

One-Bit Quantization for Random Features Models

Danil Akhtiamov ^{*†}

Reza Ghane ^{*‡}

Babak Hassibi ^{†‡}

October 21, 2025

Abstract

Recent advances in neural networks have led to significant computational and memory demands, spurring interest in one-bit weight compression to enable efficient inference on resource-constrained devices. However, the theoretical underpinnings of such compression remain poorly understood. We address this gap by analyzing one-bit quantization in the Random Features model, a simplified framework that corresponds to neural networks with random representations. We prove that, asymptotically, quantizing weights of all layers except the last incurs no loss in generalization error, compared to the full precision random features model. Our findings offer theoretical insights into neural network compression. We also demonstrate empirically that one-bit quantization leads to significant inference speed ups for the Random Features models even on a laptop GPU, confirming the practical benefits of our work. Additionally, we provide an asymptotically precise characterization of the generalization error for Random Features with an arbitrary number of layers. To the best of our knowledge, our analysis yields more general results than all previous works in the related literature.

1 Introduction

The success of deep neural networks in tasks such as image recognition, natural language processing, and reinforcement learning has come at the cost of escalating computational and memory requirements. Modern models, often comprised of billions of parameters, demand significant resources for training and inference, rendering them impractical for deployment on resource-constrained devices like mobile phones, embedded systems, or IoT devices. To address this challenge, weight quantization—reducing the precision of neural network weights—has emerged as a promising technique to lower memory footprint and accelerate inference. In particular, one-bit quantization, which restricts weights to $\{+1, -1\}$, offers extreme compression (e.g., $\sim 32\times$ memory reduction for 32-bit floats) and enables efficient hardware implementations using bitwise operations. Various works have explored the possibility of network quantization in the recent years. In particular, for Large Language Models (LLMs), some post-training have been able to reduce the model size via fine-tuning. Examples of such approach include GPTQ Frantar et al. (2022) which can quantize a 175 billion GPT model to 4 bits and QuIP which Chee et al. (2023) compresses Llama 2 70B to 2 and 3 bits. Furthermore, quantization-aware training approaches, such as Bitnet Wang et al. (2023), Bitnet 1.58b Ma et al. (2024), have been able to achieve one-bit language models with comparable performance to the models from the same weight class. For a recent survey on efficient LLMs we refer to Xu et al. (2024). Such results are desirable as they pave the way for bringing foundational models

^{*}Equal Contribution

[†]Department of Computing and Mathematical Sciences, California Institute of Technology

[‡]Department of Electrical Engineering, California Institute of Technology

to edge devices by reducing memory requirements and reducing the inference time. However, while the aforementioned empirical approaches have demonstrated practical success, the theoretical foundations of one-bit quantization remain underexplored, limiting our ability to predict its performance and design improved training algorithms.

This paper investigates the generalization properties of one-bit quantization in the Random Features model, a simplified framework that captures key properties of wide neural networks while being amenable to rigorous analysis. Introduced by Rahimi and Recht (2007), the Random Features model approximates kernel methods and corresponds to the infinite-width limit of neural networks under certain conditions (Jacot et al. (2018)). By studying quantization in this model, we aim to uncover fundamental principles that govern the trade-offs between compression and performance in neural networks, leading to memory savings and inference speed-ups. Our main contributions are twofold:

1. **Lossless Quantization of Hidden Layers:** We prove that, for sufficiently wide Random Features models, quantizing the weights of all layers except the last to one bit incurs no loss in generalization error. This surprising result is established via a Gaussian Universality (GU) and Gaussian Equivalence (GE): GU implies that the test error of the linear model trained on the outputs of the random features matches the test error of the linear model trained on Gaussians with the same covariance; GE implies that the covariance of the random features, and the necessary characteristics of the latter covariance, are the same for the quantized and unquantized weights.
2. **Precise Characterization of the Test error of Deep Random Features model:** In the proportional regime, we rigorously characterize the generalization error of Random Features model with quantized weights with multiple layers and express the generalization error in terms of a few scalar variables.

The rest of the paper is organized as follows: Section 2 introduces the Random Features model and our notation, reviews Stochastic Mirror Descent and its implicit regularization properties, and presents the Gaussian Universality and Gaussian Equivalence principles that form the foundation of our analysis. Section 3 reviews related work on Random Features models and Gaussian universality. Section 4 details our main theorems on quantization, Section 5 discusses our approach and contributions, and Section 6 presents numerical validations of our theoretical findings.

2 Preliminaries

Throughout the paper we use bold letters for vectors and matrices.

2.1 Lipschitz Concentration Property

The following definition will be necessary for presenting our main result.

Definition 1 (Lipschitz Concentration Property). *A random vector $\mathbf{z} \in \mathbb{R}^d$ satisfies the Lipschitz Concentration Property (LCP) with parameter σ if for any L -Lipschitz function $f : \mathbb{R}^d \rightarrow \mathbb{R}$, the random variable $f(\mathbf{z}) - \mathbb{E}[f(\mathbf{z})]$ is subgaussian with parameter $L\sigma$. That is, for all $t > 0$:*

$$\mathbb{P}\left(\left|f(\mathbf{z}) - \mathbb{E}[f(\mathbf{z})]\right| \geq t\right) \leq 2 \exp\left(-\frac{t^2}{2L^2\sigma^2}\right)$$

2.2 Problem Setting

We consider a Random Features model defined as follows:

- **Input and Output:** Let $\mathbf{x} \in \mathbb{R}^d$ denote the input vector, and $y \in \mathbb{R}$ denote the target. The dataset consists of n samples $(\mathbf{x}_i, y_i)$. Furthermore, we assume that the data $\mathbf{x}$ satisfies Definition 1 with $\sigma^2 = O(\frac{1}{d})$. We also assume that $\mathbf{x}$ is centered, i.e. $\mathbb{E}\mathbf{x} = 0$. Denote $\Sigma = \mathbb{E}\mathbf{x}\mathbf{x}^T$. Then the assumptions made in this bullet imply that

$$\kappa(\Sigma) = \frac{\sigma_1}{\sigma_m} = O(1) \text{ and } \text{Tr}(\Sigma) = O(1),$$

where $\sigma_1, \dots, \sigma_m$ are the eigenvalues of Σ in the decreasing order. In other words, the matrix Σ is well-conditioned and normalized so that the norms of the inputs $\|\mathbf{x}\|_2 = O(1)$.

- **Model Architecture:** The Random Features model is a neural network with L hidden layers and an activation function ϕ :

- **Hidden layers:** The input to each hidden layer is mapped to a feature vector via random weights $\mathbf{W} \in \mathbb{R}^{d_\ell \times d_{\ell-1}}$, where $d_{\ell-1}$ is the number of input features to layer ℓ and d_ℓ is the number of output features. Each entry $\mathbf{W}_{ij} \sim \mathcal{N}\left(0, \frac{1}{d_{\ell-1}}\right)$ for the non-quantized model and $\mathbf{W}_{ij} \sim \frac{1}{\sqrt{d_{\ell-1}}} \text{Unif}(-1, +1)$ for the quantized model. Note that the coefficient $\frac{1}{\sqrt{d_{\ell-1}}}$ is necessary to ensure that the quantized model has the same second order statistics as the non-quantized model. The map for the ℓ -th hidden layer is $\phi_\ell(\mathbf{x}^{\ell-1}) = \phi(\mathbf{W}^{(\ell)}\mathbf{x}^{\ell-1})$ and $\mathbf{x}^{(0)} = \mathbf{x}$ is the input distribution.
- **Last layer:** The output is a linear combination of features,

$$f(\mathbf{x}, \mathbf{a}, \mathbf{W}^1, \dots, \mathbf{W}^L) = \mathbf{a}^\top \mathbf{x}^{(L)},$$

where $\mathbf{a} \in \mathbb{R}^{d_L}$ is the output layer weights.

- **Quantization:** One-bit quantization maps weights in the hidden layer ℓ to $\frac{1}{\sqrt{d_{\ell-1}}} \{+1, -1\}$ by preserving the normalization and taking the sign of each entry. Since the hidden layers for the non-quantized model are gaussian, this means that for the quantized model. Note that we quantize only the hidden layers and do not quantize the last layer $\mathbf{a}$ and the data $\mathbf{x}$. As a motivation, note that the majority of the memory is taken up by the weights in the hidden layers for our model and, therefore, we reduce memory requirements almost by a factor of 32 assuming the non-quantized model has 32-bit weights. Moreover, we demonstrate empirically that the presented scheme leads to almost 4X inference speed ups for sufficiently wide hidden layers.
- **Training procedure:** We assume that the last layer is trained to minimize an arbitrary differentiable convex loss function satisfying $\min_t \mathcal{L}(t) = \mathcal{L}(0) = 0$, i.e. the following optimization is performed, via *Stochastic Mirror Descent*:

$$\min_{\mathbf{a}} \sum_{i=1}^n \mathcal{L}(y_i - f(\mathbf{x}_i, \mathbf{a}, \mathbf{W}^1, \dots, \mathbf{W}^L))$$

Moreover, we assume that the model is over-parametrized, i.e. the number of parameters in the last layer exceeds the number of data points. Over-parametrization is a common assumption in modern machine learning.

- **Ground truth:** We assume that the labels are generated according to

$$y = f(\mathbf{x}, \mathbf{a}_*, \mathbf{W}^1, \dots, \mathbf{W}^L) \quad (1)$$

here, $\mathbf{a}_*$ is a ground truth parameter that we take to be $\mathbf{a}_* \sim \mathcal{N}(0, \frac{\mathbf{I}}{d_L})$, as it is natural to assume the ground truth is a "generic" vector.

- **Performance Metric:** We measure performance of a trained model via the MSE loss

$$\mathbb{E}_{\mathbf{x}}[(f(\mathbf{x}_i, \mathbf{a}, \mathbf{W}^1, \dots, \mathbf{W}^L) - y)^2]$$

- **Scaling:** We assume $d \rightarrow \infty$ and the hidden layer dimensions grow proportionally, i.e. $\gamma_\ell = \frac{d_\ell}{d}$ is constant for $\ell = 1, \dots, L$.

2.3 Stochastic Mirror Descent and Implicit Regularization

Stochastic Mirror Descent (SMD) generalizes Stochastic Gradient Descent (SGD) by employing a strictly convex, differentiable mirror map ψ . For a loss function $\mathcal{L}(\mathbf{w}; \mathbf{x}, y)$ and data $\{(\mathbf{x}_1, y_1), \dots, (\mathbf{x}_n, y_n)\}$, the SMD update at step t is

$$\nabla\psi(\mathbf{w}_{t+1}) = \nabla\psi(\mathbf{w}_t) - \eta \nabla \sum_{i=1}^n \mathcal{L}(\mathbf{w}_t; \mathbf{x}_i, y_i),$$

Note that taking $\psi(\mathbf{w}) = \frac{\|\mathbf{w}\|^2}{2}$ corresponds to the usual gradient descent:

$$\mathbf{w}_{t+1} = \mathbf{w}_t - \eta \nabla \sum_{i=1}^n \mathcal{L}(\mathbf{w}_t; \mathbf{x}_i, y_i).$$

Implicit regularization refers to the phenomenon where optimization algorithms naturally favor solutions minimizing certain characteristics of the weights without explicit regularization terms in the objective function. In overparameterized linear models, where the number of parameters exceeds the number of samples ($d > n$), SMD exhibits a crucial implicit bias property Azizan et al. (2021): among all interpolating solutions, i.e., solutions satisfying $\mathbf{X}\mathbf{w} = \mathbf{y}$, it chooses the solution that minimizes the Bregman divergence from the initialization $\mathbf{w}_0$. In other words, the following holds:

$$\lim_{t \rightarrow \infty} \mathbf{w}_t = \arg \min_{\mathbf{w}} D_\psi(\mathbf{w}, \mathbf{w}_0) \text{ subject to } \mathbf{X}\mathbf{w} = \mathbf{y} \quad (2)$$

where the Bregman divergence is defined as

$$D_\psi(\mathbf{w}, \mathbf{w}') = \psi(\mathbf{w}) - \psi(\mathbf{w}') - \nabla\psi(\mathbf{w}')^T(\mathbf{w} - \mathbf{w}')$$

For the gradient descent with initialization $\mathbf{w}_0 \approx 0$, (2) takes the following simple form:

$$\lim_{t \rightarrow \infty} \mathbf{w}_t = \arg \min_{\mathbf{w}} \|\mathbf{w}\|_2^2 \text{ subject to } \mathbf{X}\mathbf{w} = \mathbf{y} \quad (3)$$

2.4 Gaussian Universality

Theorem 1, presented in Ghane et al. (2024), establishes a universality result for linear regression with implicit regularization in the overparameterized regime, where the number of features d exceeds the number of samples n . The theorem demonstrates that the test error for the linear model trained on any feature matrix satisfying certain technical conditions is asymptotically equivalent to the test error of the same linear model trained on the Gaussian distribution with matching covariance. Gaussian Universality simplifies the analysis of model performance, making it tractable to predict the generalization error using techniques for working with Gaussian data, such as Gaussian Comparison Inequalities. In this subsection, for the sake of completeness, we present Theorem 1. The following assumptions are required for Theorem 1:

Assumptions 1. 1. ***Feature Matrix** $\mathbf{X} \in \mathbb{R}^{n \times d}$: The rows of $\mathbf{X}$, denoted $\mathbf{x}_i \in \mathbb{R}^d$ for $i = 1, \dots, n$, are independently and identically distributed (i.i.d.) from a distribution $\mathbb{P}$ with mean $\boldsymbol{\mu} \in \mathbb{R}^d$ and covariance $\boldsymbol{\Sigma} \in \mathbb{R}^{d \times d}$. The distribution satisfies:*

- *Bounded moments up to the sixth order: For each row $\mathbf{x}_i$, $\mathbb{E}[\|\mathbf{x}_i - \boldsymbol{\mu}\|_2^q] = O(1)$ for $q \leq 6$.*
- *Bounded mean: $\|\boldsymbol{\mu}\|_2^2 = O(1)$.*
- *Covariance condition: For any fixed vector $\mathbf{v} \in \mathbb{R}^d$, the quadratic form $\mathbf{v}^T \boldsymbol{\Sigma} \mathbf{v}$ has vanishing variance in the sense that $\text{Var}(\mathbf{x}_i^T \mathbf{v}) = O(1/d)$ as $d \rightarrow \infty$.*
- *Minimum singular value: The smallest singular value of $\mathbf{X}\mathbf{X}^T \in \mathbb{R}^{n \times n}$, denoted $\sigma_{\min}(\mathbf{X}\mathbf{X}^T)$, satisfies $\sigma_{\min}(\mathbf{X}\mathbf{X}^T) = \Omega(1)$ with high probability, ensuring $\mathbf{X}$ is well-conditioned.*

2. ***Target Labels** ($\mathbf{y} \in \mathbb{R}^n$): The labels $\mathbf{y}$ are generated as $\mathbf{y} = \mathbf{X}\mathbf{w}^* + \boldsymbol{\epsilon}$, where $\mathbf{w}^* \in \mathbb{R}^d$ is a fixed true parameter vector with $\|\mathbf{w}^*\|_2 = O(1)$, and the noise $\boldsymbol{\epsilon} \in \mathbb{R}^n$ has i.i.d. sub-Gaussian entries with mean zero and variance $\sigma^2 = O(1)$.*

3. ***Mirror Map** ($\psi : \mathbb{R}^d \rightarrow \mathbb{R}$): The mirror map ψ is M -strongly convex (i.e., $\nabla^2 \psi \succeq M\mathbf{I}_d$ for some $M > 0$), three times differentiable with bounded third derivatives ($\|\nabla^3 \psi\| = O(1)$), and satisfies $\psi(\mathbf{0}) = O(d)$. Moreover, the gradient of the mirror map at the solution $\mathbf{w}_{\mathbf{X}}$, denoted $\nabla \psi(\mathbf{w}_{\mathbf{X}})$, satisfies $\|\nabla \psi(\mathbf{w}_{\mathbf{X}})\|_2 = O(\sqrt{d})$ with high probability.*

4. ***Overparameterization**: The dimensions d (number of parameters) and n (number of samples) tend to infinity with a fixed ratio $d/n = \kappa > 1$, ensuring an overparameterized regime where the number of parameters exceeds the number of samples.*

Theorem 1 (Ghane et al. (2024)). *Let $\mathbf{X} \in \mathbb{R}^{n \times d}$ be a feature matrix whose rows are sampled from a distribution $\mathbb{P}$ with mean $\boldsymbol{\mu} \in \mathbb{R}^d$ and covariance $\boldsymbol{\Sigma} \in \mathbb{R}^{d \times d}$ and $\mathbf{y} \in \mathbb{R}^n$ be the labels satisfying Assumptions 1. Let $\mathbf{G} \in \mathbb{R}^{n \times d}$ be a matrix with independent rows sampled from $\mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$. Define $\mathbf{w}_{\mathbf{X}}$ and $\mathbf{w}_{\mathbf{G}}$ to be the SMD solutions with a mirror ψ trained on $\mathbf{X}$ and $\mathbf{G}$ respectively for some initialization $\mathbf{w}_0$:*

$$\mathbf{w}_{\mathbf{X}} = \arg \min_{\mathbf{w}} D_{\psi}(\mathbf{w}, \mathbf{w}_0) \text{ subject to } \mathbf{X}\mathbf{w} = \mathbf{y} \quad (4)$$

$$\mathbf{w}_{\mathbf{G}} = \arg \min_{\mathbf{w}} D_{\psi}(\mathbf{w}, \mathbf{w}_0) \text{ subject to } \mathbf{G}\mathbf{w} = \mathbf{y} \quad (5)$$

Then, asymptotically, the following holds for any Lipschitz function g in probability:

$$\lim_{n \rightarrow \infty} |g(\mathbf{w}_{\mathbf{X}}) - g(\mathbf{w}_{\mathbf{G}})| = 0$$

In particular, taking $g(\mathbf{w}) = \sqrt{\mathbf{w}^T \boldsymbol{\Sigma} \mathbf{w}}$ ensures that $\mathbf{w}_{\mathbf{X}}$ and $\mathbf{w}_{\mathbf{G}}$ yield equal test MSE losses as d grows large.

2.5 Gaussian Equivalence

We utilize the **Gaussian Equivalence Principle (GEP)** to characterize the covariance matrices of the outputs of the Random Features layers. Recall that the latter outputs are defined via

$$\phi_\ell(\mathbf{x}^{\ell-1}) = \phi(\mathbf{W}^{(\ell)} \mathbf{x}^{(\ell-1)}) \text{ for } \ell = 1, \dots, L$$

Here, $\mathbf{x}^{(0)} = \mathbf{x} \in \mathbb{R}^d$ is the input, $\mathbf{W}^{(\ell)} \in \mathbb{R}^{d_\ell \times d_{\ell-1}}$ is a random weight matrix with i.i.d. entries. Namely,

$$\mathbf{W}_{ij} \sim \mathcal{N}\left(0, \frac{1}{d_{\ell-1}}\right)$$

for the full precision model and

$$\mathbf{W}_{ij} \sim \frac{1}{\sqrt{d_{\ell-1}}} \text{Unif}(-1, +1)$$

for the one-bit quantized model, $d_0 = d$ and $\phi : \mathbb{R} \rightarrow \mathbb{R}$ is an odd nonlinearity function. Define the covariance of the ℓ -th hidden layer by Σ_ℓ , i.e.

$$\Sigma_\ell = \mathbb{E} \mathbf{x}^{(\ell)} \mathbf{x}^{(\ell)T}$$

In the proportional high-dimensional limit

$$n, d, d_1, \dots, d_L \rightarrow \infty, \quad n/d = \Theta(1), \quad \frac{d_{\ell-1}}{d_\ell} = \Theta(1), \quad \ell = 1, \dots, L$$

GEP provides a recipe for finding Σ_ℓ via the following recursive relations:

$$\Sigma_\ell \approx \rho_{\ell,1}^2 \mathbf{W}_\ell \Sigma_{\ell-1} \mathbf{W}_\ell^\top + \rho_{\ell,2}^2 \mathbf{I}_{d_\ell}$$

where

$$\begin{aligned} \rho_{\ell,1} &= \frac{1}{\sigma_{\ell-1}^2} \mathbb{E}_{z \sim \mathcal{N}(0, \sigma_{\ell-1}^2)} z \phi(z) \\ \rho_{\ell,2}^2 &= \mathbb{E}_{z \sim \mathcal{N}(0, \sigma_{\ell-1}^2)} \phi(z)^2 - \sigma_{\ell-1}^2 \rho_{\ell,1}^2 \\ \sigma_\ell^2 &= \frac{\text{Tr}(\Sigma_\ell)}{d_\ell} \end{aligned}$$

with the initial conditions

$$\Sigma_0 = \Sigma \text{ and } \sigma_0^2 = \frac{\text{Tr}(\Sigma)}{d}$$

3 Related Works

In this section, we provide a brief overview of existing works relevant to our setting.

The Random Features (RF) model Rahimi and Recht (2007) has been the subject of extensive study in recent years. The generalization error of the RF model with a single hidden layer has been analyzed in many different contexts within the high-dimensional proportional regime. These include settings where the last layer is trained using a ridge regression objective Gerace et al. (2020); Dhifallah and Lu (2020); Ghorbani et al. (2021); Mei and Montanari (2022); Goldt et al. (2022), or where $\mathbf{a}$ is taken to be the minimum ℓ_2 -norm interpolating solution Hastie et al. (2022). Furthermore, for binary classification tasks, the performance of the last layer as either an ℓ_2 Montanari et al.

(2019) or ℓ_1 Liang and Sur (2022) max-margin classifier has been analyzed. Since the Random Features model resembles neural networks at initialization, one line of work Moniri et al. (2023) has considered the generalization error after taking a single step of gradient descent on the hidden layer. Other settings studied include adversarial training Hassani and Javanmard (2024), the attention mechanism as a Random Features model Fu et al. (2023), and more recently, RFs in the non-asymptotic regime Defilippis et al. (2024).

