i.d. standard Gaussian entries over \mathbb{F}_{α} and r_{α} is called the *degrees of*
 099 *freedom* of \mathbf{W}_{α} . We give an explicit expression for r_{α} based on the structure of the network, and we
 100 find that it dictates the asymptotic behavior of the loss landscape. We call this connection between
 101 the algebraic structure of a given QNN and its asymptotic Wishart process description a *Jordan*
 102 *algebraic Wishart system* (JAWS). We summarize in Table 1 how this JAWS description generalizes
 103 previous models of QNN loss landscapes which only held in specific parameter regimes.
 104

105 ¹Girardi & Palma (2024) also study instances which do not suffer from barren plateaus, but which only
 106 achieve an asymptotically-vanishing improvement over random θ in the optimization of Eq. (1).
 107

Barren Plateaus		Gaussian Process (QNTK)	
Reference	Corollary 4	Reference	Corollary 5
McClellan et al. (2018)	$\mathbb{F} = \mathbb{C}$	García-Martín et al. (2023)	$\mathbb{F} = \mathbb{C}$, \mathcal{O} Pauli
Cerezo et al. (2021)	$\{\mathbb{F}_\alpha = \mathbb{C}\}_\alpha$	Girardi & Palma (2024)	$\mathbb{F} = \mathbb{C}$, \mathcal{O} 1-local
Ragone et al. (2024)	$e^{i\mathcal{O}} \in \text{Aut}(\mathcal{A})$	García-Martín et al. (2024)	$\mathbb{F} = \mathbb{H}$, \mathcal{O} Pauli

Local Minima	
Reference	Corollary 6
Anschuetz (2022)	$\mathbb{F} = \mathbb{C}$
Anschuetz & Kiani (2022)	$\{\mathbb{F}_\alpha = \mathbb{C}\}_\alpha, p < 2r_\alpha$
You et al. (2022)	$p \gg r$

Table 1: **Previous quantum neural network loss landscape results.** We summarize (a representative subset of) previous results in quantum neural network loss landscape theory, as described in Sec. 1.1. We also state which limits of our Corollaries 4, 5, and 6 encompass these referenced results. Here, \mathcal{A} is the Jordan algebra, \mathbb{F}_α are the fields, and r_α are the degrees of freedom parameters as defined in the associated corollary statements; \mathcal{O} is the objective observable; and p is the number of trained parameters.

The JAWS framework allows us to prove many new, general properties of QNNs. First, we prove a generalized barren plateau result for the variance of the loss function $\ell(\boldsymbol{\rho}; \boldsymbol{\theta})$ of a given QNN over its initialization:

$$\text{Var}_\theta [\ell(\boldsymbol{\rho}; \boldsymbol{\theta})] = \sum_\alpha \frac{\text{Tr}((\mathcal{O}^\alpha)^2) \text{Tr}((\boldsymbol{\rho}^\alpha)^2)}{\dim_{\mathbb{R}}(\text{Aut}(\mathcal{A}_\alpha))}, \quad (4)$$

where $\mathcal{A} = \bigoplus_\alpha \mathcal{A}_\alpha$ is the Jordan algebra associated with the QNN, $\dim_{\mathbb{R}}(\text{Aut}(\mathcal{A}_\alpha))$ is the dimension of the automorphism group of \mathcal{A}_α as a real manifold, and \cdot^α denotes projection into \mathcal{A}_α (see Eq. (12) for a mathematical definition). This extends previously known barren plateau results proved in the LASA framework (Fontana et al., 2024; Ragone et al., 2024) to the setting where neither $\boldsymbol{\rho}$ nor \mathcal{O} are elements of the algebra generating $\text{Aut}(\mathcal{A}_\alpha)$. It also captures the gradient scaling in the setting of so-called matchgate networks (Diaz et al., 2023), unifying barren plateau results beyond just the LASA setting.

We are also able to calculate the asymptotic density of local minima at a loss value $\ell = z$ for a QNN with associated Jordan algebra $\mathcal{A} = \bigoplus_\alpha \mathcal{A}_\alpha$:

$$\kappa(z) = \bigstar_\alpha f_\Gamma(z; k_\alpha, s_\alpha) = \int_{\sum_\alpha z_\alpha = z} \prod_\alpha \frac{dz_\alpha}{\Gamma(k_\alpha) s_\alpha^{k_\alpha}} z_\alpha^{k_\alpha - 1} \exp\left(-\frac{z_\alpha}{s_\alpha}\right); \quad (5)$$

i.e., we show that it is a convolution of gamma distributions f_Γ with shape and scale parameters k_α and s_α , respectively, which we calculate in Sec. 5.3. The k_α and s_α are such that $\kappa(z)$ experiences a phase transition: local minima of a network with p parameters are concentrated near the global minimum if and only if the network is *overparameterized*, which occurs when (up to a constant):

$$p \gtrsim \max_\alpha \frac{\text{Tr}(\mathcal{O}^\alpha)^2}{\text{Tr}((\mathcal{O}^\alpha)^2)}. \quad (6)$$

This result generalizes previous studies of local minima in QNNs (Anschuetz, 2022; Anschuetz & Kiani, 2022) to a setting where the variational ansatz may have some sort of algebraic structure.²

Finally, we prove the necessary and sufficient conditions for a class of QNNs to asymptotically converge to a Gaussian process limit. Taken together, our results indicate that QNNs are efficiently trainable asymptotically using problem-independent optimization algorithms if and only if Eq. (4) is not exponentially vanishing with the problem size and Eq. (6) is satisfied.

²Anschuetz & Kiani (2022) consider local minima in local, shallow networks, where the Hilbert space can be decomposed into a direct sum of Hilbert spaces associated with the light cones of local observables. Here, we allow for general algebraic structures.

2 PRELIMINARIES

2.1 QUANTUM NEURAL NETWORKS

We first review *quantum neural networks* (QNNs). These are defined by a parameterized *ansatz*

$$U_{\theta} = g_{p+1} \prod_{i=p}^1 \exp(-i\theta_i A_i) g_i, \quad (7)$$

where the A_i are complex Hermitian $N \times N$ matrices and the g_i are $N \times N$ unitary matrices. Generally, the g_i and $\exp(-i\theta_i A_i)$ may be constrained to belong to a path-connected Lie subgroup G of the unitary group $U(N)$ to strengthen the inductive bias of the network (Meyer et al., 2023). Such a constraint has also been used as a theoretical model for shallow quantum networks, where G is approximately the direct product $\times_{\alpha} U(N_{\alpha})$ of unitaries acting on local patches in the network (Anschuetz & Kiani, 2022).

Training such networks involves minimizing an empirical risk of the form:

$$f(\theta) = \frac{1}{|\mathcal{R}|} \sum_{\rho \in \mathcal{R}} \ell(\rho; \theta) = \frac{1}{|\mathcal{R}|} \sum_{\rho \in \mathcal{R}} \text{Tr}(U_{\theta} \rho U_{\theta}^{\dagger} O). \quad (8)$$

Here, \mathcal{R} can be thought of a data set comprising multiple input *quantum states* ρ —that is, $\rho \succeq \mathbf{0} \in \mathbb{C}^{N \times N}$ with trace 1—and $O \in \mathbb{C}^{N \times N}$ is Hermitian. When O and the algebra elements generating U_{θ} can be expressed as a sum of O (poly $\log(N)$) Pauli matrices, such a loss can be estimated in $O(\text{poly} \log(N))$ time on a quantum computer (Nielsen & Chuang, 2010b).

Our first main result is that all losses of the form of $\ell(\rho; \theta)$ can be interpreted in terms of automorphisms of Jordan algebras. In preparation for discussing this connection, we now review Jordan algebras.

2.2 JORDAN ALGEBRAS

A *Jordan algebra* over the reals is formally a real vector space V with a commutative multiplication operation \circ acting on $u, v \in V$ satisfying the *Jordan identity*:

$$u \circ ((u \circ u) \circ v) = (u \circ u) \circ (u \circ v), \quad (9)$$

which ensures the associativity of the power. A simple example is the Jordan algebra $\mathcal{H}^N(\mathbb{C})$ of $N \times N$ complex-valued Hermitian matrices with \circ given by the anticommutator $A \circ B = AB + BA$.

The Jordan subalgebras \mathcal{A} of $\mathcal{H}^N(\mathbb{C})$ have been completely classified (Koecher, 1999d) and are known to have a semisimple decomposition:

$$\mathcal{A} \cong \bigoplus_{\alpha} \mathcal{A}_{\alpha}; \quad (10)$$

that is, \mathcal{A} is isomorphic to a direct sum of subalgebras \mathcal{A}_{α} . For our purposes, each \mathcal{A}_{α} is isomorphic to the algebra $\mathcal{H}^{N_{\alpha}}(\mathbb{F}_{\alpha})$ of $N_{\alpha} \times N_{\alpha}$ Hermitian matrices over a field \mathbb{F}_{α} that is one of the reals, complex numbers, or quaternions. We label these three cases with the integers $\beta_{\alpha} = 1, 2, 4$, respectively, and call the representation of each of these algebras as $N_{\alpha} \times N_{\alpha}$ matrices over \mathbb{F}_{α} the *defining representation*. We call the \mathcal{A}_{α} the *simple components* of \mathcal{A} and say that \mathcal{A} is *simple* when there is only a single \mathcal{A}_{α} , i.e., when $\mathcal{A} \cong \mathcal{H}^N(\mathbb{F})$ for some \mathbb{F} .

The automorphism group $\text{Aut}(\mathcal{A})$ of such a Jordan algebra \mathcal{A} is also known: its path-connected components are isomorphic to direct products of the classical Lie groups $SO(N_{\alpha})$, $SU(N_{\alpha})$, and $Sp(N_{\alpha})$, each a subgroup of the corresponding $\text{Aut}(\mathcal{A}_{\alpha})$ (Koecher, 1999b; Orliczky, 2024). As a Lie group, each $\text{Aut}(\mathcal{A}_{\alpha})$ has a well-defined dimension as a real manifold which we denote as $\dim_{\mathbb{R}}(\text{Aut}(\mathcal{A}_{\alpha}))$.

It is also known that the Frobenius inner product $\text{Tr}_R(A^{\dagger}B)$ between $A, B \in \mathcal{A}_{\alpha}$ in any representation R is always proportional to the Frobenius inner product in the defining representation (Koecher, 1999a). In particular, for any $A, B \in \mathcal{A}_{\alpha} \subseteq \mathcal{A} \subseteq \mathcal{H}^N(\mathbb{C})$, using $\text{Tr}_{\alpha}(\cdot)$ to denote the trace in the defining representation of \mathcal{A}_{α} and $\text{Tr}(\cdot)$ to denote that of $\mathcal{H}^N(\mathbb{C})$, we have:

$$\text{Tr}(AB) = I_{\alpha} \text{Tr}_{\alpha}(AB) \quad (11)$$

for some $I_\alpha > 0$ that is independent of A, B . In what follows we will continue to use $\text{Tr}(\cdot)$ to denote the trace in the defining representation of $\mathcal{H}^N(\mathbb{C})$ and $\text{Tr}_\alpha(\cdot)$ to denote that of \mathcal{A}_α . More details on Jordan algebras and some subtleties in their classification are provided in Appendix A.2.

Finally, we define the *projection* of an algebra element $A \in \mathcal{A} \cong \bigoplus_\alpha \mathcal{A}_\alpha$ into one of the \mathcal{A}_α ; to represent this, we use the notation A^α . Mathematically, letting $\text{Tr}_\alpha(\cdot)$ denote the trace in the defining representation of \mathcal{A}_α and considering an orthonormal³ basis $\{B_{\alpha,i}\}_i$ of $\mathcal{A}_\alpha \subseteq \mathcal{A}$, we have:

$$A^\alpha \equiv \sum_i \text{Tr}_\alpha(B_{\alpha,i} A) B_{\alpha,i}. \quad (12)$$

We use the same notation for the equivalent projection of a Lie group element $g \in G = \times_\alpha G_\alpha$ into one of the G_α . Using $\{B_{\alpha,i}\}_i$ to denote an orthonormal basis of the Lie algebra generating G_α , and using the observation that $g = \exp\left(\sum_{\alpha,i} c_{\alpha,i} B_{\alpha,i}\right)$ for some $\{c_{\alpha,i} \in \mathbb{R}\}_{\alpha,i}$, we define:

$$g^\alpha \equiv \exp\left(\sum_i c_{\alpha,i} B_{\alpha,i}\right). \quad (13)$$

3 JORDAN ALGEBRAIC DESCRIPTIONS OF QUANTUM NEURAL NETWORKS

Recall the general form of a QNN loss function:

$$\ell(\rho; \theta) = \text{Tr}\left(\rho \mathbf{U}_\theta^\dagger \mathbf{O} \mathbf{U}_\theta\right), \quad (14)$$

where

$$\mathbf{U}_\theta = \mathbf{g}_{p+1} \prod_{i=p}^1 \exp(-i\theta_i \mathbf{A}_i) \mathbf{g}_i \quad (15)$$

generally may be such that the \mathbf{g}_i and $\exp(-i\theta_i \mathbf{A}_i)$ belong to some path-connected Lie subgroup $G \subseteq \text{U}(N)$.

The $\mathbf{U}_\theta^\dagger \mathbf{O} \mathbf{U}_\theta$ generate under the anticommutator a Jordan subalgebra \mathcal{A} of the $N \times N$ Hermitian matrices $\mathcal{H}^N(\mathbb{C})$. Due to this fact, as well as the classification results discussed in Sec. 2.2, we may equivalently consider Eq. (14) in the following way:

1. $\mathbf{O} \in \mathcal{A} \cong \bigoplus_\alpha \mathcal{H}^{N_\alpha}(\mathbb{F}_\alpha)$.
2. $\mathbf{U}_\theta \in G = \times_\alpha G_\alpha$, where each G_α is isomorphic to a subgroup of $\text{Aut}(\mathcal{A}_\alpha)$ that is one of $\text{SO}(N_\alpha)$, $\text{SU}(N_\alpha)$, or $\text{Sp}(N_\alpha)$.
3. For I_α the proportionality constant of Eq. (11),

$$\ell(\rho; \theta) = \sum_\alpha I_\alpha \text{Tr}_\alpha\left(\rho^\alpha \mathbf{U}_\theta^{\alpha\dagger} \mathbf{O}^\alpha \mathbf{U}_\theta^\alpha\right). \quad (16)$$

We discuss the universality of this Jordan algebraic description in more detail in Appendix B.1.

This description of QNN loss landscapes in terms of Jordan algebraic properties allows us to give a simple way to classify various QNN architectures through the N_α , \mathbb{F}_α , I_α , and \mathbf{O}^α ; we call this collection of algebraic objects a *Jordan algebraic Wishart system* (JAWS). We will next tie this classification to the asymptotic properties of the loss landscape through the use of Wishart random matrices, justifying the use of ‘‘Wishart’’ in the nomenclature.

4 QUANTUM NEURAL NETWORKS ARE WISHART PROCESSES

Our main result is proving that the loss functions of wide QNNs—that is, as $\dim(\mathcal{A}) \rightarrow \infty$ —form *Wishart processes*. Such processes can be written exactly in terms of the matrix elements of Wishart-distributed positive semidefinite random matrices \mathbf{W} , which are distributed as:

$$\mathbf{W} = \mathbf{X} \mathbf{X}^\dagger \quad (17)$$

³Orthonormal with respect to the Frobenius inner product in the defining representation of \mathcal{A}_α .




Figure 1: **Loss and derivative densities.** (a) The loss density when the quantum neural network has a *pure input* (i.e., rank-1) and when it has a *mixed input* (i.e., of rank greater than 1). The distributions are centered at the mean eigenvalue of the objective observable. The mixed input density also illustrates when the input is mixed when projected into any simple component of the Jordan algebra associated with the network. (b) The gradient density conditioned on a nonzero loss function value. The distribution is centered at zero. (c) The density of local minima when the quantum neural network is *underparameterized*, *overparameterized*, and when some simple components of the associated Jordan algebra are underparameterized and some are overparameterized (*mix of sectors*).

for \mathbf{X} a rectangular matrix with i.i.d. standard normal entries over a field \mathbb{F} (assumed \mathbb{R} if not otherwise stated). The number of columns r of \mathbf{X} is called the *degrees of freedom* of \mathbf{W} in analogy with the degrees of freedom of a χ^2 -distributed random variable; indeed, the diagonal entries of \mathbf{W} are (up to a constant scaling) i.i.d. χ^2 -distributed with βr degrees of freedom.

Our main result is as follows, stated informally here with a formal statement deferred to Appendix B.2:

Theorem 1 (Quantum neural networks are Wishart processes, informal). *Consider a QNN with associated JAWS as in Sec. 3, initialized approximately uniformly at random. Let ℓ^* be the optimum of the loss and \bar{o}_α the mean eigenvalue of \mathbf{O}^α . Then, as $\dim(\mathcal{A}) \rightarrow \infty$, there is a convergence in joint distribution over $\boldsymbol{\rho}$ at any $\boldsymbol{\theta}$:*

$$\ell(\boldsymbol{\rho}; \boldsymbol{\theta}) - \ell^* \rightsquigarrow \sum_{\alpha} \frac{I_{\alpha} \bar{o}_{\alpha}}{r_{\alpha}} \text{Tr}_{\alpha}(\boldsymbol{\rho}^{\alpha} \mathbf{W}_{\alpha}). \quad (18)$$

The \mathbf{W}_{α} are each independent Wishart-distributed random matrices in the defining representation of \mathcal{A}_{α} with

$$r_{\alpha} = \left\lfloor \frac{\text{Tr}_{\alpha}(\mathbf{O}^{\alpha})^2}{\text{Tr}_{\alpha}((\mathbf{O}^{\alpha})^2)} \right\rfloor \quad (19)$$

degrees of freedom, where $\lfloor \cdot \rfloor$ denotes rounding to the nearest integer.

An illustration of this distribution is provided in Figure 1(a). The proof follows by using the fact that the marginal distributions of matrix elements of the ansatz are approximately Gaussian distributed; this is made quantitative by bounding the errors in the associated joint characteristic functions over the inputs $\boldsymbol{\rho}$. This is enough to demonstrate convergence of $\ell(\boldsymbol{\rho}; \boldsymbol{\theta}) - \ell^*$ to a convolution of Wishart processes. We then prove a Welch–Satterthwaite-like result to demonstrate an asymptotic convergence of this convolution to a Wishart process with degrees of freedom parameters given by Eq. (19). The full proof is given in Appendix D. We there also prove a strengthened result—convergence *pointwise* in the corresponding probability densities—under certain additional assumptions on the initialization of the network.

We not only prove the convergence of the QNN loss function ℓ to that of a Wishart process, but also ℓ , its gradient, and Hessian jointly. In doing so we are able to show that the *loss landscape* of QNNs—not just the landscape at a single point in parameter space, or averaged over it—converges to that of a Wishart process. We here only give an informal statement of the result in the special case when each $\boldsymbol{\rho}^{\alpha}$ is rank-1, but give the full distribution and formal statement in Appendix B.2 with proof in Appendix D. This distribution when conditioned on a nonzero value for the loss is illustrated in Figure 1(b).

Theorem 2 (Gradient distribution, informal). *Consider the setting of Theorem 1. Assume each ρ^α is rank-1 in its defining representation. Let σ_α be the standard deviation of the eigenvalues of \mathbf{O}^α . Then, as $\dim(\mathcal{A}) \rightarrow \infty$, the joint distribution of $\partial_{\theta_i} \ell(\rho; \theta)$ conditioned on all $\ell(\rho^\alpha; \theta) = z_\alpha$ asymptotically converges to the joint distribution of*

$$\hat{\ell}_i \equiv \sum_{\alpha} \frac{2I_{\alpha} \sigma_{\alpha} \text{Tr}_{\alpha}(\rho^{\alpha})}{N_{\alpha}} \sqrt{\frac{\beta_{\alpha} z_{\alpha}}{I_{\alpha} \bar{\sigma}_{\alpha}}} G_{\alpha, i} \chi_{\alpha, i}. \quad (20)$$

Here, the $\chi_{\alpha, i}$ are independent χ -distributed random variables with $\max(2, \beta_{\alpha})$ degrees of freedom and the $G_{\alpha, i}$ are i.i.d. standard normal random variables.

We also give an expression for the Hessian distribution at critical points. We informally report the result when each ρ^α is rank-1 and the spectrum of each \mathbf{O}^α is sufficiently concentrated around its mean; this latter condition is typical of sums of low-weight fermionic (Feng et al., 2019) and local spin operators (Erdős & Schröder, 2014). The full distribution and formal statement is once again given in Appendix B.2 with proof in Appendix D.

Theorem 3 (Hessian distribution, informal). *Consider the setting of Theorem 2. Assume as well that $\frac{\sigma_{\alpha}}{\bar{\sigma}_{\alpha}} \rightarrow 0$ for all α . Then, as $\dim(\mathcal{A}) \rightarrow \infty$, the joint distribution of $\partial_{\theta_i} \partial_{\theta_j} \ell(\rho; \theta)$ ($i \geq j$) conditioned on all $\ell(\rho^\alpha; \theta) = z_\alpha$ and all $\partial_{\theta_i} \ell(\rho^\alpha; \theta) = 0$ asymptotically converges to the joint distribution of*

$$\hat{\ell}_{i, j} \equiv \sum_{\alpha} \frac{2I_{\alpha} \sigma_{\alpha} \text{Tr}_{\alpha}(\rho^{\alpha})}{N_{\alpha}^2} \sqrt{\frac{z_{\alpha}}{I_{\alpha} \bar{\sigma}_{\alpha}}} G_{\alpha, i} \chi_{\alpha, j} W_{\alpha, i, j}. \quad (21)$$

Here, the $\chi_{\alpha, i}$ are independent χ -distributed random variables with $\max(2, \beta_{\alpha})$ degrees of freedom and the $G_{\alpha, i}$ are i.i.d. standard normal random variables. The \mathbf{W}_{α} are independent Wishart-distributed random matrices with $\beta_{\alpha} r_{\alpha}$ degrees of freedom.

Both of these derivative results follow by carefully considering the joint distribution of elements of a Wishart-distributed random matrix via the use of the Bartlett decomposition of Wishart-distributed matrices. Further simplifications follow by bounding the error in probability between various products of these marginal elements to simpler distributions.

5 NEW RESULTS IN LANDSCAPE THEORY FROM THE JAWS FRAMEWORK

With the full asymptotic distribution of QNN loss landscapes in hand, we are now able to prove several novel results in QNN loss landscape theory.

5.1 BARREN PLATEAUS

The first implication of our results that we will discuss is the unification of barren plateau results (Fontana et al., 2024; Ragone et al., 2024) in the limit of large Jordan algebra dimension.

Corollary 4 (General expression for the loss function variance, informal). *Let ℓ be as in Theorem 1. The variance of the loss over the initialization is:*

$$\text{Var}_{\theta} [\ell(\rho; \theta)] = \sum_{\alpha} \frac{\text{Tr} \left((\mathbf{O}^{\alpha})^2 \right) \text{Tr} \left((\rho^{\alpha})^2 \right)}{\dim_{\mathbb{R}}(\text{Aut}(\mathcal{A}_{\alpha}))}. \quad (22)$$

This follows immediately from the variance of elements of Wishart-distributed random matrices; see Appendix C.1 for an explicit calculation. Despite the simplicity of the proof (given Theorem 1), this result generalizes Theorem 1 of Ragone et al. (2024), which only considered when either $i\mathbf{O}$ or $i\rho$ was in the algebra generating $\text{Aut}(\mathcal{A}_{\alpha})$. However, our result holds for *all* variational ansatzes due to our general Jordan algebraic formulation of QNNs. Intriguingly, $\text{Tr} \left((\rho^{\alpha})^2 \right)$ can be thought of as probing the *generalized entanglement* (Barnum et al., 2004) of ρ with respect to the Jordan algebraic structure of \mathcal{A} . That is, entanglement induces barren plateaus, as previously seen in a nonalgebraic setting by Ortiz Marrero et al. (2021). A similar phenomenon was also previously noted by Ragone et al. (2024) with respect to the Lie algebraic structure of LASAs.

5.2 THE QUANTUM NEURAL TANGENT KERNEL

We now connect our results to the quantum neural tangent kernel (QNTK) literature (Liu et al., 2022; 2023; You et al., 2022; García-Martín et al., 2023; Girardi & Palma, 2024; García-Martín et al., 2024). The landmark result in this field is that, in certain settings, QNN loss functions are asymptotically Gaussian processes. However, this same body of work has noted that such a Gaussian process description cannot generally hold. For instance, if the objective observable \mathbf{O} is rank-1 and the algebra \mathcal{A} associated with the QNN is the space of complex Hermitian matrices, it is known that the loss is proportional to an exponentially-distributed random variable (Boixo et al., 2018).

Our results can be seen as a unifying model of neural network loss landscapes, including both when convergence to a Gaussian process is achieved and when it is not. Recall that Theorem 1 demonstrated the asymptotic expression for the QNN loss at any θ (left implicit) to be:

$$\hat{\ell}(\boldsymbol{\rho}) = \sum_{\alpha} \frac{I_{\alpha} \bar{o}_{\alpha}}{r_{\alpha}} \text{Tr}_{\alpha}(\boldsymbol{\rho}^{\alpha} \mathbf{W}_{\alpha}). \quad (23)$$

This correctly captures the exponential behavior when \mathcal{A} is the space of complex Hermitian matrices and $\text{rank}(\mathbf{O}) = r = 1$. This is because the diagonal entries of such a complex Wishart matrix are (up to a constant) χ^2 -distributed with two degrees of freedom each, which is identical to an exponential distribution.

Indeed, our more general result can be used to *exactly* characterize when the loss functions of QNNs asymptotically form Gaussian processes. First, we consider normalizing $\hat{\ell}$ by some $\mathcal{N} \geq 1$ such that $\mathcal{N}\hat{\ell}$ has nonvanishing variance asymptotically. Under this normalization, convergence to a Gaussian process occurs when $\mathcal{N}\hat{\ell}$ has higher-order cumulants asymptotically vanishing. In other words, using $\kappa_i(\cdot)$ to denote the i th cumulant and \mathcal{R} to denote the data set, this occurs when (see Appendix C.2 for an explicit calculation):

$$\max_{\substack{\boldsymbol{\rho} \in \mathcal{R} \\ i > 2}} \kappa_i(\mathcal{N}\hat{\ell}(\boldsymbol{\rho})) \sim \max_{\boldsymbol{\rho} \in \mathcal{R}} \frac{\mathcal{N}^3 I_{\alpha}^3 \bar{o}_{\alpha}^3 \text{Tr}_{\alpha}((\boldsymbol{\rho}^{\alpha})^3)}{r_{\alpha}^2} = o(1). \quad (24)$$

We thus can state this pair of conditions as follows.

Corollary 5 (Exact conditions for convergence to a Gaussian process, informal). *Let ℓ be as in Theorem 1. $\mathcal{N}\ell(\boldsymbol{\rho})$ is asymptotically a Gaussian process over $\boldsymbol{\rho} \in \mathcal{R}$ if and only if Eq. (24) is satisfied and $\mathcal{N}^2 \text{Var}[\ell(\boldsymbol{\rho})]$ (as in Corollary 4) is nonvanishing for all $\boldsymbol{\rho} \in \mathcal{R}$.*

When these conditions are met, the covariance is given by:

$$\begin{aligned} \mathcal{K}(\boldsymbol{\rho}, \boldsymbol{\rho}') &= \lim_{N \rightarrow \infty} \mathcal{K}_N(\boldsymbol{\rho}, \boldsymbol{\rho}') \\ &= \lim_{N \rightarrow \infty} \mathcal{N}^2 \sum_{\alpha} \frac{\beta_{\alpha} \text{Tr}((\mathbf{O}^{\alpha})^2)}{2 \dim_{\mathbb{R}}(\text{Aut}(\mathcal{A}_{\alpha})) r_{\alpha}} I_{\alpha} \mathbb{E}[\text{Tr}_{\alpha}(\boldsymbol{\rho}^{\alpha} (\mathbf{W}_{\alpha} - r_{\alpha})) \text{Tr}_{\alpha}(\boldsymbol{\rho}'^{\alpha} (\mathbf{W}_{\alpha} - r_{\alpha}))]. \end{aligned} \quad (25)$$

It is now easy to see why (for simple \mathcal{A}) the $r = 1$ case does not converge to a Gaussian process: in this setting \mathcal{N} must be $o(N)$ such that Eq. (24) is satisfied and higher-order cumulants vanish, but for such a choice the variance is vanishing.

In Appendix C.2 we are also more explicit on the form of Eq. (25). We use known results for the mixed cumulants of Wishart matrix elements to show that the covariance can be written as an explicit, quadratic expression in $\boldsymbol{\rho}^{\alpha}$ and $\boldsymbol{\rho}'^{\alpha}$, depending only on \mathcal{A} . We further show in Appendix C.2 that, when the $\boldsymbol{\rho} \in \mathcal{R}$ can be mutually diagonalized,

$$\mathcal{K}_N(\boldsymbol{\rho}, \boldsymbol{\rho}') = \mathcal{N}^2 \sum_{\alpha} \frac{\text{Tr}((\mathbf{O}^{\alpha})^2)}{\dim_{\mathbb{R}}(\text{Aut}(\mathcal{A}_{\alpha}))} \text{Tr}(\boldsymbol{\rho}^{\alpha} \boldsymbol{\rho}'^{\alpha}). \quad (26)$$

This form of the covariance can then be immediately fed into neural tangent kernel results to reason about the training behavior (Neal, 1996; Choromanska et al., 2015; Chaudhari, 2018; Lee et al.,

2018) and generalization ability (Jacot et al., 2018; Wei et al., 2022) of such networks. For instance, by the results of Girardi & Palma (2024), this covariance suggests that gradient descent does not reach the global minimum in time polynomial in the model size for unstructured QNNs. This is discussed in more detail in Appendix C.2.

5.3 LOCAL MINIMA

We end by examining the distribution of local minima of the loss landscape. This has been done previously in the more restricted setting where no structure is imposed on the ansatz or loss function outside of locality (Anschuetz, 2022; Anschuetz & Kiani, 2022). In this section we consider a general QNN with p trained parameters under the conditions of Theorem 3, as well as assume that all

$$\gamma_\alpha \equiv \frac{p}{\beta_\alpha r_\alpha} \quad (27)$$

are held constant as we take the asymptotic limit $\dim(\mathcal{A}) \rightarrow \infty$. We call the γ_α the *overparameterization ratios* of a given QNN architecture, associated with the various simple components of the Jordan algebra \mathcal{A} in the JAWS description of the network.

We calculate the distribution of local minima using the *Kac–Rice formula* (Adler & Taylor, 2007), which gives the expected density of local minima of a random field ℓ at a function value z . Anschuetz (2022) demonstrated that the assumptions for the Kac–Rice formula are satisfied for variational loss landscapes, and when rotationally invariant on the p -torus the formula takes the form:

$$\mathbb{E}[\text{Crt}_0(z)] = (2\pi)^p \mathbb{E}[\det(\mathbf{H}_z) \mathbf{1}\{\mathbf{H}_z \succeq \mathbf{0}\}] \mathbb{P}[\mathbf{G}_z = \mathbf{0}] \mathbb{P}[\ell = z]. \quad (28)$$

Here, $\mathbb{E}[\text{Crt}_0(z)]$ is the expected density of local minima at a function value z , \mathbf{G}_z is the gradient conditioned on $\ell = z$, and \mathbf{H}_z is the Hessian conditioned on $\ell = z$ and $\mathbf{G} = \mathbf{0}$. In a slight abuse of notation, here $\mathbb{P}[\cdot]$ denotes the probability density associated with the event \cdot .

We can evaluate this expression using Theorems 1, 2, and 3. To simplify the expression here, we assume the addition of a regularization term of the form:

$$R_L(\boldsymbol{\theta}) = L \|\boldsymbol{\theta} - \mathbf{C}\|_2^2 \quad (29)$$

to the loss; we also present here only the relative density rather than the total number of minima at a loss function value z . The full expression counting local minima—both with and without the regularization of Eq. (29)—are described in detail in Appendix C.3.

Corollary 6 (Density of local minima, informal). *Consider the setting of Theorem 3 with an additional regularization term of the form of Eq. (29). The density of local minima at a loss function value $z > 0$ is, to multiplicative leading order, given by the convolution over all α with $\gamma_\alpha < 1$:*

$$\kappa(z) = \bigstar_{\alpha: \gamma_\alpha < 1} f_\Gamma\left(\frac{z}{\bar{\sigma}_\alpha \text{Tr}(\boldsymbol{\rho}^\alpha)}, \frac{\beta_\alpha r_\alpha}{2}, \frac{2}{\beta_\alpha r_\alpha}\right). \quad (30)$$

Here, $f_\Gamma(\cdot; k, \theta)$ denotes the gamma distribution with shape k and scale θ parameters:

$$f_\Gamma(x; k, \theta) = \frac{1}{\Gamma(k)\theta^k} x^{k-1} \exp\left(-\frac{x}{\theta}\right). \quad (31)$$

Examples of this distribution are plotted in Figure 1(c), where we give a finite width to the density in the *overparameterized regime*—where all $\gamma_\alpha \geq 1$ —to represent potential finite-size effects. Intriguingly, the variance of this distribution when all $\gamma_\alpha < 1$ —the *underparameterized regime*—corresponds to the variance of the loss function itself, i.e., as given in Corollary 4. However, due to the exponential tails of the gamma distribution, even when there are no barren plateaus in the loss landscape there is only an exponentially small fraction of local minima in the vicinity of the global minimum in this regime. That is, both barren plateaus and poor local minima are potential obstructions to efficient trainability that must be taken into account in the design of practical QNNs.

6 CONCLUSION

Taken together, our results show that the natural model for quantum neural networks is not one of Gaussian processes obeying a quantum neural tangent kernel, but rather one of *Wishart processes*.

This Wishart process model unifies all of the recent major thrusts in calculating properties of quantum neural network loss landscapes. Indeed, our results allow us to propose a simple operational definition for the “trainability” of quantum neural networks, which to date has been a term used heuristically without any formal definition.⁴

Definition 7 (Trainability of quantum neural networks). Consider a QNN in a Hilbert space of dimension N with p trained parameters. Let $\mathcal{A} \cong \bigoplus_{\alpha} \mathcal{A}_{\alpha}$ be the corresponding Jordan algebra as in Sec. 3. Define the *degrees of freedom* parameters:

$$r_{\alpha} \equiv \frac{\text{Tr}_{\alpha}(\mathbf{O}^{\alpha})^2}{\text{Tr}_{\alpha}((\mathbf{O}^{\alpha})^2)}. \quad (32)$$

We say that the QNN is *trainable* if and only if, as $N \rightarrow \infty$, the QNN satisfies:

1. *Absence of barren plateaus* (Corollary 4):

$$\sum_{\alpha} \frac{\text{Tr}((\mathbf{O}^{\alpha})^2) \text{Tr}((\boldsymbol{\rho}^{\alpha})^2)}{\dim_{\mathbb{R}}(\text{Aut}(\mathcal{A}_{\alpha}))} = \Omega\left(\text{poly}(\log(N))^{-1}\right). \quad (33)$$

2. *Absence of poor local minima* (Corollary 6):

$$p \geq \max_{\alpha} \beta_{\alpha} r_{\alpha}. \quad (34)$$

Given knowledge of \mathcal{A} , this yields a quantum algorithm for determining the asymptotic trainability of a QNN architecture: one need only measure $\text{Tr}((\mathbf{O}^{\alpha})^2)$ and $\text{Tr}((\boldsymbol{\rho}^{\alpha})^2)$ on a quantum computer. This can be done through the measurement of basis elements of \mathcal{A} given copies of $\boldsymbol{\rho}$ and, for the former, block encodings of \mathbf{O} as prepared by the standard linear combination of unitaries subroutine (Childs & Wiebe, 2012). We provide explicit error bounds for this procedure in Appendix E.

