

NONPARAMETRIC LEARNING OF TWO-LAYER RELU RESIDUAL UNITS

Anonymous authors

Paper under double-blind review

ABSTRACT


We describe an algorithm that learns two-layer residual units using rectified linear unit (ReLU) activation: suppose the input \mathbf{x} is from a distribution with support space \mathbb{R}^d and the ground-truth generative model is a residual unit of this type, given by $\mathbf{y} = \mathbf{B}^* \left[(\mathbf{A}^* \mathbf{x})^+ + \mathbf{x} \right]$, where ground-truth network parameters $\mathbf{A}^* \in \mathbb{R}^{d \times d}$ represent a nonnegative full-rank matrix and $\mathbf{B}^* \in \mathbb{R}^{m \times d}$ is full-rank with $m \geq d$ and for $\mathbf{c} \in \mathbb{R}^d$, $[\mathbf{c}^+]_i = \max\{0, c_i\}$. We design layer-wise objectives as functionals whose analytic minimizers express the exact ground-truth network in terms of its parameters and nonlinearities. Following this objective landscape, learning residual units from finite samples can be formulated using convex optimization of a nonparametric function: for each layer, we first formulate the corresponding empirical risk minimization (ERM) as a positive semi-definite quadratic program (QP), then we show the solution space of the QP can be equivalently determined by a set of linear inequalities, which can then be efficiently solved by linear programming (LP). We further prove the statistical strong consistency of our algorithm, and demonstrate its robustness and sample efficiency through experimental results.

1 INTRODUCTION

Neural networks have achieved remarkable success in various fields such as computer vision (LeCun et al., 1998; Krizhevsky et al., 2012; He et al., 2016a) and natural language processing (Kim, 2014; Sutskever et al., 2014). This success is largely due to the strong expressive power of neural networks (Bengio & Delalleau, 2011), where nonlinear activation units, such as rectified linear units (ReLU) (Nair & Hinton, 2010) and hyperbolic tangents (tanh) play a vital role to ensure the large learning capacity of the networks (Maas et al., 2013). However, the nonlinearity of neural networks makes them significantly more difficult to train than linear models (Livni et al., 2014). Therefore, with the development of neural network applications, finding efficient algorithms with provable properties to train such nontrivial neural networks has become an important and a relatively new goal.

Residual networks, or ResNets (He et al., 2016a), are a class of deep neural networks that adopt skip connections to feed values between nonadjacent layers, where skipped layers may contain nonlinearities in between. Without loss of expressivity, ResNets avoid the vanishing gradient problem by directly passing gradient information from previous layers to current layers where otherwise gradients might vanish without skipping. In practice, ResNets have shown strong learning efficiency in several tasks, e.g. achieving at least 93% test accuracy on CIFAR-10 classification, lowering single-crop error to 20.1% on the 1000-class ImageNet dataset (Russakovsky et al., 2015; He et al., 2016b).

Common ResNets are often aggregated by many repeated shallow networks, where each network acts as a minimal unit with this kind of skip propagation, named a residual unit (He et al., 2016b). Given the flexibility and simplicity of residual units, much theoretical work has been devoted to study them and develop training algorithms for them in a way that sidesteps from the standard backpropagation regime and provides guarantees on the quality of estimation (see subsection 1.1). In this paper, we propose algorithms that can learn a general class of single-skip two-layer residual units with ReLU activation as shown on the right by the equation: $\mathbf{y} = \mathbf{B} \left[(\mathbf{A}\mathbf{x})^+ + \mathbf{x} \right]$, where for a scalar c , $c^+ = \max\{0, c\}$ (for a vector, this maximization is applied coordinate-wise), \mathbf{x}



is a random vector as the network input with support space \mathbb{R}^d , and $\mathbf{A} \in \mathbb{R}^{d \times d}$ and $\mathbf{B} \in \mathbb{R}^{m \times d}$ are weight matrices of *layer 1* and *layer 2*, respectively.

Compared to previous work (Ge et al., 2018; 2017; Zhang et al., 2018; Wu et al., 2019; Tian, 2017; Du et al., 2017; Brutzkus & Globerson, 2017; Soltanolkotabi, 2017; Li & Yuan, 2017; Zhong et al., 2017), the introduction of residual connections simplifies the recovery of the network parameters by removing the permutation and scaling invariance. However, the naive mean square error minimization used in estimating the parameters remains nonconvex. Unlike most previous work, we do not assume a specific input distribution nor that the distribution is symmetric (Ge et al., 2018; Du & Goel, 2018).

We show that under regularity conditions on the weights of the residual unit, the problem of learning the unit can be formulated through quadratic programming (QP). We use nonparametric estimation (Guntuboyina et al., 2018) to estimate the ReLU function values in the networks. We further rewrite our constructed quadratic programs to linear programs (LPs). The LP formulation is simpler to optimize and has the same solution space as the QP for the network parameters.

1.1 RELATED WORK

Provable learning of neural networks has become an important topic of research. Arora et al. (2014) recover a multi-layer generative network with sparse connections and Livni et al. (2014) study the learning of multi-layer neural networks with polynomial activation. Goel et al. (2018) learns a one-layer convolution network with a perceptron-like rule. They prove the correctness of an iterative algorithm for exact recovery of the target network. Rabusseau et al. (2019) describe a spectral algorithm for two-layer linear networks. Recent work has connected optimization and two-layer network learning (Ergen & Pilanci, 2021; Sahiner et al., 2020; Ergen & Pilanci, 2020; Pilanci & Ergen, 2020) and has showed how to optimize networks layer by layer (Belilovsky et al., 2019).