RFs with multiple hidden layers have remained underexplored compared to those with single hidden layer. The paper Schröder et al. (2023) rigorously proved the Gaussian universality of the test error for the last layer trained using ridge regression on the same task as described in Subsection 2.2. A concurrent paper Bosch et al. (2023) proved a similar universality result for much more general convex losses and regularizers. Furthermore, Schröder et al. (2023) provided a conjecture for the universality of the test error for more general convex losses and regularizers, as well as for cases where the structures of the learner and the ground-truth differ. In Schröder et al. (2024), they extended these results to networks whose weights are not necessarily isotropic, imposing a general covariance structure on the weights per layer for ridge regression with squared loss. They went beyond the well-specified settings of Bosch et al. (2023); Schröder et al. (2023) and provided an expression for the test error where the ground truth and the learner features differ. They also conjectured a Gaussian equivalence model for multiple layers. To investigate the effect of the covariance structure of weights on the performance of RFs, Zavatone-Veth and Pehlevan (2023) used the non-rigorous replica method to characterize the test error of a linear Random Features model, where the last layer is trained using ridge regression to learn a linear ground truth function. In Cui et al. (2023), the authors computed the Bayes-optimal test error for estimating the target function in both classification and regression tasks for a deep Random Features model. They also provided a conjecture for the recursion on the population covariance of the layers, which was mentioned by Bosch et al. (2023); Schröder et al. (2023, 2024).

Gaussian universality plays a key role in reducing problems with non-Gaussian distributions to equivalent problems involving Gaussian distributions that match the first and second moments of the original distribution. This phenomenon has been actively investigated for various statistical inference problems, such as the universality of the test error of classifiers and regressors obtained through ridge regression or gradient descent. For an incomplete list, see Montanari and Nguyen (2017); Panahi and Hassibi (2017); Oymak and Tropp (2018); Abbasi et al. (2019); Montanari and Saeed (2022); Han and Shen (2023); Lahiry and Sur (2023); Dandi et al. (2023); Ghane et al. (2024, 2025). In the context of Random Features models, the universality of the test error for regression was rigorously proved by Hu and Lu (2022); Bosch et al. (2023); Montanari and Saeed (2022); Schröder et al. (2023).

The Gaussian Equivalence property is a framework used in the context of Random Features models. It allows for recursive characterization of layer-wise statistics and provides theoretical justification for analyzing neural networks through their Gaussian approximations. An interested reader can refer to Section 2 for details on Gaussian Universality and Gaussian Equivalence. This principle has been used in many recent works, such as Goldt et al. (2020); Bosch et al. (2023); Hu and Lu (2022); Schröder et al. (2023, 2024); Defilippis et al. (2024). The paper Hu and Lu (2022) was the first to provide a rigorous proof of Gaussian Equivalence for Random Features Models with one hidden layer. The subsequent papers Bosch et al. (2023); Schröder et al. (2023) have proved different forms of Gaussian Equivalence for deep RF models. It should be mentioned that all works mentioned in this paragraph operate under the assumptions that the random features are Gaussian.

Theoretical analyses of quantization and pruning are limited in the literature. The investigated topics include post-training quantization Zhang et al. (2025), training-aware quantization Askari-Hemmat et al. (2024), analysis of generalization error of linear models for binary classification

Akhtiamov et al. (2024b), multiclass classification Ghane et al. (2025) and pruning in the context of random features model Chang et al. (2021).

4 Main Results

The following theorem, which is the main result of our work, provides a precise asymptotic characterization of the test loss for the quantized and non-quantized Random Features Models. Since we obtain the same expression for both, we conclude that quantizing the hidden layers to one bit naively does not lead to any degradation of performance for the Random Features Models as long as the model and the dataset are big enough and both models are trained via SMD using the same smooth mirror function.

Theorem 2. *Let $f(\mathbf{x}, \mathbf{a}, \mathbf{W}^1, \dots, \mathbf{W}^L)$ be the Random Features Model defined in Subsection 2.2, where*

$$\mathbf{W}^1 \in \mathbb{R}^{d_1 \times d}, \dots, \mathbf{W}^L \in \mathbb{R}^{d_L \times d_{L-1}}$$

are either full precision weights sampled i.i.d. from

$$\mathbf{W}_{ij}^\ell \sim \mathcal{N}\left(0, \frac{1}{d_{\ell-1}}\right)$$

or one-bit quantized weights sampled i.i.d. from

$$\mathbf{W}_{ij}^\ell \sim \frac{1}{\sqrt{d_{\ell-1}}} \text{Unif}(-1, +1)$$

Assume that

- *The data $\mathbf{x} \in \mathbb{R}^d$ satisfies*

$$\mathbb{E}\mathbf{x} = 0$$

along with the LCP property from Definition 1 with

$$\sigma^2 = O\left(\frac{1}{d}\right)$$

- *The activation function ϕ is odd and has bounded first, third and fifth derivatives.*
- *The dimension of the last layer d_L exceeds the number of training samples n and the last layer $\mathbf{a}$ is trained to minimize the following objective using SMD with a mirror ψ satisfying Assumptions 1 initialized at $\mathbf{a}_0 \in \mathbb{R}^{d_L}$:*

$$\min_{\mathbf{a}} \sum_{i=1}^n \mathcal{L}(y_i - f(\mathbf{x}_i, \mathbf{a}, \mathbf{W}^1, \dots, \mathbf{W}^L))$$

- *The labels y are generated using a ground truth*

$$\mathbf{a}_* \sim \mathcal{N}\left(0, \frac{\mathbf{I}}{d_L}\right)$$

as defined in (1).

Then, in the asymptotic proportional regime

$$n, d, d_1, \dots, d_L \rightarrow \infty, \\ \frac{n}{d}, \frac{d}{d_1}, \dots, \frac{d_{L-1}}{d_L} = \Theta(1),$$

the test loss satisfies

$$\mathbb{E}_{\mathbf{x}}[(f(\mathbf{x}, \mathbf{a}, \mathbf{W}^1, \dots, \mathbf{W}^L) - y)^2] \rightarrow \tau = \tau^{(L)}$$

Here, convergence means convergence in probability and τ can be found by solving a system of elaborate nonlinear scalar deterministic equations, which follow from (38) for the case of general mirrors and are simplified for the case of SGD in (36). It should be noted that (38) and (36) are min-max optimization objectives and τ can be found by solving the corresponding saddle-point equations.

In particular, asymptotically, the error does not depend on the realizations of

$$\mathbf{W}^1, \dots, \mathbf{W}^L$$

and does not change if we replace

$$\mathbf{W}^1, \dots, \mathbf{W}^L$$

by

$$\frac{\text{sign}(\mathbf{W}^1)}{\sqrt{d_1}}, \dots, \frac{\text{sign}(\mathbf{W}^L)}{\sqrt{d_L}},$$

where sign is applied entry-wise.

Remark 1. Examples of data satisfying LCP with $\sigma^2 = O(\frac{1}{d})$ include $\mathbf{x} = \mathbf{g} \sim \mathcal{N}(0, \mathbf{\Sigma})$ for $\mathbf{\Sigma}$ such that

$$\text{Tr}(\mathbf{\Sigma}) = O(1)$$

and

$$\kappa(\mathbf{\Sigma}) = \frac{\sigma_{\max}}{\sigma_{\min}} = O(1)$$

as well as $\mathbf{x} = f(\mathbf{g})$ for any Lipschitz f with bounded Lipschitz constant and the same $\mathbf{g}$ defined as above.

Remark 2. We observe a close match between the performances of Gaussian and Rademacher Random Features trained to classify points from MNIST dataset with ReLU activation function in Section 6. As such, we believe that it should be possible to extend Theorem 2 to non-centered data and non-odd activation functions. The main technical obstacle for this is establishing Gaussian Equivalence results applicable to the latter scenario. We leave this as an important direction for future work.

Remark 3. While we postpone presentation of the exact non-linear equations from Theorem 2 defining τ to (38) and (36) in the Appendix C, we would like to provide the essence here. To find τ for a general smooth mirror ψ , one needs to solve a nonlinear scalar deterministic system of equations involving $2L$ scalar parameters. For the case of SGD, i.e. when the mirror

$$\psi(\cdot) = \frac{1}{2} \|\cdot\|_2^2$$

the number of unknown parameters could be reduced to L . This way, we obtain a deterministic system of equations that defines the test MSE loss implicitly for both Gaussian and appropriately normalized Rademacher weights. To the best of our knowledge, our work is the first work characterizing the test loss for normalized Rademacher Random Features via a finite number of scalar equations.

5 Our approach and contributions

Our approach is different from Bosch et al. (2023) and Schröder et al. (2023), as we start with invoking Gaussian Universality for the last layer and only afterwards do we apply Gaussian Equivalence Principle to calculate the covariance of the last layer. This approach allows us to analyze the generalization error of the solutions obtained via Stochastic Mirror Descent with smooth mirrors and arbitrary convex losses, extending results available in the literature. Indeed, to the best of our knowledge, the only examples considered in the literature previously are ridge regression Gerace et al. (2020); Dhifallah and Lu (2020); Ghorbani et al. (2021); Mei and Montanari (2022); Goldt et al. (2022), SGD initialized at 0 Hastie et al. (2022) and the ℓ_2 Montanari et al. (2019) and ℓ_1 Liang and Sur (2022) max-margin classifiers. In addition, we prove Gaussian Equivalence for deep Random Features Models ($L > 1$), while Schröder et al. (2023) leaves the case $L > 1$ as a conjecture for objectives other than ridge regression and Bosch et al. (2023) takes an additional expectation with respect to the weights in the Gaussian Equivalence part.

Other works Montanari and Saeed (2022); Defilippis et al. (2024); Schröder et al. (2024) are more similar to the present paper, as they apply similar universality results to the output of the last layer as well. The main differences between Montanari and Saeed (2022); Defilippis et al. (2024); Schröder et al. (2024) and our work is that we extend their results to the case of normalized Rademacher features to capture the one-bit quantization of weights as well as apply additional steps to show that the test error converges to a deterministic quantity independent of the realizations of the weights.

After combining Gaussian Universality with Gaussian Equivalence, we proceed to apply Convex Gaussian Min-Max Theorem (CGMT) Thrampoulidis et al. (2014); Akhtiamov et al. (2024a) to each hidden layer one by one to prove that the error concentrates with respect to the randomness in each $\mathbf{W}^\ell$ as well. For Gaussian weights, this application is more straightforward, while for normalized Rademacher weights we have to employ an additional step and apply another result Han and Shen (2023) that says that CGMT can be applied to many other i.i.d. subgaussian designs. This allows us to derive identical expressions for the test losses for both Gaussian and Rademacher models and conclude that one-bit quantization does not lead to any deterioration in performance for Random Features Models. This aspect of our work is novel as well: to the best of our knowledge, our work is the first to derive expressions for deep non-Gaussian Random Features.

6 Numerical Experiments

We validate our theoretical results through experiments on Random Features models with Gaussian Weights and with one-bit quantized weights with last layer trained on synthetic Gaussian data and MNIST Deng (2012). For the Gaussian data, we used tanh activation function and trained the last layer with SGD as well as with negative entropy mirror. For MNIST, we use ReLU activations and trained the last layer using SGD.

6.1 One-Bit Quantization

6.1.1 Synthetic data

We verify that one-bit quantization incurs no loss by comparing test MSE between Gaussian and Rademacher weights across depths (defined as the number of hidden layers)

$$L \in \{1, 2, 3, 4, 5\}$$

Specifically, we compare two Random Features variants, Gaussian

$$\mathbf{W}_{ij}^{(\ell)} \sim \mathcal{N}\left(0, \frac{1}{d_{\ell-1}}\right)$$

and Rademacher

$$\mathbf{W}_{ij}^{(\ell)} \sim \frac{1}{\sqrt{d_{\ell-1}}} \text{Unif}\{-1, +1\}$$

We generate synthetic data with $\mathbf{x}_i \sim \mathcal{N}(0, \frac{\mathbf{I}_d}{d})$ and labels $y_i = \phi_L(\mathbf{x}_i)^\top \mathbf{a}_*$, where ϕ_L is the L -layer random features map with tanh activation, $\mathbf{a}_* \sim \mathcal{N}(0, \frac{1}{d_L} \mathbf{I})$. This is the data-generation procedure that will be used for demonstrating inference speedup. We use $n = 1000$ training samples, input dimension $d = 8192$ and hidden dimensions $d_1 = \dots = d_L = 4096$ for each hidden layer. Following the overparametrized regime, we fix random features and train only the last layer via minimum ℓ_2 norm solution, which can be recovered analytically as:

$$\mathbf{a} = \Phi^\top (\Phi \Phi^\top)^{-1} \mathbf{y}, \text{ where } \Phi \in \mathbb{R}^{n \times d_L}$$

As can be seen in Figure 1, we observe a close match between the test error of the RF model with Gaussian weights and the RF model with Rademacher weights. To illustrate a more general case of our theorem, we also consider the negative Shannon entropy

$$\psi(\mathbf{w}) = \sum |w_i| \log(|w_i|)$$

under the same setting in Figure 2. For both scenarios, we use $n_{\text{test}} = 5000$ test samples for estimating the test MSE loss.




Figure 1: Random Features with varying depth for Synthetic Data for SGD

6.1.2 MNIST

We run the following two experiments for MNIST:

- We train random feature networks with Gaussian and Rademacher weights and with ReLU activations on MNIST, varying the number of layers

$$L \in \{1, 2, 3, 4, 5\}$$




Figure 2: Random Features with varying depth for Synthetic Data for Negative Entropy Mirror

while fixing the hidden dimensions

$$d_1 = d_2 = \dots d_L = 512$$

Since this is a classification task, we report test accuracy rather than test MSE. For each layer count, we use 20 samples per class and average results over 20 trials. We use one-hot encoding of the classes in the optimization objective. The final layer uses minimum ℓ_2 -norm interpolation. The results are presented in Figure 3 and demonstrate a close match, despite not being covered by our theory. We use a total of 200 test samples for estimating the resulting test accuracy.

- We also train $L = 2$ - layer random feature networks with Gaussian and Rademacher weights on MNIST, varying the hidden dimension

$$d_1 = d_2 \in \{256, 512, 1024, 2048, 4196\}$$

while fixing $L = 2$. Since this is a classification task, we report test accuracy rather than test MSE. For each width, we perform evaluation in the same way as in the experiment from the previous bullet point.

6.2 Inference speedup

For investigating the potential speedup of employing one-bit weights during inference, we consider the setting in Section 6.1.1 for $n = 1000$ training samples. We proceed to load the model using PyTorch with CUDA acceleration, on an RTX 2060 laptop GPU with 6GB VRAM in FP32 precision. Furthermore, we leverage the Gemlite Badri and Mobius Labs (2024), a triton-based kernel library with 1 bit weights and group size set to 64 with 500 warmup runs and 50,000 timed iterations on batch size 1. We present the results in Figure 5 for the Random Features model with one hidden layer. We observe a 4 times speed-up on average.




Figure 3: Random Features with varying depth for MNIST




Figure 4: Random Features with varying hidden width for MNIST

7 Conclusion and Future Works

The present paper leverages a combination of Gaussian Universality and Gaussian Equivalence principles followed by an application of Gaussian Comparison Inequalities to analyze one-bit compression of the weights for Random Features Models. We demonstrate that for random features the naive one-bit compression is lossless and results in a $\sim 4X$ inference speed up assuming the hidden layer dimension is sufficiently wide. It is worth mentioning that quantizing the last layer under the same setting would neither be lossless nor result in a noticeable inference speed up. As such, we suggest that the last layer should never be quantized in practice.

Our experiments suggest that one-bit quantization might be lossless for Random Features with ReLU trained on classification tasks as well. This calls for extending our methods to classification instead of regression and to non-odd activation functions. Both former and latter extensions would require a Gaussian Equivalence Principle for non-centered data.

Finally, while we believe that the setting of random features considered in the present work sheds light on one-bit quantization, it would be interesting to study the more nuanced picture of learnable representations. While performing the full analysis might be too challenging in general, we suggest starting with the simpler case when the features are learned via one-step Gradient Descent Moniri




Figure 5: Inference speed up vs Hidden Dimension

et al. (2023). In the latter setting, it would be interesting to see the effects of more sophisticated one-bit compression techniques as well.

References

- Abbasi, E., Salehi, F., and Hassibi, B. (2019). Universality in learning from linear measurements. *Advances in Neural Information Processing Systems*, 32.
- Akhtiamov, D., Bosch, D., Ghane, R., Varma, K. N., and Hassibi, B. (2024a). A novel gaussian min-max theorem and its applications. *arXiv preprint arXiv:2402.07356*.
- Akhtiamov, D., Ghane, R., and Hassibi, B. (2024b). Regularized linear regression for binary classification. In *2024 IEEE International Symposium on Information Theory (ISIT)*, pages 202–207. IEEE.
- AskariHemmat, M., Jeddi, A., Hemmat, R. A., Lazarevich, I., Hoffman, A., Sah, S., Saboori, E., Savaria, Y., and David, J.-P. (2024). Qgen: On the ability to generalize in quantization aware training. *arXiv preprint arXiv:2404.11769*.
- Azizan, N., Lale, S., and Hassibi, B. (2021). Stochastic mirror descent on overparameterized non-linear models. *IEEE Transactions on Neural Networks and Learning Systems*, 33(12):7717–7727.
- Badri, H. and Mobius Labs (2024). Gemlite: Fast low-bit matmul kernels in triton. GitHub repository.
- Bosch, D., Panahi, A., and Hassibi, B. (2023). Precise asymptotic analysis of deep random feature models. In *The Thirty Sixth Annual Conference on Learning Theory*, pages 4132–4179. PMLR.
- Chang, X., Li, Y., Oymak, S., and Thrampoulidis, C. (2021). Provable benefits of overparameterization in model compression: From double descent to pruning neural networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 6974–6983.
- Chee, J., Cai, Y., Kuleshov, V., and De Sa, C. M. (2023). Quip: 2-bit quantization of large language models with guarantees. *Advances in Neural Information Processing Systems*, 36:4396–4429.

- Cui, H., Krzakala, F., and Zdeborová, L. (2023). Bayes-optimal learning of deep random networks of extensive-width. In *International Conference on Machine Learning*, pages 6468–6521. PMLR.
- Dandi, Y., Stephan, L., Krzakala, F., Loureiro, B., and Zdeborová, L. (2023). Universality laws for gaussian mixtures in generalized linear models. *Advances in Neural Information Processing Systems*, 36:54754–54768.
- Defilippis, L., Loureiro, B., and Misiakiewicz, T. (2024). Dimension-free deterministic equivalents for random feature regression. *arXiv preprint arXiv:2405.15699*.
- Deng, L. (2012). The mnist database of handwritten digit images for machine learning research. *IEEE Signal Processing Magazine*, 29(6):141–142.
- Dhifallah, O. and Lu, Y. M. (2020). A precise performance analysis of learning with random features. *arXiv preprint arXiv:2008.11904*.
- Frantar, E., Ashkboos, S., Hoefler, T., and Alistarh, D. (2022). Gptq: Accurate post-training quantization for generative pre-trained transformers. *arXiv preprint arXiv:2210.17323*.
- Fu, H., Guo, T., Bai, Y., and Mei, S. (2023). What can a single attention layer learn? a study through the random features lens. *Advances in Neural Information Processing Systems*, 36:11912–11951.
- Gerace, F., Loureiro, B., Krzakala, F., Mézard, M., and Zdeborová, L. (2020). Generalisation error in learning with random features and the hidden manifold model. In *International Conference on Machine Learning*, pages 3452–3462. PMLR.
- Ghane, R., Akhtiamov, D., and Hassibi, B. (2024). Universality in transfer learning for linear models. *arXiv preprint arXiv:2410.02164*.
- Ghane, R., Bao, A., Akhtiamov, D., and Hassibi, B. (2025). Concentration of measure for distributions generated via diffusion models. *arXiv preprint arXiv:2501.07741*.
- Ghorbani, B., Mei, S., Theodor, M., and Montanari, A. (2021). Linearized two-layers neural networks in high dimension. *The Annals of Statistics*, 49(2):1029–1054.
- Goldt, S., Loureiro, B., Reeves, G., Krzakala, F., Mézard, M., and Zdeborová, L. (2022). The gaussian equivalence of generative models for learning with shallow neural networks. In *Mathematical and Scientific Machine Learning*, pages 426–471. PMLR.
- Goldt, S., Mézard, M., Krzakala, F., and Zdeborová, L. (2020). Modeling the influence of data structure on learning in neural networks: The hidden manifold model. *Physical Review X*, 10(4):041044.
- Han, Q. and Shen, Y. (2023). Universality of regularized regression estimators in high dimensions. *The Annals of Statistics*, 51(4):1799–1823.
- Hassani, H. and Javanmard, A. (2024). The curse of overparametrization in adversarial training: Precise analysis of robust generalization for random features regression. *The Annals of Statistics*, 52(2):441–465.
- Hastie, T., Montanari, A., Rosset, S., and Tibshirani, R. J. (2022). Surprises in high-dimensional ridgeless least squares interpolation. *Annals of statistics*, 50(2):949.

- Hu, H. and Lu, Y. M. (2022). Universality laws for high-dimensional learning with random features. *IEEE Transactions on Information Theory*, 69(3):1932–1964.
- Jacot, A., Gabriel, F., and Hongler, C. (2018). Neural tangent kernel: Convergence and generalization in neural networks. *Advances in neural information processing systems*, 31.
- Lahiry, S. and Sur, P. (2023). Universality in block dependent linear models with applications to nonparametric regression. *arXiv preprint arXiv:2401.00344*.
- Liang, T. and Sur, P. (2022). A precise high-dimensional asymptotic theory for boosting and minimum- ℓ_1 1-norm interpolated classifiers. *The Annals of Statistics*, 50(3):1669–1695.
- Ma, S., Wang, H., Ma, L., Wang, L., Wang, W., Huang, S., Dong, L., Wang, R., Xue, J., and Wei, F. (2024). The era of 1-bit llms: All large language models are in 1.58 bits. *arXiv preprint arXiv:2402.17764*, 1(4).
- Mei, S. and Montanari, A. (2022). The generalization error of random features regression: Precise asymptotics and the double descent curve. *Communications on Pure and Applied Mathematics*, 75(4):667–766.
- Moniri, B., Lee, D., Hassani, H., and Dobriban, E. (2023). A theory of non-linear feature learning with one gradient step in two-layer neural networks. *arXiv preprint arXiv:2310.07891*.
- Montanari, A. and Nguyen, P.-M. (2017). Universality of the elastic net error. In *2017 IEEE International Symposium on Information Theory (ISIT)*, pages 2338–2342. IEEE.
- Montanari, A., Ruan, F., Sohn, Y., and Yan, J. (2019). The generalization error of max-margin linear classifiers: Benign overfitting and high dimensional asymptotics in the overparametrized regime. *arXiv preprint arXiv:1911.01544*.
- Montanari, A. and Saeed, B. N. (2022). Universality of empirical risk minimization. In *Conference on Learning Theory*, pages 4310–4312. PMLR.
- Oymak, S. and Tropp, J. A. (2018). Universality laws for randomized dimension reduction, with applications. *Information and Inference: A Journal of the IMA*, 7(3):337–446.
- Panahi, A. and Hassibi, B. (2017). A universal analysis of large-scale regularized least squares solutions. *Advances in Neural Information Processing Systems*, 30.
- Rahimi, A. and Recht, B. (2007). Random features for large-scale kernel machines. *Advances in neural information processing systems*, 20.
- Schröder, D., Cui, H., Dmitriev, D., and Loureiro, B. (2023). Deterministic equivalent and error universality of deep random features learning. In *International Conference on Machine Learning*, pages 30285–30320. PMLR.
- Schröder, D., Dmitriev, D., Cui, H., and Loureiro, B. (2024). Asymptotics of learning with deep structured (random) features. *arXiv preprint arXiv:2402.13999*.
- Seddik, M. E. A., Louart, C., Tamaazousti, M., and Couillet, R. (2020). Random matrix theory proves that deep learning representations of GAN-data behave as Gaussian mixtures. In III, H. D. and Singh, A., editors, *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 8573–8582. PMLR.