Practically, where does this leave quantum neural networks? For one, it seems unlikely that there exists any computational quantum advantage during the training of QNNs outside of HHL-like speedups (Harrow et al., 2009; Biamonte et al., 2017) and training algorithms that leverage some existing knowledge of the data to be learned (Liu et al., 2021; Hastings & O’Donnell, 2022; Huang et al., 2024). This is due to efficient classical simulation algorithms for quantum systems algebraically constrained to explore a low-dimensional subspace (Anschuetz et al., 2023a; Goh et al., 2023), as is required for efficient trainability per Definition 7. This was noted in Cerezo et al. (2023) for deep QNNs, where the authors postulated that deep QNNs exhibiting no barren plateaus are classically simulable (outside of certain special cases). Our results demonstrate a similar phenomenon for *shallow* QNNs: poor local minima in polynomially-sized circuits can only be avoided when the effective Hilbert space dimension grows at most polynomially quickly with the system size. This leaves the space for a practical, superpolynomial quantum advantage when training with a problem-agnostic algorithm such as gradient descent even narrower than previously believed: a veritable Amity Island in a sea of negative results. One ray of hope is the known existence of *polynomial* quantum advantages during inference in such a setting (Anschuetz et al., 2023b; Anschuetz & Gao, 2024). Intelligent warm-starting of the optimization procedure may be another way to circumvent poor training (Puig-i-Valls et al., 2024), though more must be done to fully understand how precise warm-starting must be for training to be efficient. Given the exact asymptotic form of the loss landscape we give here, such an analysis may be possible in the future.

Our work gives a unified understanding of quantum neural networks as Wishart processes. Great strides have been made in the classical machine learning literature in understanding the training dynamics (Choromanska et al., 2015; Chaudhari, 2018; Jacot et al., 2018; Allen-Zhu et al., 2019) and generalization behavior (Jacot et al., 2018; Wei et al., 2022) of classical neural networks via their connections to Gaussian processes, which unfortunately only port over in the specific settings where the Wishart process itself approaches a Gaussian process. Our results encourage an understanding of how specific properties of Wishart processes, not just Gaussian processes, influence the learning behavior of quantum networks in order to more fully grasp how quantum neural networks learn.

⁴The concept of “efficient learning” has been previously studied (Gil-Fuster et al., 2024), but we give the first *operational* definition in terms of the structural properties of a given QNN architecture.

540
541
542
543
544
545
546
547
548
549
550
551
552
553
554
555
556
557
558
559
560
561
562
563
564
565
566
567
568
569
570
571
572
573
574
575
576
577
578
579
580
581
582
583
584
585
586
587
588
589
590
591
592
593

REPRODUCIBILITY STATEMENT

The main results discussed in Secs. 3 and 4 are formally stated—with an explicit listing of formal assumptions—in Appendix B. Definitions and background necessary for these formal theorem and assumption statements are laid out in Appendix A. Formal proofs of the main results are given in Appendix D. Formal statements and proofs of the corollaries discussed in Sec. 5 are given in Appendix C. A proof of the claim made in Sec. 6 that one can efficiently estimate the quantities in Definition 7 using an LCU block encoding is given in Appendix E. A discussion of the settings where our results do not hold due to a violation of our formal assumptions is given in Appendix F. The remainder of the Appendices prove claims and helper lemmas used in Appendices B, C, and D.

REFERENCES

- Robert J. Adler and Jonathan E. Taylor. *Random Fields on Manifolds*, pp. 301–330. Springer New York, New York, USA, 2007. ISBN 978-0-387-48116-6. doi: 10.1007/978-0-387-48116-6_12.
- Zeyuan Allen-Zhu, Yuanzhi Li, and Zhao Song. A convergence theory for deep learning via over-parameterization. In Kamalika Chaudhuri and Ruslan Salakhutdinov (eds.), *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pp. 242–252, Long Beach, USA, 6 2019. PMLR. URL <https://proceedings.mlr.press/v97/allen-zhu19a.html>.
- Eric R. Anschuetz and Xun Gao. Arbitrary polynomial separations in trainable quantum machine learning. *CoRR*, abs/2402.08606, 2024. doi: 10.48550/ARXIV.2402.08606.
- Eric R. Anschuetz and Bobak T. Kiani. Quantum variational algorithms are swamped with traps. *Nat. Commun.*, 13:7760, 2022. doi: 10.1038/s41467-022-35364-5.
- Eric R. Anschuetz, Andreas Bauer, Bobak T. Kiani, and Seth Lloyd. Efficient classical algorithms for simulating symmetric quantum systems. *Quantum*, 7:1189, 11 2023a. ISSN 2521-327X. doi: 10.22331/q-2023-11-28-1189.
- Eric R. Anschuetz, Hong-Ye Hu, Jin-Long Huang, and Xun Gao. Interpretable quantum advantage in neural sequence learning. *PRX Quantum*, 4:020338, 2023b. doi: 10.1103/PRXQuantum.4.020338.
- Eric Ricardo Anschuetz. Critical points in quantum generative models. In Katja Hofmann, Alexander Rush, Yan Liu, Chelsea Finn, Yejin Choi, and Marc Deisenroth (eds.), *The Tenth International Conference on Learning Representations*. OpenReview, 2022. URL <https://openreview.net/forum?id=2f1z55GVQN>.
- Howard Barnum, Emanuel Knill, Gerardo Ortiz, Rolando Somma, and Lorenza Viola. A subsystem-independent generalization of entanglement. *Phys. Rev. Lett.*, 92:107902, Mar 2004. doi: 10.1103/PhysRevLett.92.107902.
- Istvan Berkes and Walter Philipp. Approximation theorems for independent and weakly dependent random vectors. *Ann. Probab.*, 7(1):29–54, 1979. doi: 10.1214/aop/1176995146.
- Jacob Biamonte, Peter Wittek, Nicola Pancotti, Patrick Rebentrost, Nathan Wiebe, and Seth Lloyd. Quantum machine learning. *Nature*, 549(7671):195–202, 2017. doi: 10.1038/nature23474.
- Sergio Boixo, Sergei V Isakov, Vadim N Smelyanskiy, Ryan Babbush, Nan Ding, Zhang Jiang, Michael J Bremner, John M Martinis, and Hartmut Neven. Characterizing quantum supremacy in near-term devices. *Nat. Phys.*, 14(6):595–600, 2018. doi: 10.1038/s41567-018-0124-x.
- Dennis D. Boos. A converse to Scheffe’s theorem. *Ann. Stat.*, 13(1):423 – 427, 1985. doi: 10.1214/aos/1176346604.
- Marco Cerezo, Akira Sone, Tyler Volkoff, Lukasz Cincio, and Patrick J Coles. Cost function dependent barren plateaus in shallow parametrized quantum circuits. *Nat. Commun.*, 12(1):1791–1802, 2021. doi: 10.1038/s41467-021-21728-w.

- 594 Marco Cerezo, Martin Larocca, Diego García-Martín, Nelson L. Diaz, Paolo Braccia, Enrico
595 Fontana, Manuel S. Rudolph, Pablo Bermejo, Aroosa Ijaz, Supanut Thanasilp, Eric R. Anschuetz,
596 and Zoë Holmes. Does provable absence of barren plateaus imply classical simulability? Or,
597 why we need to rethink variational quantum computing. *CoRR*, abs/2312.09121, 2023. doi:
598 10.48550/ARXIV.2312.09121.
- 599 Shantanav Chakraborty. Implementing any linear combination of unitaries on intermediate-term
600 quantum computers. *Quantum*, 8:1496, October 2024. ISSN 2521-327X. doi: 10.22331/
601 q-2024-10-10-1496.
- 602 Pratik Anil Chaudhari. *A Picture of the Energy Landscape of Deep Neural Networks*. PhD the-
603 sis, University of California, Los Angeles, 2018. URL [https://escholarship.org/uc/
604 item/26h5787r](https://escholarship.org/uc/item/26h5787r).
- 605 Chi-Fang Chen, Alexander M. Dalzell, Mario Berta, Fernando G. S. L. Brandão, and Joel A. Tropp.
606 Sparse random Hamiltonians are quantumly easy. *Phys. Rev. X*, 14:011014, 2 2024. doi: 10.1103/
607 PhysRevX.14.011014.
- 608 Andrew M. Childs and Nathan Wiebe. Hamiltonian simulation using linear combinations of unitary
609 operations. *Quantum Info. Comput.*, 12(11–12):901–924, November 2012. ISSN 1533-7146. doi:
610 10.26421/QIC12.11-12-1.
- 611 Anna Choromanska, Mikael Henaff, Michael Mathieu, Gerard Ben Arous, and Yann LeCun. The
612 loss surfaces of multilayer networks. In Guy Lebanon and S. V. N. Vishwanathan (eds.), *Pro-
613 ceedings of the Eighteenth International Conference on Artificial Intelligence and Statistics*, vol-
614 ume 38 of *Proceedings of Machine Learning Research*, pp. 192–204, San Diego, USA, 2015.
615 PMLR. URL <http://proceedings.mlr.press/v38/choromanska15.html>.
- 616 Benoît Collins, James A Mingo, Piotr Śniady, and Roland Speicher. Second order freeness and
617 fluctuations of random matrices. III: Higher order freeness and free cumulants. *Doc. Math.*, 12:
618 1–70, 2007. doi: 10.4171/DM/220.
- 619 Amir Dembo and Ofer Zeitouni. *General Principles*, pp. 115–174. Springer Berlin Hei-
620 delberg, Berlin, Heidelberg, Germany, 2010. ISBN 978-3-642-03311-7. doi: 10.1007/
621 978-3-642-03311-7_4.
- 622 N. L. Diaz, Diego García-Martín, Sujay Kazi, Martin Larocca, and M. Cerezo. Showcasing a barren
623 plateau theory beyond the dynamical Lie algebra. *arXiv*, abs/2310.11505, 2023. doi: 10.48550/
624 arXiv.2310.11505.
- 625 Alexander Dubbs, Alan Edelman, Plamen Koev, and Praveen Venkataramana. The beta-Wishart
626 ensemble. *J. Math. Phys.*, 54(8):083507, 08 2013. ISSN 0022-2488. doi: 10.1063/1.4818304.
- 627 László Erdős and Dominik Schröder. Phase transition in the density of states of quantum spin
628 glasses. *Math. Phys. Anal. Geom.*, 17(3):441–464, 2014. doi: 10.1007/s11040-014-9164-3.
- 629 Renjie Feng, Gang Tian, and Dongyi Wei. Spectrum of SYK model. *Peking Math. J.*, 2(1):41–70,
630 2019. doi: 10.1007/s42543-018-0007-1.
- 631 Enrico Fontana, Dylan Herman, Shouvanik Chakraborti, Niraj Kumar, Romina Yalovetzky, Jamie
632 Heredge, Shree Hari Sureshababu, and Marco Pistoia. Characterizing barren plateaus in quan-
633 tum ansätze with the adjoint representation. *Nat. Commun.*, 15(1):7171, 2024. doi: 10.1038/
634 s41467-024-49910-w.
- 635 Diego García-Martín, Martin Larocca, and Marco Cerezo. Deep quantum neural networks form
636 Gaussian processes. *CoRR*, abs/2305.09957, 2023. doi: 10.48550/ARXIV.2305.09957.
- 637 Diego García-Martín, Paolo Braccia, and Marco Cerezo. Architectures and random properties of
638 symplectic quantum circuits. *CoRR*, abs/2405.10264, 2024. doi: 10.48550/ARXIV.2405.10264.
- 639 Elies Gil-Fuster, Casper Gyurik, Adrián Pérez-Salinas, and Vedran Dunjko. On the relation between
640 trainability and dequantization of variational quantum learning models. *CoRR*, abs/2406.07072,
641 2024. doi: 10.48550/ARXIV.2406.07072.

- 648 Filippo Girardi and Giacomo De Palma. Trained quantum neural networks are Gaussian processes.
649 *CoRR*, abs/2402.08726, 2024. doi: 10.48550/ARXIV.2402.08726.
- 650
- 651 Matthew L. Goh, Martin Larocca, Lukasz Cincio, M. Cerezo, and Frédéric Sauvage. Lie-algebraic
652 classical simulations for variational quantum computing. *arXiv*, abs/2308.01432, 2023. doi:
653 10.48550/arXiv.2308.01432.
- 654 Alfred Haar. Der Massbegriff in der Theorie der kontinuierlichen Gruppen. *Ann. Math.*, 34(1):
655 147–169, 1933. ISSN 0003486X. doi: 10.2307/1968346.
- 656
- 657 Aram W Harrow and Saeed Mehraban. Approximate unitary t -designs by short random quantum
658 circuits using nearest-neighbor and long-range gates. *Commun. Math. Phys.*, 401(2):1531–1626,
659 2023. doi: 10.1007/s00220-023-04675-z.
- 660 Aram W. Harrow, Avinatan Hassidim, and Seth Lloyd. Quantum algorithm for linear systems of
661 equations. *Phys. Rev. Lett.*, 103:150502, 10 2009. doi: 10.1103/PhysRevLett.103.150502.
- 662
- 663 Matthew B. Hastings and Ryan O’Donnell. Optimizing strongly interacting fermionic Hamilto-
664 nians. In *Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing*,
665 STOC 2022, pp. 776–789, New York, USA, 2022. Association for Computing Machinery. ISBN
666 9781450392648. doi: 10.1145/3519935.3519960.
- 667 Fumio Hiai and Dénes Petz. Eigenvalue density of the Wishart matrix and large deviations.
668 *Infin. Dimens. Anal. Quantum Probab. Relat. Top.*, 01(04):633–646, 1998. doi: 10.1142/
669 S021902579800034X.
- 670 Brice Huang and Mark Sellke. Algorithmic threshold for multi-species spherical spin glasses. *CoRR*,
671 abs/2303.12172, 2023. doi: 10.48550/ARXIV.2303.12172.
- 672
- 673 Hsin-Yuan Huang, Yunchao Liu, Michael Broughton, Isaac Kim, Anurag Anshu, Zeph Landau,
674 and Jarrod R. McClean. Learning shallow quantum circuits. In *Proceedings of the 56th Annual*
675 *ACM Symposium on Theory of Computing*, STOC 2024, pp. 1343–1351, New York, USA, 2024.
676 Association for Computing Machinery. ISBN 9798400703836. doi: 10.1145/3618260.3649722.
- 677 Arthur Jacot, Franck Gabriel, and Clément Hongler. Neural tangent kernel: convergence and gen-
678 eralization in neural networks. In *Proceedings of the 32nd International Conference on Neural*
679 *Information Processing Systems*, NIPS’18, pp. 8580–8589, Red Hook, USA, 2018. Curran Asso-
680 ciates Inc. URL <https://dl.acm.org/doi/abs/10.5555/3327757.3327948>.
- 681
- 682 Tiefeng Jiang. The entries of Haar-invariant matrices from the classical compact groups. *J. Theor.*
683 *Probab.*, 23(4):1227–1243, 2010. doi: 10.1007/s10959-009-0241-7.
- 684 Eytan Katzav and Isaac Pérez Castillo. Large deviations of the smallest eigenvalue of the Wishart-
685 Laguerre ensemble. *Phys. Rev. E*, 82:040104, 10 2010. doi: 10.1103/PhysRevE.82.040104.
- 686
- 687 A.I. Khuri, T. Mathew, and D.G. Nel. A test to determine closeness of multivariate Satterthwaite’s
688 approximation. *J. Multivar. Anal.*, 51(1):201–209, 1994. ISSN 0047-259X. doi: 10.1006/jmva.
689 1994.1057.
- 690 Max Koecher. *Jordan algebras*, pp. 53–72. Springer Berlin Heidelberg, Berlin, Heidelberg, Ger-
691 many, 1999a. ISBN 978-3-540-48402-8. doi: 10.1007/BFb0096288.
- 692
- 693 Max Koecher. *Real and complex Jordan algebras*, pp. 73–92. Springer Berlin Heidelberg, Berlin,
694 Heidelberg, Germany, 1999b. ISBN 978-3-540-48402-8. doi: 10.1007/BFb0096289.
- 695
- 696 Max Koecher. *Jordan algebras and omega domains*, pp. 109–126. Springer Berlin Heidelberg,
697 Berlin, Heidelberg, Germany, 1999c. ISBN 978-3-540-48402-8. doi: 10.1007/BFb0096291.
- 698
- 699 Max Koecher. *Half-spaces*, pp. 127–155. Springer Berlin Heidelberg, Berlin, Heidelberg, Germany,
700 1999d. ISBN 978-3-540-48402-8. doi: 10.1007/BFb0096292.
- 701
- 700 Martin Larocca, Piotr Czarnik, Kunal Sharma, Gopikrishnan Muraleedharan, Patrick J. Coles, and
701 M. Cerezo. Diagnosing barren plateaus with tools from quantum optimal control. *Quantum*, 6:
824, 2022. ISSN 2521-327X. doi: 10.22331/q-2022-09-29-824.

- 702 B. Laurent and P. Massart. Adaptive estimation of a quadratic functional by model selection. *Ann.*
703 *Stat.*, 28(5):1302–1338, 2000. doi: 10.1214/aos/1015957395.
704
- 705 Jaehoon Lee, Yasaman Bahri, Roman Novak, Sam Schoenholz, Jeffrey Pennington, and Jascha
706 Sohl-Dickstein. Deep neural networks as Gaussian processes. In Yoshua Bengio and Yann Le-
707 Cun (eds.), *The Sixth International Conference on Learning Representations*, Vancouver, Canada,
708 2018. OpenReview. URL <https://openreview.net/pdf?id=B1EA-M-0Z>.
- 709 Andy C. Y. Li and Imanol Hernandez. Noise-induced transition in optimal solutions of variational
710 quantum algorithms. *arXiv*, abs/2403.02762, 2024. doi: 10.48550/arXiv.2403.02762.
711
- 712 Fei Li and Yifeng Xue. The density functions of the singular quaternion normal matrix and the
713 singular quaternion Wishart matrix. *Commun. Stat. - Theory Methods*, 39(18):3316–3331, 2010.
714 doi: 10.1080/03610920903259666.
- 715 Junyu Liu, Francesco Tacchino, Jennifer R. Glick, Liang Jiang, and Antonio Mezzacapo. Repre-
716 sentation learning via quantum neural tangent kernels. *PRX Quantum*, 3:030323, 8 2022. doi:
717 10.1103/PRXQuantum.3.030323.
718
- 719 Junyu Liu, Khadijeh Najafi, Kunal Sharma, Francesco Tacchino, Liang Jiang, and Antonio Mezza-
720 capo. Analytic theory for the dynamics of wide quantum neural networks. *Phys. Rev. Lett.*, 130:
721 150601, 4 2023. doi: 10.1103/PhysRevLett.130.150601.
- 722 Yunchao Liu, Srinivasan Arunachalam, and Kristan Temme. A rigorous and robust quantum
723 speed-up in supervised machine learning. *Nat. Phys.*, 17(9):1013–1017, 2021. doi: 10.1038/
724 s41567-021-01287-z.
725
- 726 V A Marčenko and L A Pastur. Distribution of eigenvalues for some sets of random matrices. *Math.*
727 *USSR Sb.*, 1(4):457, 4 1967. doi: 10.1070/SM1967v001n04ABEH001994.
728
- 729 Jarrod R McClean, Sergio Boixo, Vadim N Smelyanskiy, Ryan Babbush, and Hartmut Neven. Bar-
730 ren plateaus in quantum neural network training landscapes. *Nat. Commun.*, 9(1):4812, 2018.
731 doi: 10.1038/s41467-018-07090-4.
- 732 Elizabeth S. Meckes. *Concentration of Measure*, pp. 131–160. Cambridge Tracts in Mathematics.
733 Cambridge University Press, Cambridge, UK, 2019. doi: 10.1017/9781108303453.006.
734
- 735 Antonio Anna Mele, Armando Angrisani, Soumik Ghosh, Sumeet Khatri, Jens Eisert, Daniel Stilck
736 França, and Yihui Quek. Noise-induced shallow circuits and absence of barren plateaus. *arXiv*,
737 abs/2403.13927, 2024. doi: 10.48550/arXiv.2403.13927.
- 738 Johannes Jakob Meyer, Marian Mularski, Elies Gil-Fuster, Antonio Anna Mele, Francesco Arzani,
739 Alissa Wilms, and Jens Eisert. Exploiting symmetry in variational quantum machine learning.
740 *PRX Quantum*, 4:010328, 3 2023. doi: 10.1103/PRXQuantum.4.010328.
741
- 742 James A. Mingo and Mihai Popa. Real second order freeness and Haar orthogonal matrices. *J.*
743 *Math. Phys.*, 54(5):051701, 5 2013. ISSN 0022-2488. doi: 10.1063/1.4804168.
- 744 James A. Mingo and Roland Speicher. Second order freeness and fluctuations of random matrices:
745 I. Gaussian and Wishart matrices and cyclic Fock spaces. *J. Funct. Anal.*, 235(1):226–270, 2006.
746 ISSN 0022-1236. doi: 10.1016/j.jfa.2005.10.007.
747
- 748 Radford M. Neal. *Priors for Infinite Networks*, pp. 29–53. Springer New York, New York, USA,
749 1996. ISBN 978-1-4612-0745-0. doi: 10.1007/978-1-4612-0745-0_2.
- 750 Michael A. Nielsen and Isaac L. Chuang. *The quantum Fourier transform and its applica-*
751 *tions*, pp. 216–247. Cambridge University Press, Cambridge, UK, 2010a. doi: 10.1017/
752 CBO9780511976667.009.
753
- 754 Michael A. Nielsen and Isaac L. Chuang. *The quantum Fourier transform and its applica-*
755 *tions*, pp. 216–247. Cambridge University Press, Cambridge, UK, 2010b. doi: 10.1017/
CBO9780511976667.009.

- 756 Michael Orlitzky. Jordan automorphisms and derivatives of symmetric cones. *Linear Algebra Its*
757 *Appl.*, 2024. ISSN 0024-3795. doi: 10.1016/j.laa.2024.04.024.
758
- 759 Carlos Ortiz Marrero, Mária Kieferová, and Nathan Wiebe. Entanglement-induced barren plateaus.
760 *PRX Quantum*, 2:040316, Oct 2021. doi: 10.1103/PRXQuantum.2.040316.
761
- 762 Alberto Peruzzo, Jarrod McClean, Peter Shadbolt, Man-Hong Yung, Xiao-Qi Zhou, Peter J. Love,
763 Alán Aspuru-Guzik, and Jeremy L. O’Brien. A variational eigenvalue solver on a photonic quan-
764 tum processor. *Nat. Commun.*, 5(1):4213, 2014. doi: 10.1038/ncomms5213.
- 765 Ricard Puig-i-Valls, Marc Drudis, Supanut Thanasilp, and Zoë Holmes. Variational quantum
766 simulation: a case study for understanding warm starts. *CoRR*, abs/2404.10044, 2024. doi:
767 10.48550/ARXIV.2404.10044.
- 768 Michael Ragone, Bojko N Bakalov, Frédéric Sauvage, Alexander F Kemper, Carlos Ortiz Marrero,
769 Martín Larocca, and M Cerezo. A Lie algebraic theory of barren plateaus for deep parameterized
770 quantum circuits. *Nat. Commun.*, 15(1):7172, 2024. doi: 10.1038/s41467-024-49909-3.
771
- 772 C. E. I. Redelmeier. Quaternionic second-order freeness and the fluctuations of large symplectically
773 invariant random matrices. *Random Matrices: Theory Appl.*, 10(02):2150017, 2021. doi: 10.
774 1142/S2010326321500179.
- 775 C. Emily I. Redelmeier. Genus expansion for real Wishart matrices. *J. Theor. Probab.*, 24(4):1044–
776 1062, 2011. doi: 10.1007/s10959-010-0278-7.
777
- 778 Catherine Emily Iska Redelmeier. Real second-order freeness and the asymptotic real second-order
779 freeness of several real matrix models. *Int. Math. Res. Not.*, 2014(12):3353–3395, 3 2012. ISSN
780 1073-7928. doi: 10.1093/imrn/rnt043.
- 781 Alain Rouault. Asymptotic behavior of random determinants in the Laguerre, Gram and Jacobi
782 ensembles. *ALEA Lat. Am. J. Probab. Math. Stat.*, 3:181–230, 2007. URL [https://hal.
783 science/hal-00088079](https://hal.science/hal-00088079).
784
- 785 F. E. Satterthwaite. An approximate distribution of estimates of variance components. *Biometrics*
786 *Bull.*, 2(6):110–114, 1946. ISSN 00994987. doi: 10.2307/3002019.
787
- 788 Henry Scheffé. A useful convergence theorem for probability distributions. *Ann. Math. Stat.*, 18(3):
789 434–438, 1947. doi: 10.1214/aoms/1177730390.
- 790 M.S. Srivastava. Singular Wishart and multivariate beta distributions. *Ann. Stat.*, 31(5):1537–1560,
791 2003. doi: 10.1214/aos/1065705118.
792
- 793 V. Strassen. The existence of probability measures with given marginals. *Ann. Math. Stat.*, 36(2):
794 423–439, 1965. doi: 10.1214/aoms/1177700153.
- 795 Eliran Subag. The complexity of spherical p -spin models—a second moment approach. *Ann.*
796 *Probab.*, 45(5):3385 – 3450, 2017. doi: 10.1214/16-AOP1139.
797
- 798 Roman Vershynin. *Concentration of Sums of Independent Random Variables*, pp. 11–37. Cambridge
799 Series in Statistical and Probabilistic Mathematics. Cambridge University Press, Cambridge, UK,
800 2018. doi: 10.1017/9781108231596.005.
- 801 Samson Wang, Enrico Fontana, Marco Cerezo, Kunal Sharma, Akira Sone, Lukasz Cincio, and
802 Patrick J Coles. Noise-induced barren plateaus in variational quantum algorithms. *Nat. Commun.*,
803 12(1):6961, 2021. doi: 10.1038/s41467-021-27045-6.
804
- 805 Alexander Wei, Wei Hu, and Jacob Steinhardt. More than a toy: Random matrix models predict
806 how real-world neural representations generalize. In Kamalika Chaudhuri, Stefanie Jegelka,
807 Le Song, Csaba Szepesvari, Gang Niu, and Sivan Sabato (eds.), *Proceedings of the 39th Inter-
808 national Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning
809 Research*, pp. 23549–23588. PMLR, 7 2022. URL [https://proceedings.mlr.press/
v162/wei22a.html](https://proceedings.mlr.press/v162/wei22a.html).

- 810 B. L. Welch. The generalization of ‘Student’s’ problem when several different population variances
811 are involved. *Biometrika*, 34(1-2):28–35, 01 1947. ISSN 0006-3444. doi: 10.1093/biomet/34.
812 1-2.28.
- 813 Xuchen You, Shouvanik Chakrabarti, and Xiaodi Wu. A convergence theory for over-parameterized
814 variational quantum eigensolvers. *CoRR*, abs/2205.12481, 2022. doi: 10.48550/ARXIV.2205.
815 12481.
- 816 Soonyu Yu, Jaena Ryu, and Kyoohong Park. A derivation of anti-Wishart distribution. *J. Multivar.*
817 *Anal.*, 131:121–125, 2014. ISSN 0047-259X. doi: 10.1016/j.jmva.2014.06.019.

820 A PRELIMINARIES FOR FORMAL DISCUSSION OF RESULTS

821 We begin by reviewing concepts that we will use in proving our results. We also give a summary of
822 the notation we use throughout in Table 2.

823 A.1 QUANTUM NEURAL NETWORKS

824 We first review *quantum neural networks* (QNNs). These are defined by a parameterized *ansatz*

$$825 \mathbf{U}(\boldsymbol{\theta}) = \prod_{i=p}^1 \exp(-i\theta_i \mathbf{A}_i), \quad (35)$$

826 which belongs to some path-connected subgroup of the unitary group $U(N)$. Though $\mathbf{U}(\boldsymbol{\theta})$ can in
827 principle parameterize all of $U(N)$, it is often taken to instead parameterize some path-connected
828 subgroup $G \subseteq U(N)$. This might be due to enforcing some global structure on the model (Meyer
829 et al., 2023), or can be a model of the finiteness of the reverse light cones of shallow quantum
830 circuits (Anschuetz & Kiani, 2022).

831 Given a choice of ansatz, the goal of a QNN is to minimize an empirical risk of the form:

$$832 f(\boldsymbol{\theta}) = \frac{1}{|\mathcal{R}|} \sum_{\rho \in \mathcal{R}} \ell(\boldsymbol{\theta}; \rho) = \frac{1}{|\mathcal{R}|} \sum_{\rho \in \mathcal{R}} \text{Tr}(\mathbf{U}(\boldsymbol{\theta}) \rho \mathbf{U}(\boldsymbol{\theta})^\dagger \mathbf{O}). \quad (36)$$

833 Here, \mathcal{R} can be thought of a data set comprising multiple input states ρ , and ℓ the loss function. His-
834 torically, when $|\mathcal{R}| = 1$ QNNs have been referred to as *variational quantum algorithms* (VQAs) (Pe-
835 ruzzo et al., 2014) due to their connection to finding variational approximations to the ground states
836 of quantum Hamiltonians. There are known quantum-classical separations for the expressivity of
837 quantum neural networks even when taking into account the requirement that the training procedure
838 is efficient (Liu et al., 2021; Hastings & O’Donnell, 2022; Huang et al., 2024), though they require
839 very specific training algorithms that take advantage of the structure of the data.

840 There has been recent hope that, with enough ansatz structure, efficient training may follow just
841 via a simple application of gradient descent. This follows from the *Lie algebra-supported ansatz*
842 (LASA) literature, where it has been shown that if the generators $i\mathbf{A}_i$ of the ansatz \mathbf{U} as well as the
843 (scaled) objective observable $i\mathbf{O}$ belong to the same Lie algebra \mathfrak{g} —called the *dynamical Lie alge-*
844 *bra* (Larocca et al., 2022)—generating G , gradients scale inversely with the dimension of \mathfrak{g} (Fontana
845 et al., 2024; Ragone et al., 2024). It is also conjectured that when $\dim(\mathfrak{g})$ scales polynomially with
846 the system size there exist polynomial-depth ansatzes that do not have poor local minima, which
847 together with the large gradients would imply efficient trainability of these loss functions via gra-
848 dient descent. Conditioned on this conjecture there have been results demonstrating expressivity
849 separations in quantum machine learning where the QNN is efficiently trainable through a simple
850 application of gradient descent (Anschuetz et al., 2023b; Anschuetz & Gao, 2024).

851 Though the LASA framework gives sufficient conditions for loss functions to have large gradients,
852 it is known that they are not necessary. One example of this was demonstrated in Diaz et al. (2023),
853 where it was shown that parameterized matchgate circuits with an objective observable given by
854 constant-degree polynomials in Majorana fermions are efficiently trainable though they are not a
855 part of the LASA setting. We claim that both settings are special cases of a Jordan algebraic under-
856 standing of variational loss functions. In preparation of discussing this connection we now review
857 Jordan algebras.

864	\mathbb{N}_0	Natural numbers including 0
865	\mathbb{N}_1	Natural numbers excluding 0
866	$[n]$	Natural numbers from 1 through n
867	$\mathcal{H}^n(\mathbb{F})$	Jordan algebra of $n \times n$ Hermitian matrices over \mathbb{F}
868	\mathbb{F}	Field, here one of $\mathbb{R}, \mathbb{C}, \mathbb{H}$
869	β	$\beta = 1, 2, 4$ when associated field $\mathbb{F} = \mathbb{R}, \mathbb{C}, \mathbb{H}$, respectively
870	$\dim_{\mathbb{F}}(\cdot)$	Dimension of \cdot as a vector space over \mathbb{F}
871	$\mathcal{N}^{\beta}(0, 1)$	Standard normal distribution over \mathbb{F} (given by β)
872	$\mathcal{W}_n^{\beta}(r, \Sigma)$	β -Wishart distribution of $n \times n$ matrices with r degrees of freedom and scale matrix Σ
873	\rightsquigarrow	Convergence in distribution
874	\xrightarrow{p}	Convergence in probability
875	\odot	Hadamard product
876	\mathcal{A}_{α}	Simple components of semisimple Jordan algebra \mathcal{A}
877	G_{α}	Lie group isomorphic to a connected component of $\text{Aut}(\mathcal{A}_{\alpha})$
878	\mathfrak{g}_{α}	Lie algebra generating G_{α}
879	\cdot^{α}	Defining representation of the projection of \cdot into \mathcal{A}_{α}
880	N_{α}	Dimension of vector space on which the defining representation of \mathcal{A} acts
881	N	Sum of N_{α}
882	$\text{Tr}_{\alpha}(\cdot)$	Trace of \cdot in the defining representation of \mathcal{A}_{α}
883	$\text{Tr}(\cdot)$	Trace of \cdot according to its representation in $\mathcal{H}^N(\mathbb{C}) \supseteq \mathcal{A}$
884	$\ \cdot\ _{\text{op}}$	Operator norm of $\cdot \in \mathcal{A}_{\alpha}$ in its defining representation
885	$\ \cdot\ _*$	Trace (nuclear) norm of $\cdot \in \mathcal{A}_{\alpha}$ in its defining representation
886	$\ \cdot\ _{\mathbb{F}}$	Frobenius norm of $\cdot \in \mathcal{A}_{\alpha}$ in its defining representation
887	$\mathcal{O}(\text{poly}(N))$	$\mathcal{O}(N^{\kappa})$ for some constant $\kappa > 0$
888	WLOG	without loss of generality
889	w.h.p.	with high probability
890		

891 **Table 2: Table of notation.** Notation used in the presentation of our results are given in the left
892 column with corresponding meaning given in the right column. We also note here that exponents
893 written as decimals throughout our results are chosen somewhat arbitrarily, i.e., can be improved
894 by any sufficiently small constant $\epsilon > 0$ in the exponent. Finally, we note that we will often forgo
895 writing “... a sequence of [ansatzes, loss functions, observables, QNNs, ...]....” when describing
896 our asymptotic convergence results for brevity.

897 A.2 JORDAN ALGEBRAS

898 A *Jordan algebra* over the reals is formally a real vector space V with a commutative multiplication
899 operation \circ acting on $u, v \in V$ satisfying the *Jordan identity*:

$$900 u \circ ((u \circ u) \circ v) = (u \circ u) \circ (u \circ v), \quad (37)$$

901 which ensures the associativity of the power. A simple example of a Jordan algebra over the reals
902 is the real algebra $\mathcal{H}^N(\mathbb{C})$ of $N \times N$ complex-valued Hermitian matrices with \circ given by half the
903 anticommutator. In particular, for any finite-dimensional quantum system both \mathcal{O} and the inputs ρ
904 in Eq. (36) belong to $\mathcal{H}^N(\mathbb{C})$. We emphasize that though this algebra is typically written in terms
905 of complex matrices it is still a *real* Jordan algebra. This is because, for instance, $i \notin \mathbb{R}$ multiplying
906 a Hermitian matrix is no longer Hermitian. In fact, the Jordan algebra of Hermitian matrices is
907 *Euclidean* (or *formally real*) as it satisfies the defining property (Koecher, 1999c):

$$908 \forall A, B \in \mathcal{A}, A \circ A + B \circ B = 0 \iff A = B = 0. \quad (38)$$

909 It is apparent that all subalgebras of a formally real algebra are also formally real.

910 The real vector space V with which a Jordan algebra is associated also has a natural linear transfor-
911 mation $L(u)$:

$$912 L(u)v \equiv u \circ v, \quad (39)$$

913 which can be viewed as the Jordan algebraic analogue of the adjoint representation of Lie algebras.
914 This linear transformation gives rise to the *canonical trace form*:

$$915 \tau(u, v) \equiv \text{Tr}(L(u \circ v)). \quad (40)$$

When τ is nonsingular we call the associated Jordan algebra \mathcal{A} *semisimple*. As an example, for the algebra $\mathcal{H}^N(\mathbb{F})$ of Hermitian matrices over the field \mathbb{F} , $\tau(u, v)$ is just the Frobenius inner product between u and v . Just as is the case for semisimple Lie algebras, all symmetric bilinear forms on a semisimple Jordan algebra are identical up to an overall scaling. Specifically, for any symmetric bilinear form σ , there exists an element z in the center of \mathcal{A} such that (Theorem 10, Koecher (1999a)):

$$\sigma(u, v) = \tau(z \circ u, v). \quad (41)$$

For $\mathcal{A} = \mathcal{H}^N(\mathbb{C})$ in the defining representation the center is just real multiples of the identity matrix, i.e., there exists a real $I \geq 0$ such that:

$$I\sigma(u, v) = \tau(u, v). \quad (42)$$

We call I the *index* of σ in analogy with Fontana et al. (2024) for Lie algebras.

Though perhaps not as famous as the classification of compact Lie groups, the semisimple Euclidean Jordan algebras have also been classified.