For learning a one-layer ReLU network, Wu et al. (2019) optimize the norm and direction of neural network weight vectors separately, and Zhang et al. (2018) use gradient descent with a specific initialization. For learning a two-layer ReLU network, Ge et al. (2017) redesign the optimization landscape such that it is more amenable to theoretical analysis and Ge et al. (2018) use a moment-based method for estimating a neural network Janzamin et al. (2015). Many others have studied ReLU networks in various settings (Tian, 2017; Du et al., 2017; Brutzkus & Globerson, 2017; Soltanolkotabi, 2017; Li & Yuan, 2017; Zhong et al., 2017; Goel & Klivans, 2017).

The study of ReLU networks with two hidden layers has also been gaining attention (Goel & Klivans, 2017; Allen-Zhu et al., 2019), with a focus on PAC learning (Valiant, 1984). In relation to our work, Allen-Zhu & Li (2019) examined the PAC-learnable function of a specific three-layer neural network with residual connections. Their work differs from ours in two aspects. First, their learnable functions include a smaller (in comparison to the student network) three-layer residual network. Second, the assumptions they make on their three-layer model are rather different than ours.

In relation to nonparametric estimation, Guntuboyina et al. (2018) treat the final output of a shape-restricted regressor as a parameter, placing some restrictions on the type of function that can be estimated (such as convexity). They provide solutions for estimation with isotonic regression (Brunk, 1955; Ayer et al., 1955; van Eeden, 1956), convex regression (Seijo et al., 2011), shape-restricted additive models (Meyer, 2013; Chen & Samworth, 2016) and shape-restricted single index models (Kakade et al., 2011; Kuchibhotla et al., 2017).

1.2 MAIN RESULTS

We design quadratic objective functionals with a linear bounded feasible domain that take network parameters and functions as variables to estimate the ground-truth network parameters and nonlinearities. The values of the objectives are moments over the input distribution. Thm. 1.1 summarizes the landscapes of the objectives.

Theorem 1.1 (objective landscape, informal). *Suppose a ground-truth residual unit has nondegenerate weights in both layers and nonnegative weights in layer 1. Then there exist quadratic functionals defined in linear-constrained domains whose minimizers a) are unique, and are the exact ground-truth, or b) are not unique, but can be adjusted to the exact ground-truth.*

In practice, the exact moments over unobserved variables are not available. We can only construct the empirical risk minimization (ERM) of the moment-valued objectives by generated samples. With functions as variables in moment-valued objectives being optimized nonparametrically, the empirical objectives become quadratic functions with linear constraints (QP). We further show the convexity of our QP which guarantees its solution in polynomial time w.r.t. sample size and dimension. With the solution to the QP available, the strong consistency of our network learner is guaranteed:

Theorem 1.2 (strong consistency, informal). *Suppose samples are generated by a ground-truth residual unit that has nondegenerate weights in both layers and nonnegative weights in layer 1. Then there exists an algorithm that learns a network $\xrightarrow{a.s.}$ the exact ground-truth as sample size grows.*

Roadmap: Assume \mathbf{A} and \mathbf{B} are student network weights of layer 1 and 2. We first give a warm-up vanilla linear regression approach only knowing \mathbf{A} is in $\mathbb{R}^{d \times d}$ and \mathbf{B} is in $\mathbb{R}^{m \times d}$ (section 3). We then move to the details of our nonparametric learning for layer 2 (section 4), and similarly, layer 1 (section 5). We formalize Thm. 1.1, showing how nonparametric learning allows us to select the values of \mathbf{A} and \mathbf{B} from reduced spaces that are seeded by \mathbf{A}^* and \mathbf{B}^* , under which LR is faster than vanilla LR. In section 6, we describe the strong consistency of our methods for respective layers, formalizing Thm. 1.2 using the continuous mapping theorem (Mann & Wald, 1943).

2 PRELIMINARIES

We describe the notations used in this paper, introduce the model and its underlying assumptions, and state conditions which simplify the problem but which can be removed without loss of learnability.

Notation The ReLU residual units we use take a vector $\mathbf{x} \in \mathbb{R}^d$ as input and return a vector $\mathbf{y} \in \mathbb{R}^m$. We use $\mathbf{A}^* \in \mathbb{R}^{d \times d}$ and $\mathbf{B}^* \in \mathbb{R}^{m \times d}$ to denote the ground-truth network parameters for layer 1 and 2, respectively. We use circumflex to denote predicted terms (e.g. an empirical objective function \hat{f} , estimated layer 1 weights $\hat{\mathbf{A}}$). We use n to denote the number of i.i.d. samples available for the training algorithm, $\{\mathbf{x}^{(i)}, \mathbf{y}^{(i)}\}$ to denote the i -th sample drawn. For an integer k , We define $[k]$ to be $\{1, 2, \dots, k\}$, and $\mathbf{e}^{(j)}$ as the standard basis vector with a 1 at position j . All scalar-based operators are element wise in the case of vectors or matrices unless specified otherwise. We use \mathbf{x} and \mathbf{y} to refer to the input and output vectors, respectively, as random variables.