- Thrampoulidis, C., Oymak, S., and Hassibi, B. (2014). A tight version of the gaussian min-max theorem in the presence of convexity. *arXiv preprint arXiv:1408.4837*.
- Wang, H., Ma, S., Dong, L., Huang, S., Wang, H., Ma, L., Yang, F., Wang, R., Wu, Y., and Wei, F. (2023). Bitnet: Scaling 1-bit transformers for large language models. *arXiv preprint arXiv:2310.11453*.
- Xu, M., Yin, W., Cai, D., Yi, R., Xu, D., Wang, Q., Wu, B., Zhao, Y., Yang, C., Wang, S., et al. (2024). A survey of resource-efficient llm and multimodal foundation models. *arXiv preprint arXiv:2401.08092*.
- Zavatone-Veth, J. and Pehlevan, C. (2023). Learning curves for deep structured gaussian feature models. *Advances in neural information processing systems*, 36:42866–42897.
- Zhang, H., Zhang, S., Colbert, I., and Saab, R. (2025). Provable post-training quantization: Theoretical analysis of optq and qronos. *arXiv preprint arXiv:2508.04853*.

A Scheme of the proof of Theorem 2

As outlined in Section 5, our approach is comprised of a consecutive application of Gaussian Universality, Gaussian Equivalence and Gaussian Comparison Inequalities. We present the omitted proofs related to Gaussian Universality and Gaussian Equivalence in Subsections B.2 and B.1 respectively, followed by the missing CGMT derivations in Section C.

B Gaussian Universality and Gaussian Equivalence

Denote the rows of $\mathbf{W}^{(\ell)}$ by $\mathbf{w}_1^{(\ell)}, \dots, \mathbf{w}_{d_\ell}^{(\ell)}$

Note that the following event holds w.h.p. with respect to randomness in $\mathbf{W}^{(1)}, \dots, \mathbf{W}^{(L)}$

$$\max_{0 \leq i, j \leq d_\ell} \left| \mathbf{w}_i^{(\ell)T} \mathbf{w}_j^{(\ell)} - \delta_{i,j} \right| \leq \frac{C}{\sqrt{d_\ell}} \text{ and } \|\mathbf{W}^{(\ell)}\|_{\text{op}} = O(1) \quad (6)$$

Note that (6) holds w.h.p. both when the weights are normalized i.i.d. Rademacher as well as standard Gaussian. For the purposes of this section, we freeze a realization of the features $\mathbf{W}^{(\ell)}$ satisfying (6) for $\ell = 1, \dots, L$ and consider the randomness w.r.t. the inputs $\mathbf{x}$ only.

B.1 Gaussian Equivalence

B.1.1 One hidden layer

We will illustrate the argument for Random Features Models with one hidden layer first. We would like to apply Theorem 1 to $\mathbf{a} = \phi(\mathbf{W}\mathbf{x})$. Denote $m = d_1$ the width of the only hidden layer in this case for ease of notation. Note that $\mathbb{E}_{\mathbf{x}} \phi(\mathbf{W}\mathbf{x}) = 0$ since ϕ is odd and denote

$$\Sigma_{\mathbf{a}} = \mathbb{E}_{\mathbf{x}} \phi(\mathbf{W}\mathbf{x}) \phi(\mathbf{W}\mathbf{x})^T$$

Theorem 1 guarantees that, even though $\mathbf{a}$ is not Gaussian, the test error remains unchanged if we train the last layer on data sampled from $\mathbf{a}' \sim \mathcal{N}(0, \Sigma_{\mathbf{a}})$ instead of $\mathbf{a}$.

Thus, according to Lemma 5 from Hu and Lu (2022), we have

$$\|\Sigma_{\mathbf{a}} - \Sigma_{\mathbf{b}}\|_{op} = O\left(\|\mathbf{W}\|_{op} \frac{\text{polylog}(m)}{m^2}\right) \quad (7)$$

Here, m is the width of the hidden layer and $\Sigma_{\mathbf{b}}$ is the covariance of the distribution defined via

$$\begin{aligned} \mathbf{b} &\sim \rho_1 \mathbf{W}\mathbf{x} + \rho_2 \mathbf{g} \\ \gamma &\sim \mathcal{N}(0, 1) \\ \rho_1 &= \mathbb{E}_{\gamma} \gamma \phi(\gamma) \\ \rho_2 &= (\mathbb{E}_{\gamma} \phi^2(\gamma) - \rho_1^2)^{\frac{1}{2}} \end{aligned}$$

Note that $\|\mathbf{W}\|_{op} = O(1)$ holds w.h.p. as well. Therefore,

$$\|\Sigma_{\mathbf{a}} - \Sigma_{\mathbf{b}}\|_{op} = o\left(\frac{1}{\sigma_{\min}(\Sigma_{\mathbf{a}})}\right) \text{ as } m \rightarrow \infty \quad (8)$$

Hence, since the test error depends continuously on the covariance for regression trained on gaussian data, we can replace $\mathbf{a}$ by $\mathcal{N}(0, \Sigma_{\mathbf{b}})$ without changing the generalization error.

Finally, note that

$$\Sigma_{\mathbf{b}} = \rho_1^2 \mathbf{W} \Sigma \mathbf{W}^T + \rho_2^2 \mathbf{I}_m$$

B.1.2 Multiple hidden layers

Denote the output of the ℓ -th hidden layer by $\mathbf{x}^{(\ell)}$, $\ell = 0, \dots, L$. Same as in the case of one hidden layer, we apply Theorem 1 to $\mathbf{x}^{(L)} = \phi(\mathbf{W}^{(L-1)} \mathbf{x}^{(L-1)})$. Again, same as in the case of one hidden layer, we have $\mathbb{E}_{\mathbf{a} \sim \mathbf{x}^{(L)}} \mathbf{a} = 0$ and denote

$$\begin{aligned} \Sigma_{\ell} &= \mathbb{E}_{\mathbf{x}^{(\ell)}} \mathbf{x}^{(\ell)} \mathbf{x}^{(\ell)T} \quad \ell = 0, \dots, L \\ \tilde{\Sigma}_{\ell} &= \rho_{\ell,1}^2 \mathbf{W}^{(\ell)} \Sigma_{\ell-1} \mathbf{W}^{(\ell)T} + \rho_{\ell,2}^2 \mathbf{I} \quad \ell = 1, \dots, L \end{aligned} \quad (9)$$

where

$$\rho_{\ell,1} = \frac{1}{\sigma_{\ell-1}^2} \mathbb{E}_{z \sim \mathcal{N}(0, \sigma_{\ell-1}^2)} z \phi(z) \quad (10)$$

$$\rho_{\ell,2}^2 = \mathbb{E}_{z \sim \mathcal{N}(0, \sigma_{\ell-1}^2)} \phi(z)^2 - \sigma_{\ell-1}^2 \rho_{\ell,1}^2 \quad (11)$$

$$\sigma_{\ell}^2 = \frac{\text{Tr}(\Sigma_{\ell})}{d_{\ell}} \quad (12)$$

Similarly to the one hidden layer case, our goal is to show that

$$\|\Sigma_{\ell} - \tilde{\Sigma}_{\ell}\|_{op} = o\left(\frac{1}{d_{\ell}}\right) \quad (13)$$

Note that (13) provides an asymptotic recurrence for finding $\Sigma^{(L)}$.

To prove (13), note that, if $\mathbf{x}^{(\ell-1)}$ were Gaussian, we would be able to obtain the desired result in the same way as for one hidden layer by appealing to the results of Schröder et al. (2023). However, $\mathbf{x}^{(\ell-1)}$ is not Gaussian in general for $\ell > 1$. As such, we outline a more general argument based on subgaussianity below.

Definition 2. A random variable s is called subgaussian if there exists a constant $\sigma > 0$ such that for all $t \geq 0$,

$$\mathbb{P}(|s| \geq t) \leq 2e^{-\frac{t^2}{2\sigma^2}}.$$

The smallest such constant σ is called the subgaussian parameter of s .

We will also need a definition of the Lipschitz Concentration Property:

Definition 3 (Lipschitz Concentration Property). A random vector $\mathbf{z} \in \mathbb{R}^n$ satisfies the Lipschitz Concentration Property (LCP) with parameter σ if for any L -Lipschitz function $f : \mathbb{R}^n \rightarrow \mathbb{R}$, the random variable $f(\mathbf{z}) - \mathbb{E}[f(\mathbf{z})]$ is subgaussian with parameter $L\sigma$. That is, for all $t > 0$:

$$\mathbb{P}(|f(\mathbf{z}) - \mathbb{E}[f(\mathbf{z})]| \geq t) \leq 2 \exp\left(-\frac{t^2}{2L^2\sigma^2}\right)$$

Remark 4. The LCP is preserved under Lipschitz mappings: if $\mathbf{z}$ satisfies LCP with parameter σ and $g : \mathbb{R}^n \rightarrow \mathbb{R}^m$ is L_g -Lipschitz, then $g(\mathbf{z})$ satisfies LCP with parameter $L_g\sigma$.

We will make use of the following lemma in the rest of the proof:

Lemma 1. Each $\mathbf{w}_i^T \mathbf{x}^{(\ell)}$ is subgaussian with parameter $\sigma = O(\frac{1}{\sqrt{d}})$ for $\ell = 1, \dots, L$ and $i = 1, \dots, m$.

Proof. Recall that the data $\mathbf{x}^{(0)} = \mathbf{x}$ is Gaussian with a well-conditioned Σ satisfying $\text{Tr}(\Sigma) = O(1)$ by assumption. Also recall that $\mathbf{x}^{(\ell)} = \phi(\mathbf{W}^{(\ell)} \mathbf{x}^{(\ell-1)})$ by definition.

Step 1: Initial data satisfies LCP. Since $\mathbf{x}^{(0)} \sim \mathcal{N}(\mathbf{0}, \Sigma)$ with $\text{Tr}(\Sigma) = O(1)$ and Σ is well-conditioned, we have:

- $\text{Tr}(\Sigma) = \sum_{i=1}^d \lambda_i = O(1)$
- Well-conditioned means $\lambda_{\max}/\lambda_{\min} = O(1)$, so all eigenvalues are of the same order
- This implies $d \cdot \lambda_{\max} = O(1)$, hence $\lambda_{\max} = O(1/d)$

Since Gaussian random vectors satisfy LCP with parameter proportional to $\sqrt{\lambda_{\max}}$, we have that $\mathbf{x}^{(0)}$ satisfies LCP with parameter $\sigma_0 = O(1/\sqrt{d})$.

Step 2: LCP is preserved through layers. We proceed by induction on ℓ . Assume $\mathbf{x}^{(\ell-1)}$ satisfies LCP with parameter $\sigma_{\ell-1}$.

Consider the mapping $\mathbf{x}^{(\ell-1)} \mapsto \mathbf{x}^{(\ell)} = \phi(\mathbf{W}^{(\ell)} \mathbf{x}^{(\ell-1)})$. Since:

- The linear map $\mathbf{x}^{(\ell-1)} \mapsto \mathbf{W}^{(\ell)} \mathbf{x}^{(\ell-1)}$ is Lipschitz with constant $\|\mathbf{W}^{(\ell)}\|_{\text{op}}$
- The activation function ϕ is assumed to be Lipschitz (typically with constant 1 for ReLU, sigmoid, tanh, etc.)

The composition $\mathbf{x}^{(\ell-1)} \mapsto \phi(\mathbf{W}^{(\ell)} \mathbf{x}^{(\ell-1)})$ is Lipschitz with constant $L_\phi \cdot \|\mathbf{W}^{(\ell)}\|_{\text{op}}$.

By the preservation of LCP under Lipschitz mappings, $\mathbf{x}^{(\ell)}$ satisfies LCP with parameter $\sigma_\ell = L_\phi \cdot \|\mathbf{W}^{(\ell)}\|_{\text{op}} \cdot \sigma_{\ell-1}$.

Step 3: Bounding the LCP parameter. Assuming (6) holds, we have $\|\mathbf{W}^{(\ell)}\|_{\text{op}} = O(1)$ with high probability. Thus:

$$\sigma_\ell = O(1) \cdot \sigma_{\ell-1} = O(1) \cdot O(1/\sqrt{d}) = O(1/\sqrt{d})$$

for all $\ell \leq L$, maintaining the $O(1/\sqrt{d})$ LCP parameter across layers.

Step 4: Linear functionals of LCP vectors. For any fixed vector $\mathbf{w}_i$ with $\|\mathbf{w}_i\|_2 = O(1)$, the linear functional $f(\mathbf{x}^{(\ell)}) = \mathbf{w}_i^T \mathbf{x}^{(\ell)}$ is $O(1)$ -Lipschitz.

Since $\mathbf{x}^{(\ell)}$ satisfies LCP with parameter $\sigma_\ell = O(1/\sqrt{d})$, we have that $\mathbf{w}_i^T \mathbf{x}^{(\ell)}$ is subgaussian with parameter $O(1) \cdot O(1/\sqrt{d}) = O(1/\sqrt{d})$.

We will need another technical lemma as well:

Lemma 2. *The following decomposition holds for any bounded odd ϕ with bounded first, third and fifth derivatives:*

$$\phi(\mathbf{w}_i^T \mathbf{x}^{(\ell)}) = \phi'(0) \mathbf{w}_i^T \mathbf{x}^{(\ell)} + \frac{\phi'''(0)}{6} (\mathbf{w}_i^T \mathbf{x}^{(\ell)})^3 + O\left(\frac{1}{d^{5/2}}\right) \quad (14)$$

Proof. Since ϕ is odd, we have $\phi(0) = 0$ and all even derivatives vanish at 0. By Taylor's theorem with remainder:

$$\phi(z) = \phi'(0)z + \frac{\phi'''(0)}{6} z^3 + \frac{\phi^{(5)}(\xi)}{120} z^5 \quad (15)$$

for some ξ between 0 and z .

Setting $z = \mathbf{w}_i^T \mathbf{x}^{(\ell)}$:

$$\phi(\mathbf{w}_i^T \mathbf{x}^{(\ell)}) = \phi'(0) \mathbf{w}_i^T \mathbf{x}^{(\ell)} + \frac{\phi'''(0)}{6} (\mathbf{w}_i^T \mathbf{x}^{(\ell)})^3 + R_i \quad (16)$$

where the remainder term is:

$$R_i = \frac{\phi^{(5)}(\xi)}{120} (\mathbf{w}_i^T \mathbf{x}^{(\ell)})^5$$

From our earlier result, $\mathbf{w}_i^T \mathbf{x}^{(\ell)}$ is subgaussian with parameter $O(1/\sqrt{d})$. Therefore, with high probability:

$$|\mathbf{w}_i^T \mathbf{x}^{(\ell)}| = O(1/\sqrt{d})$$

Thus:

$$|R_i| = \left| \frac{\phi^{(5)}(\xi)}{120} (\mathbf{w}_i^T \mathbf{x}^{(\ell)})^5 \right| = O\left(\frac{1}{(\sqrt{d})^5}\right) = O\left(\frac{1}{d^{5/2}}\right)$$

□

Lemma 3. *Let $\mathbf{g} \sim \mathcal{N}(0, \Sigma^{(\ell)})$. Then the following holds for all i, j :*

$$\left| \mathbb{E}_{\mathbf{x}^{(\ell)}} [\phi(\mathbf{w}_i^T \mathbf{x}^{(\ell)}) \phi(\mathbf{w}_j^T \mathbf{x}^{(\ell)})] - \mathbb{E}_{\mathbf{g}} [\phi(\mathbf{w}_i^T \mathbf{g}) \phi(\mathbf{w}_j^T \mathbf{g})] \right| = O(1/d^2) \quad (17)$$

Moreover, if $i \neq j$, then one has:

$$\left| \mathbb{E}_{\mathbf{x}^{(\ell)}} [\phi(\mathbf{w}_i^T \mathbf{x}^{(\ell)}) \phi(\mathbf{w}_j^T \mathbf{x}^{(\ell)})] - \mathbb{E}_{\mathbf{g}} [\phi(\mathbf{w}_i^T \mathbf{g}) \phi(\mathbf{w}_j^T \mathbf{g})] \right| = O(1/d^3) \quad (18)$$

Proof. We will apply Taylor expansion to both terms.

For the subgaussian $\mathbf{x}^{(\ell)}$:

$$\phi(\mathbf{w}_i^T \mathbf{x}^{(\ell)}) \phi(\mathbf{w}_j^T \mathbf{x}^{(\ell)}) = \left[\phi'(0) \mathbf{w}_i^T \mathbf{x}^{(\ell)} + \frac{\phi'''(0)}{6} (\mathbf{w}_i^T \mathbf{x}^{(\ell)})^3 + O(1/d^{5/2}) \right] \quad (19)$$

$$\times \left[\phi'(0) \mathbf{w}_j^T \mathbf{x}^{(\ell)} + \frac{\phi'''(0)}{6} (\mathbf{w}_j^T \mathbf{x}^{(\ell)})^3 + O(1/d^{5/2}) \right] \quad (20)$$

Expanding:

$$= [\phi'(0)]^2 (\mathbf{w}_i^T \mathbf{x}^{(\ell)}) (\mathbf{w}_j^T \mathbf{x}^{(\ell)}) + \frac{\phi'(0)\phi'''(0)}{6} [(\mathbf{w}_i^T \mathbf{x}^{(\ell)}) (\mathbf{w}_j^T \mathbf{x}^{(\ell)})^3 + (\mathbf{w}_i^T \mathbf{x}^{(\ell)})^3 (\mathbf{w}_j^T \mathbf{x}^{(\ell)})] \quad (21)$$

$$+ \frac{[\phi'''(0)]^2}{36} (\mathbf{w}_i^T \mathbf{x}^{(\ell)})^3 (\mathbf{w}_j^T \mathbf{x}^{(\ell)})^3 + O(1/d^{5/2}) \quad (22)$$

Similarly for the Gaussian $\mathbf{g}$:

$$\phi(\mathbf{w}_i^T \mathbf{g}) \phi(\mathbf{w}_j^T \mathbf{g}) = [\phi'(0)]^2 (\mathbf{w}_i^T \mathbf{g}) (\mathbf{w}_j^T \mathbf{g}) + \frac{\phi'(0)\phi'''(0)}{6} [(\mathbf{w}_i^T \mathbf{g}) (\mathbf{w}_j^T \mathbf{g})^3 + (\mathbf{w}_i^T \mathbf{g})^3 (\mathbf{w}_j^T \mathbf{g})] \quad (23)$$

$$+ \frac{[\phi'''(0)]^2}{36} (\mathbf{w}_i^T \mathbf{g})^3 (\mathbf{w}_j^T \mathbf{g})^3 + O(1/d^{5/2}) \quad (24)$$

Now, we compare expectations of each term one by one.

For the second order term:

$$\mathbb{E}[(\mathbf{w}_i^T \mathbf{x}^{(\ell)}) (\mathbf{w}_j^T \mathbf{x}^{(\ell)})] = \mathbf{w}_i^T \boldsymbol{\Sigma}^{(\ell)} \mathbf{w}_j = \mathbb{E}[(\mathbf{w}_i^T \mathbf{g}) (\mathbf{w}_j^T \mathbf{g})] \quad (25)$$

These match exactly by assumption.

For the forth-order cross-terms like $\mathbb{E}[(\mathbf{w}_i^T \mathbf{x}^{(\ell)}) (\mathbf{w}_j^T \mathbf{x}^{(\ell)})^3]$: both are $O(1/d^2)$ since $\mathbf{w}_i^T \mathbf{x}^{(\ell)} = O(1/\sqrt{d})$ and $\mathbf{w}_j^T \mathbf{x}^{(\ell)} = O(1/\sqrt{d})$, therefore the difference is at most $O(1/d^2)$.

All remaining terms are $O(1/d^3)$ for both distributions, so we are done with the first part.

For the sharper bound for $i \neq j$, note that the fourth moment expands as:

$$\mathbb{E}[(\mathbf{w}_i^T \mathbf{x}^{(\ell)}) (\mathbf{w}_j^T \mathbf{x}^{(\ell)})^3] = \sum_{k,\ell,m,n} w_{ik} w_{j\ell} w_{jm} w_{jn} \mathbb{E}[x_k^{(\ell)} x_\ell^{(\ell)} x_m^{(\ell)} x_n^{(\ell)}]$$

Assuming that $\boldsymbol{\Sigma}_\ell$ is diagonal WLOG (rotate it to become diagonal and apply the same rotation to $\mathbf{w}_i, \mathbf{w}_j$ if not), we can assume that only terms with paired indices survive. Since each corresponding expectation is $O(\frac{1}{d^2})$, each coefficient $w_{ik} w_{j\ell} w_{jm} w_{jn}$ is of order $O(\frac{1}{d^2})$, there are $O(d^2)$ of these coefficients and $\mathbf{w}_i$ is independent from $\mathbf{w}_j$, we conclude that $\mathbb{E}[(\mathbf{w}_i^T \mathbf{x}^{(\ell)}) (\mathbf{w}_j^T \mathbf{x}^{(\ell)})^3]$ is indeed of order $O(\frac{1}{d^3})$ as desired. $\square$

Lemma 4. *Given the covariance recursions:*

$$\boldsymbol{\Sigma}_\ell = \mathbb{E}_{\mathbf{x}^{(\ell)}} [\mathbf{x}^{(\ell)} (\mathbf{x}^{(\ell)})^T], \quad \ell = 0, \dots, L \quad (26)$$

$$\tilde{\boldsymbol{\Sigma}}_\ell = \rho_1^2 \mathbf{W}^{(\ell)} \boldsymbol{\Sigma}_{\ell-1} (\mathbf{W}^{(\ell)})^T + \rho_2^2 \mathbf{I} \quad (27)$$

Suppose that $\|\boldsymbol{\Sigma}_{\ell-1} - \tilde{\boldsymbol{\Sigma}}_{\ell-1}\| = O(\delta_{\ell-1})$ for some $\delta_{\ell-1}$. Then:

$$\|\boldsymbol{\Sigma}_\ell - \tilde{\boldsymbol{\Sigma}}_\ell\| = O\left(\frac{m}{d^2} + \rho_1^2 \|\mathbf{W}^{(\ell)}\|^2 \delta_{\ell-1}\right)$$

where m is the dimension of $\mathbf{x}^{(\ell)}$.