Theorem 8 (Classification of semisimple Euclidean Jordan algebras (Koecher, 1999d)). *Any semisimple Euclidean Jordan algebra is isomorphic to a direct sum of the simple Euclidean Jordan algebras:*

- \mathcal{L}_N for $N \geq 1$: the spin factor, with vector space equal to \mathbb{R}^N and \circ given by the operation

$$\mathbf{x} \circ \mathbf{y} = (\mathbf{x} \cdot \mathbf{y}, x_1 \bar{y} + y_1 \bar{x}), \quad (43)$$

where $\bar{\cdot}$ denotes \cdot with the first coordinate projected out;

- $\mathcal{H}^N(\mathbb{R})$ for $N \geq 3$: $N \times N$ symmetric matrices over \mathbb{R} , with \circ half the anticommutator;
- $\mathcal{H}^N(\mathbb{C})$ for $N \geq 3$: $N \times N$ Hermitian matrices over \mathbb{C} , with \circ half the anticommutator;
- $\mathcal{H}^N(\mathbb{H})$ for $N \geq 3$: $N \times N$ Hermitian matrices over \mathbb{H} , with \circ half the anticommutator;
- $\mathcal{H}^3(\mathbb{O})$: 3×3 Hermitian matrices over \mathbb{O} , with \circ half the anticommutator.

We use the term *defining representation* to speak of the described matrix representations of these algebras. As the Hermitian octonion case is not an infinite family it is often called *exceptional*. We here are interested in asymptotic sequences of Jordan algebras so we will not be considering the exceptional case.

Every Jordan algebra \mathcal{A} has associated with it an automorphism group $\text{Aut}(\mathcal{A})$. As an example, for $\mathcal{H}^N(\mathbb{R})$ in the defining representation this is just given by the action of conjugation (under the usual matrix multiplication) by orthogonal matrices. More generally, the nonexceptional simple cases have automorphism groups (Orlitzky, 2024):

- $\text{Aut}(\mathcal{L}_N)$: left-action of $\{1\} \times \text{O}(N-1)$;
- $\text{Aut}(\mathcal{H}^N(\mathbb{R}))$: conjugation action of $\text{PO}(N)$;
- $\text{Aut}(\mathcal{H}^N(\mathbb{C}))$: disjoint union of the conjugation action of $\text{PU}(N)$, and the transpose followed by the conjugation action of $\text{PU}(N)$;
- $\text{Aut}(\mathcal{H}^N(\mathbb{H}))$: conjugation action of $\text{PSp}(N)$.

We are here primarily interested in the connected component $\text{Aut}_1(\cdot) \subseteq \text{Aut}(\cdot)$ containing the identity transformation, so we will only concern ourselves with the path-connected automorphism subgroups:

- $\text{SO}(N-1) \cong \text{Aut}_1(\mathcal{L}_N) \subset \text{Aut}(\mathcal{L}_N)$;
- $\text{SO}(N) \cong \text{Aut}_1(\mathcal{H}^N(\mathbb{R})) \subset \text{Aut}(\mathcal{H}^N(\mathbb{R}))$;
- $\text{SU}(N) \cong \text{Aut}_1(\mathcal{H}^N(\mathbb{C})) \subset \text{Aut}(\mathcal{H}^N(\mathbb{C}))$;
- $\text{Sp}(N) \cong \text{Aut}_1(\mathcal{H}^N(\mathbb{H})) \subset \text{Aut}(\mathcal{H}^N(\mathbb{H}))$.

Surprisingly, the automorphism groups of semisimple Jordan algebras are also classified. Given a decomposition of a semisimple Euclidean Jordan algebra \mathcal{A} into simple components:

$$\mathcal{A} \cong \bigoplus_{\alpha} \mathcal{A}_{\alpha}, \quad (44)$$

the connected component of $\text{Aut}_1(\mathcal{A}) \subseteq \text{Aut}(\mathcal{A})$ containing the identity is isomorphic to the direct product (Theorem 10, Koecher (1999b)):

$$\text{Aut}_1(\mathcal{A}) \cong \prod_{\alpha} \text{Aut}_1(\mathcal{A}_{\alpha}). \quad (45)$$

A.3 ϵ -APPROXIMATE t -DESIGNS OVER $\text{Aut}_1(\mathcal{A})$

In order to discuss the structure of the loss function in any detail we will need to consider some choice of randomness over loss functions. To achieve this we will use ϵ -approximate t -designs over $\text{Aut}_1(\mathcal{A})$. Our use of these designs formalizes the notion of approximate “independence” or “uniform initialization” of the ansatz with respect to the eigenbasis of a given objective observable, as when $\epsilon \rightarrow 0$ and $t \rightarrow \infty$ the ansatz is chosen in a completely group-invariant way over $\text{Aut}_1(\mathcal{A})$. By Eq. (45) these designs are (up to isomorphism) direct products of ϵ -approximate t -designs over $\text{SO}(N)$, $\text{SU}(N)$, and $\text{Sp}(N)$.

Before defining ϵ -approximate t -designs we define the Haar measure on compact Lie groups such as $\text{Aut}_1(\mathcal{A})$. This is the unique group-invariant normalized measure on these compact Lie groups.

Definition 9 (Haar measure on G (Haar, 1933)). Let G be a compact Lie group. The unique measure satisfying:

$$\int_G d\mu = 1 \quad (46)$$

as well as

$$d\mu = d(g\mu) \quad (47)$$

for all $g \in G$ is the *Haar measure on G* . The existence and uniqueness of this measure follows from Haar (1933).

We now define ϵ -approximate t -designs. We will here use the “trace norm definition” out of convenience, though all of the commonly used definitions are roughly equivalent (Harrow & Mehraban, 2023).

Definition 10 (ϵ -approximate t -designs over G). Let G be a compact Lie group with defining representation over an N -dimensional space, and μ the Haar measure over G . A measure ν satisfying:

$$\sum_m |\mathbb{E}_{U \sim \mu} [m(U)] - \mathbb{E}_{U \sim \nu} [m(U)]| \leq \epsilon, \quad (48)$$

where the sum is over all degree- (t, t) monomials $m(U)$ in the entries of U and U^* , is an ϵ -approximate t -design over G .

A.4 WISHART MATRICES

Our main results will be given in terms of Wishart-distributed random matrices, so before proceeding we give a brief review of this distribution. We will use $\mathcal{W}_n^{\beta}(r, \Sigma)$ to denote the β -Wishart distribution (Dubbs et al., 2013), where:

- $\beta = 1, 2, 4$ indicates it is over a field $\mathbb{R}, \mathbb{C}, \mathbb{H}$, respectively;
- $r \in \mathbb{N}_1$ is the *degrees of freedom* parameter of the distribution;
- $\Sigma \in \mathcal{H}^n(\mathbb{R})$ is the symmetric, real-valued, positive-definite *scale matrix* parameter of the distribution.

When we do not specify β we implicitly are referring to the usual Wishart distribution where $\beta = 1$. The β -Wishart distribution is over positive semidefinite matrices of rank r , constructed by drawing i.i.d. standard Gaussian entries over \mathbb{F} :

$$X_{i,j} \sim \mathcal{N}^{\beta}(0, 1) \quad (49)$$




Figure 2: **Marčenko–Pastur distribution.** The density of the Marčenko–Pastur distribution—the asymptotic empirical eigenvalue distribution of normalized Wishart matrices—in the regime where $\gamma \ll 1$ and where $\gamma = 1$. At $\gamma = 1$ the associated Wishart matrix transitions from being full-rank to being low-rank.

for $i \in [n]$ and $j \in [r]$, and considering the distribution of:

$$\mathbf{W} = \sqrt{\Sigma} \mathbf{X} \mathbf{X}^\dagger \sqrt{\Sigma}. \quad (50)$$

Here there exists the so-called *Bartlett decomposition* for $\mathbf{W} \sim \mathcal{W}_n^\beta(r, \mathbf{I}_n)$:

$$\mathbf{W} \sim \sqrt{\Sigma} \mathbf{L} \mathbf{L}^\dagger \sqrt{\Sigma}, \quad (51)$$

where \mathbf{L} is lower-triangular. When $r \geq n$, the entries below the diagonal of \mathbf{L} are i.i.d. $\mathcal{N}^\beta(0, 1)$ -distributed and the diagonal entries are i.i.d. distributed as (Rouault, 2007):

$$\beta L_{i,i}^2 \sim \chi^2(\beta(r - i + 1)), \quad (52)$$

i.e., the $L_{i,i}$ are χ -distributed up to an overall scaling of $\beta^{-\frac{1}{2}}$. When $r < n$, \mathbf{L} is $n \times r$ with first r rows as above and all other entries i.i.d. according to $\mathcal{N}^\beta(0, 1)$ (Srivastava, 2003; Li & Xue, 2010; Yu et al., 2014). This completely characterizes the marginal distribution of the entries of any $\mathbf{W} \sim \mathcal{W}_n^\beta(r, \mathbf{I}_n)$.

Our final results will be written in terms of real Wishart matrices \mathbf{W} , i.e., those with $\beta = 1$. When we analyze the asymptotic density of local minima of a JAWS it will turn out we will need to consider the asymptotic spectrum of a real Wishart matrix, which we now discuss. When the scale matrix Σ of \mathbf{W} is the identity, as $n, r \rightarrow \infty$ with $\frac{n}{r} = \gamma$ held constant, the spectrum of $r^{-1}\mathbf{W}$ is known to almost surely converge weakly to a fixed distribution (Marčenko & Pastur, 1967). This fixed distribution is the *Marčenko–Pastur distribution*, given by

$$\frac{d\mu_{\text{MP}}^\gamma}{d\lambda} = (1 - \gamma^{-1}) \delta(\lambda) \mathbf{1}\{\gamma \geq 1\} + \frac{\sqrt{(\gamma_+ - \lambda)(\lambda - \gamma_-)}}{2\pi\gamma\lambda} \mathbf{1}\{\gamma_- \leq \lambda \leq \gamma_+\}, \quad (53)$$

where

$$\gamma_\pm \equiv (1 \pm \sqrt{\gamma})^2 \quad (54)$$

and $\mathbf{1}$ is the indicator function. This distribution is illustrated in Figure 2.

A.5 CONVERGENCE OF RANDOM VARIABLES

We finally discuss in more detail the various notions of convergence of random variables that we use throughout our paper. Given a sequence of a set $(X_i^N)_i$ of real-valued random variables as $N \rightarrow \infty$, we say $(X_i^N)_i$ *weakly converges* (or *converges in distribution*) to $(X_i^\infty)_i$ if the joint cumulative distribution function F^N of the X_i^N converges to the joint cumulative distribution function F^∞ of the X_i^∞ at every point at which F^∞ is continuous. That is,

$$\lim_{N \rightarrow \infty} F^N(\mathbf{x}) = F^\infty(\mathbf{x}) \quad (55)$$

for all points x at which F^∞ is continuous. We denote this at the level of random variables with the notation:

$$(X_i^N)_i \rightsquigarrow (X_i^\infty)_i. \quad (56)$$

One way to quantify the rate of this convergence is through the use of the *Lévy–Prokhorov metric*. Letting p^N and p^∞ be the densities associated with F^N and F^∞ , respectively, the Lévy–Prokhorov metric is given by:

$$\pi(p^N, p^\infty) = \inf \{ \epsilon > 0 \mid \forall A, F^N(A) \leq F^\infty(A^\epsilon) + \epsilon \}, \quad (57)$$

where here A^ϵ is an ϵ -neighborhood in ∞ -norm of A . $\pi(p^N, p^\infty) \rightarrow 0$ if and only if $(X_i^N)_i \rightsquigarrow (X_i^\infty)_i$.

Convergence in distribution is closely related to the pointwise convergence of probability densities. Indeed, the latter implies the former by Scheffé’s theorem (Scheffé, 1947). The former implies the latter if the p^N are equicontinuous and uniformly bounded as $N \rightarrow \infty$ (Boos, 1985).

We finally discuss *convergence in probability*. This is the statement that, for all $\epsilon > 0$,

$$\lim_{N \rightarrow \infty} \mathbb{P}(\|\mathbf{X}^N - \mathbf{X}^\infty\|_\infty > \epsilon) = 0. \quad (58)$$

We denote this convergence using the notation:

$$(X_i^N)_i \xrightarrow{p} (X_i^\infty)_i. \quad (59)$$

One way to quantify convergence in probability is through the *Ky Fan metric*, which is given by:

$$\alpha(p^N, p^\infty) = \inf \{ \epsilon > 0 \mid \mathbb{P}(\|\mathbf{X}^N - \mathbf{X}^\infty\|_\infty > \epsilon) \leq \epsilon \}. \quad (60)$$

$\alpha(p^N, p^\infty) \rightarrow 0$ if and only if $(X_i^N)_i \xrightarrow{p} (X_i^\infty)_i$. The Ky Fan metric upper bounds the Lévy–Prokhorov metric (Strassen, 1965), consistent with the notion of convergence in probability implying convergence in distribution.

We will in what follows often be loose with language and say two sequences of distributions converge to one another at a rate K , either weakly or in probability, rather than state convergence to a fixed distribution. This should be understood to mean that their distance in associated metric asymptotically vanishes as K^{-1} .

B FORMAL DISCUSSION OF THE MAIN RESULTS

B.1 THE JORDAN ALGEBRAIC STRUCTURE OF VARIATIONAL LOSS LANDSCAPES

With these preliminaries in place we can write the variational loss function given in Eq. (36) as the canonical trace form on $\mathcal{H}^N(\mathbb{C})$:

$$\ell(\boldsymbol{\theta}; \boldsymbol{\rho}) = \tau(\boldsymbol{\rho}, T(\boldsymbol{\theta})\mathbf{O}), \quad (61)$$

where here $\boldsymbol{\rho}, \mathbf{O} \in \mathcal{H}^N(\mathbb{C})$, and $T(\boldsymbol{\theta})$ parameterizes some path-connected compact Lie group \mathcal{T} . Note that $\{T(\boldsymbol{\theta})\mathbf{O}\}_\boldsymbol{\theta}$ generates a Jordan subalgebra $\mathcal{A} \subseteq \mathcal{H}^N(\mathbb{C})$ such that $\mathbf{O} \in \mathcal{A}$ and $\mathcal{T} \subseteq \text{Aut}_1(\mathcal{A})$. As $T(\boldsymbol{\theta})\mathbf{O} \in \mathcal{A}$ for all $\boldsymbol{\theta}$, the trace form is zero for components of $\boldsymbol{\rho}$ orthogonal to \mathcal{A} . Because of this we will often also consider $\boldsymbol{\rho}$ an element of \mathcal{A} in a slight abuse of notation.

We call a variational loss function of this form a *Jordan algebra-supported ansatz* (JASA), in analogy with the term *Lie-algebra supported ansatz* (LASA) introduced in Fontana et al. (2024). However, whereas variational loss functions are LASAs only if $i\mathbf{O}$ belongs to the dynamical Lie algebra generating the ansatz, *all* variational loss functions are JASAs (up to assuming the path-connectedness of \mathcal{T}). Indeed, in Appendix H we give a direct mapping from both LASAs and the variational matchgate formalism of Diaz et al. (2023) to JASAs and show that a mapping in the other direction is generally not possible.

We are often interested in the case where our variational loss landscape has some sort of structure, i.e., when $\mathcal{A} \cong \bigoplus_\alpha \mathcal{A}_\alpha \neq \mathcal{H}^N(\mathbb{C})$. In Appendix G we demonstrate that in the context of examining variational loss landscapes, the spin factor sectors ($\mathcal{A}_\alpha = \mathcal{L}_{N_\alpha}$) effectively reduce to real symmetric ($\mathcal{A}_\alpha = \mathcal{H}^{N_\alpha}(\mathbb{R})$) sectors. Because of this, we will from here on out focus on the case when \mathcal{A}_α is

the Jordan algebra of Hermitian matrices over a field, i.e., $\mathcal{A}_\alpha \cong \mathcal{H}^{N_\alpha}(\mathbb{F}_\alpha)$ for $\mathbb{F}_\alpha = \mathbb{R}, \mathbb{C}, \mathbb{H}$. We will also throughout use $\beta_\alpha = 1, 2, 4$ to label these three cases, respectively, as is commonly done in the physics literature. We use $\text{Tr}_\alpha(\cdot)$ to denote the trace of \cdot in the defining representation of \mathcal{A}_α and $\text{Tr}(\cdot)$ to denote the trace of \cdot according to its representation in $\mathcal{H}^N(\mathbb{C}) \supseteq \mathcal{A}$.

Recalling from Appendix A.2 the universality of the trace form and the direct product decomposition of $\text{Aut}_1(\mathcal{A})$, we can rewrite the loss function of Eq. (61) as:

$$\ell(\boldsymbol{\theta}; \boldsymbol{\rho}) = \sum_{\alpha} \ell^{\alpha}(\boldsymbol{\theta}; \boldsymbol{\rho}) \equiv \sum_{\alpha} I_{\alpha} \tau^{\alpha}(\boldsymbol{\rho}^{\alpha}, T^{\alpha}(\boldsymbol{\theta}) \mathbf{O}^{\alpha}), \quad (62)$$

where here \cdot^{α} is used to denote the component of \cdot in \mathcal{A}_{α} , τ^{α} is the canonical trace form on \mathcal{A}_{α} , and $I_{\alpha} > 0$ is the constant due to Eq. (42).

We now give an explicit form for $T^{\alpha}(\boldsymbol{\theta})$. Recall that elements of $\text{Aut}_1(\mathcal{A}_{\alpha})$ are, in the defining representation of \mathcal{A}_{α} , given by the conjugation action by elements of $\text{SO}(N_{\alpha})$, $\text{SU}(N_{\alpha})$, or $\text{Sp}(N_{\alpha})$ when $\mathbb{F}_{\alpha} = \mathbb{R}, \mathbb{C}, \mathbb{H}$, respectively. We let G_{α} denote the corresponding Lie group and \mathfrak{g}_{α} the associated Lie algebra. We also denote G as the direct product over G_{α} , and \mathfrak{g} the corresponding Lie algebra. Following the typical structure of variational quantum algorithms (Peruzzo et al., 2014) we will assume $T^{\alpha}(\boldsymbol{\theta})$ in the defining representation corresponds to conjugation by:

$$U_{g,h}^{\alpha}(\boldsymbol{\theta}) = g_0^{\alpha \dagger} \left(\prod_{i=1}^p g_i^{\alpha} \exp(\theta_i A_i^{\alpha}) g_i^{\alpha \dagger} \right) h^{\alpha} \in G_{\alpha} \quad (63)$$

for some $g_i, h \in G$ and $A_i \in \mathfrak{g}$, where \cdot^{α} denotes the projection onto G_{α} or \mathfrak{g}_{α} appropriately. With this choice of ansatz it is also easy to see the derivatives of ℓ have algebraic interpretations. For instance, at $\boldsymbol{\theta} = \mathbf{0}$, $i \geq j$,

$$\partial_i \ell(\boldsymbol{\theta}; \boldsymbol{\rho}) = \sum_{\alpha} I_{\alpha} \tau^{\alpha} \left(g_0^{\alpha} \boldsymbol{\rho}^{\alpha} g_0^{\alpha \dagger}, \left[g_i^{\alpha} A_i^{\alpha} g_i^{\alpha \dagger}, h^{\alpha} \mathbf{O}^{\alpha} h^{\alpha \dagger} \right] \right), \quad (64)$$

$$\partial_i \partial_j \ell(\boldsymbol{\theta}; \boldsymbol{\rho}) = \sum_{\alpha} I_{\alpha} \tau^{\alpha} \left(g_0^{\alpha} \boldsymbol{\rho}^{\alpha} g_0^{\alpha \dagger}, \left[g_j^{\alpha} A_j^{\alpha} g_j^{\alpha \dagger}, \left[g_i^{\alpha} A_i^{\alpha} g_i^{\alpha \dagger}, h^{\alpha} \mathbf{O}^{\alpha} h^{\alpha \dagger} \right] \right] \right), \quad (65)$$

where the commutator action of the automorphism group of a Jordan algebra is canonically defined through its representation L defined in Eq. (39) (see Lemma 7 of Koecher (1999b)).

B.2 JORDAN ALGEBRAIC WISHART SYSTEMS

We have demonstrated that all variational loss landscapes are Jordan algebra-supported ansatzes (JASAs). We are finally ready to discuss our main results, which give an explicit expression for the loss landscape of a JASA when its ansatz takes the form of Eq. (63). To do this we first define a *Jordan algebraic Wishart system* (JAWS), leaving for now ambiguous the connection to Wishart matrices.

Definition 11 (Jordan algebraic Wishart system). Let O belong to a Jordan algebra $\mathcal{A} \subseteq \mathcal{H}^N(\mathbb{C})$ and let \mathcal{T} be a path-connected subspace $\mathcal{T} \subseteq \text{Aut}_1(\mathcal{A})$. Let

$$\mathcal{A} \cong \bigoplus_{\alpha} \mathcal{A}_{\alpha} \quad (66)$$

be the decomposition of \mathcal{A} into simple Euclidean components as in Theorem 8 with associated decomposition $\mathcal{T} \cong \times_{\alpha} \mathcal{T}_{\alpha}$ as in Eq. (45). Let $\mathcal{G}_{\alpha}^{\alpha}, \mathcal{H}^{\alpha}$ be independent distributions over \mathcal{T}_{α} , and $\mathbf{A} = \{A_i\}_{i=1}^p$ elements of the Lie algebra of $\text{Aut}_1(\mathcal{A})$. We call

$$\mathcal{J} = (\mathcal{A}, \mathcal{T}, \mathcal{G}, \mathcal{H}, \mathbf{A}, O) \quad (67)$$

a *Jordan algebraic Wishart system* (JAWS).

Every JAWS has associated with it a loss function as discussed in Appendix B.1. Using \cdot^{α} to label the projection of \cdot into \mathcal{A}_{α} in its defining representation, the associated loss function takes the form:

$$\ell(\boldsymbol{\theta}; \boldsymbol{\rho}) = \sum_{\alpha} \ell^{\alpha}(\boldsymbol{\theta}; \boldsymbol{\rho}) = \sum_{\alpha} I_{\alpha} \text{Tr}_{\alpha} \left(\boldsymbol{\rho}^{\alpha} U^{\alpha}(\boldsymbol{\theta}) \mathbf{O}^{\alpha} U^{\alpha}(\boldsymbol{\theta})^{\dagger} \right), \quad (68)$$

where $U^\alpha(\boldsymbol{\theta}) \in G_\alpha$ for all $\boldsymbol{\theta}$ and is of the form:

$$U^\alpha(\boldsymbol{\theta}) \equiv g_0^{\alpha\dagger} \left(\prod_{i=1}^p g_i^\alpha \exp(\theta_i A_i^\alpha) g_i^{\alpha\dagger} \right) h^\alpha. \quad (69)$$

Here, g_i^α is the defining representation of an element drawn from the distribution \mathcal{G}_i^α , and h^α is similar for \mathcal{H}^α .

Our main result is a concise, asymptotic description of $\ell(\boldsymbol{\theta}; \boldsymbol{\rho})$ when the ansatz is chosen “sufficiently independently” from \mathcal{O} (up to respecting the algebra \mathcal{A}). More formally, the rate at which $\ell(\boldsymbol{\theta}; \boldsymbol{\rho})$ converges to its asymptotic limit depends on the parameters defined by the following assumption.

Assumption 12 (ϵ -approximate t -design of observable basis). Each \mathcal{H}^α is an ϵ -approximate t -design over $\mathcal{T}_\alpha = \text{Aut}_1(\mathcal{A}_\alpha)$.

If there is a sense of geometric locality in the system, i.e., if the reverse lightcone of \mathcal{O}^α is of the form:

$$U^\alpha(\boldsymbol{\theta}) \mathcal{O}^\alpha U^\alpha(\boldsymbol{\theta})^\dagger = \tilde{\mathcal{O}}^\alpha \otimes \mathbf{I} \quad (70)$$

for some $\tilde{\mathcal{O}}^\alpha$ acting on a space of dimension at most $D \equiv \exp(O(p^d))$ for constant geometric dimension d , one can explicitly enforce this with only a modest overhead in circuit depth. Indeed, one may embed \mathcal{A}_α into qubits and use the construction of Harrow & Mehraban (2023) to give a d -dimensional random circuit of depth $O(\text{poly}(t)p)$ that achieves this over $\text{Aut}_1(\mathcal{H}^{2^n}(\mathbb{C})) \supset \text{Aut}_1(\mathcal{A}_\alpha)$ for some $n = O(p^d)$ and $\epsilon = \exp(-\Omega(p))$.

Given Assumption 12 we are able to prove our main result on the convergence of loss functions. We are able to prove our convergence for any N -dependent overall normalization of the loss function \mathcal{N} , the only requirement being that central moments of sufficiently large (constant) order have a finite limit as $N \rightarrow \infty$. This is always true for $\mathcal{N} = 1$ and, depending on the specific choice of t -design, potentially holds for any $\mathcal{N} \leq O(N^{1-\delta})$ where $\delta > 0$ is constant.

Theorem 13 (Loss function distribution). *Let \mathcal{J} be a JAWS satisfying Assumption 12 with loss function $\ell(\boldsymbol{\theta}; \boldsymbol{\rho})$ for $\boldsymbol{\rho} \in \mathcal{R}$. Further assume that*

$$\text{Tr}((\boldsymbol{\rho}^\alpha)^2) = \Omega\left(\frac{1}{N_\alpha^{0.999}}\right). \quad (71)$$

Fix $\boldsymbol{\theta}$. Assume $\mathcal{N} \leq O(\min_\alpha N_\alpha^{0.99})$ is such that there exists a constant k^ such that the central moments of order $k > k^*$ of $(\mathcal{N}\ell(\boldsymbol{\theta}; \boldsymbol{\rho}))_{\boldsymbol{\rho} \in \mathcal{R}}$ are finite as all $\epsilon^{-1}, t, N_\alpha \rightarrow \infty$. We have the convergence in joint distributions:*

$$(\mathcal{N}\ell(\boldsymbol{\theta}; \boldsymbol{\rho}))_{\boldsymbol{\rho} \in \mathcal{R}} \rightsquigarrow (\mathcal{N}\hat{\ell}(\boldsymbol{\theta}; \boldsymbol{\rho}))_{\boldsymbol{\rho} \in \mathcal{R}} \quad (72)$$

as all $\epsilon^{-1}, t, N_\alpha \rightarrow \infty$, where

$$\hat{\ell}(\boldsymbol{\theta}; \boldsymbol{\rho}) \equiv \sum_\alpha \hat{\ell}^\alpha(\boldsymbol{\theta}; \boldsymbol{\rho}) \equiv \sum_\alpha \frac{I_\alpha \bar{\sigma}^\alpha}{r_\alpha} \text{Tr}_\alpha(\boldsymbol{\rho}^\alpha \mathbf{W}^\alpha), \quad (73)$$

$\bar{\sigma}^\alpha$ is the arithmetic mean eigenvalue of \mathcal{O}^α in the defining representation:

$$\bar{\sigma}^\alpha = N_\alpha^{-1} \text{Tr}_\alpha(\mathcal{O}^\alpha), \quad (74)$$

and the \mathbf{W}^α are independent Wishart-distributed random matrices with $\beta_\alpha r_\alpha$ degrees of freedom, where

$$r_\alpha \equiv \frac{\|\mathcal{O}^\alpha\|_*^2}{\|\mathcal{O}^\alpha\|_F^2}. \quad (75)$$

In particular, the distributions differ in Lévy–Prokhorov metric by

$$\pi = O\left(\frac{\log \log(N^{-t}\epsilon^{-1})}{\log(N^{-t}\epsilon^{-1})} + \frac{\log(t)}{\sqrt{t}} + \sum_\alpha \sqrt{\frac{\mathcal{N} \log(N_\alpha)}{N_\alpha^{1.001}}}\right) = o(1). \quad (76)$$

In Appendix I we show that Eq. (71) can effectively be taken WLOG. This is because all ρ^α for which this is not true are such that $\ell^\alpha(\theta; \rho)$ have asymptotically vanishing variance even when rescaled by the potentially exponentially large $\mathcal{N} = O(N)$. This is why we do not take Eq. (71) as a numbered Assumption. Examples of the loss density are illustrated in Figure 1(a).

We can strengthen the statement of weak convergence in Theorem 13 to one of pointwise convergence of densities given equicontinuity and boundedness of the appropriately normalized loss function density. Whether this is true depends on the specifics of the distributions \mathcal{G}, \mathcal{H} , particularly whether the loss is equicontinuous and bounded at the same $\mathcal{N} = \sqrt{r_\alpha}$ scale as $\hat{\ell}^\alpha$. We give some standard examples where this is true in Appendix K.

Corollary 14 (Convergence of loss function densities). *Let $\mathcal{J}, \ell, \hat{\ell}$ be as in Theorem 13. Assume the density of $(\sqrt{r_\alpha} \ell^\alpha(\theta; \rho))_{\rho \in \mathcal{R}, \alpha}$ is equicontinuous and bounded as all $\epsilon^{-1}, t, N_\alpha \rightarrow \infty$. Then the joint density of $(\sqrt{r_\alpha} \ell^\alpha(\theta; \rho))_{\rho \in \mathcal{R}, \alpha}$ is pointwise equal to that of $(\sqrt{r_\alpha} \hat{\ell}^\alpha(\theta; \rho))_{\rho \in \mathcal{R}, \alpha}$ up to an additive error $O(\pi)$.*

As-written our result only gives the loss function distribution at a single point θ in parameter space. However, any finite set Θ of points in parameter space can be considered by taking

$$\rho(\theta) = U^\alpha(\theta) \rho U^{\alpha\dagger}(\theta) \quad (77)$$

to be elements in an augmented set of input states:

$$\mathcal{R}' = \mathcal{R} \times \Theta. \quad (78)$$

We can give a more concrete form for the parameter-dependence of the loss landscape by considering the joint distribution of ℓ with its derivatives. We begin with the gradient. In order to consider the joint distribution over what can potentially be many gradient components—a number growing with the number of parameters p —we take the following additional assumption on the growth of t with p so we are able to fully capture correlations between the derivatives.

Assumption 15 (Scaling of parameter space dimension with t -design). The number of trained parameters p satisfies:⁵

$$p^2 \leq o\left(\min\left(\frac{\log(\epsilon^{-1})}{\log \log(\epsilon^{-1})}, \frac{\sqrt{t}}{\log(t)}\right)\right), \quad (79)$$

where the \mathcal{G}_i^α are i.i.d. ϵ -approximate t -designs over $\mathcal{T}_\alpha = \text{Aut}_1(\mathcal{A}_\alpha)$, and

$$p \leq O\left(\text{poly}\left(\min_\alpha N_\alpha\right)\right). \quad (80)$$

It is possible to weaken this assumption and instead assume that the \mathcal{G}_i^α are ϵ -approximate 2-designs rather than t -designs, but this is at the expense of requiring $\mathcal{N} \leq o\left(N^{\frac{2}{3}}\right)$ and only considering jointly at most $\sim \log(N)$ components of the gradient. We leave further details of this alternative setting to Appendix J.

It will also be convenient for the rest of our discussion to consider a concrete set of \mathbf{A}_i .

Assumption 16 (Concrete choice of \mathbf{A}_i). For $\mathbb{F}_\alpha = \mathbb{C}, \mathbb{H}$, the \mathbf{A}_i^α are rank-1. For $\mathbb{F}_\alpha = \mathbb{R}$, the \mathbf{A}_i^α are rank-2.

The distinction between the cases $\mathbb{F}_\alpha = \mathbb{C}, \mathbb{H}$ and $\mathbb{F}_\alpha = \mathbb{R}$ is due to there being no rank 1 operators in the defining representation of $\mathfrak{so}(N_\alpha)$. This choice of \mathbf{A}_i^α may seem unphysical due to their nonlocality, but one can emulate the behavior of high-rank (e.g., Pauli) rotations by considering a factor of N_α more layers p_α in a given simple sector—each, for instance, performing a rotation under each eigenvector of \mathbf{A}_i^α —and then tying the associated parameters together. This breaks no other assumptions as taking $p_\alpha \rightarrow N_\alpha p_\alpha$ maintains Eq. (80).

We then have the following theorem. It is stated assuming only a single input ρ which is rank-1 when projected to the defining representation of each simple sector to simplify the final result. However, in Appendix D.3 we give a full, exact expression of the joint distribution in terms of Wishart matrix elements for any choice of $\mathcal{R} \ni \rho$ and independent of the scaling of $N_\alpha - r_\alpha$.

⁵We here use a bound on p^2 so that later we can reuse this assumption for the Hessian; at this stage we really only need this bound for p .

Theorem 17 (Gradient distribution). *Consider the setting of Theorem 13, with the additional Assumptions 15 and 16. Assume $\frac{\sigma_o^\alpha}{\bar{o}}$ is bounded—where σ_o^α is the standard deviation of the eigenvalues of \mathbf{O}^α —and assume $|\mathcal{R}| = 1$ with element $\boldsymbol{\rho}$ such that each $\boldsymbol{\rho}^\alpha$ is rank-1 in its defining representation. We have the convergence in joint distributions:*

$$(\mathcal{N}\ell(\boldsymbol{\theta}; \boldsymbol{\rho}), \mathcal{N}\partial_i \ell(\boldsymbol{\theta}; \boldsymbol{\rho}))_{i \in [p]} \rightsquigarrow \left(\mathcal{N}\hat{\ell}(\boldsymbol{\theta}; \boldsymbol{\rho}), \mathcal{N}\hat{\ell}_{;i}(\boldsymbol{\theta}; \boldsymbol{\rho}) \right)_{i \in [p]} \quad (81)$$

as all $\epsilon^{-1}, t, N_\alpha \rightarrow \infty$, where conditioned on

$$\hat{\ell}^\alpha(\boldsymbol{\theta}; \boldsymbol{\rho}) = z^\alpha \quad (82)$$

we have that

$$\hat{\ell}_{;i|z^\alpha}^\alpha(\boldsymbol{\theta}; \boldsymbol{\rho}) = \frac{2I_\alpha \sigma_o^\alpha \text{Tr}_\alpha(\boldsymbol{\rho}^\alpha)}{N_\alpha} \sqrt{\frac{\beta_\alpha z^\alpha}{I_\alpha \bar{o}^\alpha}} \chi_i^\alpha G_i^\alpha, \quad (83)$$

where

$$\hat{\ell}_{;i}(\boldsymbol{\theta}; \boldsymbol{\rho}) \mid \left\{ \hat{\ell}^\alpha(\boldsymbol{\theta}; \boldsymbol{\rho}) = z^\alpha \right\}_\alpha \equiv \hat{\ell}_{;i|z}(\boldsymbol{\theta}; \boldsymbol{\rho}) \equiv \sum_\alpha \hat{\ell}_{;i|z^\alpha}^\alpha(\boldsymbol{\theta}; \boldsymbol{\rho}). \quad (84)$$

Here, the χ_i^α are independent χ -distributed random variables with $\max(2, \beta_\alpha)$ degrees of freedom and the G_i^α are i.i.d. standard normal random variables. In particular, the distributions differ in Lévy–Prokhorov metric by

$$\pi' = \mathcal{O} \left(\frac{p \log \log(\mathcal{N}^{-t} \epsilon^{-1})}{\log(\mathcal{N}^{-t} \epsilon^{-1})} + \frac{p \log(t)}{\sqrt{t}} + \sum_\alpha \sqrt{\frac{\mathcal{N} \log(N_\alpha)}{N_\alpha^{1.001}} + \frac{\mathcal{N}^{\frac{1}{3}} \log(N)}{N^{\frac{5}{12}}}} \right) = \mathcal{O}(1). \quad (85)$$

This result is stronger than typical barren plateau results as it gives the full distributional form of the gradient—even when conditioned on the loss contribution z^α —rather than just its variance (McClellan et al., 2018; Larocca et al., 2022; Fontana et al., 2024; Ragone et al., 2024). The conditional distribution as in Eq. (83) is illustrated in Figure 1(b).

Just as with the loss function, we can strengthen Theorem 17 to show pointwise convergence in probability densities assuming equicontinuity of the original distribution.

Corollary 18 (Convergence of gradient densities). *Let $\mathcal{J}, \ell, \hat{\ell}$ be as in Corollary 14. Assume the density of $(\sqrt{r_\alpha} \hat{\ell}_{;i}^\alpha(\boldsymbol{\theta}; \boldsymbol{\rho}))_\alpha$ is equicontinuous and bounded for all $z^\alpha \in \mathbb{R}_{\geq 0}$ as all $\epsilon^{-1}, t, N_\alpha \rightarrow \infty$.*

Then the joint density of $(\sqrt{r_\alpha} \hat{\ell}_{;i}^\alpha(\boldsymbol{\theta}; \boldsymbol{\rho}))_\alpha$ is pointwise equal to that of $(\sqrt{r_\alpha} \hat{\ell}_{;i}^\alpha(\boldsymbol{\theta}; \boldsymbol{\rho}))_\alpha$ up to an additive error $\mathcal{O}(\pi')$.