Linear Regression: Linear regression (LR) in this paper refers to the problem of estimating \mathbf{L} for unbiased model $\mathbf{y} = \mathbf{L}\mathbf{x}$. In this case, estimation is done by minimizing the empirical risk $\hat{R}(\mathbf{L}) = \frac{1}{2n} \sum_{i \in [n]} \|\mathbf{L}\mathbf{x}^{(i)} - \mathbf{y}^{(i)}\|^2$ – using linear least squares (LLS). Its solution has a closed form which we use as given. We refer the reader to Hamilton (1994) for more information.

Models, Assumptions and Conditions Following previous work (Livni et al., 2014), we assume a given neural network structure specifies a hypothesis class that contains all the networks conforming to this structure. Learning such a class means using training samples to find a set of weights such that the neural networks predictions generalize well to unseen samples, where both the training and the unseen samples are drawn from an unobserved ground-truth distribution. In this paper, the hypothesis class is given by ReLU residual units and we assume it has sufficient expressive power to fit the ground-truth model. More specifically, we discuss the *realizable* case of learning, in which the ground-truth model is set to be a residual unit taken from the hypothesis class with the form:

$$\mathbf{y} = \mathbf{B}^* \left[(\mathbf{A}^* \mathbf{x})^+ + \mathbf{x} \right], \quad (1)$$

and is used to draw samples for learning. Unlike other multi-layer ReLU-affine models which do not apply skip connections, we cannot permute the weight matrices of the residual unit and retain the same function because of skip-adding \mathbf{x} (it breaks symmetry). This helps us circumvent issues of identifiability¹, and allows us to precisely estimate the ground-truth weight matrices \mathbf{A}^* and \mathbf{B}^* .

Our general approach for residual unit layer 2 learns a scaled ground-truth weight matrix that also minimizes the layer 2 objective. The existence of such scaled equivalence of our layer 2 approach comes from what is defined below.

¹Here, we are referring to the ability to identify the true model parameters using infinite samples.

Definition 2.1 (component-wise scale transformation). A matrix $\mathbf{A} \in \mathbb{R}^{d \times d}$ is said to be a scale transformation w.r.t. the j -th component if $(\mathbf{A}_{j,:})^\top = A_{j,j} \cdot \mathbf{e}^{(j)}$.

Additionally, estimation is more complex when the layer 2 weights \mathbf{B}^* is nonsquare. For simplicity of our algorithm presentation for layer 2, we stick to Cond. 2.1 in the following sections².

Condition 2.1 (layer 2 objective minimizer unique). \mathbf{A}^* is not a scale transformation w.r.t. any components and \mathbf{B}^* is a square matrix, i.e. $m = d$.

3 WARM-UP: VANILLA LINEAR REGRESSION

Consider a ground-truth two-layer residual unit. If we assume that the inputs only contain vectors with negative entries, i.e. $\mathbf{x} < 0$, the effect of the ReLU function in the residual unit then disappears because of the nonnegativity of \mathbf{A}^* . The residual unit turns into linear model $\mathbf{y} = \mathbf{B}^* \mathbf{x}$. Thus, direct LR on samples with negative inputs can learn the exact ground-truth layer 2 parameter \mathbf{B}^* with at least d samples, when formulating the LR as a solvable full-rank linear equation system.

On the contrary, if the inputs only contain vectors with positive entries, i.e. $\mathbf{x} > 0$, all the neurons in the hidden layer are then activated by the ReLU and the nonlinearity is eliminated. The residual unit in this case turns into $\mathbf{y} = \mathbf{B}^* (\mathbf{A}^* + \mathbf{I}_d) \mathbf{x}$. Taking the value that left-multiplies \mathbf{x} as a single weight matrix \mathbf{D}^* , this is also a linear model. Direct LR on at least d samples with positive inputs by the residual unit can learn the exact \mathbf{D}^* . Since we have the access to the exact \mathbf{B}^* , solving for the exact \mathbf{A}^* can be accomplished through solving a full-rank linear equation system $\mathbf{B}^* \cdot \tilde{\mathbf{A}} = \mathbf{D}^*$, where the unique solution $\tilde{\mathbf{A}} = \mathbf{A}^* + \mathbf{I}_d$.

While simple, this vanilla LR approach requires a large number of redundant samples, since sampled inputs usually have a small proportion of fully negative/positive vectors. Taking random input vectors i.i.d. with respect to each entry as an example, the probability of sampling a vector with all negative entries is p_-^d , where p_- is the probability of sampling a negative vector entry, then the expected number of samples to get one negative vector is $1/p_-^d$. Denoting p_+ similarly, $1/p_+^d$ samples are expected for a positive vector. Besides, each LR step in this approach requires d such samples respectively to make the linear equation system full-rank, which implies the expected sample size to be d -exponential $d \cdot (1/p_-^d + 1/p_+^d)$. For other common random vectors like Gaussian samples, the proportions of fully negative/positive vectors in sampled inputs are also expected to decrease exponentially as d grows, as high-dimensional random vectors like Gaussian are essentially concentrated uniformly in a sphere (Johnstone, 2006). Technical and experimental details about sample size expectations and the vanilla LR algorithm are further discussed in Appendix F.

4 NONPARAMETRIC LEARNING: LAYER 2

We present how we learn a residual unit layer 2 under Cond. 2.1 (estimating \mathbf{B}^*): We first design an objective functional with the arguments being a matrix and a function. The objective uses expectation of a loss over the true distribution generating the data and is uniquely minimized by $[\mathbf{B}^*]^{-1}$ and a rectifier function (ReLU). We then formulate its ERM using nonparametric estimation as a standard convex QP, further simplified as an LP that has the same capability as the QP to learn layer 2.