Proof. Step 1: Express $\boldsymbol{\Sigma}_\ell$ in terms of the activation function.

Since $\mathbf{x}^{(\ell)} = \phi(\mathbf{W}^{(\ell)} \mathbf{x}^{(\ell-1)})$ applied element-wise:

$$[\boldsymbol{\Sigma}_\ell]_{ij} = \mathbb{E}[\phi(\mathbf{w}_i^T \mathbf{x}^{(\ell-1)}) \phi(\mathbf{w}_j^T \mathbf{x}^{(\ell-1)})]$$

where $\mathbf{w}_i^T$ is the i -th row of $\mathbf{W}^{(\ell)}$.

Step 2: Define intermediate Gaussian covariance.

Let $\mathbf{g}^{(\ell-1)} \sim \mathcal{N}(0, \Sigma_{\ell-1})$ and define:

$$\hat{\Sigma}_\ell = \mathbb{E}[\hat{\mathbf{x}}^{(\ell)}(\hat{\mathbf{x}}^{(\ell)})^T]$$

where $\hat{x}_i^{(\ell)} = \phi(\mathbf{w}_i^T \mathbf{g}^{(\ell-1)})$.

By the previous lemma, each entry satisfies:

$$|[\Sigma_\ell]_{ij} - [\hat{\Sigma}_\ell]_{ij}| = O(1/d^2)$$

Step 3: Relate $\hat{\Sigma}_\ell$ to the Gaussian model.

For Gaussian inputs $\mathbf{g}^{(\ell-1)}$, using Taylor expansion and Gaussian moment formulas:

$$[\hat{\Sigma}_\ell]_{ij} = \rho_1^2 \mathbf{w}_i^T \Sigma_{\ell-1} \mathbf{w}_j + \rho_2^2 \delta_{ij} + O(1/d^{3/2})$$

This can be written as:

$$\hat{\Sigma}_\ell = \rho_1^2 \mathbf{W}^{(\ell)} \Sigma_{\ell-1} (\mathbf{W}^{(\ell)})^T + \rho_2^2 \mathbf{I} + O(1/d^{3/2})$$

Step 4: Account for the approximation error from the previous layer.

Since $\|\Sigma_{\ell-1} - \tilde{\Sigma}_{\ell-1}\| = O(\delta_{\ell-1})$:

$$\|\hat{\Sigma}_\ell - \tilde{\Sigma}_\ell\| = \|\rho_1^2 \mathbf{W}^{(\ell)} \Sigma_{\ell-1} (\mathbf{W}^{(\ell)})^T - \rho_1^2 \mathbf{W}^{(\ell)} \tilde{\Sigma}_{\ell-1} (\mathbf{W}^{(\ell)})^T\| + O(1/d^{3/2}) \quad (28)$$

$$= \rho_1^2 \|\mathbf{W}^{(\ell)} (\Sigma_{\ell-1} - \tilde{\Sigma}_{\ell-1}) (\mathbf{W}^{(\ell)})^T\| + O(1/d^{3/2}) \quad (29)$$

$$\leq \rho_1^2 \|\mathbf{W}^{(\ell)}\|^2 \|\Sigma_{\ell-1} - \tilde{\Sigma}_{\ell-1}\| + O(1/d^{3/2}) \quad (30)$$

$$= O(\rho_1^2 \|\mathbf{W}^{(\ell)}\|^2 \delta_{\ell-1}) \quad (31)$$

Step 5: Combine bounds using triangle inequality.

Using the Frobenius norm argument from Step 2:

$$\|\Sigma_\ell - \hat{\Sigma}_\ell\|_F^2 = \sum_{i,j=1}^m O(1/d^4) = O(m^2/d^4)$$

Therefore $\|\Sigma_\ell - \hat{\Sigma}_\ell\| \leq \|\Sigma_\ell - \hat{\Sigma}_\ell\|_F = O(m/d^2)$.

Combining with Step 4:

$$\|\Sigma_\ell - \tilde{\Sigma}_\ell\| \leq \|\Sigma_\ell - \hat{\Sigma}_\ell\| + \|\hat{\Sigma}_\ell - \tilde{\Sigma}_\ell\| \quad (32)$$

$$= O(m/d^2) + O(\rho_1^2 \|\mathbf{W}^{(\ell)}\|^2 \delta_{\ell-1}) \quad (33)$$

□

□

B.2 Gaussian Universality

Below we verify that we can apply Theorem 1 to the outputs of the penultimate layer $\mathbf{x}^L$ of the Random Features under the assumptions made in Subsection 2.2.

Note that Assumptions 2,3 and 4 from the list of Assumptions 1 and explicitly assumed to hold in Subsection 2.2 and Theorem 2 and are inevitable if we want to apply Theorem 1.

For Assumption 1 from the list, note that

1. The mean of each row is $\boldsymbol{\mu} = 0$ because σ is assumed to be odd and the moments of $\mathbf{x}^L$ are bounded due to the subgaussianity of $\|\mathbf{x}^L\|$, which follows from the LCP property 1 of $\mathbf{x}^L$ and is proven in the Step 1 of Lemma 1.
2. In particular, $\|\boldsymbol{\mu}\| = 0 = O(1)$.
3. For any fixed vector of bounded norm, $\mathbf{v}^T \mathbf{x}^{(L)}$ is subgaussian as $\mathbf{x}^{(L)}$ with $\sigma = O(\frac{1}{\sqrt{d}})$, as it satisfies the LCP property 1, which is proven in the Step 1 of Lemma 1. This implies $\text{Var}(\mathbf{v}^T \mathbf{x}^{(L)}) = O(\frac{1}{d})$.
4. Denoting the outputs of the L -th hidden layer applied to the training samples $\mathbf{x}_1, \dots, \mathbf{x}^L$ by $\mathbf{X}^L$, it remains to verify that $\sigma_{\min}(\mathbf{X}^{(L)} \mathbf{X}^{(L)T}) = \Omega(1)$. The latter follows from the universality of the Marchenko-Pastur law for data satisfying LCP proven in Seddik et al. (2020).

C CGMT Derivations

After applying Gaussian Universality and Gaussian Equivalence, we use a framework called Convex Gaussian Min-Max Theorem Thrampoulidis et al. (2014); Akhtiamov et al. (2024a) to derive asymptotically tight expressions for the generalization error of the Random Features trained via SMD with different mirrors. For SGD, we have provided the resulting nonlinear system of scalar equations required to find the generalization error in (36). The case of general mirrors can be found in (38). Thus, (38) is the optimization referred to in Theorem 2 and (36) is its particular case corresponding to the SGD.

C.1 SGD

We consider the training datapoints $\{(\mathbf{x}_i, y_i)\}_{i=1}^n$ being generated according to the model $y_i = \mathbf{a}_*^T \phi(\mathbf{W} \mathbf{x}_i)$ where $\mathbf{x}_i \in \mathbb{R}^d$, $\mathbf{a}_* \in \mathbb{R}^D$. We denote $\mathbb{E} \mathbf{x}_i = 0$ and $\mathbb{E} \mathbf{x}_i \mathbf{x}_i^T = \boldsymbol{\Sigma}_{L-1}$. We train $\mathbf{a}$ using SGD initialized from $\mathbf{0}$ by minimizing the squared loss $\sum_{i=1}^n (\mathbf{a}^T \phi(\mathbf{W}_L \mathbf{x}_i) - y_i)^2$. Letting the input matrix $\mathbf{X} \in \mathbb{R}^{n \times d}$ with each row corresponding to $\mathbf{x}_i$, We know from the implicit bias of SGD:

$$\begin{aligned} & \min_{\mathbf{a}} \|\mathbf{a} - \mathbf{s}_0\|_2^2 \\ \text{s.t. } & \mathbf{a}^T \phi(\mathbf{W}_L \mathbf{X}^T) = \mathbf{Y}^T = \mathbf{a}_*^T \phi(\mathbf{W}_L \mathbf{X}^T) \end{aligned}$$

Where $\mathbf{Y} \in \mathbb{R}^{n \times 1}$. We define $\mathbf{a} \leftarrow \mathbf{a} - \mathbf{a}_0$. Thus we may write:

$$\begin{aligned} & \min_{\mathbf{a}} \|\mathbf{a}\|_2^2 \\ \text{s.t. } & \mathbf{G}(\mathbf{a} - \mathbf{a}_*) = \mathbf{0} \end{aligned}$$

Where for each row of $\mathbf{G} \in \mathbb{R}^{n \times D}$, $\mathbf{g}_i \in \mathbb{R}^D$, we have from the Gaussian Equivalence Principal:

$$\mathbb{E} \mathbf{G} = \mathbf{0}, \quad \boldsymbol{\Sigma}_L \approx \rho_{L,1}^2 \mathbf{W}_L \boldsymbol{\Sigma}_{L-1} \mathbf{W}_L^T + \rho_{L,2}^2 \mathbf{I}$$

The generalization error is:

$$g.e := \mathbb{E}_{\mathbf{x}} \left(\hat{\mathbf{a}}^T \phi(\mathbf{W}_L \mathbf{x}) - \mathbf{a}_*^T \phi(\mathbf{W}_L \mathbf{x}) \right)^2 = \rho_{L,1}^2 (\hat{\mathbf{a}} - \mathbf{a}_*)^T \mathbf{W}_L \Sigma_{L-1} \mathbf{W}_L^T (\hat{\mathbf{a}} - \mathbf{a}_*) + \rho_{L,2}^2 \|\hat{\mathbf{a}} - \mathbf{a}_*\|_2^2$$

Now using a Lagrange multiplier, we formulate the optimization as a min-max:

$$\min_{\mathbf{a}} \max_{\mathbf{v}_L} \|\mathbf{a}\|_2^2 + \mathbf{v}_L^T \tilde{\mathbf{G}} \Sigma_L^{1/2} (\mathbf{a} - \mathbf{a}_*)$$

Now using CGMT, we obtain:

$$\min_{\mathbf{a}} \max_{\mathbf{v}_L} \|\mathbf{v}_L\|_2 \mathbf{g}^T \Sigma_L^{1/2} (\mathbf{a} - \mathbf{a}_*) + \|\Sigma_L^{1/2} (\mathbf{a} - \mathbf{a}_*)\|_2 \mathbf{h}_o^T \mathbf{v}_L + \|\mathbf{a}\|_2^2$$

Doing the optimization over the direction of $\mathbf{v}_L$ yields:

$$\min_{\mathbf{a}} \max_{\beta > 0} \beta \mathbf{g}^T \Sigma_L^{1/2} (\mathbf{a} - \mathbf{a}_*) + \beta \|\Sigma_L^{1/2} (\mathbf{a} - \mathbf{a}_*)\|_2 \cdot \|\mathbf{h}_o\|_2 + \|\mathbf{a}\|_2^2$$

Using the square-root trick $\sqrt{t} = \frac{\tau}{2} + \frac{t}{2\tau}$, we observe:

$$\min_{\mathbf{a}} \max_{\beta > 0} \min_{\tau > 0} \beta \mathbf{g}_o^T \Sigma_L^{1/2} (\mathbf{a} - \mathbf{a}_*) + \frac{\beta \tau}{2} + \frac{\beta}{2\tau} \|\Sigma_L^{1/2} (\mathbf{a} - \mathbf{a}_*)\|_2^2 \cdot \|\mathbf{h}_o\|_2^2 + \|\mathbf{a}\|_2^2$$

Furthermore, we have that $g.e = \tau^2$. We note the convexity and concavity of the objective, hence we may exchange the order of min and max:

$$\max_{\beta > 0} \min_{\tau > 0} \frac{\beta \tau}{2} + \min_{\mathbf{a}} \beta \mathbf{g}_o^T \Sigma_L^{1/2} (\mathbf{a} - \mathbf{a}_*) + \frac{\beta}{2\tau} \|\Sigma_L^{1/2} (\mathbf{a} - \mathbf{a}_*)\|_2^2 \cdot \|\mathbf{h}_o\|_2^2 + \|\mathbf{a}\|_2^2$$

Now note that

$$\Sigma_L^{1/2} \mathbf{g}_o \sim \mathcal{N}(\mathbf{0}, \rho_{L,1}^2 \mathbf{W}_L \Sigma_{L-1} \mathbf{W}_L^T + \rho_{L,2}^2 \mathbf{I})$$

Thus we may write $\Sigma_L^{1/2} \mathbf{g}_o = \rho_{L,1} \mathbf{W}_L \Sigma_{L-1}^{1/2} \tilde{\mathbf{g}}_{L-1,1} + \rho_{L,2} \tilde{\mathbf{g}}_{L-1,2}$ with $\tilde{\mathbf{g}}_{L-1,1}$ and $\tilde{\mathbf{g}}_{L-1,2}$ being independent of each other. Therefore the optimization turns into

$$\begin{aligned} \max_{\beta > 0} \min_{\tau > 0} \frac{\beta \tau}{2} + \min_{\mathbf{a}} \rho_{L,2} \beta \tilde{\mathbf{g}}_{L-1,2}^T (\mathbf{a} - \mathbf{a}_*) + \rho_{L,1} \beta \tilde{\mathbf{g}}_{L-1,1}^T \Sigma_{L-1}^{1/2} \mathbf{W}_L^T (\mathbf{a} - \mathbf{a}_*) \\ + \frac{\rho_{L,1}^2 \beta \|\mathbf{h}_o\|_2^2}{2\tau} \left\| \Sigma_{L-1}^{1/2} \mathbf{W}_L^T (\mathbf{a} - \mathbf{a}_*) \right\|_2^2 + \frac{\rho_{L,2}^2 \beta \|\mathbf{h}_o\|_2^2}{2\tau} \|\mathbf{a} - \mathbf{a}_*\|_2^2 + \|\mathbf{a}\|_2^2 \end{aligned}$$

Now we complete the squares over $\mathbf{W}_L^T (\mathbf{a} - \mathbf{a}_*)$.

$$\begin{aligned} \max_{\beta > 0} \min_{\tau > 0} \frac{\beta \tau}{2} \left(1 - \frac{\|\tilde{\mathbf{g}}_{L-1,1}\|_2^2}{\|\mathbf{h}_o\|_2^2} \right) + \min_{\mathbf{a}} \rho_{L,2} \beta \tilde{\mathbf{g}}_{L-1,2}^T (\mathbf{a} - \mathbf{a}_*) \\ + \frac{\beta}{2\tau} \left\| \rho_{L,1} \|\mathbf{h}_o\|_2 \cdot \Sigma_{L-1}^{1/2} \mathbf{W}_L^T (\mathbf{a} - \mathbf{a}_*) + \frac{\tau}{\|\mathbf{h}_o\|_2} \tilde{\mathbf{g}}_{L-1,1} \right\|_2^2 + \frac{\rho_{L,2}^2 \beta \|\mathbf{h}_o\|_2^2}{2\tau} \|\mathbf{a} - \mathbf{a}_*\|_2^2 + \|\mathbf{a}\|_2^2 \end{aligned}$$

Now we focus on the inner optimization and we use a Fenchel dual to rewrite the quadratic term as

$$\begin{aligned} \min_{\mathbf{a}} \max_{\mathbf{u}} \frac{\beta \rho_{L,1} \|\mathbf{h}_o\|_2}{2\tau} \mathbf{u}^T \Sigma_{L-1}^{1/2} \mathbf{W}_L^T (\mathbf{a} - \mathbf{a}_*) + \frac{\beta}{2\|\mathbf{h}_o\|_2} \mathbf{u}^T \tilde{\mathbf{g}}_{L-1,1} - \frac{\beta \|\mathbf{u}\|_2^2}{8\tau} \\ + \rho_{L,2} \beta \tilde{\mathbf{g}}_{L-1,2}^T (\mathbf{a} - \mathbf{a}_*) + \frac{\rho_{L,2}^2 \beta \|\mathbf{h}_o\|_2^2}{2\tau} \|\mathbf{a} - \mathbf{a}_*\|_2^2 + \|\mathbf{a}\|_2^2 \end{aligned}$$

Swapping the min and max:

$$\begin{aligned} \max_{\mathbf{u}} \min_{\mathbf{a}} & \frac{\beta \rho_{L,1} \|\mathbf{h}_o\|_2}{2\tau} \mathbf{u}^T \Sigma_{L-1}^{1/2} \mathbf{W}_L^T (\mathbf{a} - \mathbf{a}_*) + \frac{\beta}{2\|\mathbf{h}_o\|_2} \mathbf{u}^T \tilde{\mathbf{g}}_{L-1,1} - \frac{\beta \|\mathbf{u}\|_2^2}{8\tau} \\ & + \rho_{L,2} \beta \tilde{\mathbf{g}}_{L-1,2}^T (\mathbf{a} - \mathbf{a}_*) + \frac{\rho_{L,2}^2 \beta \|\mathbf{h}_o\|_2^2}{2\tau} \|\mathbf{a} - \mathbf{a}_*\|_2^2 + \|\mathbf{a}\|_2^2 \end{aligned}$$

Employing CGMT again:

$$\begin{aligned} \max_{\mathbf{u}} \min_{\mathbf{a}} & \frac{\beta \rho_{L,1} \|\mathbf{h}_o\|_2}{2\tau} \|\Sigma_{L-1}^{1/2} \mathbf{u}\|_2 \mathbf{g}_{L-1}^T (\mathbf{a} - \mathbf{a}_*) + \frac{\beta \rho_{L,1} \|\mathbf{h}_o\|_2}{2\tau} \|\mathbf{a} - \mathbf{a}_*\|_2 \mathbf{h}_{L-1}^T \Sigma_{L-1}^{1/2} \mathbf{u} \\ & + \frac{\beta}{2\|\mathbf{h}_o\|_2} \mathbf{u}^T \tilde{\mathbf{g}}_{L-1,1} - \frac{\beta \|\mathbf{u}\|_2^2}{8\tau} + \rho_{L,2} \beta \tilde{\mathbf{g}}_{L-1,2}^T (\mathbf{a} - \mathbf{a}_*) + \frac{\rho_{L,2}^2 \beta \|\mathbf{h}_o\|_2^2}{2\tau} \|\mathbf{a} - \mathbf{a}_*\|_2^2 + \|\mathbf{a}\|_2^2 \end{aligned}$$

Now we perform the optimization over the direction of $\mathbf{a} - \mathbf{a}_*$. First we observe that

$$\|\mathbf{a}\|_2^2 = \|\mathbf{a} - \mathbf{a}_*\|_2^2 + 2\mathbf{a}_*^T (\mathbf{a} - \mathbf{a}_*) - \|\mathbf{a}_*\|_2^2$$

Dropping the constant term $\|\mathbf{a}_*\|_2^2$, we have

$$\begin{aligned} \max_{\mathbf{u}} \min_{\mathbf{a}} & \frac{\beta \rho_{L,1} \|\mathbf{h}_o\|_2}{2\tau} \|\Sigma_{L-1}^{1/2} \mathbf{u}\|_2 \mathbf{g}_{L-1}^T (\mathbf{a} - \mathbf{a}_*) + \frac{\beta \rho_{L,1} \|\mathbf{h}_o\|_2}{2\tau} \|\mathbf{a} - \mathbf{a}_*\|_2 \mathbf{h}_{L-1}^T \Sigma_{L-1}^{1/2} \mathbf{u} \\ & + \frac{\beta}{2\|\mathbf{h}_o\|_2} \mathbf{u}^T \tilde{\mathbf{g}}_{L-1,1} - \frac{\beta \|\mathbf{u}\|_2^2}{8\tau} + \rho_{L,2} \beta \tilde{\mathbf{g}}_{L-1,2}^T (\mathbf{a} - \mathbf{a}_*) + \left(\frac{\rho_{L,2}^2 \beta \|\mathbf{h}_o\|_2^2}{2\tau} + 1 \right) \|\mathbf{a} - \mathbf{a}_*\|_2^2 + 2\mathbf{a}_*^T (\mathbf{a} - \mathbf{a}_*) \end{aligned}$$

We observe that $\mathbf{a} - \mathbf{a}_*$ aligns with

$$\frac{\beta \rho_{L,1} \|\mathbf{h}_o\|_2}{2\tau} \|\Sigma_{L-1}^{1/2} \mathbf{u}\|_2 \mathbf{g}_{L-1} + \rho_{L,2} \beta \tilde{\mathbf{g}}_{L-1,2} + 2\mathbf{a}_*$$

Thus the optimization turns into

$$\begin{aligned} \max_{\mathbf{u}} \min_{\eta_{L-1} > 0} & \frac{\beta \rho_{L,1} \eta_{L-1} \|\mathbf{h}_o\|_2}{2\tau} \mathbf{h}_{L-1}^T \Sigma_{L-1}^{1/2} \mathbf{u} - \eta_{L-1} \left\| \frac{\beta \rho_{L,1} \|\mathbf{h}_o\|_2}{2\tau} \|\Sigma_{L-1}^{1/2} \mathbf{u}\|_2 \mathbf{g}_{L-1} + \rho_{L,2} \beta \tilde{\mathbf{g}}_{L-1,2} + 2\mathbf{a}_* \right\|_2 \\ & + \frac{\beta}{2\|\mathbf{h}_o\|_2} \mathbf{u}^T \tilde{\mathbf{g}}_{L-1,1} - \frac{\beta \|\mathbf{u}\|_2^2}{8\tau} + \left(\frac{\rho_{L,2}^2 \beta \|\mathbf{h}_o\|_2^2}{2\tau} + 1 \right) \eta_{L-1}^2 \end{aligned}$$

Applying the square-root trick again

$$\begin{aligned} \max_{\mathbf{u}} \min_{\eta_{L-1} > 0} \max_{\alpha_{L-1} > 0} & \frac{\beta \rho_{L,1} \eta_{L-1} \|\mathbf{h}_o\|_2}{2\tau} \mathbf{h}_{L-1}^T \Sigma_{L-1}^{1/2} \mathbf{u} - \frac{\alpha_{L-1} \eta_{L-1}}{2} + \frac{\beta}{2\|\mathbf{h}_o\|_2} \mathbf{u}^T \tilde{\mathbf{g}}_{L-1,1} - \frac{\beta \|\mathbf{u}\|_2^2}{8\tau} \\ & + \left(\frac{\rho_{L,2}^2 \beta \|\mathbf{h}_o\|_2^2}{2\tau} + 1 \right) \eta_{L-1}^2 - \frac{\eta_{L-1}}{2\alpha_{L-1}} \left(\frac{\beta^2 \rho_{L,1}^2 \|\mathbf{h}_o\|_2^2}{4\tau^2} \|\Sigma_{L-1}^{1/2} \mathbf{u}\|_2^2 \|\mathbf{g}_{L-1}\|_2^2 \right. \\ & \left. + \rho_{L,2}^2 \beta^2 \|\tilde{\mathbf{g}}_{L-1,2}\|_2^2 + 4\|\mathbf{a}_*\|_2^2 \right) \end{aligned}$$