We now give our final result, which specifies the joint distribution of not only the loss and gradient but also the Hessian at critical points. This is required to reason about the critical point distribution of the loss landscape using the so-called Kac–Rice formula (Adler & Taylor, 2007) that we will discuss in detail in Appendix C.3. We state our result assuming $\sigma_o \ll \bar{o}$ for simplicity, giving as we did for the gradient the full expression in terms of Wishart matrix elements in Appendix D.3.

Theorem 19 (Hessian distribution). *Let $\mathcal{J}, \ell, \hat{\ell}$ be as in Theorem 17. Assume all $\frac{\sigma_o^\alpha}{\bar{o}^\alpha} = \mathcal{O}(1)$ as $N_\alpha \rightarrow \infty$. We have the convergence in joint distributions:*

$$(\ell(\boldsymbol{\theta}; \boldsymbol{\rho}), \partial_i \ell(\boldsymbol{\theta}; \boldsymbol{\rho}), \partial_i \partial_j \ell(\boldsymbol{\theta}; \boldsymbol{\rho}))_{\boldsymbol{\rho} \in \mathcal{R}, i, j \in [p]} \rightsquigarrow \left(\hat{\ell}(\boldsymbol{\theta}; \boldsymbol{\rho}), \hat{\ell}_{;i}(\boldsymbol{\theta}; \boldsymbol{\rho}), \hat{\ell}_{;i,j}(\boldsymbol{\theta}; \boldsymbol{\rho}) \right)_{\boldsymbol{\rho} \in \mathcal{R}, i, j \in [p]} \quad (86)$$

as all $\epsilon^{-1}, t, N_\alpha \rightarrow \infty$, where conditioned on

$$\hat{\ell}^\alpha(\boldsymbol{\theta}; \boldsymbol{\rho}) = z^\alpha \quad (87)$$

and

$$\hat{\ell}_{;i|z^\alpha}^\alpha(\boldsymbol{\theta}; \boldsymbol{\rho}) = 0 \quad (88)$$

we have that

$$\hat{\ell}_{;i,j|z^\alpha,0}^\alpha(\boldsymbol{\theta}; \boldsymbol{\rho}) = \frac{2I_\alpha \sigma_o^\alpha \text{Tr}_\alpha(\boldsymbol{\rho}^\alpha)}{N_\alpha^2} \sqrt{\frac{z^\alpha}{I_\alpha \bar{o}^\alpha}} G_i^\alpha \chi_j^\alpha \langle i | \mathbf{W}^\alpha | j \rangle, \quad (89)$$

1350
1351
1352
1353
1354
1355
1356
1357
1358
1359
1360
1361
1362
1363
1364
1365
1366
1367
1368
1369
1370
1371
1372
1373
1374
1375
1376
1377
1378
1379
1380
1381
1382
1383
1384
1385
1386
1387
1388
1389
1390
1391
1392
1393
1394
1395
1396
1397
1398
1399
1400
1401
1402
1403

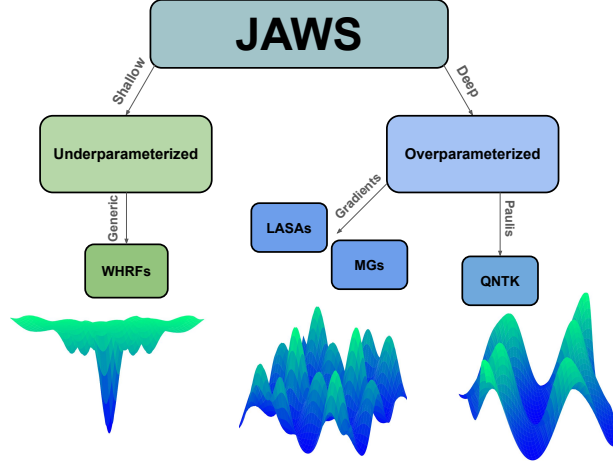


Figure 3: **Relation between models for quantum variational loss landscapes.** Our introduced theory for the loss landscapes of quantum neural networks (QNNs) are *Jordan algebraic Wishart systems* (JAWS), which relate the algebraic structure of a given QNN architecture to an asymptotic random process description of the loss landscape. This JAWS description reduces to the previously-studied Wishart hypertoroidal random fields (WHRFs) and quantum neural tangent kernel (QNTK) in different settings, and also reproduces the loss function variance known for Lie algebra-supported ansatzes (LASAs) and matchgate (MG) networks. Cartoons of the loss landscapes associated with previously studied models are shown.

where

$$\hat{\ell}_{i,j}^{\alpha}(\boldsymbol{\theta}; \boldsymbol{\rho}) \mid \left\{ \hat{\ell}^{\alpha}(\boldsymbol{\theta}, \boldsymbol{\rho}) = z^{\alpha}, \left(\hat{\ell}_{i,i}^{\alpha}(\boldsymbol{\theta}; \boldsymbol{\rho}) \right)_{i \in [p]} = \mathbf{0} \right\}_{\alpha} \equiv \hat{\ell}_{i,j|z,0}^{\alpha}(\boldsymbol{\theta}; \boldsymbol{\rho}) \equiv \sum_{\alpha} \hat{\ell}_{i,j|z^{\alpha},0}^{\alpha}(\boldsymbol{\theta}; \boldsymbol{\rho}). \quad (90)$$

Here, the G_i^{α} are i.i.d. standard normal random variables, the χ_i^{α} are independent χ -distributed random variables with $\max(2, \beta_{\alpha})$ degrees of freedom, and the \mathbf{W}^{α} are independent Wishart-distributed random matrices with $\beta_{\alpha} r_{\alpha}$ degrees of freedom. In particular, the distributions differ in Lévy–Prokhorov metric by

$$\pi'' = O\left(\frac{p^2 \log \log(\mathcal{N}^{-t} \epsilon^{-1})}{\log(\mathcal{N}^{-t} \epsilon^{-1})} + \frac{p^2 \log(t)}{\sqrt{t}} + \sum_{\alpha} \frac{\mathcal{N}^{\frac{1}{3}} \log(N_{\alpha})}{N_{\alpha}^{\frac{1}{3}}} \right) = o(1). \quad (91)$$

C FORMAL DISCUSSION OF THE CONSEQUENCES OF OUR RESULTS

Before proceeding with proofs of our results we examine their implications. These are summarized in Figure 3.

C.1 BARREN PLATEAUS

The first implication of our results that we will discuss is the unification of barren plateau results (Fontana et al., 2024; Ragone et al., 2024) in the large N_{α} limit.

Corollary 20 (General expression for the loss function variance). *Let ℓ be as in Theorem 13. The variance is:*

$$\text{Var}[\ell(\boldsymbol{\theta}; \boldsymbol{\rho})] = \sum_{\alpha} \frac{\text{Tr}((\mathbf{O}^{\alpha})^2) \text{Tr}((\boldsymbol{\rho}^{\alpha})^2)}{\dim_{\mathbb{R}}(\mathfrak{g}_{\alpha})} + O(\pi). \quad (92)$$

Here, $\dim_{\mathbb{R}}(\mathfrak{g}_{\alpha})$ is the dimension of the automorphism group of \mathcal{A}_{α} and π is as in Eq. (76).⁶

Proof. Note that

$$\dim_{\mathbb{R}}(\mathfrak{g}_{\alpha}) = (\beta_{\alpha} - 1) N_{\alpha} + \frac{\beta_{\alpha} N_{\alpha} (N_{\alpha} - 1)}{2} = \frac{\beta_{\alpha} N_{\alpha}^2}{2} + O(N_{\alpha}). \quad (93)$$

Similarly,

$$\frac{(\bar{\sigma}^{\alpha})^2}{r_{\alpha}} = \frac{\|\mathbf{O}^{\alpha}\|_{\text{F}}^2}{N_{\alpha}^2}. \quad (94)$$

By Eq. (42), we also have that

$$I_{\alpha} \|\mathbf{O}^{\alpha}\|_{\text{F}}^2 = \text{Tr} \left((\mathbf{O}^{\alpha})^2 \right), \quad (95)$$

with $\text{Tr}(\cdot)$ the trace on the full N -dimensional Hilbert space. Using the fact that the diagonal entries of $\beta_{\alpha} \mathbf{W}^{\alpha}$ are i.i.d. χ^2 -distributed with $\beta_{\alpha} r_{\alpha}$ degrees of freedom (see Appendix A.4), we then have from Theorem 13 that as $N_{\alpha} \rightarrow \infty$:

$$\begin{aligned} \text{Var}[\ell(\boldsymbol{\theta}; \boldsymbol{\rho})] &= \sum_{\alpha} \text{Var}[\hat{\ell}^{\alpha}(\boldsymbol{\theta}; \boldsymbol{\rho})] + O(\pi) \\ &= \sum_{\alpha} \frac{2I_{\alpha}^2 (\bar{\sigma}^{\alpha})^2}{\beta_{\alpha} r_{\alpha}} \text{Tr}_{\alpha} \left((\boldsymbol{\rho}^{\alpha})^2 \right) + O(\pi) \\ &= \sum_{\alpha} \frac{2I_{\alpha} (\bar{\sigma}^{\alpha})^2}{\beta_{\alpha} r_{\alpha}} \text{Tr} \left((\boldsymbol{\rho}^{\alpha})^2 \right) + O(\pi) \\ &= \sum_{\alpha} \frac{\text{Tr} \left((\mathbf{O}^{\alpha})^2 \right) \text{Tr} \left((\boldsymbol{\rho}^{\alpha})^2 \right)}{\dim_{\mathbb{R}}(\mathfrak{g}_{\alpha})} + O(\pi). \end{aligned} \quad (96)$$

□

This result immediately implies Theorem 1 of Ragone et al. (2024) in the case of Lie algebra-supported ansatzes. However, our result holds for *all* variational ansatzes, i.e., it does not rely on either $i\mathbf{O}$ or $i\rho$ being a member of the dynamical Lie algebra generating the ansatz.

In the language of Ragone et al. (2024), $\text{Tr} \left((\boldsymbol{\rho}^{\alpha})^2 \right)$ is the \mathcal{A}_{α} -purity $\mathcal{P}_{\mathcal{A}_{\alpha}}$ of \mathbf{O} , and be thought of as measuring the *generalized entanglement* (Barnum et al., 2004) of $\boldsymbol{\rho}$ with respect to the Jordan algebraic structure of \mathcal{A} . In this sense, this barren plateau result can also be thought of as generalizing the entanglement-induced barren plateaus previously studied by Ortiz Marrero et al. (2021).

C.2 THE QUANTUM NEURAL TANGENT KERNEL

We now connect our results to the quantum neural tangent kernel (QNTK) literature (Liu et al., 2022; 2023; You et al., 2022; García-Martín et al., 2023; Girardi & Palma, 2024; García-Martín et al., 2024). The landmark result in this field is that, under certain conditions on \mathcal{R} , \mathbf{O} , and \mathcal{A} , variational loss functions are asymptotically a Gaussian process when $\epsilon^{-1}, t, N \rightarrow \infty$. However, this same body of work has noted that such a Gaussian process description cannot generally be true: for instance, if \mathbf{O} is rank-1 and \mathcal{A} is the space of complex Hermitian matrices, the loss should be Porter–Thomas distributed as this reduces to a random circuit sampling setting (Boixo et al., 2018).

Our results can be seen as a unifying model of neural network loss landscapes, including both when convergence to a Gaussian process is achieved and when it is not. Recall that Theorem 13 demonstrated that the asymptotic expression for the variational loss with objective observable \mathbf{O} is distributed as a Wishart process:

$$\hat{\ell}(\boldsymbol{\rho}) = \sum_{\alpha} \frac{I_{\alpha} \bar{\sigma}^{\alpha}}{r_{\alpha}} \text{Tr}_{\alpha}(\boldsymbol{\rho}^{\alpha} \mathbf{W}^{\alpha}), \quad (97)$$

⁶Through careful examination of our proof this error is actually identically zero when $t \geq 2$ and $\epsilon = 0$ in Assumption 12 as here we are only taking a second moment.

even when rescaled by a quantity \mathcal{N} exponentially large in the problem size. This correctly captures the Porter–Thomas behavior when \mathcal{A} is the space of complex Hermitian matrices and $\text{rank}(\mathbf{O}) = r = 1$. This is because the diagonal entries of a complex Wishart matrix are χ^2 -distributed with two degrees of freedom each, which is identical to the exponential distribution.⁷

Indeed, our more general result can be used to *exactly* characterize when QNNs asymptotically form Gaussian processes. This occurs when the loss is scaled by some \mathcal{N} such that Eq. (97) has nonvanishing variance yet the higher-order cumulants vanish. We compute the third-order cumulant:

$$\kappa_3(\mathcal{N}\hat{\ell}(\boldsymbol{\rho})) = \mathcal{O}\left(\mathcal{N}^3 \sum_{\alpha} \frac{I_{\alpha}^3(\bar{\sigma}^{\alpha})^3}{r_{\alpha}^2} \text{Tr}_{\alpha}\left((\boldsymbol{\rho}^{\alpha})^3\right)\right). \quad (98)$$

We thus can state this pair of conditions formally as follows.

Corollary 21 (Exact conditions for convergence to a Gaussian process). *Let \mathcal{J}, ℓ be as in Theorem 13. $\mathcal{N}\ell(\boldsymbol{\theta}; \boldsymbol{\rho})$ is asymptotically a Gaussian process as $\epsilon^{-1}, t, N \rightarrow \infty$ if and only if*

$$\mathcal{N} \leq \mathcal{O}\left(\min_{\alpha} \frac{r_{\alpha}^{\frac{2}{3}}}{I_{\alpha} \bar{\sigma}^{\alpha} \text{Tr}_{\alpha}\left((\boldsymbol{\rho}^{\alpha})^3\right)^{\frac{1}{3}}}\right) \quad (99)$$

and $\mathcal{N}^2 \text{Var}[\ell(\boldsymbol{\theta}; \boldsymbol{\rho})]$ (with the variance as given in Corollary 20) is nonvanishing for all $\boldsymbol{\rho} \in \mathcal{R}$.

When these conditions are met, the covariance is given by:

$$\begin{aligned} \mathcal{K}(\boldsymbol{\rho}, \boldsymbol{\rho}') &= \lim_{N \rightarrow \infty} \mathcal{K}_N(\boldsymbol{\rho}, \boldsymbol{\rho}') \\ &= \lim_{N \rightarrow \infty} \mathcal{N}^2 \sum_{\alpha} \frac{\beta_{\alpha} \text{Tr}\left((\mathbf{O}^{\alpha})^2\right)}{2 \dim_{\mathbb{R}}(\mathfrak{g}_{\alpha}) r_{\alpha}} I_{\alpha} \mathbb{E}[\text{Tr}_{\alpha}(\boldsymbol{\rho}^{\alpha}(\mathbf{W}^{\alpha} - r_{\alpha})) \text{Tr}_{\alpha}(\boldsymbol{\rho}'^{\alpha}(\mathbf{W}^{\alpha} - r_{\alpha}))]. \end{aligned} \quad (100)$$

It is now easy to see why (for simple \mathcal{A}) the $r = 1$ case does not converge to a Gaussian process: when $r = 1$ we may only choose a normalization $\mathcal{N} \leq \mathcal{O}(N)$ such that higher-order cumulants vanish asymptotically, but then the variance vanishes asymptotically. In contrast, our results demonstrate convergence to a Wishart process through *any* finite number of cumulants, not just the first two.

We can also be more concrete on the form of $\mathcal{K}(\boldsymbol{\rho}, \boldsymbol{\rho}')$. Indeed, we are able to show that each algebraic sector \mathcal{A}_{α} contributes to the covariance an explicit, quadratic expression in $\boldsymbol{\rho}^{\alpha}$ and $\boldsymbol{\rho}'^{\alpha}$, depending only on the field \mathbb{F}_{α} associated with its defining representation. In other words, whether or not a given quantum neural network is asymptotically a Gaussian process just depends on how SWAP tests of the inputs scale. As the expression is complicated we do not reproduce it here, instead giving references to where it may be found for each of $\mathbb{F}_{\alpha} = \mathbb{R}, \mathbb{C}, \mathbb{H}$.

Theorem 22 (\mathcal{K} is quadratic in $\boldsymbol{\rho}, \boldsymbol{\rho}'$). $\mathcal{K}(\boldsymbol{\rho}, \boldsymbol{\rho}')$ can be written as a closed-form, explicit function of only the $\text{Tr}(\boldsymbol{\rho}^{\alpha}), \text{Tr}(\boldsymbol{\rho}'^{\alpha}), \text{Tr}\left((\boldsymbol{\rho}^{\alpha})^2\right), \text{Tr}\left((\boldsymbol{\rho}'^{\alpha})^2\right)$, and $\text{Tr}(\boldsymbol{\rho}^{\alpha} \boldsymbol{\rho}'^{\alpha})$.

Proof. By the Haar invariance of \mathbf{W}^{α} there are explicit formulas for covariances of the form of $\mathcal{K}(\boldsymbol{\rho}, \boldsymbol{\rho}')$ as degree-2 polynomials in $\boldsymbol{\rho}$ and $\boldsymbol{\rho}'$. These formulas are given in:

1. Theorem 1 of Redelmeier (2011) when $\mathbb{F} = \mathbb{R}$;
2. Theorem 3.5 of Mingo & Speicher (2006) when $\mathbb{F} = \mathbb{C}$;
3. Sec. 5.3 of Redelmeier (2021) when $\mathbb{F} = \mathbb{H}$.

The result then follows by noting that $I_{\alpha} \text{Tr}_{\alpha}(\cdot^{\alpha}) = \text{Tr}(\cdot^{\alpha})$. \square

⁷It is also easy to check when $r = 1$ and the ansatz unitaries are Haar random that the assumptions we use in proving our results hold for any $\mathcal{N} \leq \mathcal{O}(N)$.

We give a simple example of the calculation of \mathcal{K} when all inputs ρ_0 can be mutually diagonalized. By the unitary invariance of Wishart matrices we can assume WLOG that all inputs are diagonal. Inputs are then completely parameterized by the eigenvalues $\mathbf{x} \in \mathbb{R}_{\geq 0}^{N_\alpha}$ of the input states, where $\|\mathbf{x}\|_1 = 1$. By the independence of diagonal entries of a Wishart matrix we then have that (when $\mathcal{N} = 1$):

$$\begin{aligned} \mathcal{K}_N(\mathbf{x}, \mathbf{x}') &= \mathcal{N}^2 \sum_{\alpha} \frac{\beta_{\alpha} \operatorname{Tr} \left((\mathbf{O}^{\alpha})^2 \right)}{2 \dim_{\mathbb{R}}(\mathfrak{g}_{\alpha}) r_{\alpha}} \sum_{i=1}^N x_i^{\alpha} x_i'^{\alpha} \operatorname{Var} [\langle i | \mathbf{W}^{\alpha} | i \rangle] + o(1) \\ &= \mathcal{N}^2 \sum_{\alpha} \frac{\operatorname{Tr} \left((\mathbf{O}^{\alpha})^2 \right)}{\dim_{\mathbb{R}}(\mathfrak{g}_{\alpha})} \mathbf{x}^{\alpha} \cdot \mathbf{x}'^{\alpha} + o(1), \end{aligned} \quad (101)$$

where in the final line we used similar simplifications as in Appendix C.1. Noting that $\mathbf{x}^{\alpha} \cdot \mathbf{x}'^{\alpha}$ is just the overlap of ρ^{α} and ρ'^{α} yields an equivalent formulation:

$$\mathcal{K}_N(\rho, \rho') = \mathcal{N}^2 \sum_{\alpha} \frac{\operatorname{Tr} \left((\mathbf{O}^{\alpha})^2 \right)}{\dim_{\mathbb{R}}(\mathfrak{g}_{\alpha})} \operatorname{Tr}(\rho^{\alpha} \rho'^{\alpha}) + o(1) = \sum_{\alpha} \mathcal{O} \left(\frac{\mathcal{N}^2 r_{\alpha}}{N_{\alpha}^2} \right) \operatorname{Tr}(\rho^{\alpha} \rho'^{\alpha}) + o(1). \quad (102)$$

In particular, when \mathcal{A} is simple, a $\mathcal{N} = \Omega \left(\frac{N}{\sqrt{r}} \right)$ normalization is required for the network to asymptotically form a Gaussian process over pure inputs.

While previous results have shown that Gaussian processes efficiently train, we argue that this may paint an overly optimistic picture for generic QNNs. Focusing on the case where the ansatz is an $\Theta(\log(n))$ -depth, 2-dimensional circuit on n qubits, we have that the QNN forms a Gaussian process when:

$$\mathcal{N} = \exp \left(\Omega \left(\log(n)^2 \right) \right), \quad (103)$$

i.e., at a normalization superpolynomial in n . Thus while it is known that such networks can achieve a constant improvement in $\mathcal{N} \hat{\ell}$ in time polynomial in n via gradient descent (Girardi & Palma, 2024), this translates to a superpolynomially vanishing improvement in the loss function value when in the physical normalization of $\mathcal{N} = 1$. More generally, the covariance we here derive can be used with results in the classical neural tangent kernel literature to reason about the training behavior (Neal, 1996; Choromanska et al., 2015; Chaudhari, 2018; Lee et al., 2018) and generalization ability (Jacot et al., 2018; Wei et al., 2022) of this class of quantum neural networks, which we hope to analyze in more detail in the future.

C.3 LOCAL MINIMA

We end by examining the distribution of local minima of the loss landscape. This has been done previously in the more restricted setting where no structure is imposed on the loss function (Anschuetz, 2022; Anschuetz & Kiani, 2022). In this section we assume the assumptions of Theorem 19 and that as all $\epsilon^{-1}, t, N_{\alpha} \rightarrow \infty$,

$$\gamma_{\alpha} \equiv \frac{p_{\alpha}}{\beta_{\alpha} r_{\alpha}} \quad (104)$$

is held constant. Here, p_{α} is the number of \mathbf{A}_i that are nontrivial on the simple component \mathcal{A}_{α} of \mathcal{A} . γ_{α} is the so-called *overparameterization ratio* discussed in Anschuetz (2022); Anschuetz & Kiani (2022). We also assume for the simplicity of our expressions that each simple sector is fully controllable, i.e., that each ansatz generator \mathbf{A}_i is only nontrivial on a single simple component of \mathcal{A} . In principle a more general expression could be written as well, though the gradient density would be a convolution over complicated probability densities, and the Hessian distributed as a free convolution of multiple Wishart matrices.

We calculate the distribution of local minima via the *Kac-Rice formula* (Adler & Taylor, 2007), which gives the expected density of local minima of a function ℓ at a function value z . Anschuetz (2022) demonstrated that the assumptions for the Kac-Rice formula are satisfied for variational loss landscapes, and when rotationally invariant on the p -torus takes the form:

$$\mathbb{E}[\operatorname{Crt}_0(z)] = (2\pi)^p \mathbb{E}[\det(\mathbf{H}_z) \mathbf{1}\{\mathbf{H}_z \succeq \mathbf{0}\}] \mathbb{P}[\mathbf{G}_z = \mathbf{0}] \mathbb{P}[\ell = z]. \quad (105)$$

Here, $\mathbb{E}[\text{Crt}_0(z)]$ is the expected density of local minima at a function value z , \mathbf{G}_z is the gradient conditioned on $\ell = z$, and \mathbf{H}_z is the Hessian conditioned on $\ell = z$ and $\mathbf{G} = \mathbf{0}$. In a slight abuse of notation, here $\mathbb{P}[\cdot]$ denotes the probability density associated with the event \cdot .

We can evaluate this expression using Theorems 13, 17, and 19.

Corollary 23 (Density of local minima to multiplicative leading order). *Let $\mathcal{J}, \hat{\ell}$ be as in Corollary 18 and Theorem 19. Assume each ansatz generator \mathbf{A}_i is only nontrivial on a single simple component of \mathcal{A} , and further assume that ρ^α is rank-1 in its defining representation. Assume as well that the overparameterization ratios:*

$$\gamma_\alpha \equiv \frac{p_\alpha}{\beta_\alpha r_\alpha} \quad (106)$$

remain fixed as $N \rightarrow \infty$. Let $\mu_{\mathbf{H}_z^\alpha}$ be the empirical spectral measure of

$$\tilde{\mathbf{H}}_z^\alpha \equiv N^{-1} \mathbf{S}^\alpha \odot \mathbf{W}^\alpha, \quad (107)$$

where (for $i \geq j$) $S_{i,j}^\alpha = |G_i^\alpha| \chi_j^\alpha$ with random variables defined as in Theorem 19 and \odot denotes the Hadamard product. Let $\lambda_{z,\alpha}^*$ denote the infimum of the support of $\mu_{\mathbf{H}_z^\alpha}$. Then the expected density of local minima of $\hat{\ell}$ at a loss function value z is:

$$\mathbb{E}[\text{Crt}(z)] = \bigstar_{\alpha} \mathbb{E}[\text{Crt}_0^\alpha(z)], \quad (108)$$

where

$$\begin{aligned} p^{-1} \ln(\mathbb{E}[\text{Crt}_0^\alpha(z)]) &= \ln\left(\frac{\pi \max(2, \beta_\alpha)}{4\sqrt{\beta_\alpha}}\right) + \frac{1}{2\gamma_\alpha} \left(1 - \frac{z}{\text{Tr}(\rho^\alpha) \bar{o}^\alpha} + \ln\left(\frac{z}{\text{Tr}(\rho^\alpha) \bar{o}^\alpha}\right)\right) \\ &+ p^{-1} \ln \mathbb{E} \left[\exp\left(p \int d\mu_{\tilde{\mathbf{H}}_z^\alpha}(\lambda) \ln(\lambda)\right) \mathbf{1}_{\{\lambda_{z,\alpha}^* \geq 0\}} \right] + o(1). \end{aligned} \quad (109)$$

Proof. We will first consider a single simple component \mathcal{A}_α with pure ρ^α (and with α labels implicit for clarity of notation), and describe in the end how a full distribution of local minima can be determined from this via a convolution. From Corollary 14 it follows that the density of the loss is:⁸

$$\mathbb{P}[\hat{\ell} = z] = \frac{\left(\frac{\beta r}{2I\bar{o}}\right)^{\frac{\beta r}{2}-1} \exp\left(-\frac{\beta r z}{2I\bar{o}}\right)}{2\Gamma\left(\frac{\beta r}{2}\right)}. \quad (110)$$

Similarly, using Corollary 18 and recalling that β is one of 1, 2, or 4, we can evaluate the density of the gradient at zero:

$$\mathbb{P}[\mathbf{G}_z = \mathbf{0}] = \left(\frac{1}{\sqrt{2\pi}\sigma_z} \int_0^\infty dx \frac{f_{\chi^{\max(2,\beta)}(x)}}{x}\right)^p = \left(\frac{\max(2,\beta)}{4\sigma_z}\right)^p, \quad (111)$$

where $f_{\chi^{\max(2,\beta)}}$ is the density of a χ -distributed random variable with $\max(2,\beta)$ degrees of freedom and

$$\sigma_z = \frac{2I\sigma_o}{N} \sqrt{\frac{\beta z}{I\bar{o}}}. \quad (112)$$

Finally, we consider the Hessian determinant. Recall Theorem 19 for the Hessian components, which gives the Hessian as:

$$\frac{N}{2I\sigma_o} \sqrt{\frac{I\bar{o}}{z}} \mathbf{H}_z = N^{-1} \tilde{\mathbf{S}} \odot \mathbf{W}, \quad (113)$$

where $\tilde{S}_{i,j} = G_i \chi_j$. We now claim that \mathbf{H}_z can only be positive semidefinite if all $G_i \geq 0$. To see this, note that $\chi_i W_{i,i}$ is always nonnegative. When $G_i < 0$, then, the (i,i) entry of $\tilde{\mathbf{S}} \odot \mathbf{W}$ is negative and thus is not positive semidefinite. We therefore can consider $|G_i|$ rather than G_i up

⁸Technically we need to consider the density of the loss rescaled by \sqrt{N} to achieve pointwise convergence in densities, but the results are equivalent after rescaling z .

1620 to pulling out a factor of 2^{-p} from the expectation. Putting everything together, to multiplicative
1621 leading order as $N \rightarrow \infty$,

$$1622$$

$$1623 \quad p^{-1} \ln (\mathbb{E} [\text{Crt}_0(z)]) = \ln \left(\frac{\pi \max(2, \beta)}{4\sqrt{\beta}} \right) + \frac{1}{2\gamma} \left(1 - \frac{z}{I\bar{o}} + \ln \left(\frac{z}{I\bar{o}} \right) \right)$$

$$1624$$

$$1625 \quad + p^{-1} \ln \mathbb{E} \left[\exp \left(p \int d\mu_{\tilde{H}_z}(\lambda) \ln(\lambda) \right) \mathbf{1} \{ \lambda_z^* \geq 0 \} \right] + o(1).$$

$$1626$$

$$1627 \quad (114)$$

1628 This completes our proof for a single simple component \mathcal{A}_α with pure ρ^α . To calculate the loss
1629 landscape of a JAWS associated with a nonsimple Jordan algebra, note that:

$$1630 \quad \mathbb{P}[\ell = z] = \bigstar_{\alpha} \mathbb{P}[\ell^\alpha = z^\alpha], \quad (115)$$

1631 where \bigstar denotes convolution. The relative weights of the various sectors introduced by $\text{Tr}_\alpha(\rho^\alpha)$
1632 can be accounted for by taking:

$$1633 \quad I_\alpha \rightarrow I_\alpha \text{Tr}_\alpha(\rho^\alpha) = \text{Tr}(\rho^\alpha). \quad (116)$$

1634 From Eq. (105) and our simplifying assumptions, then,

$$1635 \quad \mathbb{E}[\text{Crt}_0(z)] = \bigstar_{\alpha} \mathbb{E}[\text{Crt}_0^\alpha]_{I_\alpha \rightarrow \text{Tr}(\rho^\alpha)}(z), \quad (117)$$

1636 where $\mathbb{E}[\text{Crt}_0^\alpha]$ denotes Eq. (114) associated with the algebraic sector \mathcal{A}_α . \square

1637 While Corollary 23 is exact, it is obtuse almost to the point of obscurity, particularly due to the ex-
1638 pectation over the Hessian. The obstruction to further simplification is the presence of the Hadamard
1639 product between S^α and W in \tilde{H}_z^α , which is difficult to handle analytically. To get around this, we
1640 consider a slightly modified quantity where we condition both sides of Eq. (105) on the events:

$$1641 \quad |G_i| = \chi_i. \quad (118)$$

1642 In effect this can be considered as a regularization scheme, where new parameters $\tilde{\theta}_i$ are introduced
1643 as Lagrange multipliers with associated derivatives:⁹

$$1644 \quad \tilde{\ell}_i \equiv \tilde{\theta}_i + |G_i| - \chi_i \quad (119)$$

1645 and we consider a sufficiently small neighborhood of $\tilde{\theta}_i = 0$. In this setting the nontrivial compo-
1646 nents of \tilde{H}_z^α take the much more manageable form:

$$1647 \quad \tilde{H}_z^\alpha = N_\alpha^{-1} \sqrt{\Sigma^\alpha} W^\alpha \sqrt{\Sigma^\alpha}, \quad (120)$$

1648 where Σ^α is a diagonal matrix with entries i.i.d. χ^2 -distributed with $\max(2, \beta_\alpha)$ degrees of free-
1649 dom. Analyzing the expected determinant of this random matrix leads us to prove the following.

1650 **Corollary 24** (Density of local minima to multiplicative leading order, regularized). *Consider the*
1651 *setting of Corollary 23 conditioned on $|G_i| = \chi_i$ such that \tilde{H}_z^α is as in Eq. (120). Let μ_{MP}^γ be the*
1652 *Marčenko–Pastur distribution with parameter γ . Then the expected density of local minima of $\hat{\ell}$ at*
1653 *a loss function value z is:*

$$1654 \quad \mathbb{E}[\text{Crt}(z)] = \bigstar_{\alpha: \gamma_\alpha < 1} \mathbb{E}[\text{Crt}_0^\alpha](z), \quad (121)$$

1655 where:

$$1656 \quad p^{-1} \ln (\mathbb{E} [\text{Crt}_0^\alpha(z)]) = \ln \left(\frac{\pi \max(2, \beta_\alpha)}{2\sqrt{\beta_\alpha}} \right) + \frac{1}{2\gamma_\alpha} \left(1 - \frac{z}{\text{Tr}(\rho^\alpha) \bar{o}_\alpha} + \ln \left(\frac{z}{\text{Tr}(\rho^\alpha) \bar{o}_\alpha} \right) \right)$$

$$1657$$

$$1658 \quad + \frac{\max(2, \beta_\alpha)}{2} - 1 - \gamma + \int d\mu_{MP}^{\gamma_\alpha}(\lambda) \ln(\lambda) + o(1).$$

$$1659$$

$$1660 \quad (122)$$

1661 ⁹The Kac–Rice formula as stated in Adler & Taylor (2007) allows one to consider this modified gradient
1662 jointly with the original loss.

1674 *Proof.* As in Corollary 23 we focus on a single simple component \mathcal{A}_α . Σ is positive definite with
 1675 probability 1. As it is diagonal with i.i.d. χ^2 -distributed random variables each with $\max(2, \beta)$
 1676 degrees of freedom,

$$1677 \exp\left(p \int d\mu_\Sigma(\lambda) \ln(\lambda)\right) = \exp\left(p \left(\frac{\max(2, \beta)}{2} - 1 - \gamma + \ln(2)\right)\right), \quad (123)$$

1679 where $d\mu_\Sigma$ is the empirical spectral distribution of Σ and γ is the Euler–Mascheroni constant. What
 1680 remains to be considered is the spectrum of

$$1682 \mathbf{D}_z \equiv \Sigma^{-\frac{1}{2}} \tilde{\mathbf{H}}_z \Sigma^{-\frac{1}{2}} = N^{-1} \mathbf{W} = \left(\frac{\beta r}{N}\right) (\beta r)^{-1} \mathbf{W}. \quad (124)$$

1684 Recall that by assumption $\frac{\sigma_\alpha}{\bar{\sigma}} = o(1)$ so by, e.g., Lemma 25 (proved in Appendix D.1):

$$1686 \ln\left(\frac{\beta r}{N}\right) = o(1). \quad (125)$$

1688 We need only focus on $(\beta r)^{-1} \mathbf{W}$, then. As discussed in Appendix A.4, this random matrix has
 1689 empirical eigenvalue spectrum weakly converging almost surely to the Marčenko–Pastur distribu-
 1690 tion with parameter γ . In principle, large deviations in this convergence—even if they occur with
 1691 exponentially small probability—can contribute corrections to the expected determinant due to its
 1692 exponential sensitivity on the eigenvalues of \mathbf{D}_z . However, we show in Appendix L that these large
 1693 deviations are dominated in the expectation by the Marčenko–Pastur distribution. Noting that the
 1694 Marčenko–Pastur distribution with parameter γ has support at the origin if and only if $\gamma > 1$ and
 1695 taking convolutions over simple algebraic sectors as in the proof of Corollary 23 then yields the final
 1696 result. \square

1697 Dropping multiplicatively subleading factors from Eq. (122), we effectively have demonstrated that
 1698 the density $\kappa_\alpha(z)$ of local minima for a given simple component \mathcal{A}_α is asymptotically given by:

$$1700 \kappa_\alpha(z) = f_\Gamma\left(\frac{z}{\text{Tr}(\rho^\alpha) \bar{\sigma}^\alpha}; \frac{p_\alpha}{2\gamma_\alpha}, \frac{2\gamma_\alpha}{p_\alpha}\right) \quad (126)$$

1702 if $\gamma_\alpha < 1$, and otherwise:

$$1703 \kappa_\alpha(z) = \delta(z). \quad (127)$$

1704 We here have used the expression for the density of the gamma distribution:

$$1705 f_\Gamma(x; k, \theta) = \frac{1}{\Gamma(k) \theta^k} x^{k-1} \exp\left(-\frac{x}{\theta}\right). \quad (128)$$

1707 Convoluting κ_α over many simple sectors thus yields the final density for $z > 0$:

$$1709 \kappa(z) = \bigstar_{\alpha: \gamma_\alpha < 1} f_\Gamma\left(\frac{z}{\bar{\sigma}^\alpha \text{Tr}(\rho^\alpha)}; \frac{\beta_\alpha r_\alpha}{2}, \frac{2}{\beta_\alpha r_\alpha}\right). \quad (129)$$

1711 This distribution is illustrated in Figure 1(C) for various parameter regimes. To multiplicative lead-
 1712 ing order in γ this agrees exactly with the asymptotic local minima distribution studied in Anschuetz
 1713 (2022); Anschuetz & Kiani (2022). See for instance Eq. (1) of Anschuetz (2022), which studies the
 1714 case $\mathcal{A} = \mathcal{H}^N(\mathbb{C})$; i.e., one takes $\beta_\alpha \rightarrow 2$, $r_\alpha \rightarrow m$, and $\frac{z}{\text{Tr}(\rho^\alpha) \bar{\sigma}^\alpha} \rightarrow x$ to translate from our setting
 1715 to their setting.