4.1 OBJECTIVE DESIGN AND LANDSCAPE

Consider the formulation of a residual unit as in Eq. 1. It is possible to rewrite the model as equation: $\mathbf{C}^* \mathbf{y} = (\mathbf{A}^* \mathbf{x})^+ + \mathbf{x}$, where the output of the hidden neuron with skip addition is on both sides of the equation, and $\mathbf{C}^* \mathbf{B}^* = \mathbf{I}_d$. We aim to estimate the inverse of \mathbf{B}^* by matrix variable \mathbf{C} and the nonlinearity $\mathbf{x} \mapsto (\mathbf{A}^* \mathbf{x})^+$ by a function variable h . The objective is formulated as risk functional by the L^2 error between values respectively computed by \mathbf{C} and h

$$G_2(\mathbf{C}, h) = \frac{1}{2} \mathbb{E}_{\mathbf{x}} \left[\|h(\mathbf{x}) + \mathbf{x} - \mathbf{C} \mathbf{y}\|^2 \right], \quad (2)$$

²In Appendix G, we show that estimation of \mathbf{B}^* remains solvable without satisfying Cond. 2.1.

where the estimator $\mathbf{C} \in \mathbb{R}^{d \times m}$, the domain of h is the nonnegative³ continuous⁴ $\mathbb{R}^d \rightarrow \mathbb{R}^d$ function space, written as $\mathbb{C}_{\geq 0}^0$ in shorthand. This objective is quadratic because the forward mapping $\mathbf{x} \mapsto h(\mathbf{x}) + \mathbf{x}$ and the backward mapping $\mathbf{y} \mapsto \mathbf{C}\mathbf{y}$ are both linear w.r.t. \mathbf{C} and h , and the two are linearly combined in a L^2 norm. The objective in Eq. 2 is minimized by the ground-truth, i.e. \mathbf{C}^* and $\mathbf{x} \mapsto (\mathbf{A}^* \mathbf{x})^+$, is one of its minimizers. However, it is not simple to describe other variable values that minimize the objective if any exist. We give that detail under Cond. 2.1, the minimizer of G_2 is unique, as the exact ground-truth in the given domain.

Theorem 4.1 (objective minimizer, layer 2). *Define $G_2(\mathbf{C}, h)$ as Eq. 2, where $\mathbf{C} \in \mathbb{R}^{d \times m}$, $h \in \mathbb{C}_{\geq 0}^0$. Then under Cond. 2.1, $G_2(\mathbf{C}, h)$ reaches its zero minimum iff $\mathbf{C} = [\mathbf{B}^*]^{-1}$ and $h : \mathbf{x} \mapsto (\mathbf{A}^* \mathbf{x})^+$.*

Technical details are in Appendix H, where we use Lem. 4.2 (in subsection 4.3) to prove a more general theorem which does not require Cond. 2.1 and is sufficient for Thm. 4.1. In the next subsection, we construct the ERM of G_2 and present our convex QP formulation.

4.2 ERM WITH NONPARAMETRIC ESTIMATION IS CONVEX QP

Consider the second layer objective (Eq. 2) with nonnegative continuous function space as the domain of h . We follow Vapnik (1992) and define its standard empirical risk functional:

$$\hat{G}_2(\mathbf{C}, h) = \frac{1}{2n} \sum_{i \in [n]} \left\| h(\mathbf{x}^{(i)}) + \mathbf{x}^{(i)} - \mathbf{C}\mathbf{y}^{(i)} \right\|^2. \quad (3)$$

The function variable $h \in \mathbb{C}_{\geq 0}^0$ can be optimized either parametrically or nonparametrically. If we parameterize h , and show that the nonlinearity w.r.t. its parameters would make \hat{G}_2 lose its quadratic form. Instead, we estimate h nonparametrically: for each sample input $\mathbf{x}^{(i)}$, we introduce variables $\boldsymbol{\xi}^{(i)}$ that estimates mapped values by h . This avoids introducing nonlinearity to the objective and keeps \hat{G}_2 quadratic. On the other hand, the domain of h , i.e. nonnegative continuous function space, turns into a set of linear inequalities as constraints when optimizing nonparametrically. In this sense, learning the second layer of the residual unit can be formulated as the following QP:

$$\min_{\mathbf{C}, \boldsymbol{\Xi}} \hat{G}_2^{\text{NPE}}(\mathbf{C}, \boldsymbol{\Xi}) := \frac{1}{2n} \sum_{i \in [n]} \left\| \boldsymbol{\xi}^{(i)} + \mathbf{x}^{(i)} - \mathbf{C}\mathbf{y}^{(i)} \right\|^2, \text{ s.t. } \boldsymbol{\xi}^{(i)} \geq \mathbf{0}, \forall i \in [n]. \quad (4)$$

where $\boldsymbol{\Xi} = \{\boldsymbol{\xi}^{(i)}\}_{i=1}^n$ is the nonparametric estimator of $\mathbf{x} \mapsto (\mathbf{A}^* \mathbf{x})^+$.

Nonparametric Estimation Validation: A solution to the ERM with nonparametric estimation, i.e. the QP, is guaranteed to be sufficient for minimizing the standard empirical risk functional (Eq. 3). More specifically, assuming \mathbf{C} and $\boldsymbol{\Xi} = \{\boldsymbol{\xi}^{(i)}\}_{i=1}^n$ are a solution to layer 2 QP (Eq. 4), it is clear that \mathbf{C} and $h \in \mathbb{C}_{\geq 0}^0$ such that $h(\mathbf{x}^{(i)}) = \boldsymbol{\xi}^{(i)}$ minimize the empirical risk functional Eq. 3. Conversely, a minimizer of the standard empirical risk Eq. 3, \mathbf{C} and h , corresponds to a solution to layer 2 QP as we set $\boldsymbol{\xi}^{(i)} = h(\mathbf{x}^{(i)})$. Therefore, minimizing \hat{G}_2 and solving layer 2 QP are empirically equivalent.