Using convexity-concavity, we exchange the order of optimizations:

$$\begin{aligned} \min_{\eta_{L-1} > 0} \max_{\alpha_{L-1} > 0} & \left(\frac{\rho_{L,2}^2 \beta \|\mathbf{h}_o\|_2^2}{2\tau} + 1 \right) \eta_{L-1}^2 - \frac{\alpha_{L-1} \eta_{L-1}}{2} - \frac{\eta_{L-1}}{2\alpha_{L-1}} \left(\rho_{L,2}^2 \beta^2 \|\tilde{\mathbf{g}}_{L-1,2}\|_2^2 + 4\|\mathbf{a}_*\|_2^2 \right) \\ & + \max_{\mathbf{u}} \frac{\beta}{2\|\mathbf{h}_o\|_2} \mathbf{u}^T \tilde{\mathbf{g}}_{L-1,1} - \frac{\eta_{L-1} \beta^2 \rho_{L,1}^2 \|\mathbf{h}_o\|_2^2}{8\alpha_{L-1} \tau^2} \|\Sigma_{L-1}^{1/2} \mathbf{u}\|_2^2 \|\mathbf{g}_{L-1}\|_2^2 - \frac{\beta \|\mathbf{u}\|_2^2}{8\tau} + \frac{\beta \rho_{L,1} \eta_{L-1} \|\mathbf{h}_o\|_2}{2\tau} \mathbf{h}_{L-1}^T \Sigma_{L-1}^{1/2} \mathbf{u} \end{aligned}$$

We know from the recursion $\mathbf{\Sigma}_{L-1} = \rho_{L-1,1}^2 \mathbf{W}_{L-1} \mathbf{\Sigma}_{L-2} \mathbf{W}_{L-1}^T + \rho_{L-1,2}^2 \mathbf{I}$. Applying the same technique as before, we take $\mathbf{\Sigma}_{L-1}^{1/2} \mathbf{h}_{L-1} = \rho_{L-1,1} \mathbf{W}_{L-1} \mathbf{\Sigma}_{L-2}^{1/2} \tilde{\mathbf{h}}_{L-2,1} + \rho_{L-1,2} \tilde{\mathbf{h}}_{L-2,2}$ and consider the inner optimization

$$\begin{aligned} \max_{\mathbf{u}} & \frac{\beta}{2\|\mathbf{h}_o\|_2} \mathbf{u}^T \tilde{\mathbf{g}}_{L-1,1} - \frac{\eta_{L-1} \beta^2 \rho_{L,1}^2 \rho_{L-1,1}^2 \|\mathbf{h}_o\|_2^2}{8\alpha_{L-1} \tau^2} \|\mathbf{\Sigma}_{L-2}^{1/2} \mathbf{W}_{L-1}^T \mathbf{u}\|_2^2 \|\mathbf{g}_{L-1}\|_2^2 - \frac{\eta_{L-1} \beta^2 \rho_{L,1}^2 \rho_{L-1,2}^2 \|\mathbf{h}_o\|_2^2}{8\alpha_{L-1} \tau^2} \|\mathbf{u}\|_2^2 \|\mathbf{g}_{L-1}\|_2^2 \\ & - \frac{\beta \|\mathbf{u}\|_2^2}{8\tau} + \frac{\beta \rho_{L,1} \rho_{L-1,1} \eta_{L-1} \|\mathbf{h}_o\|_2}{2\tau} \tilde{\mathbf{h}}_{L-2,1}^T \mathbf{\Sigma}_{L-2}^{1/2} \mathbf{W}_{L-1}^T \mathbf{u} + \frac{\beta \rho_{L,1} \rho_{L-1,2} \eta_{L-1} \|\mathbf{h}_o\|_2}{2\tau} \tilde{\mathbf{h}}_{L-2,2}^T \mathbf{u} \end{aligned}$$

Completing the squares:

$$\begin{aligned} \max_{\mathbf{u}} & \frac{\beta}{2\|\mathbf{h}_o\|_2} \mathbf{u}^T \tilde{\mathbf{g}}_{L-1,1} - \frac{\eta_{L-1}}{2\alpha_{L-1}} \left\| \frac{\beta \rho_{L,1} \rho_{L-1,1} \|\mathbf{h}_o\|_2 \|\mathbf{g}_{L-1}\|_2}{2\tau} \mathbf{\Sigma}_{L-2}^{1/2} \mathbf{W}_{L-1}^T \mathbf{u} - \frac{\alpha_{L-1}}{\|\mathbf{g}_{L-1}\|_2} \tilde{\mathbf{h}}_{L-2,1} \right\|_2^2 \\ & + \frac{\|\tilde{\mathbf{g}}_{L-2,1}\|_2^2}{2\|\mathbf{g}_{L-1}\|_2^2} \alpha_{L-1} \eta_{L-1} - \frac{\eta_{L-1} \beta^2 \rho_{L,1}^2 \rho_{L-1,2}^2 \|\mathbf{h}_o\|_2^2}{8\alpha_{L-1} \tau^2} \|\mathbf{u}\|_2^2 \|\mathbf{g}_{L-1}\|_2^2 - \frac{\beta \|\mathbf{u}\|_2^2}{8\tau} + \frac{\beta \rho_{L,1} \rho_{L-1,2} \eta_{L-1} \|\mathbf{h}_o\|_2}{2\tau} \tilde{\mathbf{h}}_{L-2,2}^T \mathbf{u} \end{aligned}$$

We drop the term $\frac{\|\tilde{\mathbf{g}}_{L-2,1}\|_2^2}{2\|\mathbf{g}_{L-1}\|_2^2} \alpha_{L-1} \eta_{L-1}$ from the optimization as it does not depend on $\mathbf{u}$. Now introducing $\mathbf{v}_{L-2}$ as the Fenchel dual:

$$\begin{aligned} \max_{\mathbf{u}} \min_{\mathbf{v}_{L-2}} & \frac{\beta}{2\|\mathbf{h}_o\|_2} \mathbf{u}^T \tilde{\mathbf{g}}_{L-1,1} - \frac{\eta_{L-1}}{2\alpha_{L-1}} \frac{\beta \rho_{L,1} \rho_{L-1,1} \|\mathbf{h}_o\|_2 \|\mathbf{g}_{L-1}\|_2}{2\tau} \mathbf{v}_{L-2}^T \mathbf{\Sigma}_{L-2}^{1/2} \mathbf{W}_{L-1}^T \mathbf{u} + \frac{\eta_{L-1}}{2\|\mathbf{g}_{L-1}\|_2} \mathbf{v}_{L-2}^T \tilde{\mathbf{h}}_{L-2,1} \\ & + \frac{\eta_{L-1} \|\mathbf{v}_{L-2}\|_2^2}{8\alpha_{L-1}} - \frac{\eta_{L-1} \beta^2 \rho_{L,1}^2 \rho_{L-1,2}^2 \|\mathbf{h}_o\|_2^2}{8\alpha_{L-1} \tau^2} \|\mathbf{u}\|_2^2 \|\mathbf{g}_{L-1}\|_2^2 - \frac{\beta \|\mathbf{u}\|_2^2}{8\tau} + \frac{\beta \rho_{L,1} \rho_{L-1,2} \eta_{L-1} \|\mathbf{h}_o\|_2}{2\tau} \tilde{\mathbf{h}}_{L-2,2}^T \mathbf{u} \end{aligned}$$

Exchanging the order of min and max, we then apply CGMT w.r.t $\mathbf{W}$, obtaining:

$$\begin{aligned} \min_{\mathbf{v}_{L-2}} \max_{\mathbf{u}} & \frac{\beta}{2\|\mathbf{h}_o\|_2} \mathbf{u}^T \tilde{\mathbf{g}}_{L-1,1} - \frac{\eta_{L-1} \beta \rho_{L,1} \rho_{L-1,1} \|\mathbf{h}_o\|_2 \|\mathbf{g}_{L-1}\|_2}{4\alpha_{L-1} \tau} \left(\|\mathbf{\Sigma}_{L-2}^{1/2} \mathbf{v}_{L-2}\|_2 \mathbf{h}_{L-2}^T \mathbf{u} + \|\mathbf{u}\|_2 \mathbf{g}_{L-2}^T \mathbf{\Sigma}_{L-2}^{1/2} \mathbf{v}_{L-2} \right) \\ & + \frac{\eta_{L-1}}{2\|\mathbf{g}_{L-1}\|_2} \mathbf{v}_{L-2}^T \tilde{\mathbf{h}}_{L-2,1} + \frac{\eta_{L-1} \|\mathbf{v}_{L-2}\|_2^2}{8\alpha_{L-1}} - \frac{\eta_{L-1} \beta^2 \rho_{L,1}^2 \rho_{L-1,2}^2 \|\mathbf{h}_o\|_2^2}{8\alpha_{L-1} \tau^2} \|\mathbf{u}\|_2^2 \|\mathbf{g}_{L-1}\|_2^2 - \frac{\beta \|\mathbf{u}\|_2^2}{8\tau} \\ & + \frac{\beta \rho_{L,1} \rho_{L-1,2} \eta_{L-1} \|\mathbf{h}_o\|_2}{2\tau} \tilde{\mathbf{h}}_{L-2,2}^T \mathbf{u} \end{aligned}$$

Doing the optimization over the direction of $\mathbf{u}$, yields

$$\begin{aligned} \min_{\mathbf{v}_{L-2}} \max_{\eta_{L-2} > 0} & - \frac{\eta_{L-1} \beta \rho_{L,1} \rho_{L-1,1} \eta_{L-1} \|\mathbf{h}_o\|_2 \|\mathbf{g}_{L-1}\|_2}{4\alpha_{L-1} \tau} \eta_{L-2} \mathbf{g}_{L-2}^T \mathbf{\Sigma}_{L-2}^{1/2} \mathbf{v}_{L-2} + \frac{\eta_{L-1}}{2\|\mathbf{g}_{L-1}\|_2} \mathbf{v}_{L-2}^T \tilde{\mathbf{h}}_{L-2,1} \\ & + \eta_{L-2} \left\| \frac{\beta}{2\|\mathbf{h}_o\|_2} \tilde{\mathbf{g}}_{L-1,1} - \frac{\eta_{L-1} \beta \rho_{L,1} \rho_{L-1,1} \|\mathbf{h}_o\|_2 \|\mathbf{g}_{L-1}\|_2}{4\alpha_{L-1} \tau} \|\mathbf{\Sigma}_{L-2}^{1/2} \mathbf{v}_{L-2}\|_2 \mathbf{h}_{L-2} \right. \\ & \left. + \frac{\beta \rho_{L,1} \rho_{L-1,2} \eta_{L-1} \|\mathbf{h}_o\|_2}{2\tau} \tilde{\mathbf{h}}_{L-2,2} \right\|_2 + \frac{\eta_{L-1} \|\mathbf{v}_{L-2}\|_2^2}{8\alpha_{L-1}} - \left(\frac{\eta_{L-1} \beta^2 \rho_{L,1}^2 \rho_{L-1,2}^2 \|\mathbf{h}_o\|_2^2 \|\mathbf{g}_{L-1}\|_2^2}{8\alpha_{L-1} \tau^2} + \frac{\beta}{8\tau} \right) \eta_{L-2}^2 \end{aligned}$$

Applying the square-root trick again, we obtain:

$$\begin{aligned} \min_{\mathbf{v}_{L-2}} \max_{\eta_{L-2} > 0} \min_{\alpha_{L-2} > 0} & \frac{\eta_{L-2} \alpha_{L-2}}{2} - \frac{\eta_{L-1} \beta \rho_{L,1} \rho_{L-1,1} \|\mathbf{h}_o\|_2 \|\mathbf{g}_{L-1}\|_2}{4\alpha_{L-1} \tau} \eta_{L-2} \mathbf{g}_{L-2}^T \mathbf{\Sigma}_{L-2}^{1/2} \mathbf{v}_{L-2} \\ & + \frac{\eta_{L-2}}{2\alpha_{L-2}} \left(\frac{\beta^2}{4\|\mathbf{h}_o\|_2^2} \|\tilde{\mathbf{g}}_{L-1,1}\|_2^2 + \frac{\eta_{L-1}^2 \beta^2 \rho_{L,1}^2 \rho_{L-1,1}^2 \|\mathbf{h}_o\|_2^2 \|\mathbf{g}_{L-1}\|_2^2}{16\alpha_{L-1}^2 \tau^2} \|\mathbf{\Sigma}_{L-2}^{1/2} \mathbf{v}_{L-2}\|_2^2 \|\mathbf{h}_{L-2}\|_2^2 \right. \\ & \left. + \frac{\beta^2 \rho_{L,1}^2 \rho_{L-1,2}^2 \eta_{L-1}^2 \|\mathbf{h}_o\|_2^2}{4\tau^2} \|\tilde{\mathbf{h}}_{L-2,2}\|_2^2 \right) + \frac{\eta_{L-1}}{2\|\mathbf{g}_{L-1}\|_2} \mathbf{v}_{L-2}^T \tilde{\mathbf{h}}_{L-2,1} + \frac{\eta_{L-1} \|\mathbf{v}_{L-2}\|_2^2}{8\alpha_{L-1}} \\ & - \left(\frac{\eta_{L-1} \beta^2 \rho_{L,1}^2 \rho_{L-1,2}^2 \|\mathbf{h}_o\|_2^2 \|\mathbf{g}_{L-1}\|_2^2}{8\alpha_{L-1} \tau^2} + \frac{\beta}{8\tau} \right) \eta_{L-2}^2 \end{aligned}$$

We exchange the orders of min and max because of convexity and concavity

$$\begin{aligned}
& \max_{\eta_{L-2} > 0} \min_{\alpha_{L-2} > 0} \frac{\eta_{L-2} \alpha_{L-2}}{2} - \left(\frac{\eta_{L-1} \beta^2 \rho_{L,1}^2 \rho_{L-1,2}^2 \|\mathbf{h}_o\|_2^2 \|\mathbf{g}_{L-1}\|_2^2}{8\alpha_{L-1} \tau^2} + \frac{\beta}{8\tau} \right) \eta_{L-2}^2 \\
& + \frac{\eta_{L-2}}{2\alpha_{L-2}} \left(\frac{\beta^2 \|\tilde{\mathbf{g}}_{L-1,1}\|_2^2}{4\|\mathbf{h}_o\|_2^2} + \frac{\beta^2 \rho_{L,1}^2 \rho_{L-1,2}^2 \eta_{L-1}^2 \|\mathbf{h}_o\|_2^2 \|\tilde{\mathbf{h}}_{L-2,2}\|_2^2}{4\tau^2} \right) \\
& + \min_{\mathbf{v}_{L-2}} \frac{\eta_{L-2}}{2\alpha_{L-2}} \frac{\eta_{L-1}^2 \beta^2 \rho_{L,1}^2 \rho_{L-1,1}^2 \|\mathbf{h}_o\|_2^2 \|\mathbf{g}_{L-1}\|_2^2 \|\mathbf{h}_{L-2}\|_2^2}{16\alpha_{L-1}^2 \tau^2} \|\Sigma_{L-2}^{1/2} \mathbf{v}_{L-2}\|_2^2 + \frac{\eta_{L-1} \|\mathbf{v}_{L-2}\|_2^2}{8\alpha_{L-1}} \\
& - \frac{\eta_{L-1} \beta \rho_{L,1} \rho_{L-1,1} \|\mathbf{h}_o\|_2 \|\mathbf{g}_{L-1}\|_2}{4\alpha_{L-1} \tau} \eta_{L-2} \mathbf{g}_{L-2}^T \Sigma_{L-2}^{1/2} \mathbf{v}_{L-2} + \frac{\eta_{L-1}}{2\|\mathbf{g}_{L-1}\|_2} \mathbf{v}_{L-2}^T \tilde{\mathbf{h}}_{L-2,1}
\end{aligned}$$

Now we use the recursion step $\Sigma_{L-2} = \rho_{L-2,1}^2 \mathbf{W}_{L-2} \Sigma_{L-3} \mathbf{W}_{L-2}^T + \rho_{L-2,2}^2 \mathbf{I}$ and write

$$\begin{aligned}
& \min_{\mathbf{v}_{L-2}} \frac{\eta_{L-2}}{2\alpha_{L-2}} \frac{\eta_{L-1}^2 \beta^2 \rho_{L,1}^2 \rho_{L-1,1}^2 \rho_{L-2,1}^2 \|\mathbf{h}_o\|_2^2 \|\mathbf{g}_{L-1}\|_2^2 \|\mathbf{h}_{L-2}\|_2^2}{16\alpha_{L-1}^2 \tau^2} \|\Sigma_{L-3}^{1/2} \mathbf{W}_{L-2}^T \mathbf{v}_{L-2}\|_2^2 \\
& + \frac{\eta_{L-2}}{2\alpha_{L-2}} \frac{\eta_{L-1}^2 \beta^2 \rho_{L,1}^2 \rho_{L-1,1}^2 \rho_{L-2,2}^2 \|\mathbf{h}_o\|_2^2 \|\mathbf{g}_{L-1}\|_2^2 \|\mathbf{h}_{L-2}\|_2^2}{16\alpha_{L-1}^2 \tau^2} \|\mathbf{v}_{L-2}\|_2^2 \\
& - \frac{\eta_{L-1} \beta \rho_{L,1} \rho_{L-1,1} \rho_{L-2,1} \|\mathbf{h}_o\|_2 \|\mathbf{g}_{L-1}\|_2}{4\alpha_{L-1} \tau} \eta_{L-2} \tilde{\mathbf{g}}_{L-3,1}^T \Sigma_{L-3}^{1/2} \mathbf{W}_{L-2}^T \mathbf{v}_{L-2} \\
& - \frac{\eta_{L-1} \beta \rho_{L,1} \rho_{L-1,1} \rho_{L-2,2} \|\mathbf{h}_o\|_2 \|\mathbf{g}_{L-1}\|_2}{4\alpha_{L-1} \tau} \eta_{L-2} \tilde{\mathbf{g}}_{L-3,2}^T \mathbf{v}_{L-2} + \frac{\eta_{L-1}}{2\|\mathbf{g}_{L-1}\|_2} \mathbf{v}_{L-2}^T \tilde{\mathbf{h}}_{L-2,1} + \frac{\eta_{L-1} \|\mathbf{v}_{L-2}\|_2^2}{8\alpha_{L-1}}
\end{aligned}$$

Completing the squares yields

$$\begin{aligned}
& \min_{\mathbf{v}_{L-2}} \frac{\eta_{L-2}}{2\alpha_{L-2}} \left\| \frac{\eta_{L-1} \beta \rho_{L,1} \rho_{L-1,1} \rho_{L-2,1} \|\mathbf{h}_o\|_2 \|\mathbf{g}_{L-1}\|_2 \|\mathbf{h}_{L-2}\|_2}{4\alpha_{L-1} \tau} \Sigma_{L-3}^{1/2} \mathbf{W}_{L-2}^T \mathbf{v}_{L-2} - \frac{\alpha_{L-2}}{\|\mathbf{h}_{L-2}\|_2} \tilde{\mathbf{g}}_{L-3,1} \right\|_2^2 \\
& - \frac{\alpha_{L-2} \eta_{L-2}}{2} \frac{\|\tilde{\mathbf{g}}_{L-3,1}\|_2^2}{\|\mathbf{h}_{L-2}\|_2^2} + \left(\frac{\eta_{L-2}}{2\alpha_{L-2}} \frac{\eta_{L-1}^2 \beta^2 \rho_{L,1}^2 \rho_{L-1,1}^2 \rho_{L-2,2}^2 \|\mathbf{h}_o\|_2^2 \|\mathbf{g}_{L-1}\|_2^2 \|\mathbf{h}_{L-2}\|_2^2}{16\alpha_{L-1}^2 \tau^2} + \frac{\eta_{L-1}}{8\alpha_{L-1}} \right) \|\mathbf{v}_{L-2}\|_2^2 \\
& - \frac{\beta \eta_{L-1} \rho_{L,1} \rho_{L-1,1} \rho_{L-2,2} \|\mathbf{h}_o\|_2 \|\mathbf{g}_{L-1}\|_2}{4\alpha_{L-1} \tau} \eta_{L-2} \tilde{\mathbf{g}}_{L-3,2}^T \mathbf{v}_{L-2} + \frac{\eta_{L-1}}{2\|\mathbf{g}_{L-1}\|_2} \mathbf{v}_{L-2}^T \tilde{\mathbf{h}}_{L-2,1}
\end{aligned}$$