1716 We can simplify this expression even further by noting that, asymptotically, the convolution of many
 1717 gamma distributions is also gamma-distributed by the Welch–Satterthwaite equation (Satterthwaite,
 1718 1946; Welch, 1947). This yields:

$$1720 \kappa(z) = f_\Gamma(z; k_{\text{eff}}, \theta_{\text{eff}}), \quad (130)$$

1721 where

$$1722 k_{\text{eff}} = \frac{\left(\sum_{\alpha: \gamma_\alpha < 1} \bar{\sigma}^\alpha \text{Tr}(\rho^\alpha)\right)^2}{\sum_{\alpha: \gamma_\alpha < 1} \frac{2(\bar{\sigma}^\alpha)^2 \text{Tr}(\rho^\alpha)^2}{\beta_\alpha r_\alpha}} \xrightarrow{\{N_\alpha \rightarrow \infty\}_\alpha} \bar{\ell}_{\text{u.p.}}^2 \left(\sum_{\alpha: \gamma_\alpha < 1} \frac{\text{Tr}((\mathbf{O}^\alpha)^2) \text{Tr}((\rho^\alpha)^2)}{\dim_{\mathbb{R}}(\mathfrak{g}_\alpha)}\right)^{-1}, \quad (131)$$

$$1726 \theta_{\text{eff}} = \frac{\sum_{\alpha: \gamma_\alpha < 1} \frac{2(\bar{\sigma}^\alpha)^2 \text{Tr}(\rho^\alpha)^2}{\beta_\alpha r_\alpha}}{\sum_{\alpha: \gamma_\alpha < 1} \bar{\sigma}^\alpha \text{Tr}(\rho^\alpha)} \xrightarrow{\{N_\alpha \rightarrow \infty\}_\alpha} \bar{\ell}_{\text{u.p.}}^{-1} \sum_{\alpha: \gamma_\alpha < 1} \frac{\text{Tr}((\mathbf{O}^\alpha)^2) \text{Tr}((\rho^\alpha)^2)}{\dim_{\mathbb{R}}(\mathfrak{g}_\alpha)}. \quad (132)$$

Here,

$$\bar{\ell}_{\text{u.p.}} \equiv \sum_{\alpha: \gamma_\alpha < 1} \bar{\sigma}^\alpha \text{Tr}(\rho^\alpha) \quad (133)$$

is the mean loss function value over the underparameterized sectors, and the limit in each line is due to the identities (recalling that here we assume that ρ^α is rank-1 in its defining representation):

$$(\bar{\sigma}^\alpha)^2 \text{Tr}(\rho^\alpha)^2 = I_\alpha^2 (\bar{\sigma}^\alpha)^2 \text{Tr}_\alpha \left((\rho^\alpha)^2 \right) = I_\alpha (\bar{\sigma}^\alpha)^2 \text{Tr} \left((\rho^\alpha)^2 \right), \quad (134)$$

$$I_\alpha (\bar{\sigma}^\alpha)^2 = \frac{I_\alpha}{N_\alpha^2} \text{Tr}_\alpha (\mathcal{O}^\alpha)^2 = \frac{I_\alpha r_\alpha}{N_\alpha^2} \text{Tr}_\alpha \left((\mathcal{O}^\alpha)^2 \right) = \frac{r_\alpha}{N_\alpha^2} \text{Tr} \left((\mathcal{O}^\alpha)^2 \right), \quad (135)$$

as well as the identities considered in Appendix C.1.

Intriguingly, the relevant features of this density in the underparameterized regime are controlled by the \mathcal{A}_α -purities of both \mathcal{O} and ρ in the sectors in which they are underparameterized. This is the same quantity which controls the variance of the loss function (see Appendix C.1). However, even when there are no barren plateaus in the loss landscape—for instance, if the variance of the loss function are polynomially vanishing in N —the density $\kappa(z)$ may still have exponentially small measure near $z = 0$ as the gamma distribution has exponential tails.

We reemphasize that our calculation of the local minima density was performed assuming the variance of the spectral distribution of each \mathcal{O}^α (in units of the mean eigenvalue) asymptotically vanishes. This was also the setting studied in previous work on the local minima of QNNs (Anschuetz, 2022; Anschuetz & Kiani, 2022). This assumption allows us to dramatically simplify the Hessian to the form given in Theorem 19. Though it holds for low-weight fermionic (Feng et al., 2019) and local spin Hamiltonians (Erdős & Schröder, 2014), it does *not* hold for the Gaussian unitary ensemble (GUE) or nonlocal spin systems; these systems are also known to have efficient quantum algorithms that prepare their low-energy states, unlike their local cousins (Chen et al., 2024). We hope in the future to analyze whether this property also has an impact on the behavior of local minima of QNNs.

We finally note an interesting connection between our results and algorithmic hardness. We here only calculate the expected density of local minima at a given function value. If the second moment of the local minima density is also sufficiently well-behaved, then it could be the case that the asymptotic density of local minima has a fixed distributional form; this is the case in (classical) spherical spin glass models (Subag, 2017). The function value at which these local minima proliferate w.h.p. is conjectured to also hold as a *general* (i.e., beyond gradient descent) *algorithmic threshold*, that is, it is generically believed to be the function value at which better approximations to the ground state become algorithmically intractable to find. This conjecture is known to be true for the specific case of pure spherical spin glass models (Huang & Sellke, 2023). Studying second moments of quantum spin glass local minima distributions may thus be a tractable avenue for studying the algorithmic hardness of quantum problems.

D PROOFS OF THE MAIN RESULTS

D.1 PRELIMINARIES

We now give in full detail the proofs of the main results discussed in Appendix B.2. We begin by giving definitions and notational conventions that we will use throughout our proofs. Recall that, given a JAWS $\mathcal{J} = (\bigoplus_\alpha \mathcal{A}_\alpha, \mathcal{T}, \mathcal{G}, \mathcal{H}, \mathbf{A}, \mathcal{O})$, we are interested in the joint distribution of the loss $\ell(\theta; \rho)$ and its first two derivatives over a set of input states $\rho \in \mathcal{R}$, where:

$$\ell(\theta; \rho) = \sum_\alpha \ell^\alpha(\theta; \rho) \equiv \sum_\alpha I_\alpha \text{Tr}_\alpha \left(\rho^\alpha U^\alpha(\theta) \mathcal{O}^\alpha U^{\alpha\dagger}(\theta) \right) \quad (136)$$

and

$$U^\alpha(\theta) \equiv \mathbf{g}_0^{\alpha\dagger} \left(\prod_{i=1}^p \mathbf{g}_i^\alpha \exp(\theta_i \mathbf{A}_i^\alpha) \mathbf{g}_i^{\alpha\dagger} \right) \mathbf{h}^\alpha \quad (137)$$

with $\mathbf{g}_i^\alpha, \mathbf{h}^\alpha \in G_\alpha$, and $\mathbf{A}_i^\alpha \in \mathfrak{g}_\alpha$. As previously discussed we will use \cdot^α to denote the defining representation of the projection of \cdot into \mathcal{A}_α . As each sector labeled by α is independent we will here only consider a single α WLOG, with nonzero ρ^α and \mathcal{O}^α . This will also allow us to remove

1782 the cumbersome notation of labeling all objects with the index α for the remainder of this section.
 1783 To further simplify the language, we will use the term “unitary” to refer to “orthogonal,” “unitary,”
 1784 or “hyperunitary” in the context of $\mathbb{F} = \mathbb{R}, \mathbb{C}, \mathbb{H}$, respectively, unless otherwise explicitly stated.
 1785 Similarly, we will later see that our results hold for any θ ; we will thus leave the θ -dependence of ℓ
 1786 implicit from here on out to save on notation. In the following we will use $|\mu\rangle$ (i.e., Greek letters),
 1787 $0 \leq \mu \leq N_\alpha - 1$ to denote basis vectors in the vector space on which the defining representation of
 1788 $\mathcal{H}^{N_\alpha}(\mathbb{F}_\alpha)$ acts. We will later use $|i\rangle$ (i.e., Latin letters), $0 \leq i \leq p$ to denote vectors in the vector
 1789 space on which the defining representation of $\mathcal{H}^{p+1}(\mathbb{F}_\alpha)$ acts.

1790 We now detail our choice of \mathbf{A}_i given Assumption 16. Assumption 16 has a nice interpretation as
 1791 taking the \mathbf{A}_i to be low-rank (representations of) basis elements of the Cartan subalgebra $\mathfrak{h} \subseteq \mathfrak{g}$. Up
 1792 to the adjoint action of G and an overall normalization (which can be absorbed into θ_i), then, we can
 1793 take WLOG:

$$1794 \quad \mathbf{A}_i = |0\rangle\langle 1| - |1\rangle\langle 0| \quad (138)$$

1795 when $\mathbb{F} = \mathbb{R}$ and

$$1796 \quad \mathbf{A}_i = i|0\rangle\langle 0| \quad (139)$$

1797 when $\mathbb{F} = \mathbb{C}$. When $\mathbb{F} = \mathbb{H}$ we will take a parameterization of the form:

$$1799 \quad \exp(\theta_i \mathbf{A}_i) = \exp\left(i\theta_i^{(i)}|0\rangle\langle 0|\right) \exp\left(j\theta_i^{(j)}|0\rangle\langle 0|\right) \exp\left(k\theta_i^{(k)}|0\rangle\langle 0|\right); \quad (140)$$

1801 that is, we will assume we have full control over the quaternionic phase. We have chosen here for
 1802 the \mathbf{A}_i to be i -independent for convenience, moving any i -dependence to the conjugating unitaries
 1803 of each layer $\mathbf{g}_i^\alpha \in G_\alpha$.

1804 Finally, before continuing we give a convenient relation between

$$1806 \quad r = \frac{\|\mathbf{O}\|_*^2}{\|\mathbf{O}\|_{\mathbb{F}}^2} \quad (141)$$

1807 and the standard deviation σ_o of the eigenvalues of \mathbf{O} . This relation will be used to simplify some
 1808 of our later expressions.

1809 **Lemma 25** (r and σ_o relation). *Let \mathbf{O} be an $N \times N$ Hermitian operator and r be as in Eq. (141).
 1810 Let σ_o be the standard deviation of the eigenvalues of \mathbf{O} , and \bar{o} the arithmetic mean. Then:*

$$1812 \quad \sqrt{\frac{N}{r} - 1} = \frac{\sigma_o}{\bar{o}}. \quad (142)$$

1813 *Proof.* Let \bar{o}_{RMS} be the root mean square of the eigenvalues of \mathbf{O} . Then:

$$1814 \quad \begin{aligned} 1815 \quad \sqrt{\frac{N}{r} - 1} &= \sqrt{\frac{N \|\mathbf{O}\|_{\mathbb{F}}^2}{\|\mathbf{O}\|_*^2} - 1} \\ 1816 \quad &= \sqrt{\frac{\bar{o}_{\text{RMS}}^2}{\bar{o}^2} - 1} \\ 1817 \quad &= \frac{\sigma_o}{\bar{o}}. \end{aligned} \quad (143)$$

1818 \square

1829 D.2 ASYMPTOTIC EXPRESSION FOR THE LOSS

1830 We now proceed with the proofs of our main results. This subsection is devoted to a series of
 1831 reductions that will allow us to consider the $\mathbf{hO}\mathbf{h}^\dagger$ as β -Wishart matrices up to a controlled error in
 1832 Lévy–Prokhorov metric, thus proving Theorem 13. Along the way we will also prove a reduction to
 1833 taking the \mathbf{g}_i to be Haar random, once again up to some bounded error in Lévy–Prokhorov metric.
 1834 In proving these results we will heavily rely on various lemmas on convergence in distribution given
 1835 in Appendix M.

1836 D.2.1 REDUCTION TO HAAR RANDOM \mathbf{h}, \mathbf{g}_i

1837 We first argue that, under Assumption 12, the \mathbf{h} in Eq. (137) can be assumed to be Haar random over
 1838 G up to some bounded error in Lévy–Prokhorov metric. We also show that this remains true when
 1839 scaling by any choice of N -dependent normalization \mathcal{N} with the only requirement being that central
 1840 moments of large (constant) order have a well-defined, finite limit as $N \rightarrow \infty$; for which maximal
 1841 choice of \mathcal{N} this is true depends on the exact form of the ϵ -approximate t -design, but is always true
 1842 for $\mathcal{N} = 1$. In the following we implicitly consider a sequence of objects as $\epsilon^{-1}, t, N \rightarrow \infty$.

1843 **Lemma 26** (Weak convergence to Haar random \mathbf{h}). *Assume \mathbf{h} forms an ϵ -approximate t -design.
 1844 Let $L_i, i \in [d]$ be multilinear functions of the form:*

$$1845 L_i = \text{Tr}(\mathbf{M}_i \mathbf{h} \mathbf{O} \mathbf{h}^\dagger), \quad (144)$$

1846 where the \mathbf{M}_i have bounded operator norm and $\mathbf{h} \sim \mathcal{H}$. Let \tilde{L}_i be the same with $\mathbf{h} \rightarrow \tilde{\mathbf{h}}$, where $\tilde{\mathbf{h}}$
 1847 is Haar random. Assume any N -dependent normalization $\mathcal{N} \leq O(N)$, and assume that there exists
 1848 a constant k^* such that all central moments of constant order $k > k^*$ of the $\mathcal{N}L_i$ have a finite limit
 1849 as $\epsilon^{-1}, t, N \rightarrow \infty$. Then the joint distribution of $\mathcal{N}L_i$ differs from that of $\mathcal{N}\tilde{L}_i$ by an error at most

$$1850 \pi_1 = O\left(\frac{d \log \log(\mathcal{N}^{-t} \epsilon^{-1})}{\log(\mathcal{N}^{-t} \epsilon^{-1})} + \frac{d \log(t)}{\sqrt{t}}\right) \quad (145)$$

1851 in Lévy–Prokhorov metric as $\epsilon^{-1}, t, N \rightarrow \infty$.

1852 *Proof.* As the $L_i - \mathbb{E}[L_i]$ are bounded random variables, there exists some constant $C > 0$ such
 1853 that:

$$1854 \mathbb{E}\left[\prod_{i \in s_k} (L_i - \mathbb{E}[L_i])\right] \leq (C\sqrt{k})^k \quad (146)$$

1855 for all index multisets s_k of cardinality k . By the given assumption on \mathcal{N} a similar bound also holds
 1856 for central moments of the $\mathcal{N}L_i$ of sufficiently high order. The bound $\mathcal{N} = O(N)$ ensures that this
 1857 central moment bound also holds true for the $\mathcal{N}\tilde{L}_i$ by sub-Gaussianity (Meckes, 2019; Vershynin,
 1858 2018), and also ensures that the μ of Corollary 38 is subleading to Eq. (145). The conditions of
 1859 Corollary 38 are then satisfied and the final result yielded. \square

1860 By incurring this error in Lévy–Prokhorov metric we now can assume that the \mathbf{h} (given Assump-
 1861 tion 12) are Haar random when considering just the distribution of the loss (so $d = 1$ in Lemma 26).
 1862 When considering the loss, gradient, and Hessian jointly, the addition of Assumption 15 means that
 1863 we can assume that both the \mathbf{h} and the \mathbf{g}_i are Haar random. We note, however, that up to requiring a
 1864 worse bound on \mathcal{N} (i.e., $\mathcal{N} \leq o(N^{\frac{2}{3}})$) the \mathbf{g}_i being drawn from an ϵ -approximate 2-design suffices.
 1865 This argument is given in detail in Appendix J.

1866 When the \mathbf{h} and \mathbf{g}_i are reduced to being drawn i.i.d. from the Haar distribution it is apparent that
 1867 the joint distribution of the loss with its first two derivatives is invariant under translations of the
 1868 parameters. This justifies us fixing $\boldsymbol{\theta} = \mathbf{0}$ in the sequel.

1878 D.2.2 REDUCTION TO GAUSSIAN \mathbf{h}, \mathbf{g}_i

1879 We now show that certain marginal distributions of the the entries of $\sqrt{N}\mathbf{h}, \sqrt{N}\mathbf{g}_i$ can be ap-
 1880 proximated as random Gaussian matrices over \mathbb{F} up to a small error in Ky Fan metric whenever
 1881 $\mathcal{N} \leq O(N)$. This will allow us to replace $\sqrt{N}\mathbf{h}, \sqrt{N}\mathbf{g}_i$ with random matrices with i.i.d. Gaussian
 1882 entries, simplifying our results further.

1883 **Lemma 27** (Convergence in probability of Haar marginals to Gaussian matrices). *Let $L_i, i \in [d]$ be
 1884 uniformly bounded multilinear functions of the form:*

$$1885 L_i = \text{Tr}(\mathbf{M}_i \mathbf{h} \boldsymbol{\rho} \mathbf{h}^\dagger), \quad (147)$$

1886 where there exists some D such that

$$1887 \frac{1}{\text{Tr}(\boldsymbol{\rho}^2)} \leq D = o\left(\frac{N}{\log(N)}\right). \quad (148)$$

1890 Assume as well that $d \leq O(\text{poly}(N))$ and the M_i are independent from the \mathbf{h} . Let \tilde{L}_i be the
 1891 same, where now the \mathbf{h} are random matrices \mathbf{G} with i.i.d. standard Gaussian entries over \mathbb{F} . For
 1892 sufficiently large N , the joint distribution of $\mathcal{N}L_i$ differs from the joint distribution of $\frac{\mathcal{N}}{N}\tilde{L}_i$ by an
 1893 error in Ky Fan metric of at most:

$$1894 \alpha_3 = O\left(\frac{\sqrt{\mathcal{N}D \log(N)}}{N} + \frac{\mathcal{N}}{D}\right) = o(1) \quad (149)$$

1895 for all $\mathcal{N} \leq O(N)$. Alternatively, if $\text{rank}(\boldsymbol{\rho}) \leq D$,

$$1896 \alpha_3 = O\left(\frac{\sqrt{\mathcal{N}D \log(N)}}{N}\right) = o(1). \quad (150)$$

1897 *Proof.* We assume $\boldsymbol{\rho}$ has trace norm 1 by absorbing $\text{Tr}(\boldsymbol{\rho})$ into M_i . We first argue that if the
 1898 high-purity condition is assumed then $\boldsymbol{\rho}$ has an approximation to small error in trace norm that is
 1899 low-rank. To see this, let r_i be the eigenvalues of $\boldsymbol{\rho}$ in non-increasing order. If

$$1900 \sum_{i=D+1}^{N_\alpha} r_i \geq \frac{1}{D+1} \quad (151)$$

1901 then it must be that $\boldsymbol{\rho}$ has low purity, i.e.,

$$1902 \text{Tr}(\boldsymbol{\rho}^2) \leq \frac{D+1}{(D+1)^2} = \frac{1}{D+1} < \frac{1}{D}, \quad (152)$$

1903 breaking our assumption on the purities of the inputs (Eq. (148)). Thus,

$$1904 \sum_{i=D+1}^{N_\alpha} r_i < \frac{1}{D+1}. \quad (153)$$

1905 In particular, $\boldsymbol{\rho}$ has a rank- D approximation that agrees up to an $O(\frac{1}{D})$ additive error in trace
 1906 distance.

1907 Consider now the low-rank case. Let

$$1908 \delta \equiv \frac{D \ln(N)}{N}. \quad (154)$$

1909 $\delta < \frac{1}{4}$ for sufficiently large N by the assumed scaling of D . Following the proof of Corollary 1.1 in
 1910 Jiang (2010), then, for sufficiently large N and $t > 10\sqrt{\frac{\mathcal{N}\delta}{N}}$,

$$1911 \mathbb{P}\left[\max_{\substack{i \in [N] \\ j \in [D]}} \left| \sqrt{N}h_{i,j} - G_{i,j} \right| \geq \sqrt{\frac{N}{\mathcal{N}}}t\right] = \exp\left(-\Omega\left(\log(N)^{\frac{3}{2}}\right)\right) + \exp\left(-\Omega\left(\frac{N^2}{\mathcal{N}D}\right)\right). \quad (155)$$

1912 The right-hand side decays superpolynomially with N whenever $\mathcal{N} \leq O(N)$. This implies the
 1913 error in Ky Fan metric is dominated by the cutoff at $t = \Theta\left(\sqrt{\frac{\mathcal{N}\delta}{N}}\right)$. This in conjunction with the
 1914 error from the initial low-rank approximation yields the final result. \square

1915 We will mostly be concerned with the specific case when $D = N^{0.999}$, yielding for $\mathcal{N} \leq O(N^{0.99})$:

$$1916 \alpha_3 = O\left(\sqrt{\frac{\mathcal{N} \log(N)}{N^{1.001}}}\right) = o(1). \quad (156)$$

1944 D.2.3 REDUCTION TO SEMI-ISOTROPIC \mathbf{O}

1945
1946 We end with a reduction to a more convenient form for the spectrum of \mathbf{O} . Let $|\mu\rangle$ be the eigenbasis
1947 of \mathbf{O} . Consider the Hermitian $\tilde{\mathbf{O}}$:

$$1948 \quad \tilde{\mathbf{O}} \equiv k \sum_{\mu=0}^{r-1} |\mu\rangle \langle \mu|, \quad (157)$$

1950 where

$$1951 \quad r \equiv \frac{\|\mathbf{O}\|_*^2}{\|\mathbf{O}\|_F^2} \quad (158)$$

1952 is assumed to be an integer and (for \bar{o} the mean eigenvalue of o)

$$1953 \quad k \equiv \frac{\|\mathbf{O}\|_*}{r} = \frac{N\bar{o}}{r}. \quad (159)$$

1954 By construction

$$1955 \quad \text{Tr}(\mathbf{O}) = \text{Tr}(\tilde{\mathbf{O}}), \quad (160)$$

$$1956 \quad \text{Tr}(\mathbf{O}^2) = \text{Tr}(\tilde{\mathbf{O}}^2). \quad (161)$$

1957 We claim that replacing \mathbf{O} with $\tilde{\mathbf{O}}$ incurs only a vanishingly small error in Ky Fan metric. Intuitively
1958 this follows from the Welch–Satterthwaite approximation of a weighted sum of Wishart random
1959 matrices as a single Wishart random matrix (Khuri et al., 1994). We formally state this as the
1960 following lemma.

1961 **Lemma 28** (Reduction to semi-isotropic \mathbf{O}). *Let $L_i, i \in [d]$ be of the form:*

$$1962 \quad L_i = \text{Tr}(\mathbf{M}_i \mathbf{X} (\mathbf{O} - \tilde{\mathbf{O}}) \mathbf{X}^\dagger), \quad (162)$$

1963 where $\sqrt{N}\mathbf{X}$ has i.i.d. standard Gaussian entries over \mathbb{F} , \mathbf{M}_i is independent from \mathbf{X} with bounded
1964 trace norm, and $d = O(\text{poly}(N))$. Let \tilde{L}_i be the same multilinear functions, where instead of \mathbf{O}
1965 one has $\tilde{\mathbf{O}}$ as defined in Eq. (157). The joint distribution of NL_i over \mathbf{X} differs from $\mathbf{0}$ by an error
1966 at most

$$1967 \quad \alpha_4 = O\left(\frac{\mathcal{N}\Delta \log\left(\frac{dN}{N\Delta}\right)}{N}\right) \quad (163)$$

1968 in Ky Fan metric, where

$$1969 \quad \Delta \equiv 1 - \frac{\bar{o}}{\|\mathbf{O}\|_{op}} - \frac{\sigma_o^2}{\|\mathbf{O}\|_{op} \bar{o}}. \quad (164)$$

1970 Here, o_i is the i th eigenvalue of \mathbf{O} in non-increasing order, σ_o the standard deviation of the eigen-
1971 values of \mathbf{O} , and \bar{o} the mean.

1972 *Proof.* \mathbf{O} and $\tilde{\mathbf{O}}$ can be assumed to be diagonal WLOG as \mathbf{X} is unitarily invariant. Therefore,
1973 $\mathbf{X}(\mathbf{O} - \tilde{\mathbf{O}})\mathbf{X}^\dagger$ can be written as a weighted sum of standard β -Wishart matrices each with a
1974 single degree of freedom:

$$1975 \quad \mathbf{X}(\mathbf{O} - \tilde{\mathbf{O}})\mathbf{X}^\dagger = \frac{1}{N} \sum_{i=1}^N (o_i - \tilde{o}_i) |x_i\rangle \langle x_i|, \quad (165)$$

1976 where $|x_i\rangle$ is the (unnormalized) i th column of $\sqrt{N}\mathbf{X}$ and o_i, \tilde{o}_i are the eigenvalues of $\mathbf{O}, \tilde{\mathbf{O}}$, res-
1977 pectively, in nonincreasing order. Note that for any normalized $|\mu\rangle$, $|\langle x_i|\mu\rangle|$ is gamma-distributed.
1978 Thus, from a generalization of Bernstein’s inequality (see, e.g., Theorem 2.8.1 of Vershynin (2018)),
1979 for any $|\mu\rangle$,

$$1980 \quad \mathbb{P}\left[\left|\frac{1}{N} \sum_{i=1}^N (o_i - \tilde{o}_i) |\langle x_i|\mu\rangle|^2\right| \geq \epsilon\right] \quad (166)$$

$$1981 \quad \leq 2 \exp\left(-cN \min\left(\frac{\epsilon}{\max_i |o_i - \tilde{o}_i|}, \frac{\epsilon^2}{\frac{1}{N} \sum_{i=1}^N (o_i - \tilde{o}_i)^2}\right)\right)$$

for some constant $c > 0$. Due to the equality of traces of \mathbf{O} and $\tilde{\mathbf{O}}$ we calculate:

$$\begin{aligned}
\frac{1}{N} \sum_{i=1}^N (o_i - \tilde{o}_i)^2 &= \frac{1}{N} \text{Tr} \left((\mathbf{O} - \tilde{\mathbf{O}})^2 \right) \\
&= \frac{2}{N} \left(\text{Tr}(\mathbf{O}^2) - \text{Tr}(\mathbf{O}\tilde{\mathbf{O}}) \right) \\
&\leq \frac{2}{N} \left(\text{Tr}(\mathbf{O}^2) - \text{Tr}(\mathbf{O})k + ko_{r+1} \right) \\
&= \frac{2}{r} \bar{o} o_{r+1} \\
&= \mathcal{O} \left(\frac{o_{r+1}}{N} \right).
\end{aligned} \tag{167}$$

Furthermore,

$$\begin{aligned}
\max_i |o_i - \tilde{o}_i| &= \mathcal{O}(o_1 - k) + \mathcal{O}(o_{r+1}) \\
&= \mathcal{O} \left(o_1 - \bar{o} - \frac{\sigma_o^2}{\bar{o}} + o_{r+1} \right) \\
&= \mathcal{O} \left(1 - \frac{\bar{o}}{\|\mathbf{O}\|_{\text{op}}} - \frac{\sigma_o^2}{\|\mathbf{O}\|_{\text{op}} \bar{o}} \right).
\end{aligned} \tag{168}$$

The desired convergence follows by the union bound and taking $\epsilon \rightarrow \frac{\epsilon}{N}$. \square

We now claim these suffice to prove Theorem 13.

Proof of Theorem 13. For \mathbf{X} with i.i.d. standard Gaussian entries over \mathbb{F} and $\tilde{\mathbf{O}}$ as in Lemma 28, $k^{-1} \mathbf{X} \tilde{\mathbf{O}} \mathbf{X}^\dagger$ is β -Wishart-distributed with r degrees of freedom. The result then follows from Lemmas 26, 27, and 28 applied in sequence, with a final error in Lévy–Prokhorov metric of:

$$\pi = \pi_1 + \alpha_3 + \alpha_4 = \mathcal{O} \left(\frac{\log \log (\mathcal{N}^{-t} \epsilon^{-1})}{\log (\mathcal{N}^{-t} \epsilon^{-1})} + \frac{\log(t)}{\sqrt{t}} + \sqrt{\frac{\mathcal{N} \log(N)}{N^{1.001}}} \right). \tag{169}$$

\square

D.3 ASYMPTOTIC EXPRESSIONS FOR THE FIRST TWO DERIVATIVES

We will now move on and construct an explicit asymptotic expression for the first two derivatives of the loss function. As previously mentioned, we need only consider $\boldsymbol{\theta} = \mathbf{0}$ by Lemma 26. Here we have that the first derivatives are distributed as:

$$\partial_i \ell = I \text{Tr} \left(\boldsymbol{\rho} \left[\mathbf{g}_i \mathbf{A}_i \mathbf{g}_i^\dagger, \mathbf{h} \mathbf{O} \mathbf{h}^\dagger \right] \right). \tag{170}$$

The second derivatives are of the form (for $i \geq j$):

$$\partial_i \partial_j \ell = I \text{Tr} \left(\boldsymbol{\rho} \left[\mathbf{g}_j \mathbf{A}_j \mathbf{g}_j^\dagger \left[\mathbf{g}_i \mathbf{A}_i \mathbf{g}_i^\dagger, \mathbf{h} \mathbf{O} \mathbf{h}^\dagger \right] \right] \right). \tag{171}$$

Furthermore, using the lemmas from Appendix D.2, up to a cost:

$$\pi_1 + \alpha_3 + \alpha_4 = \mathcal{O} \left(\frac{d \log \log (\mathcal{N}^{-t} \epsilon^{-1})}{\log (\mathcal{N}^{-t} \epsilon^{-1})} + \frac{d \log(t)}{\sqrt{t}} + \sqrt{\frac{\mathcal{N} \log(N)}{N}} \right) \tag{172}$$

in Lévy–Prokhorov metric—where d is either p or p^2 —we may reduce to considering:

$$\hat{\ell} = \frac{I}{N} \text{Tr} \left(\boldsymbol{\rho} \mathbf{X}_0^\dagger \mathbf{O} \mathbf{X}_0 \right), \tag{173}$$

$$\hat{\ell}_{;i} = \frac{I}{N^2} \text{Tr} \left(\boldsymbol{\rho} \mathbf{X}_0^\dagger \left[\mathbf{X}_i \mathbf{A}_i \mathbf{X}_i^\dagger, \mathbf{O} \right] \mathbf{X}_0 \right), \tag{174}$$

$$\hat{\ell}_{;i,j} = \frac{I}{N^3} \text{Tr} \left(\boldsymbol{\rho} \mathbf{X}_0^\dagger \left[\mathbf{X}_j \mathbf{A}_j \mathbf{X}_j^\dagger, \left[\mathbf{X}_i \mathbf{A}_i \mathbf{X}_i^\dagger, \mathbf{O} \right] \right] \mathbf{X}_0 \right), \tag{175}$$

where:

1. \mathbf{X}_0 is an $N \times \text{rank}(\rho)$ random matrix with i.i.d. standard Gaussian entries over \mathbb{F} , and the \mathbf{X}_i are $N \times p$ or $N \times 2p$ (depending on the rank of the \mathbf{A}_i) random matrices with i.i.d. standard Gaussian entries over \mathbb{F} .
2. The initial states ρ are of rank at most $O(N^{0.999})$.

Note that these quantities depend only on $M = O(\text{rank}(\rho))$ rows of \mathbf{X}_0 and either 1 ($\mathbb{F} = \mathbb{C}, \mathbb{H}$) or 2 ($\mathbb{F} = \mathbb{R}$) rows of \mathbf{X}_i . We can therefore further simplify this by collecting all of the relevant rows of $\mathbf{X}_0, \mathbf{X}_i$ into a single \mathbf{X} of dimensions $N \times (p + M)$ (or $N \times (2p + M)$ when $\mathbb{F} = \mathbb{R}$), writing these expressions as:¹⁰

$$\hat{\ell} = \frac{I\bar{o}}{r} \text{Tr} \left(\rho \mathbf{X}^\dagger \tilde{\mathbf{O}} \mathbf{X} \right), \quad (176)$$

$$\hat{\ell}_{;i} = \frac{I\bar{o}}{rN} \text{Tr} \left(\rho \mathbf{X}^\dagger \left[\mathbf{X} |i\rangle \langle i| \mathbf{X}^\dagger, \tilde{\mathbf{O}} \right] \mathbf{X} \right), \quad (177)$$

$$\hat{\ell}_{;i,j} = \frac{I\bar{o}}{rN^2} \text{Tr} \left(\rho \mathbf{X}^\dagger \left[\mathbf{X} |j\rangle \langle j| \mathbf{X}^\dagger, \left[\mathbf{X} |i\rangle \langle i| \mathbf{X}^\dagger, \tilde{\mathbf{O}} \right] \right] \mathbf{X} \right). \quad (178)$$

Here, we took another error of α_4 in Lévy–Prokhorov metric to change \mathbf{O} to

$$\tilde{\mathbf{O}} = \sum_{\mu=0}^{r-1} |\mu\rangle \langle \mu| \quad (179)$$

as in Eq. (157). Though this expression for the full distribution is unwieldy, it completely characterizes the joint distribution of the loss and first two derivatives in terms of a single Gaussian random matrix \mathbf{X} . To simplify further we will assume ρ is a single state which is rank-1 when projected into the simple sector we are considering, taken WLOG to be $|0\rangle \langle 0|$. We will also assume in the following that $\frac{\sigma_o}{\sigma}$ is asymptotically bounded.

Details of further simplifications vary depending on whether \mathbb{F} is algebraically closed or not (i.e., whether or not we are working in \mathbb{R}). We will thus consider the two cases separately.

D.3.1 $\mathbb{F} = \mathbb{C}, \mathbb{H}$

When $\mathbb{F} = \mathbb{C}, \mathbb{H}$, we are considering the joint distribution of:

$$\hat{\ell} = \frac{I\bar{o}}{r} \text{Tr} \left(|0\rangle \langle 0| \mathbf{X}^\dagger \tilde{\mathbf{O}} \mathbf{X} \right), \quad (180)$$

$$\hat{\ell}_{;i} = \frac{iI\bar{o}}{rN} \text{Tr} \left(|0\rangle \langle 0| \mathbf{X}^\dagger \left[\mathbf{X} |i\rangle \langle i| \mathbf{X}^\dagger, \tilde{\mathbf{O}} \right] \mathbf{X} \right), \quad (181)$$

$$\hat{\ell}_{;i,j} = -\frac{I\bar{o}}{rN^2} \text{Tr} \left(|0\rangle \langle 0| \mathbf{X}^\dagger \left[\mathbf{X} |j\rangle \langle j| \mathbf{X}^\dagger, \left[\mathbf{X} |i\rangle \langle i| \mathbf{X}^\dagger, \tilde{\mathbf{O}} \right] \right] \mathbf{X} \right), \quad (182)$$

where now \mathbf{X} is a $N \times (p + 1)$ random matrix with i.i.d. standard Gaussian entries over \mathbb{F} . Here, when $\mathbb{F} = \mathbb{H}$, “ i ” should be taken to mean one of i, j, k , depending on which parameter of Eq. (140) the derivative is being taken with respect to. We take the two contributions of quaternionic phase to be equal in the second derivative expression as, by the anticommutivity of quaternions, this expression is otherwise identically zero. In particular, this abuse of notation will not change any of the following analysis unless specifically stated otherwise.

Recall that

$$\tilde{\mathbf{O}} = \sum_{\mu=0}^{r-1} |\mu\rangle \langle \mu|. \quad (183)$$

Let \mathbf{C} be the complement of $\tilde{\mathbf{O}}$, i.e.,

$$\mathbf{C} = \sum_{\mu=r}^{N-1} |\mu\rangle \langle \mu|. \quad (184)$$

¹⁰We have slightly abused notation here; the $|i\rangle, |j\rangle$ here are taken to be orthogonal to the nonzero eigenvectors of ρ .