Convexity: The convexity of the QP: Eq. 4 is also guaranteed. First, constraints are linear. Second, for each sample with index $i \in [n]$, the L^2 norm wraps linearity w.r.t. \mathbf{C} and $\boldsymbol{\xi}^{(i)}$. Such formulation ensures the quadratic coefficient matrix is positive semidefinite. Thus, with the sum of convex functions still being convex, the QP objective (Eq. 4) becomes convex. Even without the knowledge of how samples are generated, this QP would be a convex program. Strict proofs are in Appendix I.

4.3 LP SIMPLIFICATION

Consider single-sample error written as $g_2(\mathbf{C}, h; \mathbf{x}, \mathbf{y}) = \frac{1}{2} \|h(\mathbf{x}) + \mathbf{x} - \mathbf{C}\mathbf{y}\|^2$. If there is a feasible \mathbf{C} such that $\mathbf{C}\mathbf{y} - \mathbf{x} \geq 0$ holds for all $\mathbf{x} \in \mathbb{R}^d$, then \mathbf{C} and $h : \mathbf{x} \mapsto \mathbf{C}\mathbf{y} - \mathbf{x}$ always minimize g_2 as zero, and thereby minimize the layer 2 objective (Eq. 2). Thus, we obtain a condition that is equivalent to G_2 reaching the minimum in Thm. 4.1 and avoids randomness (see Lem. 4.2).

³Setting h as nonnegative ensures that a) only ReLU nonlinearity minimizes G_2 (see Thm. 4.1), and b) h 's nonparametric estimator is linearly constrained (explained in subsection 4.2).

⁴If h is not a continuous function, only a null set of discontinuities is possible to make G_2 reach zero as its minimum. Setting h as continuous simplifies our theoretical results that still strictly support empirical discussion.

Algorithm 1 Learn a ReLU residual unit, layer 2.

- 1: **Input:** $\{(\mathbf{x}^{(i)}, \mathbf{y}^{(i)})\}_{i=1}^n$, samples drawn by Eq. 1.
- 2: **Output:** $\hat{\mathbf{B}}, \hat{\Xi}$, a layer 2 and $\mathbf{x} \mapsto (\mathbf{A}^* \mathbf{x})^+$ estimate.
- 3: Go to line 4 if QP, line 5 if LP.
- 4: Solve QP: Eq. 4 and obtain a \hat{G}_2^{NPE} minimizer, denoted by $\hat{\mathbf{C}}, \hat{\Xi}$. Go to line 6.
- 5: Solve LP: Eq. 5 and obtain a minimizer $\hat{\mathbf{C}}$, then assign $\hat{\xi}^{(i)} \leftarrow \hat{\mathbf{C}} \mathbf{y}^{(i)} - \mathbf{x}^{(i)}$.
- 6: **return** $\hat{\mathbf{C}}^{-1}, \hat{\Xi}$.

Algorithm 2 Learn a ReLU residual unit, layer 1.

- 1: **Input:** $\{(\mathbf{x}^{(i)}, \mathbf{h}^{(i)})\}_{i=1}^n$, layer 1 samples.
- 2: **Output:** $\hat{\mathbf{A}}$, a layer 1 estimate.
- 3: Solve QP: Eq. 7 or LP: Eq. 8 and obtain a \hat{G}_1^{NPE} minimizer, denoted by $\hat{\mathbf{A}}, \hat{\Phi}$. $\{\hat{\Phi}$ is no longer needed. $\}$
- 4: **for all** $j \in [d]$ **do**
- 5: $\hat{k}_j \leftarrow \text{LR}\{h_j^{(i)}, \hat{\mathbf{A}}_{j,:} \mathbf{x}^{(i)}\}_{h_j^{(i)} > 0}$.
- 6: $\hat{\mathbf{A}}_{j,:} \leftarrow \hat{\mathbf{A}}_{j,:} / \hat{k}_j$. $\{\text{Rescale } \hat{\mathbf{A}}.\}$
- 7: **end for**
- 8: **return** $\hat{\mathbf{A}}$.

Lemma 4.2. $G_2(\mathbf{C}, h)$ reaches its zero minimum iff $\mathbf{C}\mathbf{y} - \mathbf{x} \geq 0$ holds for any $\mathbf{x} \in \mathbb{R}^d$ and its corresponding residual unit output \mathbf{y} , and $h : \mathbf{x} \mapsto \mathbf{C}\mathbf{y} - \mathbf{x}$.

The pointwise satisfaction of the inequality in Lem. 4.2 describes the solution space for the minimization of G_2 . The sufficiency of satisfying this inequality to minimize G_2 is directly obtained by assigning $h : \mathbf{x} \mapsto \mathbf{C}\mathbf{y} - \mathbf{x}$ where \mathbf{C} complies with $\mathbf{C}\mathbf{y} - \mathbf{x} \geq 0$ for all \mathbf{x} in the support space \mathbb{R}^d . The necessity of satisfying this inequality comes from its contraposition: If a \mathbf{C} violates the inequality, there must be a non-null set of \mathbf{x} that yield the violation due to the continuity of $\mathbf{C}\mathbf{y} - \mathbf{x}$. In this sense, the resulting G_2 value becomes nonzero.