Now we consider the optimization over $\mathbf{v}_{L-2}$, we observe that, the inner optimization takes a similar

form to that what was obtained earlier. So far, we have:

$$\begin{aligned}
& \max_{\beta > 0} \min_{\tau > 0} \frac{\beta \tau}{2} \left(1 - \frac{\|\tilde{\mathbf{g}}_{L-1,1}\|_2^2}{\|\mathbf{h}_o\|_2^2} \right) \\
& + \min_{\eta_{L-1} > 0} \max_{\alpha_{L-1} > 0} - \frac{\alpha_{L-1} \eta_{L-1}}{2} \left(1 - \frac{\|\tilde{\mathbf{g}}_{L-2,1}\|_2^2}{\|\mathbf{g}_{L-1}\|_2^2} \right) + \left(\frac{\rho_{L,2}^2 \beta \|\mathbf{h}_o\|_2^2}{2\tau} + 1 \right) \eta_{L-1}^2 - \frac{\eta_{L-1}}{2\alpha_{L-1}} \left(\beta^2 \rho_{L,2}^2 \|\tilde{\mathbf{g}}_{L-1,2}\|_2^2 + 4 \|\mathbf{a}_*\|_2^2 \right) \\
& + \max_{\eta_{L-2} > 0} \min_{\alpha_{L-2} > 0} \frac{\eta_{L-2} \alpha_{L-2}}{2} \left(1 - \frac{\|\tilde{\mathbf{g}}_{L-3,1}\|_2^2}{\|\mathbf{h}_{L-2}\|_2^2} \right) - \left(\frac{\eta_{L-1} \beta^2 \rho_{L,1}^2 \rho_{L-1,2}^2 \|\mathbf{h}_o\|_2^2 \|\mathbf{g}_{L-1}\|_2^2}{8\alpha_{L-1} \tau^2} + \frac{\beta}{8\tau} \right) \eta_{L-2}^2 \\
& \quad + \frac{\eta_{L-2}}{2\alpha_{L-2}} \left(\frac{\beta^2 \|\tilde{\mathbf{g}}_{L-1,1}\|_2^2}{4\|\mathbf{h}_o\|_2^2} + \frac{\beta^2 \rho_{L,1}^2 \rho_{L-1,2}^2 \eta_{L-1}^2 \|\mathbf{h}_o\|_2^2 \|\mathbf{h}_{L-2,2}\|_2^2}{4\tau^2} \right) \\
& \quad + \min_{\mathbf{v}_{L-2}} \frac{\eta_{L-2}}{2\alpha_{L-2}} \left\| \frac{\eta_{L-1} \beta \rho_{L,1} \rho_{L-1,1} \rho_{L-2,1} \|\mathbf{h}_o\|_2 \|\mathbf{g}_{L-1}\|_2 \|\mathbf{h}_{L-2}\|_2}{4\alpha_{L-1} \tau} \boldsymbol{\Sigma}_{L-3}^{1/2} \mathbf{W}_{L-2}^T \mathbf{v}_{L-2} - \frac{\alpha_{L-2}}{\|\mathbf{h}_{L-2}\|_2} \tilde{\mathbf{g}}_{L-3,1} \right\|_2^2 \\
& \quad + \left(\frac{\eta_{L-2}}{2\alpha_{L-2}} \frac{\eta_{L-1}^2 \beta^2 \rho_{L,1}^2 \rho_{L-1,1}^2 \rho_{L-2,2}^2 \|\mathbf{h}_o\|_2^2 \|\mathbf{g}_{L-1}\|_2^2 \|\mathbf{h}_{L-2}\|_2^2}{16\alpha_{L-1}^2 \tau^2} + \frac{\eta_{L-1}}{8\alpha_{L-1}} \right) \|\mathbf{v}_{L-2}\|_2^2 \\
& \quad - \frac{\beta \eta_{L-1} \rho_{L,1} \rho_{L-1,1} \rho_{L-2,2} \|\mathbf{h}_o\|_2 \|\mathbf{g}_{L-1}\|_2}{4\alpha_{L-1} \tau} \eta_{L-2} \tilde{\mathbf{g}}_{L-3,2}^T \mathbf{v}_{L-2} + \frac{\eta_{L-1}}{2\|\mathbf{g}_{L-1}\|_2} \mathbf{v}_{L-2}^T \tilde{\mathbf{h}}_{L-2,1}
\end{aligned}$$

Continuing this process, for the final optimization we have:

$$\min_{\mathbf{v}_1} \frac{\eta_2}{2\alpha_2} \left\| c_L \boldsymbol{\Sigma}_0^{1/2} \mathbf{W}_1^T \mathbf{v}_1 - \frac{\alpha_2}{\|\mathbf{h}_1\|_2} \tilde{\mathbf{g}}_{0,1} \right\|_2^2 + \frac{\eta_2}{2\alpha_2} \frac{c_L^2 \rho_{1,2}^2}{\rho_{1,1}^2} \|\mathbf{v}_1\|_2^2 - \frac{c_L \rho_{1,2}}{\rho_{1,1} \|\mathbf{g}_1\|_2} \eta_2 \tilde{\mathbf{g}}_{0,2}^T \mathbf{v}_1 + \frac{\eta_3}{2\|\mathbf{g}_1\|_2} \mathbf{v}_1^T \tilde{\mathbf{h}}_{1,1} + \frac{\eta_3 \|\mathbf{v}_1\|_2^2}{8\alpha_3}$$

Using Fenchel dual and introducing $\mathbf{v}_0$, yields

$$\begin{aligned}
& \min_{\mathbf{v}_1} \max_{\mathbf{v}_0} \frac{\eta_2}{2\alpha_2} c_L \mathbf{v}_0^T \boldsymbol{\Sigma}_0^{1/2} \mathbf{W}_1^T \mathbf{v}_1 - \frac{\eta_2}{2\|\mathbf{h}_1\|_2} \mathbf{v}_0^T \tilde{\mathbf{g}}_{0,1} - \frac{\eta_2 \|\mathbf{v}_1\|_2^2}{8\alpha_2} \\
& \quad + \frac{\eta_2}{2\alpha_2} \frac{c_L^2 \rho_{1,2}^2}{\rho_{1,1}^2} \|\mathbf{v}_1\|_2^2 - \frac{c_L \rho_{1,2}}{\rho_{1,1} \|\mathbf{g}_1\|_2} \eta_2 \tilde{\mathbf{g}}_{0,2}^T \mathbf{v}_1 + \frac{\eta_3}{2\|\mathbf{g}_1\|_2} \mathbf{v}_1^T \tilde{\mathbf{h}}_{1,1} + \frac{\eta_3 \|\mathbf{v}_1\|_2^2}{8\alpha_3}
\end{aligned}$$

Applying CGMT

$$\begin{aligned}
& \min_{\mathbf{v}_1} \max_{\mathbf{v}_0} \frac{\eta_2}{2\alpha_2} c_L \|\boldsymbol{\Sigma}_0^{1/2} \mathbf{v}_0\|_2 \mathbf{g}_0^T \mathbf{v}_1 + \frac{\eta_2}{2\alpha_2} c_L \|\mathbf{v}_1\|_2 \mathbf{h}_0^T \boldsymbol{\Sigma}_0^{1/2} \mathbf{v}_0 - \frac{\eta_2}{2\|\mathbf{h}_1\|_2} \mathbf{v}_0^T \tilde{\mathbf{g}}_{0,1} - \frac{\eta_2 \|\mathbf{v}_0\|_2^2}{8\alpha_2} \\
& \quad + \frac{\eta_2}{2\alpha_2} \frac{c_L^2 \rho_{1,2}^2}{\rho_{1,1}^2} \|\mathbf{v}_1\|_2^2 - \frac{c_L \rho_{1,2}}{\rho_{1,1} \|\mathbf{g}_1\|_2} \eta_2 \tilde{\mathbf{g}}_{0,2}^T \mathbf{v}_1 + \frac{\eta_3}{2\|\mathbf{g}_1\|_2} \mathbf{v}_1^T \tilde{\mathbf{h}}_{1,1} + \frac{\eta_3 \|\mathbf{v}_1\|_2^2}{8\alpha_3}
\end{aligned}$$

Doing the optimization over $\mathbf{v}_1$:

$$\begin{aligned}
& \max_{\mathbf{v}_0} \min_{\eta_1 > 0} - \eta_1 \left\| \frac{c_L \eta_2}{2\alpha_2} \|\boldsymbol{\Sigma}_0^{1/2} \mathbf{v}_0\|_2 \mathbf{g}_1 - \frac{c_L \rho_{1,2}}{\rho_{1,1} \|\mathbf{g}_1\|_2} \eta_2 \tilde{\mathbf{g}}_{0,2} + \frac{\eta_3}{2\|\mathbf{g}_1\|_2} \tilde{\mathbf{h}}_{1,1} \right\|_2^2 - \frac{\eta_2}{2\|\mathbf{h}_1\|_2} \mathbf{v}_0^T \tilde{\mathbf{g}}_{0,1} - \frac{\eta_2 \|\mathbf{v}_0\|_2^2}{8\alpha_2} \\
& \quad + \frac{\eta_2}{2\alpha_2} c_L \eta_1 \mathbf{h}_1^T \boldsymbol{\Sigma}_0^{1/2} \mathbf{v}_0 + \frac{\eta_2}{2\alpha_2} \frac{c_L^2 \rho_{1,2}^2}{\rho_{1,1}^2} \eta_1^2 + \frac{\eta_3 \eta_1^2}{8\alpha_3}
\end{aligned}$$

Now using the square-root trick, we have

$$\begin{aligned}
& \max_{\mathbf{v}_0} \min_{\eta_1 > 0} \max_{\alpha_1 > 0} - \frac{\alpha_1 \eta_1}{2} - \frac{\eta_1}{2\alpha_1} \left(\frac{c_L^2 \eta_2^2}{4\alpha_2^2} \|\boldsymbol{\Sigma}_0^{1/2} \mathbf{v}_0\|_2^2 \|\mathbf{g}_1\|_2^2 + \frac{c_L^2 \rho_{1,2}^2}{\rho_{1,1}^2 \|\mathbf{g}_1\|_2^2} \eta_2^2 \|\tilde{\mathbf{g}}_{0,2}\|_2^2 + \frac{\eta_3^2}{4\|\mathbf{g}_1\|_2^2} \|\tilde{\mathbf{h}}_{1,1}\|_2^2 \right) \\
& \quad - \frac{\eta_2}{2\|\mathbf{h}_1\|_2} \mathbf{v}_0^T \tilde{\mathbf{g}}_{0,1} - \frac{\eta_2 \|\mathbf{v}_0\|_2^2}{8\alpha_2} + \frac{\eta_2}{2\alpha_2} c_L \eta_1 \mathbf{h}_1^T \boldsymbol{\Sigma}_0^{1/2} \mathbf{v}_0 + \frac{\eta_2}{2\alpha_2} \frac{c_L^2 \rho_{1,2}^2}{\rho_{1,1}^2} \eta_1^2 + \frac{\eta_3 \eta_1^2}{8\alpha_3}
\end{aligned}$$

Swapping min, max, we complete the squares over $\mathbf{v}_0$:

$$\begin{pmatrix} \mathbf{v}_0^T & 1 \end{pmatrix} \begin{pmatrix} \frac{\eta_1}{2\alpha_1} \frac{c_L^2 \eta_2^2 \|\mathbf{g}_1\|_2^2}{4\alpha_2^2} \mathbf{\Sigma}_0 + \frac{\eta_2}{8\alpha_2} \mathbf{I} & -\frac{\eta_2}{4\|\mathbf{h}_1\|_2} \tilde{\mathbf{g}}_{0,1} + \frac{\eta_2}{4\alpha_2} c_L \eta_1 \mathbf{\Sigma}_0^{1/2} \mathbf{h}_1 \\ -\frac{\eta_2}{4\|\mathbf{h}_1\|_2} \tilde{\mathbf{g}}_{0,1}^T + \frac{\eta_2}{4\alpha_2} c_L \eta_1 \mathbf{h}_1^T \mathbf{\Sigma}_0^{1/2} & 0 \end{pmatrix} \begin{pmatrix} \mathbf{v}_0 \\ 1 \end{pmatrix}$$

Which yields the scalar optimization

$$\begin{aligned} \min_{\eta_1 > 0} \max_{\alpha_1 > 0} & -\frac{\alpha_1 \eta_1}{2} - \frac{\eta_1}{2\alpha_1} \left(\frac{c_L^2 \rho_{1,2}^2}{\rho_{1,1}^2 \|\mathbf{g}_1\|_2^2} \eta_2^2 \|\tilde{\mathbf{g}}_{0,2}\|_2^2 + \frac{\eta_3^2}{4\|\mathbf{g}_1\|_2^2} \|\tilde{\mathbf{h}}_{1,1}\|_2^2 \right) + \left(\frac{\eta_3}{8\alpha_3} + \frac{\eta_2}{2\alpha_2} \frac{c_L^2 \rho_{1,2}^2}{\rho_{1,1}^2} \right) \eta_1^2 \\ & + \left(-\frac{\eta_2}{4\|\mathbf{h}_1\|_2} \tilde{\mathbf{g}}_{0,1} + \frac{\eta_2}{4\alpha_2} c_L \eta_1 \mathbf{\Sigma}_0^{1/2} \mathbf{h}_1 \right)^T \left(\frac{\eta_1}{2\alpha_1} \frac{c_L^2 \eta_2^2 \|\mathbf{g}_1\|_2^2}{4\alpha_2^2} \mathbf{\Sigma}_0 + \frac{\eta_2}{8\alpha_2} \mathbf{I} \right)^{-1} \left(-\frac{\eta_2}{4\|\mathbf{h}_1\|_2} \tilde{\mathbf{g}}_{0,1} + \frac{\eta_2}{4\alpha_2} c_L \eta_1 \mathbf{\Sigma}_0^{1/2} \mathbf{h}_1 \right) \end{aligned}$$

By Hanson-Wright's inequality, we have that

$$\begin{aligned} & \left(-\frac{\eta_2}{4\|\mathbf{h}_1\|_2} \tilde{\mathbf{g}}_{0,1} + \frac{\eta_2}{4\alpha_2} c_L \eta_1 \mathbf{\Sigma}_0^{1/2} \mathbf{h}_1 \right)^T \left(\frac{\eta_1}{2\alpha_1} \frac{c_L^2 \eta_2^2 \|\mathbf{g}_1\|_2^2}{4\alpha_2^2} \mathbf{\Sigma}_0 + \frac{\eta_2}{8\alpha_2} \mathbf{I} \right)^{-1} \left(-\frac{\eta_2}{4\|\mathbf{h}_1\|_2} \tilde{\mathbf{g}}_{0,1} + \frac{\eta_2}{4\alpha_2} c_L \eta_1 \mathbf{\Sigma}_0^{1/2} \mathbf{h}_1 \right) \\ & \xrightarrow{\mathbb{P}} \frac{\eta_2^2}{16\|\mathbf{h}_1\|_2^2} \text{Tr} \left(\frac{c_L^2 \eta_2^2 \|\mathbf{g}_1\|_2^2}{4\alpha_2^2} \mathbf{\Sigma}_0 + \frac{\eta_2}{8\alpha_2} \mathbf{I} \right)^{-1} + \frac{\eta_2^2 c_L^2 \eta_1^2}{16\alpha_2^2} \text{Tr} \left(\frac{c_L^2 \eta_2^2 \|\mathbf{g}_1\|_2^2}{4\alpha_2^2} \mathbf{\Sigma}_0 + \frac{\eta_2}{8\alpha_2} \mathbf{I} \right)^{-1} \mathbf{\Sigma}_0 \end{aligned}$$

Thus the last scalar optimization would be:

$$\begin{aligned} \min_{\eta_1 > 0} \max_{\alpha_1 > 0} & -\frac{\alpha_1 \eta_1}{2} - \frac{\eta_1}{2\alpha_1} \left(\frac{c_L^2 \rho_{1,2}^2}{\rho_{1,1}^2 \|\mathbf{g}_1\|_2^2} \eta_2^2 \|\tilde{\mathbf{g}}_{0,2}\|_2^2 + \frac{\eta_3^2}{4\|\mathbf{g}_1\|_2^2} \|\tilde{\mathbf{h}}_{1,1}\|_2^2 \right) + \left(\frac{\eta_3}{8\alpha_3} + \frac{\eta_2}{2\alpha_2} \frac{c_L^2 \rho_{1,2}^2}{\rho_{1,1}^2} \right) \eta_1^2 \\ & + \frac{\eta_2^2}{16\|\mathbf{h}_1\|_2^2} \text{Tr} \left(\frac{\eta_1}{2\alpha_1} \frac{c_L^2 \eta_2^2 \|\mathbf{g}_1\|_2^2}{4\alpha_2^2} \mathbf{\Sigma}_0 + \frac{\eta_2}{8\alpha_2} \mathbf{I} \right)^{-1} + \frac{\eta_2^2 c_L^2 \eta_1^2}{16\alpha_2^2} \text{Tr} \left(\frac{\eta_1}{2\alpha_1} \frac{c_L^2 \eta_2^2 \|\mathbf{g}_1\|_2^2}{4\alpha_2^2} \mathbf{\Sigma}_0 + \frac{\eta_2}{8\alpha_2} \mathbf{I} \right)^{-1} \mathbf{\Sigma}_0 \end{aligned}$$

The final optimization is as follows:

$$\begin{aligned} & \max_{\beta > 0} \min_{\tau > 0} \frac{\beta \tau}{2} \left(1 - \frac{\|\tilde{\mathbf{g}}_{L-1,1}\|_2^2}{\|\mathbf{h}_o\|_2^2} \right) \\ & + \min_{\eta_{L-1} > 0} \max_{\alpha_{L-1} > 0} -\frac{\alpha_{L-1} \eta_{L-1}}{2} \left(1 - \frac{\|\tilde{\mathbf{g}}_{L-2,1}\|_2^2}{\|\mathbf{g}_{L-1}\|_2^2} \right) + \left(\frac{\rho_{L,2}^2 \beta \|\mathbf{h}_o\|_2^2}{2\tau} + 1 \right) \eta_{L-1}^2 - \frac{\eta_{L-1}}{2\alpha_{L-1}} \left(\beta^2 \rho_{L,2}^2 \|\tilde{\mathbf{g}}_{L-1,2}\|_2^2 + 4\|\mathbf{a}_*\|_2^2 \right) \\ & + \max_{\eta_{L-2} > 0} \min_{\alpha_{L-2} > 0} \frac{\eta_{L-2} \alpha_{L-2}}{2} \left(1 - \frac{\|\tilde{\mathbf{g}}_{L-3,1}\|_2^2}{\|\mathbf{h}_{L-2}\|_2^2} \right) - \left(\frac{\eta_{L-1} \beta^2 \rho_{L,1}^2 \rho_{L-1,2}^2 \|\mathbf{h}_o\|_2^2 \|\mathbf{g}_{L-1}\|_2^2}{8\alpha_{L-1} \tau^2} + \frac{\beta}{8\tau} \right) \eta_{L-2}^2 \\ & + \frac{\eta_{L-2}}{2\alpha_{L-2}} \left(\frac{\beta^2 \|\tilde{\mathbf{g}}_{L-1,1}\|_2^2}{4\|\mathbf{h}_o\|_2^2} + \frac{\beta^2 \rho_{L,1}^2 \rho_{L-1,2}^2 \eta_{L-1}^2 \|\mathbf{h}_o\|_2^2 \|\mathbf{h}_{L-2,2}\|_2^2}{4\tau^2} \right) \\ & \dots \\ & + \min_{\eta_1 > 0} \max_{\alpha_1 > 0} -\frac{\alpha_1 \eta_1}{2} - \frac{\eta_1}{2\alpha_1} \left(\frac{c_L^2 \rho_{1,2}^2}{\rho_{1,1}^2 \|\mathbf{g}_1\|_2^2} \eta_2^2 \|\tilde{\mathbf{g}}_{0,2}\|_2^2 + \frac{\eta_3^2}{4\|\mathbf{g}_1\|_2^2} \|\tilde{\mathbf{h}}_{1,1}\|_2^2 \right) + \left(\frac{\eta_3}{8\alpha_3} + \frac{\eta_2}{2\alpha_2} \frac{c_L^2 \rho_{1,2}^2}{\rho_{1,1}^2} \right) \eta_1^2 \\ & + \frac{\eta_2^2}{16\|\mathbf{h}_1\|_2^2} \text{Tr} \left(\frac{\eta_1}{2\alpha_1} \frac{c_L^2 \eta_2^2 \|\mathbf{g}_1\|_2^2}{4\alpha_2^2} \mathbf{\Sigma}_0 + \frac{\eta_2}{8\alpha_2} \mathbf{I} \right)^{-1} + \frac{\eta_2^2 c_L^2 \eta_1^2}{16\alpha_2^2} \text{Tr} \left(\frac{\eta_1}{2\alpha_1} \frac{c_L^2 \eta_2^2 \|\mathbf{g}_1\|_2^2}{4\alpha_2^2} \mathbf{\Sigma}_0 + \frac{\eta_2}{8\alpha_2} \mathbf{I} \right)^{-1} \mathbf{\Sigma}_0 \end{aligned}$$

We take derivative with respect to α_1 :

$$\begin{aligned} \frac{\partial}{\partial \alpha_1} = 0 \Rightarrow 0 = & -\frac{\eta_1}{2} + \frac{\eta_1}{\alpha_1^2} \left(\frac{c_L^2 \rho_{1,2}^2}{\rho_{1,1}^2 \|\mathbf{g}_1\|_2^2} \eta_2^2 \|\tilde{\mathbf{g}}_{0,2}\|_2^2 + \frac{\eta_3^2}{4 \|\mathbf{g}_1\|_2^2} \|\tilde{\mathbf{h}}_{1,1}\|_2^2 \right) \\ & + \frac{\eta_2^2}{16 \|\mathbf{h}_1\|_2^2} \frac{\eta_1}{2\alpha_1^2} \frac{c_L^2 \eta_2^2 \|\mathbf{g}_1\|_2^2}{4\alpha_2^2} \text{Tr} \left(\frac{\eta_1}{2\alpha_1} \frac{c_L^2 \eta_2^2 \|\mathbf{g}_1\|_2^2}{4\alpha_2^2} \Sigma_0 + \frac{\eta_2}{8\alpha_2} \mathbf{I} \right)^{-2} \Sigma_0 \\ & + \frac{\eta_2^2 c_L^2 \eta_1^2}{16\alpha_2^2} \frac{\eta_1}{2\alpha_1^2} \frac{c_L^2 \eta_2^2 \|\mathbf{g}_1\|_2^2}{4\alpha_2^2} \text{Tr} \left(\frac{\eta_1}{2\alpha_1} \frac{c_L^2 \eta_2^2 \|\mathbf{g}_1\|_2^2}{4\alpha_2^2} \Sigma_0 + \frac{\eta_2}{8\alpha_2} \mathbf{I} \right)^{-2} \Sigma_0^2 \end{aligned}$$

For η_1 and α_1 , after taking derivatives we observe that:

$$0 = -\alpha_1 + 2 \left(\frac{\eta_3}{8\alpha_3} + 2 \frac{\eta_2}{2\alpha_2} \frac{c_L^2 \rho_{1,2}^2}{\rho_{1,1}^2} \right) \eta_1 + 2\eta_1 \frac{\eta_2^2 c_L^2}{16\alpha_2^2} \text{Tr} \left(\frac{\eta_1}{2\alpha_1} \frac{c_L^2 \eta_2^2 \|\mathbf{g}_1\|_2^2}{4\alpha_2^2} \Sigma_0 + \frac{\eta_3}{8\alpha_3} \mathbf{I} \right)^{-1} \Sigma_0$$

Hence, denoting $\zeta_1 = \frac{\eta_1}{\alpha_1}$, we observe that

$$\zeta_1 = \frac{1 - \zeta_1 \zeta_2^2 c_L^2 \text{Tr} \left(c_L^2 \zeta_1 \zeta_2^2 \|\mathbf{g}_1\|_2^2 \Sigma_0 + \zeta_2 \mathbf{I} \right)^{-1}}{\frac{\zeta_3}{4} + \frac{2c_L^2 \zeta_2 \rho_{1,2}^2}{\rho_{1,1}^2}}$$