2106 In particular, \mathbf{C} and $\tilde{\mathbf{O}}$ sum to the identity:

$$2108 \quad \mathbf{C} + \tilde{\mathbf{O}} = \mathbf{I}_N. \quad (185)$$

2109 Furthermore,

$$2111 \quad \mathbf{W} \equiv \mathbf{X}^\dagger \tilde{\mathbf{O}} \mathbf{X}, \quad (186)$$

$$2113 \quad \mathbf{V} \equiv \mathbf{X}^\dagger \mathbf{C} \mathbf{X}, \quad (187)$$

2114 are each β -Wishart-distributed matrices; the former with r degrees of freedom, and the latter with
 2115 $N - r$. The independence of \mathbf{W} and \mathbf{V} can be checked from the associated Bartlett decompositions.
 2116 We thus have that the distribution we are interested in is the joint distribution of:

$$2118 \quad \hat{\ell} = \frac{I\bar{o}}{r} \langle 0 | \mathbf{W} | 0 \rangle, \quad (188)$$

$$2120 \quad \hat{\ell}_{;i} = \frac{2I\bar{o}}{rN} \operatorname{Re} \{i \langle 0 | \mathbf{W} | i \rangle \langle i | \mathbf{X}^\dagger \mathbf{X} | 0 \rangle\} = \frac{2I\bar{o}}{rN} \operatorname{Re} \{i \langle 0 | \mathbf{W} | i \rangle \langle i | \mathbf{V} | 0 \rangle\}, \quad (189)$$

2122 and:

$$2123 \quad \begin{aligned} \hat{\ell}_{;i,j} &= -\frac{I\bar{o}}{rN^2} \langle 0 | \mathbf{X}^\dagger [\mathbf{X} | j \rangle \langle j | \mathbf{X}^\dagger, [\mathbf{X} | i \rangle \langle i | \mathbf{X}^\dagger, \mathbf{O}]] \mathbf{X} | 0 \rangle \\ &= -\frac{I\bar{o}}{rN^2} \langle 0 | \mathbf{X}^\dagger \mathbf{X} | j \rangle \langle j | \mathbf{X}^\dagger \mathbf{X} | i \rangle \langle i | \mathbf{X}^\dagger \mathbf{O} \mathbf{X} | 0 \rangle \\ &\quad - \frac{I\bar{o}}{rN^2} \langle 0 | \mathbf{X}^\dagger \mathbf{O} \mathbf{X} | i \rangle \langle i | \mathbf{X}^\dagger \mathbf{X} | j \rangle \langle j | \mathbf{X}^\dagger \mathbf{X} | 0 \rangle \\ &\quad + \frac{I\bar{o}}{rN^2} \langle 0 | \mathbf{X}^\dagger \mathbf{X} | i \rangle \langle i | \mathbf{X}^\dagger \mathbf{O} \mathbf{X} | j \rangle \langle j | \mathbf{X}^\dagger \mathbf{X} | 0 \rangle \\ &\quad + \frac{I\bar{o}}{rN^2} \langle 0 | \mathbf{X}^\dagger \mathbf{X} | j \rangle \langle j | \mathbf{X}^\dagger \mathbf{O} \mathbf{X} | i \rangle \langle i | \mathbf{X}^\dagger \mathbf{X} | 0 \rangle \\ &= -\frac{I\bar{o}}{rN^2} \langle 0 | (\mathbf{W} + \mathbf{V}) | j \rangle \langle j | (\mathbf{W} + \mathbf{V}) | i \rangle \langle i | \mathbf{W} | 0 \rangle \\ &\quad - \frac{I\bar{o}}{rN^2} \langle 0 | \mathbf{W} | i \rangle \langle i | (\mathbf{W} + \mathbf{V}) | j \rangle \langle j | (\mathbf{W} + \mathbf{V}) | 0 \rangle \\ &\quad + \frac{I\bar{o}}{rN^2} \langle 0 | (\mathbf{W} + \mathbf{V}) | i \rangle \langle i | \mathbf{W} | j \rangle \langle j | (\mathbf{W} + \mathbf{V}) | 0 \rangle \\ &\quad + \frac{I\bar{o}}{rN^2} \langle 0 | (\mathbf{W} + \mathbf{V}) | j \rangle \langle j | \mathbf{W} | i \rangle \langle i | (\mathbf{W} + \mathbf{V}) | 0 \rangle. \end{aligned} \quad (190)$$

2142 We can simplify the expression for the second derivative by grouping terms by their order in \mathbf{W} .
 2143 Note first that all terms cubic in matrix elements of \mathbf{W} sum to zero. For the quadratic terms, note
 2144 that for the outer two terms:

$$2146 \quad \begin{aligned} &-\langle 0 | \mathbf{V} | j \rangle \langle j | \mathbf{W} | i \rangle \langle i | \mathbf{W} | 0 \rangle - \langle 0 | \mathbf{W} | j \rangle \langle j | \mathbf{V} | i \rangle \langle i | \mathbf{W} | 0 \rangle \\ &+ \langle 0 | \mathbf{V} | j \rangle \langle j | \mathbf{W} | i \rangle \langle i | \mathbf{W} | 0 \rangle + \langle 0 | \mathbf{W} | j \rangle \langle j | \mathbf{W} | i \rangle \langle i | \mathbf{V} | 0 \rangle \\ &= -\langle 0 | \mathbf{W} | j \rangle \langle j | \mathbf{V} | i \rangle \langle i | \mathbf{W} | 0 \rangle + \langle 0 | \mathbf{W} | j \rangle \langle j | \mathbf{W} | i \rangle \langle i | \mathbf{V} | 0 \rangle, \end{aligned} \quad (191)$$

2149 and for the inner two terms:

$$2151 \quad \begin{aligned} &-\langle 0 | \mathbf{W} | i \rangle \langle i | \mathbf{V} | j \rangle \langle j | \mathbf{W} | 0 \rangle - \langle 0 | \mathbf{W} | i \rangle \langle i | \mathbf{W} | j \rangle \langle j | \mathbf{V} | 0 \rangle \\ &+ \langle 0 | \mathbf{V} | i \rangle \langle i | \mathbf{W} | j \rangle \langle j | \mathbf{W} | 0 \rangle + \langle 0 | \mathbf{W} | i \rangle \langle i | \mathbf{W} | j \rangle \langle j | \mathbf{V} | 0 \rangle \\ &= -\langle 0 | \mathbf{W} | i \rangle \langle i | \mathbf{V} | j \rangle \langle j | \mathbf{W} | 0 \rangle + \langle 0 | \mathbf{V} | i \rangle \langle i | \mathbf{W} | j \rangle \langle j | \mathbf{W} | 0 \rangle. \end{aligned} \quad (192)$$

2155 We can then combine terms (along with the terms linear in \mathbf{W}) to see that:

$$2157 \quad \begin{aligned} \hat{\ell}_{;i,j} &= \frac{2I\bar{o}}{rN^2} \operatorname{Re} \{ \langle 0 | \mathbf{V} | i \rangle \langle i | \mathbf{W} | j \rangle \langle j | (\mathbf{W} + \mathbf{V}) | 0 \rangle \} \\ &\quad - \frac{2I\bar{o}}{rN^2} \operatorname{Re} \{ \langle 0 | \mathbf{W} | i \rangle \langle i | \mathbf{V} | j \rangle \langle j | (\mathbf{W} + \mathbf{V}) | 0 \rangle \}. \end{aligned} \quad (193)$$

We can simplify the Hessian even further if we assume that $\frac{\sigma_o}{\sigma} = o(1)$ by dropping terms that are w.h.p. subleading in $\frac{\sigma_o}{\sigma}$. In particular, consider the following terms:

$$\begin{aligned} & \frac{1}{rN} \operatorname{Re} \{ \langle 0 | \mathbf{V} | i \rangle \langle i | \mathbf{W} | j \rangle \langle j | \mathbf{V} | 0 \rangle \} \\ & \leq \frac{N-r}{N} \left| \frac{1}{\sqrt{N-r}} \langle 0 | \mathbf{V} | i \rangle \right| \left| \frac{1}{r} \langle i | \mathbf{W} | j \rangle \right| \left| \frac{1}{\sqrt{N-r}} \langle j | \mathbf{V} | 0 \rangle \right|, \end{aligned} \quad (194)$$

$$\begin{aligned} & \frac{1}{rN} \operatorname{Re} \{ \langle 0 | \mathbf{W} | i \rangle \langle i | \mathbf{V} | j \rangle \langle j | \mathbf{W} | 0 \rangle \} \\ & \leq \frac{N-r}{N} \left| \frac{1}{\sqrt{r}} \langle 0 | \mathbf{W} | i \rangle \right| \left| \frac{1}{N-r} \langle i | \mathbf{V} | j \rangle \right| \left| \frac{1}{\sqrt{r}} \langle j | \mathbf{W} | 0 \rangle \right|, \end{aligned} \quad (195)$$

$$\begin{aligned} & \frac{1}{rN} \operatorname{Re} \{ \langle 0 | \mathbf{W} | i \rangle \langle i | \mathbf{V} | j \rangle \langle j | \mathbf{V} | 0 \rangle \} \\ & \leq \frac{(N-r)^{\frac{3}{2}}}{\sqrt{r}N} \left| \frac{1}{\sqrt{r}} \langle 0 | \mathbf{W} | i \rangle \right| \left| \frac{1}{N-r} \langle i | \mathbf{V} | j \rangle \right| \left| \frac{1}{\sqrt{N-r}} \langle j | \mathbf{V} | 0 \rangle \right|. \end{aligned} \quad (196)$$

Recall that the diagonal elements of a β -Wishart matrix are χ^2 -distributed with βr degrees of freedom (see Appendix A.4) and thus have exponential tails. In Lemma 43—proven in Appendix N—we prove that the off-diagonal elements also have exponential tails. In particular, all three of these terms are of the form $\kappa R_1 R_2 R_3$, where $0 < \kappa < 1$ is some overall prefactor and the R_i are random variables obeying tail bounds of the form:

$$\mathbb{P}[R_i \geq 1+t] \leq K \exp(-C \min(t, t^2)) \quad (197)$$

for some universal constants $C, K > 0$. We thus have by the union bound:

$$\begin{aligned} & \mathbb{P} \left[\kappa R_1 R_2 R_3 \geq \left(\kappa^{\frac{1}{3}} + t \right)^3 \right] \\ & \leq \mathbb{P} \left[\kappa \max(R_1, R_2, R_3)^3 \geq \left(\kappa^{\frac{1}{3}} + t \right)^3 \right] \\ & \leq 3 \max \left(\mathbb{P} \left[\kappa^{\frac{1}{3}} R_1 \geq \kappa^{\frac{1}{3}} + t \right], \mathbb{P} \left[\kappa^{\frac{1}{3}} R_2 \geq \kappa^{\frac{1}{3}} + t \right], \mathbb{P} \left[\kappa^{\frac{1}{3}} R_3 \geq \kappa^{\frac{1}{3}} + t \right] \right) \\ & \leq 3K \exp \left(-C \min \left(\kappa^{-\frac{1}{3}} t, \kappa^{-\frac{2}{3}} t^2 \right) \right). \end{aligned} \quad (198)$$

Taking into account the p^2 matrix elements via the union bound, these terms thus converge to 0 in Ky Fan metric at a rate $\Omega \left(\kappa^{\frac{1}{3}} \log \left(\kappa^{-\frac{1}{3}} p \right) \right)$. We can similarly use Lemma 41 to justify replacing \mathbf{W} with a low-rank perturbation $\tilde{\mathbf{W}}$, which differs only in the removal of the first column of its Bartlett factor \mathbf{L}_W . Altogether, we have that we can take up to an order

$$\alpha_5 = O \left(\frac{\mathcal{N}^{\frac{1}{3}} (N-r)^{\frac{1}{3}} \log \left(\frac{p^2 N^{\frac{2}{3}}}{\mathcal{N}^{\frac{1}{3}} (N-r)^{\frac{1}{3}}} \right)}{N^{\frac{2}{3}}} \right) \leq O \left(\frac{\mathcal{N}^{\frac{1}{3}} \sigma_o^{\frac{2}{3}} \log \left(\frac{p^2 N^{\frac{1}{3}} \sigma_o^{\frac{2}{3}}}{\mathcal{N}^{\frac{1}{3}} \sigma_o^{\frac{2}{3}}} \right)}{N^{\frac{1}{3}} \sigma_o^{\frac{2}{3}}} \right) \quad (199)$$

error in Ky Fan metric:

$$\hat{\ell}_{i,j} = \frac{2I\bar{\sigma}}{rN^2} \operatorname{Re} \left\{ \langle 0 | \mathbf{V} | i \rangle \langle i | \tilde{\mathbf{W}} | j \rangle \langle j | \mathbf{W} | 0 \rangle \right\}. \quad (200)$$

At this point it is instructive to consider this joint distribution in terms of the marginal distributions of elements in the Bartlett decompositions of \mathbf{W} and \mathbf{V} . Let $N_{W,i,j}$ be the i.i.d. standard normally distributed random variables over \mathbb{F} in the off-diagonal elements of \mathbf{L}_W (with the convention $i < j$), and similarly for $N_{V,i,j}$ and the Bartlett factor \mathbf{L}_V of \mathbf{V} . Furthermore, let $\beta^{-\frac{1}{2}} \chi_{W,i}$ be the i.i.d. χ -distributed random variables along the diagonal of \mathbf{L}_W , and similarly for $\beta^{-\frac{1}{2}} \chi_{V,i}$ and \mathbf{L}_V .

We first condition on $\hat{\ell} = z$. From Eq. (188) this is equivalent to conditioning on $\chi_{W,0} = \sqrt{\frac{\beta r}{I\bar{\sigma}}} z$. Conditioned on this event, the first derivative is distributed as:

$$\hat{\ell}_{i|z} = \frac{2I\beta\bar{\sigma}}{rN} \sqrt{\frac{r}{I\bar{\sigma}}} z \chi_{V,0} \operatorname{Re} \left\{ i N_{W,0,i}^* N_{V,0,i} \right\}. \quad (201)$$

Similarly, the second derivative conditioned on this event is distributed as:

$$\hat{\ell}_{;i,j|z} = \frac{2I\beta\bar{o}}{rN^2} \sqrt{\frac{r}{I\bar{o}}} z \chi_{V,0} \operatorname{Re} \left\{ N_{V,0,i}^* \langle i | \tilde{\mathbf{W}} | j \rangle N_{W,0,j} \right\}. \quad (202)$$

Note that $\tilde{\mathbf{W}}$ is independent from $N_{W,0,i}$ exactly due to the removal of the first column of the Bartlett factor of \mathbf{W} .

We will next simplify further using the fact that $\chi_{V,0}$ concentrates around its mean, in particular using the χ -distribution tail bound (for $N - r > 0$, some universal constant $C > 0$, and bounded t):

$$\mathbb{P} \left[\frac{1}{(N-r)^{\frac{1}{4}}} \left| \chi - \sqrt{\beta(N-r)} \right| \geq t \right] \leq 2 \exp \left(-C\sqrt{N-r} \min(t, t^2) \right) \quad (203)$$

which follows from Lemma 42, proved in Appendix N. Just as we accounted for α_5 , up to an order

$$\alpha_6 = \mathcal{O} \left(\frac{\mathcal{N}^{\frac{1}{3}} r^{\frac{1}{6}} (N-r)^{\frac{1}{12}} \log \left(\frac{p^2 N^{\frac{2}{3}}}{\mathcal{N}^{\frac{1}{3}} r^{\frac{1}{6}} (N-r)^{\frac{1}{12}}} \right)}{N^{\frac{2}{3}}} \right) \leq \mathcal{O} \left(\frac{\mathcal{N}^{\frac{1}{3}} \sigma_o^{\frac{1}{6}} \log \left(\frac{p^2 N^{\frac{5}{12}} \bar{o}^{\frac{1}{6}}}{\mathcal{N}^{\frac{1}{3}} \sigma_o^{\frac{1}{6}}} \right)}{N^{\frac{5}{12}} \bar{o}^{\frac{1}{6}}} \right) \quad (204)$$

error in Ky Fan metric,

$$\hat{\ell}_{;i|z} = \frac{2I\beta\bar{o}}{N} \sqrt{\frac{\beta}{I\bar{o}}} \left(\frac{N}{r} - 1 \right) z \operatorname{Re} \left\{ i N_{W,0,i}^* N_{V,0,i} \right\}, \quad (205)$$

$$\hat{\ell}_{;i,j|z} = \frac{2I\beta\bar{o}}{N^2} \sqrt{\frac{\beta}{I\bar{o}}} \left(\frac{N}{r} - 1 \right) z \operatorname{Re} \left\{ N_{V,0,i}^* \langle i | \tilde{\mathbf{W}} | j \rangle N_{W,0,j} \right\}. \quad (206)$$

Using Lemma 25 we can simplify this further as:

$$\hat{\ell}_{;i|z} = \frac{2I\beta\sigma_o}{N} \sqrt{\frac{\beta z}{I\bar{o}}} \operatorname{Re} \left\{ i N_{W,0,i}^* N_{V,0,i} \right\}, \quad (207)$$

$$\hat{\ell}_{;i,j|z} = \frac{2I\beta\sigma_o}{N^2} \sqrt{\frac{\beta z}{I\bar{o}}} \operatorname{Re} \left\{ N_{V,0,i}^* \langle i | \tilde{\mathbf{W}} | j \rangle N_{W,0,j} \right\}. \quad (208)$$

We can simplify this even further by realizing that the distribution of $N_{V,0,i}$ is invariant under multiplication by an overall phase, i.e., it is circularly symmetric. We can thus absorb the argument of $N_{W,0,i}^*$ into $N_{V,0,i}$; we simultaneously conjugate $\tilde{\mathbf{W}}$ by unitary operators corresponding to this phase, so we need only consider the joint distribution of:

$$\hat{\ell}_{;i|z} = \frac{2I\beta\sigma_o}{N} \sqrt{\frac{z}{I\bar{o}}} \chi_i \operatorname{Re} \left\{ i N_{V,0,i} \right\}, \quad (209)$$

$$\hat{\ell}_{;i,j|z} = \frac{2I\beta\sigma_o}{N^2} \sqrt{\frac{z}{I\bar{o}}} \chi_j \langle i | \operatorname{Re} \left\{ N_{V,0,i}^* \tilde{\mathbf{W}} \right\} | j \rangle, \quad (210)$$

where here the χ_i are i.i.d. χ -distributed random variables with β degrees of freedom given by:

$$\chi_i \equiv \sqrt{\beta} |N_{W,0,i}|. \quad (211)$$

Noting $\sqrt{\beta} \operatorname{Re} \left\{ i N_{V,0,i} \right\}$ is distributed as a standard normal random variable proves Theorem 17 for sectors with $\mathbb{F} = \mathbb{C}, \mathbb{H}$, with accumulated error in Lévy-Prokhorov distance:

$$\pi' = \pi_1 + \alpha_3 + 2\alpha_4 + \alpha_6 = \mathcal{O} \left(\frac{p \log \log (\mathcal{N}^{-t} \epsilon^{-1})}{\log (\mathcal{N}^{-t} \epsilon^{-1})} + \frac{p \log (t)}{\sqrt{t}} + \sqrt{\frac{\mathcal{N} \log (N)}{N^{1.001}}} + \frac{\mathcal{N}^{\frac{1}{3}} \log (N)}{N^{\frac{5}{12}}} \right). \quad (212)$$

We finally consider the second derivative conditioned on both $\hat{\ell} = z$ and all $\hat{\ell}_{;i|z} = 0$ (for $z > 0$). Note that the χ distribution density is zero at 0. In particular, this conditioning is equivalent to conditioning on $\operatorname{Re} \left\{ i N_{V,0,i} \right\} = 0$. Furthermore, when $\mathbb{F} = \mathbb{H}$, recall that “i” could be any one

of i, j, k due to Eq. (140). In particular, this is conditioning on *all* nonreal components of $N_{V,0,i}$ equaling zero. Taking this all together yields:

$$\hat{\ell}_{;i,j|z,\mathbf{0}} = \frac{2I\beta\sigma_o}{N^2} \sqrt{\frac{z}{I\bar{o}}} \chi_j \operatorname{Re} \{N_{V,0,i}^*\} \langle i | \operatorname{Re} \{ \tilde{\mathbf{W}} \} | j \rangle. \quad (213)$$

As the real parts of complex or symplectic Wishart matrices are real Wishart with a factor of β more degrees of freedom and rescaled by a factor of β^{-1} , and similarly the real parts of standard complex or symplectic Gaussian random variables are standard real Gaussian random variables rescaled by a factor of $\beta^{-\frac{1}{2}}$, we can more conveniently rewrite this as

$$\hat{\ell}_{;i,j|z,\mathbf{0}} = \frac{2I\sigma_o}{N^2} \sqrt{\frac{z}{I\beta\bar{o}}} \hat{N}_i \chi_j \langle i | \hat{\mathbf{W}} | j \rangle, \quad (214)$$

where now $\hat{\mathbf{W}}$ is a real Wishart matrix with βr degrees of freedom and the \hat{N}_i are i.i.d. standard normal random variables. This proves Theorem 19 for sectors with $\mathbb{F} = \mathbb{C}, \mathbb{H}$, with accumulated error in Lévy–Prokhorov distance:

$$\pi'' = \pi_1 + \alpha_3 + 2\alpha_4 + \alpha_5 + \alpha_6 = O\left(\frac{p^2 \log \log (\mathcal{N}^{-t} \epsilon^{-1})}{\log (\mathcal{N}^{-t} \epsilon^{-1})} + \frac{p^2 \log (t)}{\sqrt{t}} + \frac{\mathcal{N}^{\frac{1}{3}} \log (N)}{N^{\frac{1}{3}}}\right). \quad (215)$$

D.3.2 $\mathbb{F} = \mathbb{R}$

We now consider when $\mathbb{F} = \mathbb{R}$, where we are interested in the joint distribution of:

$$\hat{\ell} = \frac{I\bar{o}}{r} \operatorname{Tr} \left(\rho \mathbf{X}^\top \tilde{\mathbf{O}} \mathbf{X} \right), \quad (216)$$

$$\hat{\ell}_{;i} = \frac{I\bar{o}}{rN} \operatorname{Tr} \left(\rho \mathbf{X}^\top \left[\mathbf{X} (|2i\rangle \langle 2i+1| - |2i+1\rangle \langle 2i|) \mathbf{X}^\top, \tilde{\mathbf{O}} \right] \mathbf{X} \right), \quad (217)$$

$$\hat{\ell}_{;i,j} = \frac{I\bar{o}}{rN^2} \operatorname{Tr} \left(\rho \mathbf{X}^\top \left[\mathbf{X} (|2j\rangle \langle 2j+1| - |2j+1\rangle \langle 2j|) \mathbf{X}^\top, \right. \right. \\ \left. \left. \left[\mathbf{X} (|2i\rangle \langle 2i+1| - |2i+1\rangle \langle 2i|) \mathbf{X}^\top, \tilde{\mathbf{O}} \right] \mathbf{X} \right] \right), \quad (218)$$

where now \mathbf{X} is a $N \times (2p+1)$ random matrix with i.i.d. standard Gaussian entries over \mathbb{R} . As the commutator of symmetric matrices is antisymmetric, and taking note of the identity:

$$\operatorname{Tr} (A [B, C]) = \operatorname{Tr} ([C, A] B), \quad (219)$$

we can rewrite this as:

$$\hat{\ell} = \frac{I\bar{o}}{r} \operatorname{Tr} \left(\rho \mathbf{X}^\top \tilde{\mathbf{O}} \mathbf{X} \right), \quad (220)$$

$$\hat{\ell}_{;i} = \frac{2I\bar{o}}{rN} \operatorname{Tr} \left(\rho \mathbf{X}^\top \left[\mathbf{X} |2i\rangle \langle 2i+1| \mathbf{X}^\top, \tilde{\mathbf{O}} \right] \mathbf{X} \right), \quad (221)$$

$$\hat{\ell}_{;i,j} = \frac{2I\bar{o}}{rN^2} \operatorname{Tr} \left(\rho \mathbf{X}^\top \left[\mathbf{X} |2j\rangle \langle 2j+1| \mathbf{X}^\top, \left[\mathbf{X} (|2i\rangle \langle 2i+1| - |2i+1\rangle \langle 2i|) \mathbf{X}^\top, \tilde{\mathbf{O}} \right] \mathbf{X} \right] \right). \quad (222)$$

Recall that

$$\tilde{\mathbf{O}} = \sum_{\mu=0}^{r-1} |\mu\rangle \langle \mu|. \quad (223)$$

Let \mathbf{C} be the complement of $\tilde{\mathbf{O}}$, i.e.,

$$\mathbf{C} = \sum_{\mu=r}^{N-1} |\mu\rangle \langle \mu|. \quad (224)$$

In particular, \mathbf{C} and $\tilde{\mathbf{O}}$ sum to the identity:

$$\mathbf{C} + \tilde{\mathbf{O}} = \mathbf{I}_N. \quad (225)$$

2322 Furthermore,

$$2323 \mathbf{W} \equiv \mathbf{X}^\top \tilde{\mathbf{O}} \mathbf{X}, \quad (226)$$

$$2325 \mathbf{V} \equiv \mathbf{X}^\top \mathbf{C} \mathbf{X}, \quad (227)$$

2326 are each Wishart-distributed matrices; the former with r degrees of freedom, and the latter with
2327 $N - r$. The independence of \mathbf{W} and \mathbf{V} can be checked from the associated Bartlett decompositions.
2328 We thus have that the distribution we are interested in is the joint distribution of:

$$2330 \hat{\ell} = \frac{I\bar{o}}{r} \langle 0 | \mathbf{W} | 0 \rangle, \quad (228)$$

$$2331 \hat{\ell}_{;i} = \frac{2I\bar{o}}{rN} \langle 0 | (\mathbf{X}^\top \mathbf{X} | 2i) \langle 2i + 1 | \mathbf{W} - \mathbf{W} | 2i) \langle 2i + 1 | \mathbf{X}^\top \mathbf{X} | 0 \rangle$$

$$2333 = -\frac{2I\bar{o}}{rN} \langle 0 | \mathbf{W} (|2i) \langle 2i + 1 | - |2i + 1) \langle 2i |) \mathbf{V} | 0 \rangle, \quad (229)$$

2336 and:

$$2338 \hat{\ell}_{;i,j} = \frac{2I\bar{o}}{rN^2} \langle 0 | \mathbf{X}^\top [\mathbf{X} | 2j) \langle 2j + 1 | \mathbf{X}^\top, [\mathbf{X} (|2i) \langle 2i + 1 | - |2i + 1) \langle 2i |) \mathbf{X}^\top, \mathbf{O}] \mathbf{X} | 0 \rangle$$

$$2339 = \frac{2I\bar{o}}{rN^2} \langle 0 | \mathbf{X}^\top \mathbf{X} | 2j) \langle 2j + 1 | \mathbf{X}^\top \mathbf{X} (|2i) \langle 2i + 1 | - |2i + 1) \langle 2i |) \mathbf{X}^\top \mathbf{O} \mathbf{X} | 0 \rangle$$

$$2341 + \frac{2I\bar{o}}{rN^2} \langle 0 | \mathbf{X}^\top \mathbf{O} \mathbf{X} (|2i) \langle 2i + 1 | - |2i + 1) \langle 2i |) \mathbf{X}^\top \mathbf{X} | 2j) \langle 2j + 1 | \mathbf{X}^\top \mathbf{X} | 0 \rangle$$

$$2342 - \frac{2I\bar{o}}{rN^2} \langle 0 | \mathbf{X}^\top \mathbf{X} (|2i) \langle 2i + 1 | - |2i + 1) \langle 2i |) \mathbf{X}^\top \mathbf{O} \mathbf{X} | 2j) \langle 2j + 1 | \mathbf{X}^\top \mathbf{X} | 0 \rangle$$

$$2343 - \frac{2I\bar{o}}{rN^2} \langle 0 | \mathbf{X}^\top \mathbf{X} | 2j) \langle 2j + 1 | \mathbf{X}^\top \mathbf{O} \mathbf{X} (|2i) \langle 2i + 1 | - |2i + 1) \langle 2i |) \mathbf{X}^\top \mathbf{X} | 0 \rangle$$

$$2344 = \frac{2I\bar{o}}{rN^2} \langle 0 | (\mathbf{W} + \mathbf{V}) | 2j) \langle 2j + 1 | (\mathbf{W} + \mathbf{V}) (|2i) \langle 2i + 1 | - |2i + 1) \langle 2i |) \mathbf{W} | 0 \rangle$$

$$2345 + \frac{2I\bar{o}}{rN^2} \langle 0 | \mathbf{W} (|2i) \langle 2i + 1 | - |2i + 1) \langle 2i |) (\mathbf{W} + \mathbf{V}) | 2j) \langle 2j + 1 | (\mathbf{W} + \mathbf{V}) | 0 \rangle$$

$$2346 - \frac{2I\bar{o}}{rN^2} \langle 0 | (\mathbf{W} + \mathbf{V}) (|2i) \langle 2i + 1 | - |2i + 1) \langle 2i |) \mathbf{W} | 2j) \langle 2j + 1 | (\mathbf{W} + \mathbf{V}) | 0 \rangle$$

$$2347 - \frac{2I\bar{o}}{rN^2} \langle 0 | (\mathbf{W} + \mathbf{V}) | 2j) \langle 2j + 1 | \mathbf{W} (|2i) \langle 2i + 1 | - |2i + 1) \langle 2i |) (\mathbf{W} + \mathbf{V}) | 0 \rangle. \quad (230)$$

2357 We can simplify the expression for the second derivative by grouping terms by their order in \mathbf{W} .
2358 Note first that all terms cubic in matrix elements of \mathbf{W} sum to zero. For the quadratic terms, note
2359 that for the outer two terms:

$$2361 + \langle 0 | \mathbf{V} | 2j) \langle 2j + 1 | \mathbf{W} (|2i) \langle 2i + 1 | - |2i + 1) \langle 2i |) \mathbf{W} | 0 \rangle$$

$$2362 + \langle 0 | \mathbf{W} | 2j) \langle 2j + 1 | \mathbf{V} (|2i) \langle 2i + 1 | - |2i + 1) \langle 2i |) \mathbf{W} | 0 \rangle$$

$$2363 - \langle 0 | \mathbf{V} | 2j) \langle 2j + 1 | \mathbf{W} (|2i) \langle 2i + 1 | - |2i + 1) \langle 2i |) \mathbf{W} | 0 \rangle$$

$$2364 - \langle 0 | \mathbf{W} | 2j) \langle 2j + 1 | \mathbf{W} (|2i) \langle 2i + 1 | - |2i + 1) \langle 2i |) \mathbf{V} | 0 \rangle$$

$$2365 = + \langle 0 | \mathbf{W} | 2j) \langle 2j + 1 | \mathbf{V} (|2i) \langle 2i + 1 | - |2i + 1) \langle 2i |) \mathbf{W} | 0 \rangle$$

$$2366 - \langle 0 | \mathbf{W} | 2j) \langle 2j + 1 | \mathbf{W} (|2i) \langle 2i + 1 | - |2i + 1) \langle 2i |) \mathbf{V} | 0 \rangle, \quad (231)$$

2368 and for the inner two terms:

$$2369 + \langle 0 | \mathbf{W} (|2i) \langle 2i + 1 | - |2i + 1) \langle 2i |) \mathbf{V} | 2j) \langle 2j + 1 | \mathbf{W} | 0 \rangle$$

$$2370 + \langle 0 | \mathbf{W} (|2i) \langle 2i + 1 | - |2i + 1) \langle 2i |) \mathbf{W} | 2j) \langle 2j + 1 | \mathbf{V} | 0 \rangle$$

$$2371 - \langle 0 | \mathbf{V} (|2i) \langle 2i + 1 | - |2i + 1) \langle 2i |) \mathbf{W} | 2j) \langle 2j + 1 | \mathbf{W} | 0 \rangle$$

$$2372 - \langle 0 | \mathbf{W} (|2i) \langle 2i + 1 | - |2i + 1) \langle 2i |) \mathbf{W} | 2j) \langle 2j + 1 | \mathbf{V} | 0 \rangle$$

$$2373 = + \langle 0 | \mathbf{W} (|2i) \langle 2i + 1 | - |2i + 1) \langle 2i |) \mathbf{V} | 2j) \langle 2j + 1 | \mathbf{W} | 0 \rangle$$

$$2374 - \langle 0 | \mathbf{V} (|2i) \langle 2i + 1 | - |2i + 1) \langle 2i |) \mathbf{W} | 2j) \langle 2j + 1 | \mathbf{W} | 0 \rangle. \quad (232)$$

2376
2377
2378
2379
2380
2381
2382
2383
2384
2385
2386
2387
2388
2389
2390
2391
2392
2393
2394
2395
2396
2397
2398
2399
2400
2401
2402
2403
2404
2405
2406
2407
2408
2409
2410
2411
2412
2413
2414
2415
2416
2417
2418
2419
2420
2421
2422
2423
2424
2425
2426
2427
2428
2429

We can, in turn, combine these terms to obtain:

$$\begin{aligned}
& + \langle 0 | \mathbf{W} | 2j \rangle \langle 2j + 1 | \mathbf{V} (|2i\rangle \langle 2i + 1| - |2i + 1\rangle \langle 2i|) \mathbf{W} | 0 \rangle \\
& - \langle 0 | \mathbf{W} | 2j \rangle \langle 2j + 1 | \mathbf{W} (|2i\rangle \langle 2i + 1| - |2i + 1\rangle \langle 2i|) \mathbf{V} | 0 \rangle \\
& + \langle 0 | \mathbf{W} (|2i\rangle \langle 2i + 1| - |2i + 1\rangle \langle 2i|) \mathbf{V} | 2j \rangle \langle 2j + 1 | \mathbf{W} | 0 \rangle \\
& - \langle 0 | \mathbf{V} (|2i\rangle \langle 2i + 1| - |2i + 1\rangle \langle 2i|) \mathbf{W} | 2j \rangle \langle 2j + 1 | \mathbf{W} | 0 \rangle \\
& = + \langle 0 | \mathbf{W} | 2j \rangle \langle 2j + 1 | \mathbf{V} (|2i\rangle \langle 2i + 1| - |2i + 1\rangle \langle 2i|) \mathbf{W} | 0 \rangle \\
& - \langle 0 | \mathbf{W} | 2j \rangle \langle 2j + 1 | \mathbf{W} (|2i\rangle \langle 2i + 1| - |2i + 1\rangle \langle 2i|) \mathbf{V} | 0 \rangle \\
& - \langle 0 | \mathbf{W} | 2j + 1 \rangle \langle 2j | \mathbf{V} (|2i\rangle \langle 2i + 1| - |2i + 1\rangle \langle 2i|) \mathbf{W} | 0 \rangle \\
& + \langle 0 | \mathbf{W} | 2j + 1 \rangle \langle 2j | \mathbf{W} (|2i\rangle \langle 2i + 1| - |2i + 1\rangle \langle 2i|) \mathbf{V} | 0 \rangle \\
& = + \langle 0 | \mathbf{W} (|2j\rangle \langle 2j + 1| - |2j + 1\rangle \langle 2j|) \mathbf{V} (|2i\rangle \langle 2i + 1| - |2i + 1\rangle \langle 2i|) \mathbf{W} | 0 \rangle \\
& - \langle 0 | \mathbf{W} (|2j\rangle \langle 2j + 1| - |2j + 1\rangle \langle 2j|) \mathbf{W} (|2i\rangle \langle 2i + 1| - |2i + 1\rangle \langle 2i|) \mathbf{V} | 0 \rangle.
\end{aligned} \tag{233}$$

We can then combine terms with those linear in \mathbf{W} to see that:

$$\begin{aligned}
\hat{\ell}_{i,j} &= \frac{2I\bar{o}}{rN^2} \langle 0 | \mathbf{W} (|2i\rangle \langle 2i + 1| - |2i + 1\rangle \langle 2i|) \mathbf{V} (|2j\rangle \langle 2j + 1| - |2j + 1\rangle \langle 2j|) (\mathbf{W} + \mathbf{V}) | 0 \rangle \\
& - \frac{2I\bar{o}}{rN^2} \langle 0 | \mathbf{V} (|2i\rangle \langle 2i + 1| - |2i + 1\rangle \langle 2i|) \mathbf{W} (|2j\rangle \langle 2j + 1| - |2j + 1\rangle \langle 2j|) (\mathbf{W} + \mathbf{V}) | 0 \rangle.
\end{aligned} \tag{234}$$

We now use two lemmas—proved in Appendix N—that will simplify the form of our answer by allowing us to justify the dropping of terms that are subleading w.h.p. In particular, we use Lemma 43 followed by Lemma 41 just as we did in Eqs. (194), (195), and (196). By the same logic, we have that we can take up to an order

$$\alpha_5 = O\left(\frac{\mathcal{N}^{\frac{1}{3}}(N-r)^{\frac{1}{3}} \log\left(\frac{p^2 N^{\frac{2}{3}}}{\mathcal{N}^{\frac{1}{3}}(N-r)^{\frac{1}{3}}}\right)}{N^{\frac{2}{3}}}\right) \leq O\left(\frac{\mathcal{N}^{\frac{1}{3}}\sigma_o^{\frac{2}{3}} \log\left(\frac{p^2 N^{\frac{1}{3}}\bar{o}^{\frac{2}{3}}}{\mathcal{N}^{\frac{1}{3}}\sigma_o^{\frac{2}{3}}}\right)}{N^{\frac{1}{3}}\bar{o}^{\frac{2}{3}}}\right) \tag{235}$$

error in Ky Fan metric:

$$\hat{\ell}_{i,j} = -\frac{2I\bar{o}}{rN^2} \langle 0 | \mathbf{V} (|2i\rangle \langle 2i + 1| - |2i + 1\rangle \langle 2i|) \tilde{\mathbf{W}} (|2j\rangle \langle 2j + 1| - |2j + 1\rangle \langle 2j|) \mathbf{W} | 0 \rangle, \tag{236}$$

where $\tilde{\mathbf{W}}$ is as \mathbf{W} , but with the first column of its Bartlett factor \mathbf{L}_W removed.