Empirically speaking, we cannot solve an inequality that holds w.r.t. to every point in the support space if we only observe finite samples. We can only estimate \mathbf{C} by solving the inequality that holds w.r.t. each sample. Following this, we formulate such estimation as to find a feasible point in the space defined by a set of linear inequalities, each of which corresponds to a sample: $\mathbf{C}\mathbf{y}^{(i)} - \mathbf{x}^{(i)} \geq 0$. Each point in the feasibility defined by the inequalities has a one-to-one correspondence in the solution space to layer 2 QP (Eq. 4): $\mathbf{C} \leftrightarrow (\mathbf{C}, \{\mathbf{C}\mathbf{y}^{(i)} - \mathbf{x}^{(i)}\}_{i \in [n]})$. The set of inequalities can be solved by a standard LP with a constant objective and with constraints that are inequalities, i.e.

$$\min_{\mathbf{C}} \text{const, s.t. } \mathbf{C}\mathbf{y}^{(i)} - \mathbf{x}^{(i)} \geq 0, \forall i \in [n]. \quad (5)$$

With the one-to-one correspondences, LP: Eq. 5 and QP: Eq. 4 have equivalent solution spaces. Moreover, LP: Eq. 5 is also a convex program. Alg. 1 summarizes the layer 2 estimator: Simply solve the QP/LP⁵ and return the inverse of $\hat{\mathbf{C}}$ as the estimate of layer 2 weights and $\hat{\Xi}$ as $\mathbf{x} \mapsto (\mathbf{A}^* \mathbf{x})^+$ estimate. Regardless of time complexity, QP and LP in Alg. 1 work equivalently since their solution spaces are equivalent to each other.

Our nonparametric learning directly finds a unique layer 2 estimate under Cond. 2.1. In the general case without Cond. 2.1 (discussed in Appendix G), nonparametric learning essentially reduces the possible values of \mathbf{B} from $\mathbb{R}^{m \times d}$ to a \mathbf{B}^* scale-equivalent matrix space, where LR uses sampled data much more efficiently than vanilla LR on layer 2 (section 3).

5 NONPARAMETRIC LEARNING: LAYER 1

With layer 2 learned, outputs by hidden neurons become observable. The two-layer problem is thereby reduced to single-layer. Consider a ground-truth single-layer model: $\mathbf{h} = (\mathbf{A}^* \mathbf{x})^+$. To construct a learning objective for this model, we rewrite the model as a nonlinearity plus a linear mapping by $\mathbf{A}^* : \mathbf{h} = (-\mathbf{A}^* \mathbf{x})^+ + \mathbf{A}^* \mathbf{x}$, where on both sides of the equation is the output of layer 1, and the nonlinearity is $\mathbf{x} \mapsto (-\mathbf{A}^* \mathbf{x})^+$. The objective of layer 1 is formulated as:

$$G_1(\mathbf{A}, r) = \frac{1}{2} \mathbb{E}_{\mathbf{x}} \left[\|r(\mathbf{x}) + \mathbf{A}\mathbf{x} - \mathbf{h}\|^2 \right], \quad (6)$$

⁵We use CVX (Grant & Boyd, 2014; 2008) that calls SDPT3 (Toh et al., 1999) (a free solver under GPLv3 license) and solves our convex QP/LP in polynomial time. See Appendix C for technical details.

where $\mathbf{A} \in \mathbb{R}^{d \times d}$ is of the layer 1 weights estimator, the domain of r is also $\mathbb{C}_{\geq 0}^0$. The minimizer \mathbf{A} of the risk G_1 falls into a matrix space such that for any matrix \mathbf{A} in the space, each row of \mathbf{A} is a scaled-down version of the same row of \mathbf{A}^* without changing the direction (Thm. 5.1).

Theorem 5.1 (objective minimizer space, layer 1). *Define $G_1(\mathbf{A}, r)$ as Eq. 6, where $\mathbf{A} \in \mathbb{R}^{d \times d}$, $r \in \mathbb{C}_{\geq 0}^0$. Then $G_1(\mathbf{A}, r)$ reaches its zero minimum iff for each $j \in [d]$, $\mathbf{A}_{j,:} = k_j \mathbf{A}_{j,:}^*$ where $0 \leq k_j \leq 1$, and $r : \mathbf{x} \mapsto (\mathbf{A}^* \mathbf{x})^+ - \text{diag}(\mathbf{k}) \cdot \mathbf{A}^* \mathbf{x}$.*

The scale equivalence in the solution space is derived from a ReLU inequality: $(x)^+ \geq kx$ where $0 \leq k \leq 1$. Let the j -th row of \mathbf{A} be a scaled-down version of \mathbf{A}^* , i.e. $\mathbf{A}_{j,:} = k_j \mathbf{A}_{j,:}^*$ where $0 \leq k_j \leq 1$. According to the inequality, we have $(\mathbf{A}_{j,:}^* \mathbf{x})^+ \geq k_j \mathbf{A}_{j,:}^* \mathbf{x}$, which indicates $(\mathbf{A}^* \mathbf{x})^+ \geq \text{diag}(\mathbf{k}) \cdot \mathbf{A}^* \mathbf{x}$. Thus, r minimizing G_1 in Thm. 5.1 lies in feasibility $\mathbb{C}_{\geq 0}^0$ when $\mathbf{A} = \text{diag}(\mathbf{k}) \cdot \mathbf{A}^*$.