Now we can find α_1 by:

$$\alpha_1^2 = \frac{\left(\frac{2c_L^2 \rho_{1,2}^2}{\rho_{1,1}^2 \|\mathbf{g}_1\|_2^2} \|\tilde{\mathbf{g}}_{0,2}\|_2^2 + \frac{c_L^2 \zeta_2^2 \|\mathbf{g}_1\|_2^2}{\|\mathbf{h}_1\|_2^2} \text{Tr} \left(\zeta_1 c_L^2 \zeta_2^2 \|\mathbf{g}_1\|_2^2 \Sigma_0 + \zeta_2 \mathbf{I} \right)^{-2} \Sigma_0 \right) \zeta_2^2 \alpha_2^2 + \frac{\|\tilde{\mathbf{h}}_{1,1}\|_2^2}{2 \|\mathbf{g}_1\|_2^2} \zeta_3^2 \alpha_3^2}{1 - \zeta_2^4 c_L^4 \zeta_1^2 \|\mathbf{g}_1\|_2^2 \text{Tr} \left(\zeta_1 c_L^2 \zeta_2^2 \|\mathbf{g}_1\|_2^2 \Sigma_0 + \zeta_2 \mathbf{I} \right)^{-2} \Sigma_0^2}$$

For η_2, α_2 we observe that

$$\begin{aligned} & \max_{\eta_2 > 0} \min_{\alpha_2 > 0} \frac{\eta_2 \alpha_2}{2} \left(1 - \frac{\|\tilde{\mathbf{g}}_{1,1}\|_2^2}{\|\mathbf{h}_2\|_2^2} \right) - \left(\frac{\eta_4}{8\alpha_4} + \frac{\eta_3}{2\alpha_3} \frac{c_{L-1}^2 \rho_{2,2}^2}{\rho_{2,1}^2} \right) \eta_2^2 + \frac{\eta_2}{2\alpha_2} \left(\frac{\eta_4^2 \|\tilde{\mathbf{g}}_{L-1,1}\|_2^2}{4 \|\mathbf{h}_o\|_2^2} + \eta_3^2 \frac{c_{L-1}^2 \rho_{2,2}^2}{\rho_{2,1}^2} \right) \\ & + \min_{\eta_1 > 0} \max_{\alpha_1 > 0} -\frac{\alpha_1 \eta_1}{2} - \frac{\eta_1}{2\alpha_1} \left(\frac{c_L^2 \rho_{1,2}^2}{\rho_{1,1}^2 \|\mathbf{g}_1\|_2^2} \eta_2^2 \|\tilde{\mathbf{g}}_{0,2}\|_2^2 + \frac{\eta_3^2}{4 \|\mathbf{g}_1\|_2^2} \|\tilde{\mathbf{h}}_{1,1}\|_2^2 \right) + \left(\frac{\eta_3}{8\alpha_3} + \frac{\eta_2}{2\alpha_2} \frac{c_L^2 \rho_{1,2}^2}{\rho_{1,1}^2} \right) \eta_1^2 \\ & + \frac{\eta_2^2}{16 \|\mathbf{h}_1\|_2^2} \text{Tr} \left(\frac{\eta_1}{2\alpha_1} \frac{c_L^2 \eta_2^2 \|\mathbf{g}_1\|_2^2}{4\alpha_2^2} \Sigma_0 + \frac{\eta_2}{8\alpha_2} \mathbf{I} \right)^{-1} + \frac{\eta_2^2 c_L^2 \eta_1^2}{16\alpha_2^2} \text{Tr} \left(\frac{\eta_1}{2\alpha_1} \frac{c_L^2 \eta_2^2 \|\mathbf{g}_1\|_2^2}{4\alpha_2^2} \Sigma_0 + \frac{\eta_2}{8\alpha_2} \mathbf{I} \right)^{-1} \Sigma_0 \end{aligned}$$

Which yields:

$$\begin{aligned} \zeta_2 &= \frac{1 - \frac{1}{\|\mathbf{h}_1\|_2^2} \text{Tr} \left(c_L^2 \zeta_2 \|\mathbf{g}_1\|_2^2 \zeta_1 \Sigma_0 + \mathbf{I} \right)^{-1}}{\frac{2\zeta_3 c_{L-1}^2 \rho_{2,2}^2}{2\rho_{2,1}^2} + \frac{\zeta_4}{4}} \\ \alpha_2 &= \frac{\frac{c_L^2 \rho_{1,2}^2}{\rho_{1,1}^2} \zeta_1^2 \alpha_1^2 + F_1}{1 - \frac{d_2}{d_1}} \end{aligned}$$

Where

$$\begin{aligned}
F_1 := & -\frac{\zeta_2^2 \alpha_2^2}{\|\mathbf{h}_1\|_2^2} \left[c_L^2 \zeta_2^2 \zeta_1 \|\mathbf{g}_1\|_2^2 \text{Tr} \left(c_L^2 \zeta_2^2 \zeta_1 \|\mathbf{g}_1\|_2^2 \mathbf{\Sigma}_0 + \zeta_2 \mathbf{I} \right)^{-2} \mathbf{\Sigma}_0 + \text{Tr} \left(c_L^2 \eta_2^2 \zeta_1 \|\mathbf{g}_1\|_2^2 \mathbf{\Sigma}_0 + \zeta_2 \mathbf{I} \right)^{-2} \right] \\
& + \zeta_2 c_L^2 \zeta_1^2 \alpha_1^2 \text{Tr} \left(c_L^2 \zeta_2^2 \zeta_1 \|\mathbf{g}_1\|_2^2 \mathbf{\Sigma}_0 + \zeta_2 \mathbf{I} \right)^{-1} \mathbf{\Sigma}_0 \\
& - \zeta_2^2 c_L^2 \zeta_1^2 \alpha_1^2 \left[c_L^2 \zeta_2^2 \zeta_1 \|\mathbf{g}_1\|_2^2 \text{Tr} \left(c_L^2 \zeta_2^2 \zeta_1 \|\mathbf{g}_1\|_2^2 \mathbf{\Sigma}_0 + \zeta_2 \mathbf{I} \right)^{-2} \mathbf{\Sigma}_0^2 + \text{Tr} \left(c_L^2 \zeta_2^2 \zeta_1 \|\mathbf{g}_1\|_2^2 \mathbf{\Sigma}_0 + \zeta_2 \mathbf{I} \right)^{-2} \mathbf{\Sigma}^2 \right]
\end{aligned} \tag{34}$$

Overall, we observe that, for each pair of (η_i, α_i) the inner optimization is a function of $\frac{\eta_i}{\alpha_i}$ for $i > 2$ and also η_i^2 for $i = 2$. As demonstrated above, to find the generalization error, we find ζ_i 's first and then solve the linear system of equations in terms of α_i^2 . This implies that for $i > 2$:

$$\begin{aligned}
\alpha_i^2 &= \frac{\frac{c_{L-i+1}^2 \rho_{i,2}^2}{\rho_{i,1}^2 \|\mathbf{g}_i\|_2^2} \zeta_{i+1}^2 \alpha_{i+1}^2 \|\tilde{\mathbf{g}}_{i-1,2}\|_2^2 + \frac{\eta_{i+2}^2}{4 \|\mathbf{g}_i\|_2^2} \|\tilde{\mathbf{h}}_{i,1}\|_2^2 + \frac{F'}{\zeta_i^2} + \sum_{j=1}^i \frac{c_{L-j}^2 \rho_{j,2}^2}{\rho_{j,1}^2} \zeta_j^2 \alpha_j^2}{1 - \frac{d_{i+1}}{d_i}} \\
\zeta_i &= \frac{1 - \frac{d_{i+1}}{d_i}}{\frac{c_{L-i}^2 \rho_{i,2}^2 \zeta_{i+1}}{\rho_{i,1}^2} + \frac{\zeta_{i+2}}{8}}
\end{aligned}$$

With

$$F' := \frac{\partial}{\partial x} \left[\frac{\zeta_2^2 \alpha_2^2}{\|\mathbf{h}_1\|_2^2} \text{Tr} \left(c_L^2 x^2 \zeta_1 \|\mathbf{g}_1\|_2^2 \mathbf{\Sigma}_0 + \zeta_2 \mathbf{I} \right)^{-2} \mathbf{\Sigma}_0 + \zeta_2^2 c_L^2 \zeta_1^2 \alpha_1^2 \text{Tr} \left(c_L^2 x^2 \zeta_1 \|\mathbf{g}_1\|_2^2 \mathbf{\Sigma}_0 + \zeta_2 \mathbf{I} \right)^{-2} \right] \tag{35}$$

Finally, for β, τ , we would have:

$$\begin{aligned}
\theta &= \frac{1 - \frac{d_L}{d_{L-1}}}{\zeta_{L-2} \frac{d_1}{d_0} + \zeta_{L-1} \rho_{L,2}^2 \|\tilde{\mathbf{g}}_{L-1,2}\|_2^2} \\
\tau^2 &= \frac{\frac{F'}{\zeta_i^2} + \sum_{j=1}^{L-1} \frac{c_{L-j}^2 \rho_{j,2}^2}{\rho_{j,1}^2} \zeta_j^2 \alpha_j^2}{1 - \frac{d_L}{d_{L-1}}}
\end{aligned}$$

Summarizing, we have

$$\begin{aligned}
\zeta_1 &= \frac{1 - \zeta_1 \zeta_2^2 c_L^2 \text{Tr} \left(c_L^2 \zeta_1 \zeta_2^2 \mathbf{\Sigma}_0 + \zeta_2 \mathbf{I} \right)^{-1}}{\frac{\zeta_3}{4} + \frac{2c_L^2 \zeta_2 \rho_{1,2}^2}{\rho_{1,1}^2}}, & \zeta_2 &= \frac{1 - \frac{1}{d_1} \text{Tr} \left(c_L^2 \zeta_2 \zeta_1 \mathbf{\Sigma}_0 + \mathbf{I} \right)^{-1}}{\frac{2\zeta_3 c_{L-1}^2 \rho_{2,2}^2}{2\rho_{2,1}^2} + \frac{\zeta_4}{4}} \\
\zeta_i &= \frac{1 - \frac{d_{i+1}}{d_i}}{\frac{c_{L-i}^2 \rho_{i,2}^2 \zeta_{i+1}}{\rho_{i,1}^2} + \frac{\zeta_{i+2}}{8}}, & i &= 3, \dots, L-1, & \theta &= \frac{1 - \frac{d_L}{d_{L-1}}}{\zeta_{L-2} \frac{d_1}{d_0} + \zeta_{L-1} \rho_{L,2}^2 d_{L-1}} \\
\alpha_1^2 &= \frac{\left(\frac{2c_L^2 \rho_{1,2}^2 d_0}{\rho_{1,1}^2 d_1} + \frac{c_L^2 \zeta_2^2}{d_1} \text{Tr} \left(\zeta_1 c_L^2 \zeta_2^2 \mathbf{\Sigma}_0 + \zeta_2 \mathbf{I} \right)^{-2} \mathbf{\Sigma}_0 \right) \zeta_2^2 \alpha_2^2 + \frac{d_1}{2d_0} \zeta_3^2 \alpha_3^2}{1 - \zeta_2^4 c_L^4 \zeta_1^2 \text{Tr} \left(\zeta_1 c_L^2 \zeta_2^2 \mathbf{\Sigma}_0 + \zeta_2 \mathbf{I} \right)^{-2} \mathbf{\Sigma}_0^2} \\
\alpha_2^2 &= \frac{\frac{c_L^2 \rho_{1,2}^2}{\rho_{1,1}^2} \zeta_1^2 \alpha_1^2 + F_1}{1 - \frac{d_2}{d_1}} \\
\alpha_i^2 &= \frac{\frac{c_{L-i+1}^2 \rho_{i,2}^2 d_{i-1}}{\rho_{i,1}^2 d_{i+1}} \zeta_{i+1}^2 \alpha_{i+1}^2 + \frac{\eta_{i+2}^2 d_i}{4d_{i+1}} + \frac{F'}{\zeta_i^2} + \sum_{j=1}^i \frac{c_{L-j}^2 \rho_{j,2}^2}{\rho_{j,1}^2} \zeta_j^2 \alpha_j^2}{1 - \frac{d_{i+1}}{d_i}}, & i &= 3, \dots, L-1 \\
\tau^2 &= \frac{\frac{F'}{\zeta_i^2} + \sum_{j=1}^{L-1} \frac{c_{L-j}^2 \rho_{j,2}^2}{\rho_{j,1}^2} \zeta_j^2 \alpha_j^2}{1 - \frac{d_L}{d_{L-1}}}
\end{aligned} \tag{36}$$

Where F and F' are defined in (34) and (35), respectively. And we let $c_j := \prod_{\ell=L-j+1}^L d_\ell \zeta_\ell$. To find τ^2 , the generalization error, first we find ζ_i 's and θ through the nonlinear equations described above. Note that these equations are only in terms of ζ_i 's and θ . Then we proceed to solve the linear system of equations in α_i^2 and τ^2 to find the generalization error.

C.2 General mirrors

After completing the analysis of the case of $\psi = \|\cdot\|_2^2$ in the previous section, we show that for the general case, we can use the results from the previous section. For that, consider the following optimization:

$$\begin{aligned}
&\min_{\mathbf{a}} D_\psi(\mathbf{a}, \mathbf{a}_0) \\
&s.t \quad \mathbf{G}(\mathbf{a} - \mathbf{a}_*) = \mathbf{0}
\end{aligned}$$

We have Now using a Lagrange multiplier, we the optimization as min-max:

$$\min_{\mathbf{a}} \max_{\mathbf{v}_L} D_\psi(\mathbf{a}, \mathbf{a}_0) + \mathbf{v}_L^T \tilde{\mathbf{G}} \mathbf{\Sigma}_L^{1/2} (\mathbf{a} - \mathbf{a}_*)$$

Now using CGMT, we obtain:

$$\min_{\mathbf{a}} \max_{\mathbf{v}_L} \|\mathbf{v}_L\|_2 \mathbf{g}^T \mathbf{\Sigma}_L^{1/2} (\mathbf{a} - \mathbf{a}_*) + \|\mathbf{\Sigma}_L^{1/2} (\mathbf{a} - \mathbf{a}_*)\|_2 \mathbf{h}_o^T \mathbf{v}_L + D_\psi(\mathbf{a}, \mathbf{a}_0)$$

Doing the optimization over the direction of $\mathbf{v}_L$ yields:

$$\min_{\mathbf{a}} \max_{\beta > 0} \beta \mathbf{g}^T \mathbf{\Sigma}_L^{1/2} (\mathbf{a} - \mathbf{a}_*) + \beta \|\mathbf{\Sigma}_L^{1/2} (\mathbf{a} - \mathbf{a}_*)\|_2 \cdot \|\mathbf{h}_o\|_2 + D_\psi(\mathbf{a}, \mathbf{a}_0)$$

Using the square-root trick $\sqrt{t} = \frac{\tau}{2} + \frac{t}{2\tau}$, we observe:

$$\min_{\mathbf{a}} \max_{\beta > 0} \min_{\tau > 0} \beta \mathbf{g}_o^T \Sigma_L^{1/2} (\mathbf{a} - \mathbf{a}_*) + \frac{\beta\tau}{2} + \frac{\beta}{2\tau} \|\Sigma_L^{1/2} (\mathbf{a} - \mathbf{a}_*)\|_2^2 \cdot \|\mathbf{h}_o\|_2^2 + D_\psi(\mathbf{a}, \mathbf{a}_0)$$

Furthermore, we have that $g.e = \tau^2$. We note the convexity and concavity of the objective, hence we may exchange the order of min and max:

$$\max_{\beta > 0} \min_{\tau > 0} \frac{\beta\tau}{2} + \min_{\mathbf{a}} \beta \mathbf{g}_o^T \Sigma_L^{1/2} (\mathbf{a} - \mathbf{a}_*) + \frac{\beta}{2\tau} \|\Sigma_L^{1/2} (\mathbf{a} - \mathbf{a}_*)\|_2^2 \cdot \|\mathbf{h}_o\|_2^2 + D_\psi(\mathbf{a}, \mathbf{a}_0)$$

Now note that

$$\Sigma_L^{1/2} \mathbf{g}_o \sim \mathcal{N}(\mathbf{0}, \rho_{L,1}^2 \mathbf{W}_L \Sigma_{L-1} \mathbf{W}_L^T + \rho_{L,2}^2 \mathbf{I})$$

Thus we may write $\Sigma_L^{1/2} \mathbf{g}_o = \rho_{L,1} \mathbf{W}_L \Sigma_{L-1}^{1/2} \tilde{\mathbf{g}}_{L-1,1} + \rho_{L,2} \tilde{\mathbf{g}}_{L-1,2}$ with $\tilde{\mathbf{g}}_{L-1,1}$ and $\tilde{\mathbf{g}}_{L-1,2}$ being independent of each other. Therefore the optimization turns into

$$\begin{aligned} \max_{\beta > 0} \min_{\tau > 0} \frac{\beta\tau}{2} + \min_{\mathbf{a}} \rho_{L,2} \beta \tilde{\mathbf{g}}_{L-1,2}^T (\mathbf{a} - \mathbf{a}_*) + \rho_{L,1} \beta \tilde{\mathbf{g}}_{L-1,1}^T \Sigma_{L-1}^{1/2} \mathbf{W}_L^T (\mathbf{a} - \mathbf{a}_*) \\ + \frac{\rho_{L,1}^2 \beta \|\mathbf{h}_o\|_2^2}{2\tau} \left\| \Sigma_{L-1}^{1/2} \mathbf{W}_L^T (\mathbf{a} - \mathbf{a}_*) \right\|_2^2 + \frac{\rho_{L,2}^2 \beta \|\mathbf{h}_o\|_2^2}{2\tau} \|\mathbf{a} - \mathbf{a}_*\|_2^2 + D_\psi(\mathbf{a}, \mathbf{a}_0) \end{aligned}$$

Now we complete the square over $\mathbf{W}_L^T (\mathbf{a} - \mathbf{a}_*)$.

$$\begin{aligned} \max_{\beta > 0} \min_{\tau > 0} \frac{\beta\tau}{2} \left(1 - \frac{\|\tilde{\mathbf{g}}_{L-1,1}\|_2^2}{\|\mathbf{h}_o\|_2^2} \right) + \min_{\mathbf{a}} \rho_{L,2} \beta \tilde{\mathbf{g}}_{L-1,2}^T (\mathbf{a} - \mathbf{a}_*) \\ + \frac{\beta}{2\tau} \left\| \rho_{L,1} \|\mathbf{h}_o\|_2 \cdot \Sigma_{L-1}^{1/2} \mathbf{W}_L^T (\mathbf{a} - \mathbf{a}_*) + \frac{\tau}{\|\mathbf{h}_o\|_2} \tilde{\mathbf{g}}_{L-1,1} \right\|_2^2 + \frac{\rho_{L,2}^2 \beta \|\mathbf{h}_o\|_2^2}{2\tau} \|\mathbf{a} - \mathbf{a}_*\|_2^2 + D_\psi(\mathbf{a}, \mathbf{a}_0) \end{aligned}$$

Now we focus on the inner optimization and we use a Fenchel dual to rewrite the quadratic term as

$$\begin{aligned} \min_{\mathbf{a}} \max_{\mathbf{u}} \frac{\beta \rho_{L,1} \|\mathbf{h}_o\|_2}{2\tau} \mathbf{u}^T \Sigma_{L-1}^{1/2} \mathbf{W}_L^T (\mathbf{a} - \mathbf{a}_*) + \frac{\beta}{2\|\mathbf{h}_o\|_2} \mathbf{u}^T \tilde{\mathbf{g}}_{L-1,1} - \frac{\beta \|\mathbf{u}\|_2^2}{8\tau} \\ + \rho_{L,2} \beta \tilde{\mathbf{g}}_{L-1,2}^T (\mathbf{a} - \mathbf{a}_*) + \frac{\rho_{L,2}^2 \beta \|\mathbf{h}_o\|_2^2}{2\tau} \|\mathbf{a} - \mathbf{a}_*\|_2^2 + D_\psi(\mathbf{a}, \mathbf{a}_0) \end{aligned}$$

Employing CGMT again:

$$\begin{aligned} \min_{\mathbf{a}} \max_{\mathbf{u}} \frac{\beta \rho_{L,1} \|\mathbf{h}_o\|_2}{2\tau} \|\Sigma_{L-1}^{1/2} \mathbf{u}\|_2 \mathbf{g}_{L-1}^T (\mathbf{a} - \mathbf{a}_*) + \frac{\beta \rho_{L,1} \|\mathbf{h}_o\|_2}{2\tau} \|\mathbf{a} - \mathbf{a}_*\|_2 \mathbf{h}_{L-1}^T \Sigma_{L-1}^{1/2} \mathbf{u} \\ + \frac{\beta}{2\|\mathbf{h}_o\|_2} \mathbf{u}^T \tilde{\mathbf{g}}_{L-1,1} - \frac{\beta \|\mathbf{u}\|_2^2}{8\tau} + \rho_{L,2} \beta \tilde{\mathbf{g}}_{L-1,2}^T (\mathbf{a} - \mathbf{a}_*) + \frac{\rho_{L,2}^2 \beta \|\mathbf{h}_o\|_2^2}{2\tau} \|\mathbf{a} - \mathbf{a}_*\|_2^2 + D_\psi(\mathbf{a}, \mathbf{a}_0) \end{aligned}$$

We use the change of variable $\tilde{\mathbf{u}} := \Sigma_{L-1}^{1/2} \mathbf{u}$ and use the Lagrange multiplier to bring in the constraints:

$$\begin{aligned} \min_{\mathbf{a}} \min_{\mathbf{v}_L} \mathbf{v}_L^T (\tilde{\mathbf{u}} - \Sigma_{L-1}^{1/2} \mathbf{u}) + \frac{\beta \rho_{L,1} \|\mathbf{h}_o\|_2}{2\tau} \|\tilde{\mathbf{u}}\|_2 \mathbf{g}_{L-1}^T (\mathbf{a} - \mathbf{a}_*) + \frac{\beta \rho_{L,1} \|\mathbf{h}_o\|_2}{2\tau} \|\mathbf{a} - \mathbf{a}_*\|_2 \mathbf{h}_{L-1}^T \tilde{\mathbf{u}} \\ + \frac{\beta}{2\|\mathbf{h}_o\|_2} \mathbf{u}^T \tilde{\mathbf{g}}_{L-1,1} - \frac{\beta \|\mathbf{u}\|_2^2}{8\tau} + \rho_{L,2} \beta \tilde{\mathbf{g}}_{L-1,2}^T (\mathbf{a} - \mathbf{a}_*) + \frac{\rho_{L,2}^2 \beta \|\mathbf{h}_o\|_2^2}{2\tau} \|\mathbf{a} - \mathbf{a}_*\|_2^2 + D_\psi(\mathbf{a}, \mathbf{a}_0) \end{aligned}$$