At this point it is instructive to consider this joint distribution in terms of the marginal distributions of elements in the Bartlett decompositions of \mathbf{W} and \mathbf{V} . Let $N_{W,i,j}$ be the i.i.d. standard normally distributed random variables over \mathbb{R} in the off-diagonal elements of \mathbf{L}_W (with the convention $i < j$), and similarly for $N_{V,i,j}$ and the Bartlett factor \mathbf{L}_V of \mathbf{V} . Furthermore, let $\chi_{W,i}$ be the i.i.d. χ -distributed random variables along the diagonal of \mathbf{L}_W , and similarly for $\chi_{V,i}$ and \mathbf{L}_V .

We first condition on $\hat{\ell} = z$. From Eq. (188) this is equivalent to conditioning on $\chi_{W,0} = \sqrt{\frac{r}{I\bar{o}}}z$. Conditioned on this event, the first derivative is distributed as:

$$\hat{\ell}_{i,j|z} = \frac{2I\bar{o}}{rN} \sqrt{\frac{r}{I\bar{o}}} z \chi_{V,0} (N_{W,0,2i+1} N_{V,0,2i} - N_{W,0,2i} N_{V,0,2i+1}). \tag{237}$$

Similarly, the second derivative conditioned on this event is distributed as:

$$\begin{aligned}
\hat{\ell}_{i,j|z} &= \frac{2I\bar{o}}{rN^2} \sqrt{\frac{r}{I\bar{o}}} z \chi_{V,0} \\
&\times (N_{V,0,2i+1} \langle 2i | - N_{V,0,2i} \langle 2i + 1 |) \tilde{\mathbf{W}} (|2j\rangle N_{W,0,2j+1} - |2j + 1\rangle N_{W,0,2j}).
\end{aligned} \tag{238}$$

Note that $\tilde{\mathbf{W}}$ is independent from $N_{W,0,i}$ exactly due to the removal of the first column of the Bartlett factor of \mathbf{W} .

We will next simplify further using the fact that $\chi_{V,0}$ concentrates around its mean, in particular using the χ -distribution tail bound (for $N - r > 0$, some universal constant $C > 0$, and bounded t):

$$\mathbb{P}\left[\frac{1}{(N-r)^{\frac{1}{4}}}\left|\chi - \sqrt{\beta(N-r)}\right| \geq t\right] \leq 2 \exp\left(-C\sqrt{N-r} \min(t, t^2)\right) \tag{239}$$

which follows from Lemma 42, proved in Appendix N. Just as we accounted for α_5 , up to an order

$$\alpha_6 = O\left(\frac{\mathcal{N}^{\frac{1}{3}} r^{\frac{1}{6}} (N-r)^{\frac{1}{12}} \log\left(\frac{p^2 N^{\frac{2}{3}}}{\mathcal{N}^{\frac{1}{3}} r^{\frac{1}{6}} (N-r)^{\frac{1}{12}}}\right)}{N^{\frac{2}{3}}}\right) \leq O\left(\frac{\mathcal{N}^{\frac{1}{3}} \sigma_o^{\frac{1}{6}} \log\left(\frac{p^2 N^{\frac{5}{12}} \bar{\sigma}^{\frac{1}{6}}}{\mathcal{N}^{\frac{1}{3}} \sigma_o^{\frac{1}{6}}}\right)}{N^{\frac{5}{12}} \bar{\sigma}^{\frac{1}{6}}}\right) \quad (240)$$

error in Ky Fan metric,

$$\hat{\ell}_{;i|z} = \frac{2I\bar{\sigma}}{N} \sqrt{\left(\frac{N}{r} - 1\right)} \frac{z}{I\bar{\sigma}} (N_{W,0,2i+1} N_{V,0,2i} - N_{W,0,2i} N_{V,0,2i+1}), \quad (241)$$

$$\hat{\ell}_{;i,j|z} = \frac{2I\bar{\sigma}}{N^2} \sqrt{\left(\frac{N}{r} - 1\right)} \frac{z}{I\bar{\sigma}} \quad (242)$$

$$\times (N_{V,0,2i+1} \langle 2i | - N_{V,0,2i} \langle 2i+1 |) \tilde{\mathbf{W}} (|2j\rangle N_{W,0,2j+1} - |2j+1\rangle N_{W,0,2j}).$$

Using Lemma 25 we can simplify this further as:

$$\hat{\ell}_{;i|z} = \frac{2I\sigma_o}{N} \sqrt{\frac{z}{I\bar{\sigma}}} (N_{W,0,2i+1} N_{V,0,2i} - N_{W,0,2i} N_{V,0,2i+1}), \quad (243)$$

$$\hat{\ell}_{;i,j|z} = \frac{2I\sigma_o}{N^2} \sqrt{\frac{z}{I\bar{\sigma}}} (N_{V,0,2i+1} \langle 2i | - N_{V,0,2i} \langle 2i+1 |) \tilde{\mathbf{W}} (|2j\rangle N_{W,0,2j+1} - |2j+1\rangle N_{W,0,2j}). \quad (244)$$

We can simplify this even further by realizing that the joint distribution of $(N_{V,0,2i}, N_{V,0,2i+1})$ is invariant under orthogonal transformations, i.e., it is circularly symmetric. We can thus simultaneously transform $(N_{V,0,2i}, N_{V,0,2i+1})$ and conjugate $\tilde{\mathbf{W}}$ by orthogonal operators such that we need only consider the joint distribution of:

$$\hat{\ell}_{;i|z} = \frac{2I\sigma_o}{N} \sqrt{\frac{z}{I\bar{\sigma}}} \chi_i N_{V,0,2i+1}, \quad (245)$$

$$\hat{\ell}_{;i,j|z} = \frac{2I\sigma_o}{N^2} \sqrt{\frac{z}{I\bar{\sigma}}} \chi_j (N_{V,0,2i} \langle 2i | - N_{V,0,2i+1} \langle 2i+1 |) \tilde{\mathbf{W}} |2j\rangle, \quad (246)$$

where here the χ_i are i.i.d. χ -distributed random variables with 2 degrees of freedom given by:

$$\chi_i \equiv \sqrt{N_{W,0,2i}^2 + N_{W,0,2i+1}^2}. \quad (247)$$

This proves Theorem 17 for sectors with $\mathbb{F} = \mathbb{R}$, with accumulated error in Lévy–Prokhorov distance:

$$\pi' = \pi_1 + \alpha_3 + 2\alpha_4 + \alpha_6 = O\left(\frac{p \log \log (\mathcal{N}^{-t} \epsilon^{-1})}{\log (\mathcal{N}^{-t} \epsilon^{-1})} + \frac{p \log (t)}{\sqrt{t}} + \sqrt{\frac{\mathcal{N} \log (N)}{N^{1.001}}} + \frac{\mathcal{N}^{\frac{1}{3}} \log (N)}{N^{\frac{5}{12}}}\right). \quad (248)$$

We finally consider the second derivative conditioned on both $\hat{\ell} = z$ and all $\hat{\ell}_{;i|z} = 0$ (for $z > 0$). Note that the χ distribution density is zero at 0. In particular, this conditioning is equivalent to conditioning on $N_{V,0,2i+1} = 0$, yielding:

$$\hat{\ell}_{;i,j|z,0} = \frac{2I\sigma_o}{N^2} \sqrt{\frac{z}{I\bar{\sigma}}} N_{V,0,2i} \chi_j \langle 2i | \tilde{\mathbf{W}} |2j\rangle. \quad (249)$$

As submatrices of Wishart matrices are also Wishart, we can more conveniently rewrite this by taking \tilde{W} to be $p \times p$:

$$\hat{\ell}_{;i,j|z,0} = \frac{2I\sigma_o}{N^2} \sqrt{\frac{z}{I\bar{\sigma}}} \hat{N}_i \chi_j \langle i | \tilde{\mathbf{W}} |j\rangle, \quad (250)$$

where the \hat{N}_i are i.i.d. standard normal random variables. This proves Theorem 19 for sectors with $\mathbb{F} = \mathbb{R}$, with accumulated error in Lévy–Prokhorov distance:

$$\pi'' = \pi_1 + \alpha_3 + 2\alpha_4 + \alpha_5 + \alpha_6 = O\left(\frac{p^2 \log \log (\mathcal{N}^{-t} \epsilon^{-1})}{\log (\mathcal{N}^{-t} \epsilon^{-1})} + \frac{p^2 \log (t)}{\sqrt{t}} + \frac{\mathcal{N}^{\frac{1}{3}} \log (N)}{N^{\frac{1}{3}}}\right). \quad (251)$$

E QUANTUM ALGORITHM FOR DETERMINING ASYMPTOTIC TRAINABILITY

We elaborate here on a claim made in Sec. 6 that one can efficiently estimate $\text{Tr} \left((O^\alpha)^2 \right)$ and $\text{Tr} \left((\rho^\alpha)^2 \right)$ on a quantum computer and thus, by Definition 7, determine whether a given quantum neural network architecture is asymptotically trainable or not. We assume as input the Jordan algebra $\mathcal{A} \cong \bigoplus_\alpha \mathcal{A}_\alpha$ associated with the model, given as orthonormal bases $\{B_{\alpha,i}\}_i$ spanning each \mathcal{A}_α . Note that the corresponding traces $\text{Tr}_\alpha \left((O^\alpha)^2 \right)$ and $\text{Tr}_\alpha \left((\rho^\alpha)^2 \right)$ in the defining representation of \mathcal{A}_α can then also be evaluated using Eq. (11):

$$\text{Tr}(\cdot) = I_\alpha \text{Tr}_\alpha(\cdot); \quad (252)$$

in particular, I_α can be calculated from the ratio of traces of any $B_{\alpha,i}$ in the defining representation of $\mathcal{H}^N(\mathbb{C}) \supseteq \mathcal{A}_\alpha$ and that of \mathcal{A}_α .

We begin with $\text{Tr} \left((\rho^\alpha)^2 \right)$. It is immediate from the orthonormality of the $B_{\alpha,i}$ that:

$$\text{Tr} \left((\rho^\alpha)^2 \right) = \sum_i \text{Tr} \left(B_{\alpha,i} \rho^2 \right); \quad (253)$$

in particular, $\text{Tr} \left((\rho^\alpha)^2 \right)$ is calculable via, for instance, phase estimation (Nielsen & Chuang, 2010a) of the $B_{\alpha,i}$.

This leaves only $\text{Tr} \left((O^\alpha)^2 \right)$. If O were an efficiently preparable, valid quantum state, one could immediately calculate $\text{Tr} \left((O^\alpha)^2 \right)$ in the same way as $\text{Tr} \left((\rho^\alpha)^2 \right)$. In the more general case where, for instance, O is given via its Pauli decomposition:

$$O = \sum_{i=1}^S o_i P_i, \quad (254)$$

one can proceed similarly via a block encoding using the linear combination of unitaries (LCU) subroutine (Childs & Wiebe, 2012).

To outline this procedure, we assume O is an observable on n qubits. Let:

$$f(O) \equiv \sum_{i=1}^S \frac{o_i}{2} P_i \otimes (I_2 + Z), \quad (255)$$

where I_N is the $N \times N$ identity matrix and Z the 2×2 single-qubit Pauli Z matrix, and consider the $(n+1)$ -qubit quantum state:

$$\omega_0 \equiv \frac{I_{2^n}}{2^n} \otimes |0\rangle\langle 0|. \quad (256)$$

We have by construction that:

$$O^2 \otimes |0\rangle\langle 0| = f(O) \omega_0 f(O). \quad (257)$$

Thus,

$$\begin{aligned} \text{Tr} \left((O^\alpha)^2 \right) &= \sum_i \text{Tr} \left(B_{\alpha,i} O^2 B_{\alpha,i} \right) \\ &= \sum_i \text{Tr} \left((B_{\alpha,i})^2 f(O) \omega_0 f(O) \right). \end{aligned} \quad (258)$$

Of course, $f(O)$ is not necessarily unitary, so naive application of $f(O)$ is impossible on a quantum computer. However, one can proceed using the standard LCU subroutine (Childs & Wiebe, 2012). Assuming O is normalized such that its root mean square eigenvalue is $\Theta(1)$,

$$\text{Tr} \left(f(O) \omega_0 f(O) \right) = 2^{-n} \text{Tr} \left(O^2 \right) = \Theta(1). \quad (259)$$

Following the protocol described by Theorem 1 of Chakraborty (2024), then, each $\text{Tr} \left((B_{\alpha,i})^2 f(O) \omega_0 f(O) \right)$ can be estimated to an additive error ϵ using one ancilla qubit and $O \left(\frac{S^4}{\epsilon^2} \right)$ samples through an LCU block encoding.

F OPEN QUESTIONS

As our results are formal, there is room for violation of their conclusions in the settings where their assumptions do not hold. One example of this is the local minima distribution of quantum neural networks discussed in Appendix C.3. The results given there assume that the eigenvalue spectrum of the objective observable is concentrated as is typical of low-weight fermionic (Feng et al., 2019) and local spin Hamiltonians (Erdős & Schröder, 2014). However, the behavior of the local minima density differs when the spectrum is not concentrated, as is typical of observables drawn from the Gaussian unitary ensemble (GUE) and nonlocal spin systems. Observables from this latter class are known to be “quantumly easy” in the sense that phase estimation performed on the maximally mixed state efficiently prepares the ground state (Chen et al., 2024). Based on this intuition these nonlocal systems may also yield local minima distributions for quantum neural networks more amenable to variational optimization even at shallow depth. If this is true, this may give a natural class of optimization problems which are easy to solve given access to a quantum computer.

Furthermore, we nowhere discuss the impact of noise: we assume all operations implemented by the quantum neural network are unitary, not general quantum channels. It is already known that noise can affect the presence of barren plateaus (Wang et al., 2021) as well as local minima (Li & Hernandez, 2024; Mele et al., 2024) in variational loss landscapes. Heuristically the former can be understood as an additional channel mixing the variational ansatz over Hilbert space, and the latter as noise effectively limiting the number of parameters influencing the variational ansatz and thus always keeping the network in the underparameterized regime. We hope in the future to make this heuristic understanding rigorous.

G REDUCTION OF SPIN FACTOR SECTORS TO REAL SYMMETRIC SECTORS

We elaborate here on the claim made in Appendix B.1 that the \mathcal{A}_α where $\mathcal{A}_\alpha \cong \mathcal{L}_{N_\alpha}$ —i.e., the spin factor sectors—are conceptually equivalent to the case $\mathcal{A}_\alpha \cong \mathcal{H}^{N_\alpha-1}(\mathbb{R})$. To see this, consider a parameterization of $T^\alpha(\boldsymbol{\theta})$ in the spin factor case:

$$T(\boldsymbol{\theta}) = g_0^{\alpha\top} \prod_{i=1}^p g_i^\alpha \exp(\theta_i A_i^\alpha) g_i^{\alpha\top}, \quad (260)$$

where $g_i^\alpha \in \text{SO}(N-1)$ and $A_i^\alpha \in \mathfrak{so}(N-1)$. As stated in Appendix B.1, g_i^α acts trivially on the first component in the defining representation. Because of this (assuming the loss is shifted by a constant to always be nonnegative), the associated $\ell^\alpha(\boldsymbol{\theta}; \boldsymbol{\rho})$ is just equal to the square root of an instance of the $\mathcal{H}^{N_\alpha-1}(\mathbb{R})$ case with initial state:

$$\tilde{\boldsymbol{\rho}}^\alpha \equiv \boldsymbol{\rho}^{\alpha\top} \otimes \boldsymbol{\rho}^\alpha \quad (261)$$

and objective observable:

$$\tilde{\boldsymbol{O}}^\alpha \equiv \boldsymbol{O}^\alpha \otimes \boldsymbol{O}^{\alpha\top}. \quad (262)$$

H REDUCTION OF PREVIOUSLY STUDIED ANSATZ CLASSES TO JORDAN ALGEBRA-SUPPORTED ANSATZES

We here show that Jordan algebra-supported ansatzes (JASAs) are reduced to by the algebraically defined classes of variational ansatzes previously introduced in the literature, with a primary focus on the Lie algebra-supported ansatzes (LASAs) (Fontana et al., 2024). A reduction in the other direction—i.e., given a Jordan algebra-defined ansatz, define an equivalent LASA—is not possible. As a simple example, consider the case when the objective observable \boldsymbol{O} is a real symmetric matrix and the variational ansatz is PT symmetric (i.e., real) but has no other symmetries. This fits our framework with Jordan algebra $\mathcal{A} \cong \mathcal{H}^N(\mathbb{R})$ and automorphism group given by the action of conjugation by $\text{SO}(N)$. However, $\mathfrak{so}(N)$ in the defining representation consists of antisymmetric matrices, so this is not a LASA. We also give a reduction from the variational matchgate circuits of Diaz et al. (2023) to our Jordan algebraic framework, and will later demonstrate that such a reduction to LASAs is not possible.

2592 We first discuss how JASAs are reduced to by LASAs. The loss function of a LASA is of the form:

$$2593 \ell(\boldsymbol{\theta}; \rho) = \tau(\rho, T(\boldsymbol{\theta})O), \quad (263)$$

2594 where O belongs to some dynamical Lie algebra \mathfrak{g} , $T(\boldsymbol{\theta})$ is the adjoint action of some element of the
2595 compact Lie group G associated with \mathfrak{g} , and ρ is arbitrary. As demonstrated in Fontana et al. (2024),
2596 only the projection of ρ onto the \mathfrak{g} contributes to the loss, so we consider when $\rho \in \mathfrak{g}$ WLOG.
2597

2598 G is a compact Lie group. Thus, up to an Abelian sector that does not contribute to the loss (as
2599 $O \in \mathfrak{g}$), G can be written as the direct sum of simple compact Lie groups. By the well-known
2600 classification of simple compact Lie groups these are either the classical $\mathfrak{so}(N)$, $\mathfrak{su}(N)$, $\mathfrak{sp}(N)$, or
2601 one of the exceptional Lie groups. The exceptional Lie groups are of fixed dimension—i.e., there is
2602 no sense of an asymptotic limit—so we consider here the classical cases only.

2603 We begin with the case $\mathfrak{g} \cong \mathfrak{su}(N)$, with defining representation $d\phi$ over \mathbb{C} . It is easy to see that
2604 $i d\phi(O)$ and $i d\phi(\rho)$ are in the defining representation of the Jordan algebra $\mathcal{H}(\mathbb{C})$, with identical
2605 trace form. Such a simple component is thus equivalent to a JASA with simple subalgebra isomor-
2606 phic to $\mathcal{H}^N(\mathbb{C})$.

2607 We now consider when $\mathfrak{g} \cong \mathfrak{su}(N)$, with defining representation $d\phi$ over \mathbb{H} . As the loss is real-
2608 valued we have the expansion in this sector:

$$2609 \ell(\boldsymbol{\theta}; \rho) \propto \text{Tr} \left(\text{Re} \{i d\phi(\rho)\} d\phi(U(\boldsymbol{\theta})) \text{Re} \{i d\phi(O)\} d\phi(U(\boldsymbol{\theta}))^\dagger \right) \\ 2610 + \text{Tr} \left(\text{Re} \{j d\phi(\rho)\} d\phi(U(\boldsymbol{\theta})) \text{Re} \{j d\phi(O)\} d\phi(U(\boldsymbol{\theta}))^\dagger \right) \quad (264) \\ 2611 + \text{Tr} \left(\text{Re} \{k d\phi(\rho)\} d\phi(U(\boldsymbol{\theta})) \text{Re} \{k d\phi(O)\} d\phi(U(\boldsymbol{\theta}))^\dagger \right). \\ 2612 \\ 2613 \\ 2614$$

2615 Each of these real parts is Hermitian and can be written as elements of the defining representation
2616 of $\mathcal{H}^N(\mathbb{H})$, reducing this simple component to that of a JASA.

2617 We finally consider when $\mathfrak{g} \cong \mathfrak{so}(N)$. There are more subtleties compared with the other classical
2618 Lie algebras as, in the defining representation over \mathbb{R} , there is no simple relation between Hermitian
2619 and anti-Hermitian operators. To make the following concrete, we here consider the representation
2620 of $\mathfrak{so}(N)$ defined by Majorana operators γ_i , with generators:

$$2621 M_{ij} = \frac{1}{2} [\gamma_i, \gamma_j], \quad (265)$$

2622 i.e., polynomials quadratic in Majorana operators. These are the generators of the spin algebra \mathfrak{spin}
2623 associated with the vector space spanned by the γ_i and thus is isomorphic to $\mathfrak{so}(N)$. Furthermore,
2624 the gates generated by the M_{ij} are just matchgates. The loss function variance of variational match-
2625 gate circuits was studied in detail in Diaz et al. (2023). The authors there also consider when $O \notin \mathfrak{g}$,
2626 making this an example of a barren plateau result beyond the LASA formalism.
2627

2628 We now demonstrate how this setting fits into the Jordan algebra picture. Let \mathcal{A} be the Jordan algebra
2629 of Majorana fermions, with Jordan multiplication \circ given by half the anticommutator. As in the Lie
2630 algebra picture, the dimension of this algebra grows exponentially with n ; similarly, the Jordan al-
2631 gebra generated by k -body ($2 < k < n$) Majorana terms has dimension growing exponentially with
2632 n , as does the generated algebra when acting on these k -body terms via conjugation by a matchgate.
2633 However, though these operators are not closed as an algebra under matchgate conjugation, they *are*
2634 closed as a vector space. Furthermore, the trace form of a k -body term ψ and a k' -body term ψ' is
2635 always zero when $k \neq k'$. This allows us to instead consider O as a member of a different Jordan
2636 algebra \mathcal{A}' , with the same underlying vector space as \mathcal{A} and new Jordan multiplication operation \bullet
2637 defined as:

$$2638 \psi \bullet \psi' = \begin{cases} 1 & \text{if } \psi = \psi', \\ \psi & \text{if } \psi' = 1, \\ \psi' & \text{if } \psi = 1, \\ 0 & \text{otherwise.} \end{cases} \quad (266)$$

2642 Note that this new multiplication operation still satisfies the Jordan identity:

$$2643 (\psi \bullet \psi') (\psi \bullet \psi) = \psi \bullet \psi' \\ 2644 = \psi \bullet (\psi' \bullet 1) \\ 2645 = \psi \bullet (\psi' \bullet (\psi \bullet \psi)). \quad (267)$$

2646 It also preserves the canonical trace form, as:
2647

$$2648 \quad \text{Tr}(L_\circ(\psi \circ \psi')) = 2^n \delta_{\psi, \psi'} = \text{Tr}(L_\bullet(\psi \bullet \psi')). \quad (268)$$

2649 Here, L_\bullet is the linear transformation associated with the Jordan multiplication \bullet :
2650

$$2651 \quad L_\bullet(u)v = u \bullet v, \quad (269)$$

2652 and similarly for L_\circ . The loss function in this sector is thus a JASA under the Jordan algebra \mathcal{A}' ,
2653 and thus both general matchgate ansatzes or the LASA setting with $\mathfrak{g} \cong \mathfrak{so}(N)$ reduce to JASAs.

2654 Interestingly, one can show that the matchgate setting *cannot* generally be described by a LASA
2655 while maintaining the trace form, even when the Lie bracket is redefined. To see this, we consider
2656 the adjoint endomorphism

$$2657 \quad \text{ad}_{[\cdot, \cdot]}(u)v = [u, v] \quad (270)$$

2658 associated with the Lie bracket $[\cdot, \cdot]$. The canonical trace form for Lie algebras is the *Killing form*,
2659 given by:

$$2660 \quad K(u, v) = \text{Tr}(\text{ad}_{[\cdot, \cdot]}(u) \text{ad}_{[\cdot, \cdot]}(v)). \quad (271)$$

2661 Note that (as $\text{ad}_{[\cdot, \cdot]}(u)$ is antisymmetric):
2662

$$2663 \quad K(u, u) = -\|\text{ad}_{[\cdot, \cdot]}(u)\|_F^2. \quad (272)$$

2664 For u a Majorana monomial under the usual Lie algebra of Majoranas, this is just (minus) the
2665 number of terms anticommuting with the monomial. For u a degree-1 monomial this will grow
2666 exponentially in the number of Majoranas n . If instead one considered defining a new Lie bracket
2667 $[\cdot, \cdot]'$ that is zero when mapping from k -degree Majorana monomials to those of different degree,
2668 $K(\cdot, \cdot)$ would always be zero on degree-1 Majorana monomials—in particular, the trace form is not
2669 preserved when modifying the Lie bracket.
2670

2671

2672 I REDUCTION OF LOW-PURITY INPUTS TO MAXIMALLY MIXED INPUTS

2673

2674 We elaborate here on the claim made in Appendix B.2 that if the input state ρ projected into a simple
2675 sector \mathcal{A}_α has sufficiently low purity, the loss contribution from this sector does not contribute in any
2676 meaningful way. This is true even when the loss is rescaled by some $\mathcal{N} \leq O(N^{0.99})$ as we consider
2677 elsewhere. Intuitively this is due to ρ^α being close to maximally mixed and thus the contribution to
2678 the loss in this algebraic sector is effectively constant. We formally state this claim as the following
2679 lemma.

2680 **Lemma 29** (Low purity inputs are approximately trivial). *Let \mathcal{J}, ℓ be as in Theorem 13. For a*
2681 *$\rho \in \mathcal{R}$ such that for some α :*

$$2682 \quad \text{Tr}((\rho^\alpha)^2) \leq N_\alpha^{-0.999}, \quad (273)$$

2683 *we have for sufficiently large ϵ^{-1}, t that:*
2684

$$2685 \quad \text{Var}[\mathcal{N}\ell^\alpha(\theta; \rho)] \leq O\left(\frac{\mathcal{N}^2}{N_\alpha^{1.997}}\right) + O(\pi) = o(1) \quad (274)$$

2686 *for all $\mathcal{N} \leq O(N^{0.99})$.*
2687

2688

2689 *Proof.* Assume $\text{Tr}(\rho^\alpha) = 1$ WLOG by otherwise absorbing a factor of $\text{Tr}(\rho^\alpha)$ into O^α . Consider
2690 the decomposition into eigenspaces:

$$2691 \quad \rho^\alpha = \sigma^\alpha + \tau^\alpha, \quad (275)$$

2692 where the nonzero eigenvalues of σ^α are the largest $\lfloor N_\alpha^{0.999} \rfloor$ eigenvalues of ρ^α . Iterate this proce-
2693 dure

$$2694 \quad M_\alpha \equiv \left\lceil \frac{N_\alpha}{\lfloor N_\alpha^{0.999} \rfloor} \right\rceil \quad (276)$$

2695 total times to yield the decomposition:
2696

$$2697 \quad \rho^\alpha = \sigma^\alpha + \sum_{i=1}^{M_\alpha-1} \tau_i^\alpha. \quad (277)$$

2698

2699

The standard deviation of a sum of random variables is upper bounded by the sums of the standard deviations. Furthermore, from Eq. (277) we have that

$$\mathrm{Tr} \left((\boldsymbol{\rho}^\alpha)^2 \right) \geq \mathrm{Tr} \left((\boldsymbol{\sigma}^\alpha)^2 \right). \quad (278)$$

We can thus upper bound the standard deviation of the loss by applying Corollary 20 to the terms in Eq. (277):

$$\begin{aligned} \mathrm{Var} [\mathcal{N}\ell^\alpha(\boldsymbol{\theta}; \boldsymbol{\rho})] &\leq \mathcal{N}^2 M_\alpha^2 \frac{\mathrm{Tr} \left((\mathbf{O}^\alpha)^2 \right) \mathrm{Tr} \left((\boldsymbol{\sigma}^\alpha)^2 \right)}{\dim_{\mathbb{R}}(\mathfrak{g}_\alpha)} + \mathcal{O}(\pi) \\ &\leq \mathcal{N}^2 M_\alpha^2 \frac{\mathrm{Tr} \left((\mathbf{O}^\alpha)^2 \right) \mathrm{Tr} \left((\boldsymbol{\rho}^\alpha)^2 \right)}{\dim_{\mathbb{R}}(\mathfrak{g}_\alpha)} + \mathcal{O}(\pi) \\ &\leq \mathcal{O} \left(\frac{\mathcal{N}^2}{N^{1.997}} \right) + \mathcal{O}(\pi), \end{aligned} \quad (279)$$

yielding the final result. \square

J ϵ -APPROXIMATE 2-DESIGNS SUFFICE FOR THE \mathcal{G}_i

We here argue that one need only assume that the \mathcal{G}_i form ϵ -approximate 2-designs—not large- t t -designs—if one takes the overall normalization \mathcal{N} to be at most $\mathcal{O} \left(N^{\frac{2}{3}} \right)$ and if p grows sufficiently slowly with the N_α . More concretely, instead of taking Assumption 15 we may take the following assumption:

Assumption 30 (Scaling of parameter space with ansatz moments). The number of trained parameters p satisfies:¹¹

$$p^2 \leq \mathcal{O} \left(\min_{\alpha} \left(\frac{\log \left(\mathcal{N}^{-1} N_\alpha^{\frac{2}{3}} \right)}{\log \log \left(\mathcal{N}^{-1} N_\alpha^{\frac{2}{3}} \right)} \right) \right), \quad (280)$$

where the \mathcal{G}_i^α are i.i.d. ϵ -approximate 2-designs over $\mathcal{T}_\alpha = \mathrm{Aut}_1(\mathcal{A}_\alpha)$, where

$$\epsilon \leq \mathcal{O} \left(N_\alpha^{-\frac{2}{3}} \right). \quad (281)$$

To achieve this we will rely on tools from random matrix theory to bound the error of various mixed cumulants of our expressions under such a change. In particular, we will use the existence of so-called *limit distributions of all orders*.

Definition 31 (Limit distribution of all orders (Collins et al., 2007)). A sequence of sets of $N \times N$ random matrices $\{\mathbf{A}_i\}_{i=1}^d$ has a *limit distribution of all orders* if there exist functions $\alpha_{\mathbf{k}}$ such that the cumulants $\kappa_{\mathbf{k}}$ obey as $N \rightarrow \infty$:

$$\begin{aligned} \lim_{N \rightarrow \infty} N^{\|\mathbf{k}\|_1 - 2} \kappa_{\mathbf{k}} \left(\mathrm{Tr} \left(p_1(\mathbf{A}_1, \dots, \mathbf{A}_s) \right), \dots, \mathrm{Tr} \left(p_r(\mathbf{A}_1, \dots, \mathbf{A}_s) \right) \right) \\ = \alpha_{\mathbf{k}} \left(p_1(\mathbf{A}_1, \dots, \mathbf{A}_s), \dots, p_r(\mathbf{A}_1, \dots, \mathbf{A}_s) \right) \end{aligned} \quad (282)$$

for all fixed polynomials p_i . Here, $r \equiv \|\mathbf{k}\|_0$.

Simple classes of distributions which satisfy this property include constant matrices of bounded operator norm, Gaussian random matrices, Wishart random matrices, and Haar random unitaries. We here will be interested in a weaker notion of the existence of a limit distribution of all orders, requiring only that higher-order moments vanish sufficiently quickly and not necessarily that they have well-defined limits.

Definition 32 (Weak limit distribution of all orders). A sequence of sets of $N \times N$ random matrices $\{\mathbf{A}_i\}_{i=1}^d$ has a *weak limit distribution of all orders* if the cumulants $\kappa_{\mathbf{k}}$ obey as $N \rightarrow \infty$:

$$\kappa_{\mathbf{k}} \left(\mathrm{Tr} \left(p_1(\mathbf{A}_1, \dots, \mathbf{A}_s) \right), \dots, \mathrm{Tr} \left(p_r(\mathbf{A}_1, \dots, \mathbf{A}_s) \right) \right) = \mathcal{O} \left(N^{2 - \|\mathbf{k}\|_1} \right) \quad (283)$$

for all fixed polynomials p_i . Here, $r \equiv \|\mathbf{k}\|_0$.

¹¹As in Assumption 15, this can be weakened to a bound on p if one is only interested in the gradient.

Products of matrices with limit distributions of all orders over $\mathbb{F} = \mathbb{R}, \mathbb{C}, \mathbb{H}$ were considered in Mingo & Popa (2013); Collins et al. (2007); Redelmeier (2021) respectively, allowing us to prove the following lemma.

Lemma 33 (Existence of weak limit distributions of all orders). *Let $L_i, i \in [d]$ be multilinear functions of the form:*

$$L_i = \text{Tr}(\mathbf{M}_i \mathbf{h} \mathbf{O} \mathbf{h}^\dagger), \quad (284)$$

where \mathbf{h} is Haar random and independent from the \mathbf{M}_i . Assume all \mathbf{M}_i have weak limit distributions of all orders. For any cumulant $\kappa_{\mathbf{k}}(L_{i_1}, \dots, L_{i_r})$ with $\|\mathbf{k}\|_1 \geq 3$,

$$\kappa_{\mathbf{k}}(L_{i_1}, \dots, L_{i_r}) = O(N^{2-\|\mathbf{k}\|_1}). \quad (285)$$

Proof. Following:

1. Lemma 5.12 of Redelmeier (2012) when $\mathbb{F} = \mathbb{R}$;
2. Remark 4.8 of Collins et al. (2007) when $\mathbb{F} = \mathbb{C}$;
3. and Corollary 3.2 of Redelmeier (2021) when $\mathbb{F} = \mathbb{H}$,

the product of a matrix from a unitarily-invariant ensemble with any other independent matrix with a weak limit distribution of all orders has a weak limit distribution of all orders. This implies that the $\mathbf{M}_i \mathbf{h} \mathbf{O} \mathbf{h}^\dagger$ have a weak limit distribution of all orders. The vanishing of higher-order cumulants as in Eq. (285) then follows from Definition 32. \square

We thus have the following lemma.

Lemma 34 (Weak convergence to Haar random \mathbf{g}_i). *Assume Assumption 30. Let $L_i, i \in [d]$ be uniformly bounded multilinear functions of the form:*

$$L_i = \text{Tr} \left(\boldsymbol{\rho} \mathbf{h} \mathbf{O} \mathbf{h}^\dagger \prod_{j \in \mathcal{I}_i} \mathbf{g}_j \mathbf{M}_j \mathbf{g}_j^\dagger \right), \quad (286)$$

where \mathcal{I}_i is some index set associated with i of constant cardinality, \mathbf{h} is Haar random, $\mathbf{g}_j \sim \mathcal{G}_j$, $\boldsymbol{\rho}$ and the \mathbf{M}_i are fixed with bounded trace norm, and $d = o\left(\frac{\log(\mathcal{N}^{-1} N^{\frac{2}{3}})}{\log \log(\mathcal{N}^{-1} N^{\frac{2}{3}})}\right)$. Let \hat{L}_i be the same multilinear functions, where now the \mathbf{g}_j are Haar random. The joint distribution of $\mathcal{N} L_i$ differs from the joint distribution of $\mathcal{N} \hat{L}_i$ by an error at most

$$\pi_2 = O \left(\frac{d \log(\mathcal{N}^{-1} N^{\frac{2}{3}})}{(\mathcal{N}^{-1} N^{\frac{2}{3}})^{0.99}} \right) \quad (287)$$

in Lévy–Prokhorov metric for any $\mathcal{N} \leq O(N)$.