Due to the existence of scale equivalence, we must compute the scale factor to obtain the ground-truth weights \mathbf{A}^* . The scale factor k_j is sufficiently obtainable with $(\mathbf{A}_{j,:} \mathbf{x})^+$ and $(\mathbf{A}_{j,:}^* \mathbf{x})^+$ observable: Conditioned on nonnegative ReLU input, we have a linear model $\mathbf{A}_{j,:} \mathbf{x} = k_j \mathbf{A}_{j,:}^* \mathbf{x}$ where $\mathbf{A}_{j,:} \mathbf{x}$ and $\mathbf{A}_{j,:}^* \mathbf{x}$ are observed. Thm. 5.2 summarizes layer 1 scale factor property, allowing us to correct a minimizer of G_1 to the ground-truth weights \mathbf{A}^* by computing a scalar for each row.

Theorem 5.2 (scale factor, layer 1). *Assume \mathbf{A} is a minimizer of G_1 . Then for any $j \in [d]$, $(\mathbf{A}_{j,:} \mathbf{x})^+ / (\mathbf{A}_{j,:}^* \mathbf{x})^+$ is always equal to the scale factor k_j given that $(\mathbf{A}_{j,:} \mathbf{x})^+ > 0$.*

The elimination of randomness for G_1 follows the same pattern as layer 2. If there is a feasible \mathbf{A} such that $(\mathbf{A}^* \mathbf{x})^+ - \mathbf{A} \mathbf{x} \geq 0$ holds for all $\mathbf{x} \in \mathbb{R}^d$, such \mathbf{A} is a solution to our objective by Eq. 6. We can also have the following proposition that avoids randomness.

Lemma 5.3. $G_1(\mathbf{A}, r)$ reaches its zero minimum iff $\mathbf{h} - \mathbf{A} \mathbf{x} \geq 0$ holds for any $\mathbf{x} \in \mathbb{R}^d$ and its corresponding hidden output \mathbf{h} , and $r : \mathbf{x} \mapsto \mathbf{h} - \mathbf{A} \mathbf{x}$.

We now turn into empirical discussion. Similar to layer 2, we formulate layer 1 QP by its empirical objective with nonparametric estimation and linear constraints representing the nonnegativity of r :

$$\min_{\mathbf{A}, \Phi} \hat{G}_1^{\text{NPE}}(\mathbf{A}, \Phi) := \frac{1}{2n} \sum_{i \in [n]} \left\| \phi^{(i)} + \mathbf{A} \mathbf{x}^{(i)} - \mathbf{h}^{(i)} \right\|^2, \text{ s.t. } \phi^{(i)} \geq \mathbf{0}, \forall i \in [n]. \quad (7)$$

where $\Phi = \{\phi^{(i)}\}_{i=1}^n$ is the function estimator of $\mathbf{x} \mapsto (\mathbf{A}^* \mathbf{x})^+ - \text{diag}(\mathbf{k}) \cdot \mathbf{A}^* \mathbf{x}$ where \mathbf{k} refers to the scaling equivalence. Similarly, the solution space of QP: Eq. 7 can be represented by a set of linear inequalities as constraints of the following efficiently solvable LP

$$\min_{\mathbf{A}} \text{const, s.t. } \mathbf{h}^{(i)} - \mathbf{A} \mathbf{x}^{(i)} \geq \mathbf{0}, \forall i \in [n]. \quad (8)$$

Alg. 2 describes how a layer 1 is learned: First, a G_1 minimizer estimate $\hat{\mathbf{A}}$ is obtained by solving the QP/LP. Then for the j -th row, the scale factor k_j is estimated by running LR on $h_j^{(i)}$ and $\hat{\mathbf{A}}_{j,:} \mathbf{x}^{(i)}$ s.t. $h_j^{(i)} > 0$ to correct $\hat{\mathbf{A}}$, as $\mathbf{A}_{j,:} \mathbf{x} = k_j h_j$ is an unbiased linear model given that $h_j > 0$. By nonparametric learning, the value space of \mathbf{A} is reduced from $\mathbb{R}^{d \times d}$ to $\{\text{diag}(\mathbf{k}) \cdot \mathbf{A}^* \mid 0 \leq k_j \leq 1\}$, where LR uses sampled data much more efficiently than vanilla LR layer 1 (section 3).

6 FULL ALGORITHM AND ANALYSIS

The full algorithm concatenates Alg. 1 and 2 in a layerwise fashion, with observations of input/output by the ground-truth network: *a)* Estimates layer 2 and nonlinearity: $\mathbf{x} \mapsto (\mathbf{A}^* \mathbf{x})^+$ by Alg. 1. *b)* Estimates layer 1 by running Alg. 2 on input samples and the nonlinearity estimate. It has provable guarantees. For empirical analysis, we use $\hat{\mathbf{C}}_n$ to denote the estimation of \mathbf{C}^* from n random samples. Similar notations are applied to other estimations. First of all, our methods to solve respective layers, Alg. 1 and 2, are strongly consistent if any convex QPs/LPs involved can be solved exactly.