Now we perform the optimization over $\mathbf{u}, \tilde{\mathbf{u}}$ and obtain

$$\begin{aligned} \min_{\mathbf{a}, \mathbf{v}_L} \max_{\eta_L, \tilde{\eta}_L} & \left\| \mathbf{v}_L + \frac{\beta \rho_{L,1} \|\mathbf{h}_o\|_2}{2\tau} \|\mathbf{a} - \mathbf{a}_*\|_2 \mathbf{h}_{L-1} \right\|_2 + \eta_L \left\| \Sigma_{L-1}^{1/2} \mathbf{v}_L + \frac{\beta}{2\|\mathbf{h}_o\|_2} \tilde{\mathbf{g}}_{L-1,1} \right\|_2 + D_\psi(\mathbf{a}, \mathbf{a}_0) \\ & + \frac{\beta \rho_{L,1} \|\mathbf{h}_o\|_2}{2\tau} \tilde{\eta}_L \mathbf{g}_{L-1}^T (\mathbf{a} - \mathbf{a}_*) - \frac{\beta \eta_L^2}{8\tau} + \rho_{L,2} \beta \tilde{\mathbf{g}}_{L-1,2}^T (\mathbf{a} - \mathbf{a}_*) + \frac{\rho_{L,2}^2 \beta \|\mathbf{h}_o\|_2^2}{2\tau} \|\mathbf{a} - \mathbf{a}_*\|_2^2 \end{aligned}$$

Now using the square-root trick again:

$$\begin{aligned} \min_{\mathbf{a}, \mathbf{v}_L} \max_{\eta_L, \tilde{\eta}_L} \min_{\alpha_L > 0, \tilde{\alpha}_L > 0} & \frac{\tilde{\eta}_L \tilde{\alpha}_L}{2} + \frac{\tilde{\eta}_L}{2\tilde{\alpha}_L} \left\| \mathbf{v}_L + \frac{\beta \rho_{L,1} \|\mathbf{h}_o\|_2}{2\tau} \|\mathbf{a} - \mathbf{a}_*\|_2 \mathbf{h}_{L-1} \right\|_2^2 \\ & + \frac{\alpha_L \eta_L}{2} + \frac{\eta_L}{2\alpha_L} \left\| \Sigma_{L-1}^{1/2} \mathbf{v}_L + \frac{\beta}{2\|\mathbf{h}_o\|_2} \tilde{\mathbf{g}}_{L-1,1} \right\|_2^2 + \frac{\beta \rho_{L,1} \|\mathbf{h}_o\|_2}{2\tau} \tilde{\eta}_L \mathbf{g}_{L-1}^T (\mathbf{a} - \mathbf{a}_*) \\ & - \frac{\beta \eta_L^2}{8\tau} + \rho_{L,2} \beta \tilde{\mathbf{g}}_{L-1,2}^T (\mathbf{a} - \mathbf{a}_*) + \frac{\rho_{L,2}^2 \beta \|\mathbf{h}_o\|_2^2}{2\tau} \|\mathbf{a} - \mathbf{a}_*\|_2^2 + D_\psi(\mathbf{a}, \mathbf{a}_0) \end{aligned}$$

Using the recursion, $\Sigma_{L-1} = \rho_{L-1,1}^2 \mathbf{W}_{L-1} \Sigma_{L-2} \mathbf{W}_{L-1}^T + \rho_{L-1,2}^2 \mathbf{I}$. Applying the same technique as before, we take $\Sigma_{L-1}^{1/2} \tilde{\mathbf{g}}_{L-1,1} = \rho_{L-1,1} \mathbf{W} \Sigma_{L-2}^{1/2} \tilde{\mathbf{g}}_{L-2,1} + \rho_{L-1,2} \tilde{\mathbf{g}}_{L-2,2}$, we can write:

$$\begin{aligned} \left\| \Sigma_{L-1}^{1/2} \mathbf{v}_L + \frac{\beta}{2\|\mathbf{h}_o\|_2} \tilde{\mathbf{g}}_{L-1,1} \right\|_2^2 &= \left\| \rho_{L-1,1} \Sigma_{L-2}^{1/2} \mathbf{W}_{L-2}^T \mathbf{v}_L + \frac{\beta}{2\|\mathbf{h}_o\|_2} \tilde{\mathbf{g}}_{L-2,1} \right\|_2^2 \\ &+ \rho_{L-1,2}^2 \|\mathbf{v}_L\|_2^2 + \frac{\beta^2 \rho_{L-1,2}^2}{4\|\mathbf{h}_o\|_2^2} \|\tilde{\mathbf{g}}_{L-2,2}\|_2^2 + \rho_{L-1,2} \tilde{\mathbf{g}}_{L-2,2}^T \mathbf{v}_L \end{aligned}$$

Plugging back in, we only consider the optimization over $\mathbf{a}, \mathbf{v}_L$:

$$\begin{aligned} \min_{\mathbf{a}} \max_{\tilde{\alpha}_L} & \frac{\tilde{\eta}_L}{2\tilde{\alpha}_L} \frac{\beta \rho_{L,1} \|\mathbf{h}_o\|_2}{2\tau} \tilde{\eta}_L \mathbf{g}_{L-1}^T (\mathbf{a} - \mathbf{a}_*) + \rho_{L,2} \beta \tilde{\mathbf{g}}_{L-1,2}^T (\mathbf{a} - \mathbf{a}_*) + \frac{\rho_{L,2}^2 \beta \|\mathbf{h}_o\|_2^2}{2\tau} \|\mathbf{a} - \mathbf{a}_*\|_2^2 + D_\psi(\mathbf{a}, \mathbf{a}_0) \\ & + \min_{\mathbf{v}_L} \frac{\eta_L}{2\alpha_L} \left\| \rho_{L-1,1} \Sigma_{L-2}^{1/2} \mathbf{W}_{L-2}^T \mathbf{v}_L + \frac{\beta}{2\|\mathbf{h}_o\|_2} \tilde{\mathbf{g}}_{L-2,1} \right\|_2^2 + \frac{\eta_L}{2\alpha_L} \rho_{L-1,2} \tilde{\mathbf{g}}_{L-2,2}^T \mathbf{v}_L \\ & + \frac{\eta_L}{2\alpha_L} \rho_{L-1,2}^2 \|\mathbf{v}_L\|_2^2 + \frac{\tilde{\eta}_L}{2\tilde{\alpha}_L} \left\| \mathbf{v}_L + \frac{\beta \rho_{L,1} \|\mathbf{h}_o\|_2}{2\tau} \|\mathbf{a} - \mathbf{a}_*\|_2 \mathbf{h}_{L-1} \right\|_2^2 \end{aligned}$$

Focusing on the inner optimization, we introduce the Fenchel dual variable $\mathbf{u}_{L-1}$. We note that from here on, the procedure is similar to that of SGD analysis ($\psi = \|\cdot\|_2^2$)

$$\begin{aligned} \min_{\mathbf{v}_L} \max_{\mathbf{u}_{L-1}} & \frac{\eta_L}{2\alpha_L} \rho_{L-1,1} \mathbf{u}_{L-1}^T \Sigma_{L-2}^{1/2} \mathbf{W}_{L-2}^T \mathbf{v}_L - \frac{\eta_L \|\mathbf{u}_{L-1}\|_2^2}{8\alpha_L} + \frac{\beta}{2\|\mathbf{h}_o\|_2} \mathbf{u}_{L-1}^T \tilde{\mathbf{g}}_{L-2,1} + \frac{\eta_L}{2\alpha_L} \rho_{L-1,2} \tilde{\mathbf{g}}_{L-2,2}^T \mathbf{v}_L \\ & + \frac{\eta_L}{2\alpha_L} \rho_{L-1,2}^2 \|\mathbf{v}_L\|_2^2 + \frac{\tilde{\eta}_L}{2\tilde{\alpha}_L} \left\| \mathbf{v}_L + \frac{\beta \rho_{L,1} \|\mathbf{h}_o\|_2}{2\tau} \|\mathbf{a} - \mathbf{a}_*\|_2 \mathbf{h}_{L-1} \right\|_2^2 \end{aligned}$$

Applying CGMT yields,

$$\begin{aligned} \min_{\mathbf{v}_L} \max_{\mathbf{u}_{L-1}} & \frac{\eta_L \rho_{L-1,1}}{2\alpha_L} \|\mathbf{v}_L\|_2 \mathbf{g}_{L-2}^T \Sigma_{L-2}^{1/2} \mathbf{u}_{L-1} + \frac{\eta_L \rho_{L-1,1}}{2\alpha_L} \|\Sigma_{L-2}^{1/2} \mathbf{u}_{L-1}\|_2 \mathbf{h}_{L-2}^T \mathbf{v}_L \\ & - \frac{\eta_L \|\mathbf{u}_{L-1}\|_2^2}{8\alpha_L} + \frac{\beta}{2\|\mathbf{h}_o\|_2} \mathbf{u}_{L-1}^T \tilde{\mathbf{g}}_{L-2,1} + \frac{\eta_L}{2\alpha_L} \rho_{L-1,2} \tilde{\mathbf{g}}_{L-2,2}^T \mathbf{v}_L \\ & + \frac{\eta_L}{2\alpha_L} \rho_{L-1,2}^2 \|\mathbf{v}_L\|_2^2 + \frac{\tilde{\eta}_L}{2\tilde{\alpha}_L} \left\| \mathbf{v}_L + \frac{\beta \rho_{L,1} \|\mathbf{h}_o\|_2}{2\tau} \|\mathbf{a} - \mathbf{a}_*\|_2 \mathbf{h}_{L-1} \right\|_2^2 \end{aligned}$$

We perform the optimization over the direction of $\mathbf{v}_L$, we have

$$\begin{aligned} \max_{\mathbf{u}_{L-1}} \min_{\eta_{L-1}} & \frac{\eta_L \rho_{L-1,1}}{2\alpha_L} \eta_{L-1} \mathbf{g}_{L-2}^T \boldsymbol{\Sigma}_{L-2}^{1/2} \mathbf{u}_{L-1} \\ & - \eta_{L-1} \left\| \frac{\eta_L \rho_{L-1,1}}{2\alpha_L} \boldsymbol{\Sigma}_{L-2}^{1/2} \mathbf{u}_{L-1} \right\|_2 \|\mathbf{h}_{L-2}\|_2 + \frac{\eta_L}{2\alpha_L} \rho_{L-1,2} \tilde{\mathbf{g}}_{L-2,2} + \frac{\tilde{\eta}_L}{\tilde{\alpha}_L} \frac{\beta \rho_{L,1} \|\mathbf{h}_o\|_2}{2\tau} \|\mathbf{a} - \mathbf{a}_*\|_2 \|\mathbf{h}_{L-1}\|_2 \\ & - \frac{\eta_L \|\mathbf{u}_{L-1}\|_2^2}{8\alpha_L} + \frac{\beta}{2\|\mathbf{h}_o\|_2} \mathbf{u}_{L-1}^T \tilde{\mathbf{g}}_{L-2,1} + \left(\frac{\eta_L}{2\alpha_L} \rho_{L-1,2}^2 + \frac{\tilde{\eta}_L}{2\tilde{\alpha}_L} \right) \eta_{L-1}^2 + \frac{\tilde{\eta}_L}{2\tilde{\alpha}_L} \frac{\beta \|\mathbf{a} - \mathbf{a}_*\|_2^2 \rho_{L,1}^2 \|\mathbf{h}_o\|_2^2 \|\mathbf{h}_{L-1}\|_2^2}{4\tau^2} \end{aligned}$$

Using the square-root trick again

$$\begin{aligned} \max_{\mathbf{u}_{L-1}} \min_{\eta_{L-1}} \max_{\alpha_{L-1}} & \frac{\eta_L \rho_{L-1,1}}{2\alpha_L} \eta_{L-1} \mathbf{g}_{L-2}^T \boldsymbol{\Sigma}_{L-2}^{1/2} \mathbf{u}_{L-1} - \frac{\alpha_{L-1} \eta_{L-1}}{2} \\ & - \frac{\eta_{L-1}}{2\alpha_{L-1}} \left(\frac{\eta_L^2 \rho_{L-1,1}^2 \|\mathbf{h}_{L-2}\|_2^2}{4\alpha_L^2} \|\boldsymbol{\Sigma}_{L-2}^{1/2} \mathbf{u}_{L-1}\|_2^2 + \frac{\eta_L^2}{4\alpha_L^2} \rho_{L-1,2}^2 \|\tilde{\mathbf{g}}_{L-2,2}\|_2^2 \right. \\ & \left. + \frac{\tilde{\eta}_L^2}{\tilde{\alpha}_L^2} \frac{\beta^2 \|\mathbf{a} - \mathbf{a}_*\|_2^2 \rho_{L,1}^2 \|\mathbf{h}_o\|_2^2 \|\mathbf{h}_{L-1}\|_2^2}{4\tau^2} \right) - \frac{\eta_L \|\mathbf{u}_{L-1}\|_2^2}{8\alpha_L} + \frac{\beta}{2\|\mathbf{h}_o\|_2} \mathbf{u}_{L-1}^T \tilde{\mathbf{g}}_{L-2,1} \\ & + \left(\frac{\eta_L}{2\alpha_L} \rho_{L-1,2}^2 + \frac{\tilde{\eta}_L}{2\tilde{\alpha}_L} \right) \eta_{L-1}^2 + \frac{\tilde{\eta}_L}{2\tilde{\alpha}_L} \frac{\beta^2 \|\mathbf{a} - \mathbf{a}_*\|_2^2 \rho_{L,1}^2 \|\mathbf{h}_o\|_2^2 \|\mathbf{h}_{L-1}\|_2^2}{4\tau^2} \end{aligned}$$

Now similar to before, we repeatedly apply CGMT and arrive at the following optimization:

$$\begin{aligned} & \max_{\beta > 0} \min_{\tau > 0} \frac{\beta \tau}{2} \left(1 - \frac{\|\tilde{\mathbf{g}}_{L-1,1}\|_2^2}{\|\mathbf{h}_o\|_2^2} \right) \\ & + \max_{\eta_L, \tilde{\eta}_L} \min_{\alpha_L > 0, \tilde{\alpha}_L > 0} \frac{\tilde{\eta}_L \tilde{\alpha}_L}{2} + \frac{\alpha_L \eta_L}{2} - \frac{\beta \eta_L^2}{8\tau} + \frac{\eta_L}{2\alpha_L} \frac{\beta^2 \rho_{L-1,2}^2 \|\tilde{\mathbf{g}}_{L-2,2}\|_2^2}{4\|\mathbf{h}_o\|_2^2} \\ & + \min_{\mathbf{a}} (\mathbf{a} - \mathbf{a}_*)^T \left(\frac{\tilde{\eta}_L}{2\tilde{\alpha}_L} \frac{\beta \rho_{L,1} \|\mathbf{h}_o\|_2}{2\tau} \tilde{\eta}_L \mathbf{g}_{L-1} + \rho_{L,2} \beta \tilde{\mathbf{g}}_{L-1,2} \right) + \frac{\rho_{L,2}^2 \beta \|\mathbf{h}_o\|_2^2}{2\tau} \|\mathbf{a} - \mathbf{a}_*\|_2^2 \\ & + \frac{\tilde{\eta}_L}{2\tilde{\alpha}_L} \frac{\beta^2 \|\mathbf{a} - \mathbf{a}_*\|_2^2 \rho_{L,1}^2 \|\mathbf{h}_o\|_2^2 \|\mathbf{h}_{L-1}\|_2^2}{4\tau^2} + D_\psi(\mathbf{a}, \mathbf{a}_0) + F\left(\|\mathbf{a} - \mathbf{a}_*\|_2^2 \frac{\beta^2}{\tau^2}, \frac{\eta_L}{\alpha_L}, \frac{\tilde{\eta}_L}{\tilde{\alpha}_L}, \beta\right) \end{aligned}$$

Where F is defined as

$$\begin{aligned} & \max_{\eta_{L-1} > 0} \min_{\alpha_{L-1} > 0} - \left(1 - \frac{d_{L-2}}{d_{L-3}} \right) \frac{\alpha_{L-1} \eta_{L-1}}{2} + \left(\frac{\eta_L}{2\alpha_L} \rho_{L-1,2}^2 + \frac{\tilde{\eta}_L}{2\tilde{\alpha}_L} \right) \eta_{L-1}^2 \\ & - \frac{\eta_{L-1}}{2\alpha_{L-1}} \left(\frac{\eta_L^2}{4\alpha_L^2} \rho_{L-1,2}^2 d_{L-2} + \frac{\tilde{\eta}_L^2}{\tilde{\alpha}_L^2} \frac{\beta^2 \|\mathbf{a} - \mathbf{a}_*\|_2^2 \rho_{L,1}^2 n d_{L-1}}{4\tau^2} \right) \\ & \dots + \min_{\eta_1 > 0} \max_{\alpha_1 > 0} - \frac{\alpha_1 \eta_1}{2} - \frac{\eta_1}{2\alpha_1} \left(\frac{c_L^2 \rho_{1,2}^2 d_0}{\rho_{1,1}^2 d_1} \eta_2^2 + \frac{\eta_2^2 d_1}{4d_2} \right) + \left(\frac{\eta_3}{8\alpha_3} + \frac{\eta_2}{2\alpha_2} \frac{c_L^2 \rho_{1,2}^2}{\rho_{1,1}^2} \right) \eta_1^2 \\ & + \frac{\eta_2^2}{16\|\mathbf{h}_1\|_2^2} \text{Tr} \left(\frac{\eta_1}{2\alpha_1} \frac{c_L^2 \eta_2^2}{4\alpha_2^2} \boldsymbol{\Sigma}_0 + \frac{\eta_2}{8\alpha_2} \mathbf{I} \right)^{-1} + \frac{\eta_2^2 c_L^2 \eta_1^2}{16\alpha_2^2} \text{Tr} \left(\frac{\eta_1}{2\alpha_1} \frac{c_L^2 \eta_2^2}{4\alpha_2^2} \boldsymbol{\Sigma}_0 + \frac{\eta_2}{8\alpha_2} \mathbf{I} \right)^{-1} \boldsymbol{\Sigma}_0 \end{aligned} \tag{37}$$

Using the Lagrange multiplier λ , we set $\xi := \|\mathbf{a} - \mathbf{a}_*\|_2^2$ and we perform the optimization over $\mathbf{a}$ by completing the squares and obtain that

$$\begin{aligned} & \min_{\mathbf{a}} (\mathbf{a} - \mathbf{a}_*)^T \left(\frac{\tilde{\eta}_L}{2\tilde{\alpha}_L} \frac{\beta \rho_{L,1} \|\mathbf{h}_o\|_2}{2\tau} \tilde{\eta}_L \mathbf{g}_{L-1} + \rho_{L,2} \beta \tilde{\mathbf{g}}_{L-1,2} \right) + \frac{\rho_{L,2}^2 \beta \|\mathbf{h}_o\|_2^2}{2\tau} \|\mathbf{a} - \mathbf{a}_*\|_2^2 \\ & + \frac{\tilde{\eta}_L}{2\tilde{\alpha}_L} \frac{\beta^2 \|\mathbf{a} - \mathbf{a}_*\|_2^2 \rho_{L,1}^2 \|\mathbf{h}_o\|_2^2 \|\mathbf{h}_{L-1}^2\|_2}{4\tau^2} + D_\psi(\mathbf{a}, \mathbf{a}_0) - \lambda \|\mathbf{a} - \mathbf{a}_*\|_2^2 \\ & \xrightarrow{\mathbb{P}} d_L \mathbb{E} \mathcal{M}_{\psi;c} \left(a_* - \nabla \psi(a_0) - cz \right) + cd_L \mathbb{E} z^2 + d_L \mathbb{E} (\psi(a_0) - a_0 \nabla \psi(a_0)) \end{aligned}$$

Where

$$\begin{aligned} \mathcal{M}_{\psi;c}(\cdot) &:= \min_x \frac{1}{2c} (\cdot - x)^2 + \psi(x) \\ c &:= \frac{\rho_{L,2}^2 \beta \|\mathbf{h}_o\|_2^2}{\tau} + \frac{\tilde{\eta}_L}{\tilde{\alpha}_L} \frac{\beta^2 \rho_{L,1}^2 \|\mathbf{h}_o\|_2^2 \|\mathbf{h}_{L-1}^2\|_2}{4\tau^2} - \lambda \\ z &\sim \mathcal{N} \left(0, \frac{\tilde{\eta}_L^2}{16\tilde{\alpha}_L^2} \frac{\beta^2 \rho_{L,1}^2 \|\mathbf{h}_o\|_2^2}{\tau^2} \tilde{\eta}_L^2 + \rho_{L,2}^2 \beta^2 \right) \end{aligned}$$

Hence the final scalar optimization would be

$$\begin{aligned} & \max_{\beta > 0} \min_{\tau > 0} \frac{\beta \tau}{2} \left(1 - \frac{d_{L-1}}{n} \right) + \max_{\eta_L, \tilde{\eta}_L} \min_{\alpha_L > 0, \tilde{\alpha}_L > 0} \frac{\tilde{\eta}_L \tilde{\alpha}_L}{2} + \frac{\alpha_L \eta_L}{2} - \frac{\beta \eta_L^2}{8\tau} + \frac{\eta_L}{2\alpha_L} \frac{\beta^2 \rho_{L-1,2}^2 d_{L-2}}{4n} \\ & + \max_{\lambda} \min_{\xi} \lambda \xi^2 + F \left(\xi^2 \frac{\beta^2}{\tau^2}, \frac{\eta_L}{\alpha_L}, \frac{\tilde{\eta}_L}{\tilde{\alpha}_L}, \beta \right) + d_L \mathbb{E} \mathcal{M}_{\psi;c} \left(a_* - \nabla \psi(a_0) - cz \right) \\ & + d_L \left(\frac{\rho_{L,2}^2 \beta n}{\tau} + \frac{\tilde{\eta}_L}{\tilde{\alpha}_L} \frac{\beta^2 \rho_{L,1}^2 n d_{L-1}}{4\tau^2} - \lambda \right) \cdot \left(\frac{\tilde{\eta}_L^2}{16\tilde{\alpha}_L^2} \frac{\beta^2 \rho_{L,1}^2 n}{\tau^2} \tilde{\eta}_L^2 + \rho_{L,2}^2 \beta^2 \right) \end{aligned} \quad (38)$$

Where F is defined in (37).