Proof. The $\mathbf{g}_j \mathbf{M}_j \mathbf{g}_j^\dagger$ have a limit distribution of all orders due to the \mathbf{M}_j being fixed. Thus, iteratively one can show the trace arguments of the L_i have weak limit distributions of all orders by the results cited in the proof of Lemma 33. Incorporating the bounded trace norm of the \mathbf{M}_i and the normalization \mathcal{N} gives an error for the cumulants of order $k > 2$:

$$\tilde{\epsilon} = O(\mathcal{N}^k N^{1-k}) \leq O\left(\left(\mathcal{N} N^{-\frac{2}{3}}\right)^k\right). \quad (288)$$

This also bounds the $k = 1, 2$ cases whenever

$$\epsilon \leq O\left(N^{-\frac{2}{3}}\right), \quad (289)$$

which is true by Assumption 30. The result then follows from Corollary 40, where μ is subleading as in the proof of Lemma 26. \square

K EQUICONTINUITY OF PROBABILITY DENSITIES

In Appendix D we prove that the joint distribution of the loss function and its first two derivatives for a certain random class of variational quantum loss functions converges in distribution to a fairly simple expression involving Wishart-distributed random matrices. However, Corollaries 14 and 18 are stated in terms of the *probability densities* of the loss and first derivative, and weak convergence does not necessarily imply pointwise convergence in densities. However, it is known that this *is* the case when the sequence of densities are *equicontinuous* and bounded (Boos, 1985). Whether or not this is true depends on the details of the distributions \mathcal{G}, \mathbf{H} . As an example, we here show this is true for the loss when \mathcal{G}, \mathbf{H} are Haar random, under reasonable assumptions (which we conjecture are not necessary) on the spectrum of the objective observable \mathbf{O} projected into any simple component.

Lemma 35 (Equicontinuity of loss density with Haar random ansatz). *Consider a sequence of \mathbf{O} as $N \rightarrow \infty$ with eigenvalues:*

$$0 = o_1 \leq o_2 \leq \dots \leq o_{N-1} \leq o_N = 1, \quad (290)$$

with mean eigenvalue \bar{o} and $\frac{\text{Tr}(\mathbf{O}^2)}{\text{Tr}(\mathbf{O}^2)N} = \frac{r}{N} = \Theta(1)$. For $\mathbf{w} \in \mathbb{R}^{N-2}$, define:

$$E_{\mathbf{w}} \equiv \sqrt{\bar{o} \cdot (\mathbf{w} \odot \mathbf{w})}, \quad (291)$$

$$R_{\mathbf{w}} \equiv \sqrt{\mathbf{w}^\top \cdot \mathbf{w}}, \quad (292)$$

where \bar{o} is the vector (o_2, \dots, o_{N-1}) and \odot denotes the Hadamard product. Let D_x be the domain of \mathbf{w} satisfying:

$$E_{\mathbf{w}}^2 \leq \bar{o} + N^{-\frac{1}{2}}x \leq 1 + E_{\mathbf{w}}^2 - R_{\mathbf{w}}^2. \quad (293)$$

Assume

$$p_{\sqrt{N}\ell}(x) = \frac{1}{\sqrt{N}(\bar{o} + N^{-\frac{1}{2}}x)S_{N-1}} \int_{D_x} d\mathbf{w} \left(1 - \left(\bar{o} + N^{-\frac{1}{2}}x \right)^{-1} E_{\mathbf{w}}^2 \right)^{-\frac{1}{2}} \quad (294)$$

is locally Lipschitz and bounded at any x as $N \rightarrow \infty$, with Lipschitz constants independent from N . Then the density of $\sqrt{r}\ell$ is bounded by a constant independent of N and is equicontinuous as a sequence in N .

Proof. By assumption $r = \Theta(N)$ so we consider the density of $\sqrt{N}\ell$ WLOG. Let $\{|o_i\rangle\}_{i=1}^N$ be the eigenvectors of \mathbf{O} associated with the eigenvalues o_i . Let:

$$|u\rangle = \sum_{i=1}^N u_i |o_i\rangle, \quad (295)$$

and let:

$$S(\mathbf{u}) = \sum_{i=2}^N o_i |u_i|^2. \quad (296)$$

Let $p_{\sqrt{N}\ell}(x)$ be the density of $\sqrt{N}\ell = \sqrt{N}\bar{o} + x$. We have that:

$$p_{\sqrt{N}\ell}(x) = S_{\beta N-1}^{-1} \int_{\|\mathbf{u}\|_2=1} d\mathbf{u} \delta\left(\sqrt{N}(S(\mathbf{u}) - \bar{o}) - x\right). \quad (297)$$

Here, $S_{\beta N-1}$ is the surface area of the $(\beta N - 1)$ -sphere:

$$S_{\beta N-1} = \frac{2\pi^{\frac{\beta N}{2}}}{\Gamma\left(\frac{\beta N}{2}\right)}. \quad (298)$$

We can integrate out the phases of the u_i and note that the delta function has no dependence on u_1 since $o_1 = 0$. Therefore,

$$p_{\sqrt{N}\ell}(x) = S_{N-1}^{-1} 2^N \int_{\substack{\|\tilde{\mathbf{u}}\|_2 \leq 1, \\ \tilde{\mathbf{u}} \geq 0}} d\tilde{\mathbf{u}} \delta\left(\sqrt{N}(S(\tilde{\mathbf{u}}) - \bar{o}) - x\right). \quad (299)$$

2862 Here, $\tilde{\mathbf{u}} \in \mathbb{R}_{\geq 0}^{N-1}$ has indices in $2, \dots, N$ such that $\tilde{u}_i = |u_i|$. We will similarly use $\tilde{\mathbf{o}}$ to denote \mathbf{o}
 2863 with the $o_1 = 0$ component projected out.
 2864

2865 We now proceed to show equicontinuity and uniform boundedness of $p_f(x)$ by showing that the
 2866 derivative of $p_f(x)$ is bounded everywhere by some N -independent constant. By symmetry we
 2867 restrict to the components of $v_N \geq 0$, introducing an overall factor of 2^β . We also introduce new
 2868 coordinates w_2, \dots, w_N , where $w_i = \tilde{u}_i$ for $2 \leq i \leq N-1$ and $w_N = S(\tilde{\mathbf{u}})$. The constraint
 2869 $\tilde{u}_N^2 \geq 0$ becomes in these coordinates:

$$2870 \sum_{i=2}^{N-1} o_i w_i^2 \leq w_N. \quad (300)$$

2871 Only the final row of the resulting Jacobian is nontrivial and is equal to:

$$2872 \mathbf{J}_N = 2\tilde{\mathbf{o}}^\top \odot \tilde{\mathbf{u}}^\top, \quad (301)$$

2873 where \odot denotes the Hadamard product. This gives a Jacobian determinant of (recalling that $o_N =$
 2874 1):

$$2875 \det(\mathbf{J}) = 2\tilde{u}_N = 2\sqrt{w_N - \sum_{i=2}^{N-1} o_i w_i^2}. \quad (302)$$

2876 Putting everything together yields a density:

$$2877 p_{\sqrt{N}\ell}(x) = S_{N-1}^{-1} 2^{N-2} \int_D d\mathbf{w} \frac{\delta(\sqrt{N}(w_N - \bar{o}) - x)}{\sqrt{w_N - \sum_{i=2}^{N-1} o_i w_i^2}}, \quad (303)$$

2878 where D is the domain of positive w_i obeying the normalization constraint of Eq. (300) as well as:

$$2879 w_N + \sum_{i=2}^{N-1} (1 - o_i) w_i^2 \leq 1. \quad (304)$$

2880 Defining:

$$2881 E_{\tilde{\mathbf{w}}}^2 = \sum_{i=2}^{N-1} o_i \tilde{w}_i^2, \quad (305)$$

$$2882 R_{\tilde{\mathbf{w}}}^2 = \sum_{i=2}^{N-1} \tilde{w}_i^2, \quad (306)$$

2883 and rescaling w_N we can rewrite this as:

$$2884 p_{\sqrt{N}\ell}(x) = \frac{1}{\sqrt{N} S_{N-1}} \int_{B^{N-2}} d\tilde{\mathbf{w}} \int_{\sqrt{N} E_{\tilde{\mathbf{w}}}^2}^{\sqrt{N}(E_{\tilde{\mathbf{w}}}^2 + 1 - R_{\tilde{\mathbf{w}}}^2)} dw_N \frac{\delta(w_N - \sqrt{N}\bar{o} - x)}{\sqrt{N^{-\frac{1}{2}} w_N - E_{\tilde{\mathbf{w}}}^2}}, \quad (307)$$

2885 where B^{N-2} is the unit $(N-2)$ -ball. Integrating out the delta function then yields:

$$2886 p_{\sqrt{N}\ell}(x) = \frac{1}{\sqrt{N} (\bar{o} + N^{-\frac{1}{2}} x) S_{N-1}} \int_{D_x} d\tilde{\mathbf{w}} \left(1 - (\bar{o} + N^{-\frac{1}{2}} x)^{-1} E_{\tilde{\mathbf{w}}}^2\right)^{-\frac{1}{2}}, \quad (308)$$

2887 where D_x is the domain of $\tilde{\mathbf{w}}$ satisfying:

$$2888 E_{\tilde{\mathbf{w}}}^2 \leq \bar{o} + N^{-\frac{1}{2}} x \leq 1 + E_{\tilde{\mathbf{w}}}^2 - R_{\tilde{\mathbf{w}}}^2. \quad (309)$$

2889 By assumption this is locally Lipschitz and bounded as $N \rightarrow \infty$ with Lipschitz constants indepen-
 2890 dent from N , and thus is equicontinuous and bounded. \square

We now discuss the main assumption of Lemma 35 in more detail, particularly the local Lipschitz continuity of Eq. (294). Note that the defining equations of D_x imply that all $w_i^2 \leq 1$, i.e., the domain of integration is always of $O(1)$ volume. This holds true even if some o_i vanish with N ; it is apparent that these terms do not contribute to the integrand at leading order O , and only the $\Theta(1)$ eigenvalues contribute. Because of this, we will perform an example calculation of the Lipschitz constants of Eq. (294) when all $o_i = \bar{o} = \frac{1}{2}$ for $i \neq 1, N$, though due to this observation similar results hold for all reasonable spectrums.

For such O ,

$$E_{\tilde{w}}^2 = \frac{1}{2} R_{\tilde{w}}^2, \quad (310)$$

so

$$\begin{aligned} p_{\sqrt{N}\ell}(x) &= \frac{\sqrt{2}}{\sqrt{N(1+2N^{-\frac{1}{2}}x)}S_{N-1}} \int_{D_x} d\tilde{w} \left(1 - \left(1 + 2N^{-\frac{1}{2}}x\right)^{-1} R_{\tilde{w}}^2\right)^{-\frac{1}{2}} \\ &= \frac{\sqrt{2}S_{N-3}}{\sqrt{N(1+2N^{-\frac{1}{2}}x)}S_{N-1}} \int_0^{\sqrt{1-2N^{-\frac{1}{2}}|x|}} dr r^{N-3} \left(1 - \left(1 + 2N^{-\frac{1}{2}}x\right)^{-1} r^2\right)^{-\frac{1}{2}} \\ &= \frac{S_{N-3}}{\sqrt{2N(1+2N^{-\frac{1}{2}}x)}S_{N-1}} \int_0^{1-2N^{-\frac{1}{2}}|x|} dr r^{\frac{N}{2}-2} \left(1 - \left(1 + 2N^{-\frac{1}{2}}x\right)^{-1} r\right)^{-\frac{1}{2}} \\ &= \frac{\left(1 - 2N^{-\frac{1}{2}}|x|\right)^{\frac{N}{2}-1} S_{N-3}}{\sqrt{2N(1+2N^{-\frac{1}{2}}x)}S_{N-1}} \int_0^1 dr r^{\frac{N}{2}-2} \left(1 - \frac{1 - 2N^{-\frac{1}{2}}|x|}{1 + 2N^{-\frac{1}{2}}x} r\right)^{-\frac{1}{2}} \\ &= \frac{\left(1 - 2N^{-\frac{1}{2}}|x|\right)^{\frac{N}{2}-1} \Gamma\left(\frac{N}{2}\right)}{\pi\sqrt{2N(1+2N^{-\frac{1}{2}}x)}\Gamma\left(\frac{N}{2}-1\right)} \int_0^1 dr r^{\frac{N}{2}-2} \left(1 - \frac{1 - 2N^{-\frac{1}{2}}|x|}{1 + 2N^{-\frac{1}{2}}x} r\right)^{-\frac{1}{2}}. \end{aligned} \quad (311)$$

This integral is just the integral definition of the hypergeometric function ${}_2F_1$, yielding:

$$p_{\sqrt{N}\ell}(x) = \frac{\left(1 - 2N^{-\frac{1}{2}}|x|\right)^{\frac{N}{2}-1}}{\pi\sqrt{2N(1+2N^{-\frac{1}{2}}x)}} {}_2F_1\left(\frac{1}{2}, \frac{N}{2} - 1; \frac{N}{2}; \frac{1 - 2N^{-\frac{1}{2}}|x|}{1 + 2N^{-\frac{1}{2}}x}\right). \quad (312)$$

We now consider the large- N asymptotics of the derivative of this expression. By Gauss's hypergeometric theorem, for $z = 1 + O(N^{-\frac{1}{2}})$,

$${}_2F_1\left(\frac{1}{2}, \frac{N}{2} - 1; \frac{N}{2}; z\right) = \left(1 + O(N^{-\frac{1}{2}})\right) \frac{\Gamma\left(\frac{N}{2}\right)\Gamma\left(\frac{1}{2}\right)}{\Gamma\left(\frac{N-1}{2}\right)\Gamma(1)} = O(\sqrt{N}). \quad (313)$$

Similarly, we have by the derivative rule for the hypergeometric function that:

$${}_2F_1'\left(\frac{1}{2}, \frac{N}{2} - 1; \frac{N}{2}; z\right) = \frac{1}{2} (1 - 2N^{-1}) {}_2F_1\left(\frac{3}{2}, \frac{N}{2}; \frac{N}{2} + 1; z\right) = O(N^{\frac{3}{2}}). \quad (314)$$

Taking into account the $N^{-\frac{1}{2}}$ factors scaling x thus yields (where defined):

$$p'_{\sqrt{N}\ell}(x) = O(1). \quad (315)$$

In particular, $p_{\sqrt{N}\ell}(x)$ is locally Lipschitz with Lipschitz constants independent from N . These O thus satisfy the assumptions of Lemma 35, as claimed.

L BOUNDING LARGE DEVIATIONS OF THE HESSIAN DETERMINANT

In calculating the density of local minima in Appendix C.3 we require analyzing the asymptotics of:

$$\mathbb{E}[\det(\mathbf{D}_z) \mathbf{1}\{\mathbf{D}_z \succeq \mathbf{0}\}], \quad (316)$$

where \mathbf{D}_z is as in Eq. (124), i.e., it is a Wishart-distributed random matrix normalized by its degrees of freedom parameter. As the determinant is exponentially sensitive to its argument, exponentially unlikely deviations in the convergence of the spectrum of \mathbf{D}_z to its asymptotic value can, in principle, cause this expectation to not converge to the determinant of the (almost sure) asymptotic spectral measure of \mathbf{D}_z . The same is true of the minimum eigenvalue of \mathbf{D}_z . We here show that these deviations do not asymptotically contribute to Eq. (316). We will here be brief as this is not the main focus of our work, and instead refer the interested reader to Anschuetz (2022) for a more detailed account of very similar arguments. In the following we consider only the case when $\gamma \leq 1$ as otherwise \mathbf{D}_z is never full-rank and Eq. (316) is identically zero.

We begin with the empirical spectral measure of \mathbf{D}_z . The empirical spectral measures of \mathbf{W} satisfies a large deviations principle at a speed p^2 with good rate function maximized by the Marčenko–Pastur distribution $d\mu_{\text{MP}}^\gamma(\lambda)$ (Hiai & Petz, 1998). As the determinant is only sensitive at a speed p , i.e.,

$$\det(\mathbf{D}_z) = \exp\left(p \int d\mu_{\text{MP}}^\gamma(\lambda) \ln(\lambda)\right), \quad (317)$$

by Varadhan’s lemma (Dembo & Zeitouni, 2010) the expected determinant converges to the determinant of our asymptotic expression for \mathbf{D}_z .

We now consider the smallest eigenvalue of \mathbf{D}_z . For fluctuations below its asymptotic value this satisfies a large deviations principle with good rate function at a speed p , where this rate function is maximized at the infimum of the support of the Marčenko–Pastur distribution; for fluctuations above its asymptotic value it satisfies a large deviations principle at a speed p^2 (Katzav & Pérez Castillo, 2010). By Varadhan’s lemma, then, only fluctuations below the asymptotic value potentially contribute, but as these fluctuations only decrease the value of the argument of Eq. (316) they do not contribute asymptotically (Dembo & Zeitouni, 2010). These together imply that Eq. (316) converges as claimed in Appendix C.3.

M LEMMAS ON CONVERGENCE IN DISTRIBUTION

We here prove a few helper lemmas that will allow us to obtain quantitative error bounds on the rate of convergence of two sequences of distributions given bounds on the differences of their moments or cumulants. We will achieve this by bounding the Lévy–Prokhorov distance between the two distributions, which metricizes weak convergence. We first restate a result of Berkes & Philipp (1979) which bounds this distance, with a slight modification due to our use of the infinity norm rather than the Euclidean norm (see Appendix A.5).

Theorem 36 (Slightly modified version of Lemma 2.2, Berkes & Philipp (1979)). *Let $p(\mathbf{x}), q(\mathbf{x})$ be two distributions on \mathbb{R}^d with characteristic functions $\phi(\mathbf{u}), \gamma(\mathbf{u})$, respectively. Assume there exists a C such that $\int_{\|\mathbf{x}\|_\infty \geq C} p(\mathbf{x}) \leq \mu$. There exists a universal constant K such that for all $T \geq \max(2C, Kd)$, the Lévy–Prokhorov distance between p and q is bounded by:*

$$\pi(p, q) \leq \left(\frac{T}{\pi}\right)^d \int_{\|\mathbf{u}\|_\infty \leq T} |\phi(\mathbf{u}) - \gamma(\mathbf{u})| d\mathbf{u} + \frac{16 \ln(T)}{T} + \mu. \quad (318)$$

Motivated by this, we prove a general bound on the error of the characteristic functions given errors in the moments.

Lemma 37 (Bound on characteristic functions from moments). *Let $p(\mathbf{x}), q(\mathbf{x})$ be two probability densities on \mathbb{R}^d with characteristic functions $\phi(\mathbf{u}), \gamma(\mathbf{u})$. Assume each moment of order $k \leq t$ of p differs from that of q by an additive error at most $\epsilon > 0$, and assume all moments of order $k > t$ are bounded by $(C\sqrt{k})^k$ for some k -independent $C \geq 0$. Then, for all \mathbf{u} such that $\|\mathbf{u}\|_1 \leq \frac{\sqrt{t+1}}{2\epsilon C}$,*

$$|\phi(\mathbf{u}) - \gamma(\mathbf{u})| \leq \epsilon \exp(\|\mathbf{u}\|_1) + \exp_2(1-t). \quad (319)$$

3024 *Proof.* We have by the definition of the characteristic function and the triangle inequality that:

$$3025 | \phi(\mathbf{u}) - \gamma(\mathbf{u}) | \leq \epsilon \sum_{\mathbf{k} \neq \mathbf{0} \in \mathbb{N}_0^d} \prod_{j=1}^d \frac{|(iu_j)^{k_j}|}{k_j!} + 2 \sum_{\|\mathbf{k}\|_1 \geq t+1} C^{\|\mathbf{k}\|_1} \|\mathbf{k}\|_1^{\frac{\|\mathbf{k}\|_1}{2}} \prod_{j=1}^d \frac{|(iu_j)^{k_j}|}{k_j!}. \quad (320)$$

3029 The first term has the upper bound:

$$3030 \epsilon \sum_{\mathbf{k} \neq \mathbf{0} \in \mathbb{N}_0^d} \prod_{j=1}^d \frac{|(iu_j)^{k_j}|}{k_j!} \leq \epsilon \exp(\|\mathbf{u}\|_1). \quad (321)$$

3034 For the second term, we have the upper bound (for $\|\mathbf{u}\|_1 < \frac{\sqrt{t+1}}{eC}$):

$$3035 \sum_{\|\mathbf{k}\|_1 \geq t+1} C^{\|\mathbf{k}\|_1} \|\mathbf{k}\|_1^{\frac{\|\mathbf{k}\|_1}{2}} \prod_{j=1}^d \frac{|u_j|^{k_j}}{k_j!} \leq \sum_{k \geq t+1} \frac{\|\sqrt{eC}\mathbf{u}\|_1^k}{\sqrt{k!}} \\ 3036 \leq \sum_{k \geq t+1} \left\| \frac{eC\mathbf{u}}{\sqrt{t+1}} \right\|_1^k \\ 3037 = \frac{\left(\frac{eC}{\sqrt{t+1}} \|\mathbf{u}\|_1 \right)^{t+1}}{1 - \frac{eC}{\sqrt{t+1}} \|\mathbf{u}\|_1}. \quad (322)$$

3045 When $\|\mathbf{u}\|_1 \leq \frac{\sqrt{t+1}}{2eC}$ this gives the bound:

$$3046 2 \sum_{\|\mathbf{k}\|_1 \geq t+1} C^{\|\mathbf{k}\|_1} \|\mathbf{k}\|_1^{\frac{\|\mathbf{k}\|_1}{2}} \prod_{j=1}^d \frac{|(iu_j)^{k_j}|}{k_j!} \leq 4 \left(\frac{eC}{\sqrt{t+1}} \|\mathbf{u}\|_1 \right)^{t+1} \leq \exp_2(1-t). \quad (323)$$

3050 Combining these bounds yields the final result. \square

3052 This characteristic function bound combined with Theorem 36 gives us the following corollary.

3053 **Corollary 38** (Bound on Lévy–Prokhorov metric from moments). *Let $p_N(\mathbf{x}), q_N(\mathbf{x})$ be two sequences of distributions on \mathbb{R}^d with characteristic functions $\phi_N(\mathbf{u}), \gamma_N(\mathbf{u})$, respectively. Assume each moment of order $k \leq t$ of p_N differs from that of q_N by an (N -dependent) additive error at most $\epsilon > 0$, and assume all moments of order $k > t$ (for t N -dependent) are bounded by $(C\sqrt{k})^k$ for some (k, N) -independent constant C . Assume as well that $\int_{\|\mathbf{x}\|_\infty \geq \frac{1}{2} \min\left(\frac{0.99}{d} \ln(\epsilon^{-1}), \frac{\sqrt{t+1}}{2eCd}\right)} p(\mathbf{x}) \leq \mu$ for some (N -dependent) μ . The Lévy–Prokhorov distance between p_N, q_N is bounded by:*

$$3061 \pi(p_N, q_N) = O\left(\frac{d \log \log(\epsilon^{-1})}{\log(\epsilon^{-1})} + \frac{d \log(t)}{\sqrt{t}} + \mu\right). \quad (324)$$

3064 *Proof.* Let $T = \min\left(\frac{0.99}{d} \ln(\epsilon^{-1}), \frac{\sqrt{t+1}}{2eCd}\right)$. From Theorem 36,

$$3065 \pi(p_N, q_N) \leq \tilde{O}(T^{2d}) \sup_{\|\mathbf{u}\|_\infty \leq T} |\phi_N(\mathbf{u}) - \gamma_N(\mathbf{u})| + O\left(\frac{d \log(T)}{T}\right) + O(\mu). \quad (325)$$

3068 From Lemma 37,

$$3069 \sup_{\|\mathbf{u}\|_\infty \leq T} |\phi_N(\mathbf{u}) - \gamma_N(\mathbf{u})| \leq \epsilon^{0.01} + \exp_2(1-t). \quad (326)$$

3072 The convergence rate given in Eq. (324) is trivial when $d = \Omega\left(\min\left(\frac{\log(\epsilon^{-1})}{\log \log(\epsilon^{-1})}, \frac{\sqrt{t}}{\log(t)}\right)\right)$ as

3073 $\epsilon^{-1}, t, N \rightarrow \infty$. When instead $d = o\left(\min\left(\frac{\log(\epsilon^{-1})}{\log \log(\epsilon^{-1})}, \frac{\sqrt{t}}{\log(t)}\right)\right)$,

$$3074 T^{2d} = \exp\left(o\left(\min\left(\log(\epsilon^{-1}), \sqrt{t}\right)\right)\right). \quad (327)$$

3077 The second term in Eq. (325) thus dominates in this setting, implying the final bound. \square

We also prove a bound on the error of characteristic functions given bounds on the cumulants.

Lemma 39 (Bound on characteristic functions from cumulants). *Let $p(\mathbf{x}), q(\mathbf{x})$ be two probability densities on \mathbb{R}^d with characteristic functions $\phi(\mathbf{u}), \gamma(\mathbf{u})$. Assume each cumulant of order k of p differs from that of q by an additive error of at most $\epsilon^k > 0$. Then*

$$|\phi(\mathbf{u}) - \gamma(\mathbf{u})| \leq \max(\exp(\exp(\epsilon \|\mathbf{u}\|_1) - 1) - 1, 1 - \exp(1 - \exp(\epsilon \|\mathbf{u}\|_1))). \quad (328)$$

Proof. We have that:

$$\begin{aligned} |\phi(\mathbf{u}) - \gamma(\mathbf{u})| &\leq |\phi(\mathbf{u})| |1 - \exp(\ln(\gamma(\mathbf{u})) - \ln(\phi(\mathbf{u})))| \\ &= |1 - \exp(\ln(\gamma(\mathbf{u})) - \ln(\phi(\mathbf{u})))|. \end{aligned} \quad (329)$$

By the definition of the cumulant generating function,

$$|\ln(\gamma(\mathbf{u})) - \ln(\phi(\mathbf{u}))| \leq \sum_{\|\mathbf{k}\|_1 > 0} \prod_{j=1}^d \frac{|(i\epsilon u_j)^{k_j}|}{k_j!} = \exp(\epsilon \|\mathbf{u}\|_1) - 1. \quad (330)$$

Taking into account the absolute values yields the final result. \square

This characteristic function bound combined with Theorem 36 gives us the following corollary.

Corollary 40 (Bound on Lévy–Prokhorov metric from cumulants). *Let $p_N(\mathbf{x}), q_N(\mathbf{x})$ be two sequences of distributions on \mathbb{R}^d with characteristic functions $\phi_N(\mathbf{u}), \gamma_N(\mathbf{u})$, respectively. Assume each cumulant of order k of p_N differs from that of q_N by an additive error of at most $\epsilon^k > 0$. Assume as well that $\int_{\|\mathbf{x}\|_\infty \geq \frac{1}{2d}\epsilon^{-0.99}} p(\mathbf{x}) \leq \mu$ for some (N -dependent) μ . Finally, assume that*

$d = o\left(\frac{\log(\epsilon^{-1})}{\log \log(\epsilon^{-1})}\right)$ as $N \rightarrow \infty$. The Lévy–Prokhorov distance between p_N, q_N is bounded by:

$$\pi(p_N, q_N) = O\left(\frac{d \log(\epsilon^{-1})}{\epsilon^{-0.99}} + \mu\right). \quad (331)$$

Proof. Let $T = \frac{1}{d}\epsilon^{-0.99}$. From Theorem 36,

$$\pi(p_N, q_N) \leq \tilde{O}(T^{2d}) \sup_{\|\mathbf{u}\|_\infty \leq T} |\phi_N(\mathbf{u}) - \gamma_N(\mathbf{u})| + O\left(\frac{d \log(T)}{T}\right) + O(\mu). \quad (332)$$

From Lemma 39,

$$\sup_{\|\mathbf{u}\|_\infty \leq T} |\phi_N(\mathbf{u}) - \gamma_N(\mathbf{u})| \leq \Theta(\epsilon^{0.01}). \quad (333)$$

Similarly,

$$T^{2d} = \exp(o(\log(\epsilon^{-1}))). \quad (334)$$

The second term in Eq. (332) thus dominates, implying the given convergence rate. \square

N TAIL BOUND LEMMAS

We here prove some helper lemmas giving tail bounds for the marginal entries of β -Wishart random variables. We will particularly leverage the Bartlett decomposition of β -Wishart matrices (see Appendix A.4). In the following we will use $|\cdot\rangle$ to denote vectors in the vector space on which the defining representation of $\mathcal{H}^{p+1}(\mathbb{F})$ acts, with basis $|0\rangle \cup \{|i\rangle\}_{i=1}^p$.

We first prove that β -Wishart matrices are robust to certain low-rank perturbations of their Bartlett decomposition. In particular, we show that removing a single column from the Bartlett factor \mathbf{L} of a β -Wishart matrix with many degrees of freedom—effectively reducing the degrees of freedom parameter of the matrix by one—has little effect on the joint distribution of entries of the matrix. Intuitively, this allows us to argue in Appendix D.3 that such a Wishart matrix is approximately independent from these entries.

Lemma 41 (Robustness of β -Wishart Bartlett decompositions). *Let \mathbf{W} be a $(p+1) \times (p+1)$ β -Wishart matrix with r degrees of freedom and identity scale matrix. Let $\mathbf{L}\mathbf{L}^\dagger$ be the Bartlett decomposition of \mathbf{W} . Let $\tilde{\mathbf{L}}$ be \mathbf{L} with the first column removed, and define $\tilde{\mathbf{W}} \equiv \tilde{\mathbf{L}}\tilde{\mathbf{L}}^\dagger$. For every $t \geq 0$,*

$$\mathbb{P} \left[r^{-\frac{1}{2}} \max_{i,j \in [1, \dots, p]} \left| \langle i | \mathbf{W} | j \rangle - \langle i | \tilde{\mathbf{W}} | j \rangle \right| \geq t^2 + \sqrt{2}\beta^{-\frac{1}{2}}r^{-\frac{1}{4}}t + r^{-\frac{1}{2}} \right] \leq p \exp \left(-\frac{\sqrt{r}}{2}t^2 \right). \quad (335)$$

In particular, the argument converges to zero under the Ky Fan metric as $r \rightarrow \infty$ at a rate $\Omega \left(\frac{\sqrt{r}}{\log(pr)} \right)$.

Proof. Let $N_{0,i}$ be the $(i, 0)$ entry of \mathbf{L} . Note that the $N_{0,i}$ are i.i.d. The lemma statement is then equivalent to showing that

$$\mathbb{P} \left[\max_{i,j \in [1, \dots, p]} |N_{0,i}N_{0,j}^*|^2 \geq \left(\sqrt{r}t^2 + \sqrt{2}\beta^{-\frac{1}{2}}r^{\frac{1}{4}}t + 1 \right)^2 \right] \leq p \exp \left(-\frac{\sqrt{r}}{2}t^2 \right). \quad (336)$$

This in turn is implied by:

$$\mathbb{P} \left[\beta \max_{i \in [1, \dots, p]} |N_{0,i}|^2 \geq \beta\sqrt{r}t^2 + \sqrt{2}\beta r^{\frac{1}{4}}t + \beta \right] \leq p \exp \left(-\frac{\sqrt{r}}{2}t^2 \right). \quad (337)$$

$\beta|N_{0,i}|^2$ is χ^2 -distributed with β degrees of freedom. The union bound along with standard χ^2 tail bounds (see, e.g., Lemma 1 of Laurent & Massart (2000)) thus immediately imply Eq. (337) and therefore also the final result. \square

We now prove a tail bound for a χ -distributed random variable with D degrees of freedom.

Lemma 42 (Tail bounds for χ -distributed random variables). *Let χ be χ -distributed with D degrees of freedom. For every $0 \leq t \leq \frac{3}{8}D^{\frac{1}{4}}$,*

$$\mathbb{P} \left[D^{-\frac{1}{4}} \left| \chi - \sqrt{D} \right| \geq \frac{4}{3}t + D^{-\frac{1}{4}}t^2 \right] \leq 2 \exp \left(-\sqrt{D}t^2 \right). \quad (338)$$

In particular, the argument converges to zero under the Ky Fan metric as $D \rightarrow \infty$ at a rate $\Omega \left(\frac{\sqrt[4]{D}}{\sqrt{\log(D)}} \right)$.

Proof. We have from standard χ^2 tail bounds (see, e.g., Lemma 1 of Laurent & Massart (2000)) that:

$$\mathbb{P} \left[\chi^2 - D \geq 2D^{\frac{3}{4}}t + 2\sqrt{D}t^2 \right] \leq \exp \left(-\sqrt{D}t^2 \right), \quad (339)$$

$$\mathbb{P} \left[\chi^2 - D \leq -2D^{\frac{3}{4}}t \right] \leq \exp \left(-\sqrt{D}t^2 \right). \quad (340)$$

This implies that:

$$\mathbb{P} \left[\chi \geq \sqrt{D + 2D^{\frac{3}{4}}t + 2\sqrt{D}t^2} \right] \leq \exp \left(-\sqrt{D}t^2 \right) \quad (341)$$

and (for $t \leq \frac{1}{2}D^{\frac{1}{4}}$)

$$\mathbb{P} \left[\chi \leq \sqrt{D - 2D^{\frac{3}{4}}t} \right] \leq \exp \left(-\sqrt{D}t^2 \right). \quad (342)$$

Using the general inequality (for $x \geq -1$):

$$\sqrt{1+x} \leq 1 + \frac{x}{2}, \quad (343)$$

we have that:

$$\sqrt{D + 2D^{\frac{3}{4}}t + 2\sqrt{D}t^2} - \sqrt{D} \leq D^{\frac{1}{4}}t + t^2 \quad (344)$$

3186 such that

$$3187 \mathbb{P} \left[\chi - \sqrt{D} \geq D^{\frac{1}{4}}t + t^2 \right] \leq \exp \left(-\sqrt{D}t^2 \right). \quad (345)$$

3188 Similarly, using the general inequality (for $0 \leq x \leq \frac{3}{4}$):

$$3190 1 - \frac{2x}{3} \leq \sqrt{1-x}, \quad (346)$$

3191 we have under the given assumptions that:

$$3194 \mathbb{P} \left[\sqrt{D} - \chi \geq \frac{4}{3}D^{\frac{1}{4}}t \right] \leq \exp \left(-\sqrt{D}t^2 \right). \quad (347)$$

3195 Noting that both $D^{\frac{1}{4}}t + t^2$ and $\frac{4}{3}D^{\frac{1}{4}}t$ are upper bounded by $\frac{4}{3}D^{\frac{1}{4}}t + t^2$ then implies the final
3196 result. \square

3200 We end with a tail bound for the off-diagonal entries of a β -Wishart matrix.

3201 **Lemma 43** (Off-diagonal β -Wishart tail bounds). *Let \mathbf{W} be a $(p+1) \times (p+1)$ β -Wishart matrix*
3202 *with r degrees of freedom and identity scale matrix. There exist universal constants $C, K > 0$ that*
3203 *depends only on β such that, for every $t \geq 0$,*

$$3205 \mathbb{P} \left[r^{-\frac{1}{2}} \max_{i \in [1, \dots, p]} |\langle 0 | \mathbf{W} | i \rangle| \geq 1 + t \right] \leq Kp \exp \left(-C \min(t, t^2) \right). \quad (348)$$

3207 *Proof.* Let $\mathbf{L}\mathbf{L}^\dagger$ be the Bartlett decomposition of \mathbf{W} . Let χ_0 be the $(0, 0)$ entry of \mathbf{L} and $N_{0,i}$ the
3208 $(i, 0)$ entry. The lemma statement is then equivalent to showing that

$$3210 \mathbb{P} \left[\chi_0^2 \max_{i \in [1, \dots, p]} |N_{0,i}|^2 \geq r(1+t)^2 \right] \leq Kp \exp \left(-C \min(t, t^2) \right). \quad (349)$$

3211 Note that:

$$3214 \mathbb{P} \left[\chi_0^2 \max_{i \in [1, \dots, p]} |N_{0,i}|^2 \geq r(1+t)^2 \right] \leq \mathbb{P} \left[\max \left(\frac{\chi_0^2}{r}, \max_{i \in [1, \dots, p]} |N_{0,i}|^2 \right)^2 \geq (1+t)^2 \right] \quad (350)$$

$$3217 = \mathbb{P} \left[\max \left(\frac{\chi_0^2}{r}, \max_{i \in [1, \dots, p]} |N_{0,i}|^2 \right) \geq 1+t \right],$$

3218 so by the union bound we need only check whether

$$3220 \max \left(\mathbb{P} \left[\chi_0^2 \geq r(1+t) \right], \mathbb{P} \left[|N_{0,i}|^2 \geq 1+t \right] \right) \leq \frac{K}{2} \exp \left(-C \min(t, t^2) \right) \quad (351)$$

3221 to prove the final result. As $\beta |N_{0,i}|^2$ is χ^2 -distributed with β degrees of freedom the final result
3222 immediately follows from standard χ^2 tail bounds. \square

3223
3224
3225
3226
3227
3228
3229
3230
3231
3232
3233
3234
3235
3236
3237
3238
3239