Lemma 6.1 (layer 2 strong consistency). *Under Cond. 2.1, $\hat{\mathbf{C}}_n \xrightarrow{a.s.} \mathbf{C}^*$ and $\hat{\mathbf{B}}_n \xrightarrow{a.s.} \mathbf{B}^*$, $n \rightarrow \infty$.*

For $\hat{\mathbf{C}}_n$: In Appendix J, we prove its more general a.s. convergence without satisfying Cond. 2.1, where the solution space to layer 2 objective is a non-compact continuous set where all the elements

are scaling equivalences. We use Hausdorff distance (Rockafellar & Wets, 2009) as metric and prove that the empirical solution space a.s. converges to the theoretical solution space. Then the more general a.s. convergence holds with the strong consistency of layer 2 scale factor estimator where we use LR to estimate the scale factors.

For $\hat{\mathbf{B}}_n$: Using the continuous mapping theorem (Mann & Wald, 1943), we directly propagate the strong consistency of \mathbf{C}^* estimator to its inverse \mathbf{B}^* 's estimator.

According to full algorithm description, layer 1 estimation uses $\hat{\mathbf{C}}_n \mathbf{y}^{(i)} - \mathbf{x}^{(i)}$ as the outputs where $i \in [n]$. Thus, the strong consistency of the hidden neuron estimator is also guaranteed by the continuous mapping theorem. Following the proof sketch of the $\hat{\mathbf{C}}_n$ a.s. convergence, we obtain the strong consistency of layer 1 estimator (See Lem. 6.2).

Lemma 6.2 (layer 1 strong consistency). $\hat{\mathbf{A}}_n \xrightarrow{a.s.} \mathbf{A}^*, n \rightarrow \infty$.

By the continuous mapping theorem, the strong consistency of the full algorithm (Thm. 6.3, formal version of Thm. 1.2), which is commonly defined by a loss function that is continuous on network weights, is implied by the strong consistency of network weights estimators (Lem. 6.1 and 6.2). See Appendix J for proofs of strong consistency discussions in this section.

Theorem 6.3 (strong consistency, formal). Define L as L^2 output loss. $L(\hat{\mathbf{A}}_n, \hat{\mathbf{B}}_n) \xrightarrow{a.s.} 0, n \rightarrow \infty$.

7 EXPERIMENTS

We provide experimental analysis to demonstrate the effectiveness and robustness of our approach in comparison to stochastic gradient descent (SGD) on L^2 output loss: $L(\mathbf{A}, \mathbf{B}) = \frac{1}{2} \mathbb{E}_{\mathbf{x}} \|\hat{\mathbf{y}} - \mathbf{y}\|^2$, where we parameterize the output prediction by $\hat{\mathbf{y}} = \mathbf{B} \left[(\mathbf{A}\mathbf{x})^+ + \mathbf{x} \right]$. Our proposed methods outperform SGD in terms of sample efficiency and robustness to different network weights and noise strengths, which indicates a poor optimization landscape of L^2 output loss for ReLU residual units.

Setup: The ground-truth weights are generated through i.i.d. folded standard Gaussian⁶ and standard Gaussian for layer 1 and 2 respectively, i.e. $\mathbf{A}^* \stackrel{i.i.d.}{\sim} |\mathcal{N}|(0, 1)$, $\mathbf{B}^* \stackrel{i.i.d.}{\sim} \mathcal{N}(0, 1)$. The input distribution is set to be an i.i.d. zero mean Gaussian-uniform equal mixture $\mathcal{N}(-0.1, 1) - \mathcal{U}(-0.9, 1.1)$. SGD is conducted on mini-batch empirical losses of $L(\mathbf{A}, \mathbf{B})$ with batch size 32 for 256 epochs in each learning trial. We apply time-based learning rate decay $\eta = \eta_0 / (1 + \gamma \cdot T)$ with initial rate $\eta_0 = 10^{-3}$ and decay rate $\gamma = 10^{-5}$, where T is the epoch number. The above hyperparameters are tuned to outperform other hyperparameters in learning ReLU residual units in terms of output errors.

Evaluation: We use relative errors to measure the accuracy of our vector/matrix estimates: For a network with weights \mathbf{A} and \mathbf{B} and its teacher network with weights \mathbf{A}^* and \mathbf{B}^* , a) layer 1 error refers to $\|\mathbf{A} - \mathbf{A}^*\| / \|\mathbf{A}^*\|$, similar to layer 2. b) output error refers to $\hat{\mathbb{E}}[\|\hat{\mathbf{y}} - \mathbf{y}\| / \|\mathbf{y}\|]$ by test data. Due to the equivalence between the solution spaces of our QP and LP without label noise, we choose LP in noiseless experiments, referred to as ‘‘ours’’. In addition, to reduce variance, the results of learning the same ground-truths are computed as means across 16 trials.

Sample Efficiency: Consider Fig. 1. It depicts the variation in the prediction errors (warmer color, larger error) as a function of d (input dimension) and the number of samples the learning algorithm is using. We compare SGD against our algorithm. We observe that our approach to the estimating the neural network is more sample efficient. For SGD, the estimation is relatively easy with only up to 10 dimensions. As expected, once the dimension grows, the sample size required for the same level of error as our method is larger. Still, overall, our method is capable of learning robustly with small sample sizes and more efficiently than SGD even for larger sample sizes.

Network Weight Robustness: This experiment aims to verify whether our method can learn a broader class of residual units. In Tab. 1, our method shows a light-tailed distribution with nearly zero means and standard deviations for layers 1 and 2, and output errors across various ground-truth networks, whereas SGD is less robust in the same context. Our method shows strong robustness to network weight changes, indicating its applicability across the whole hypothesis class.

⁶A folded Gaussian is the absolute value of a Gaussian, with p.d.f. $p(|x|)$ where $x \sim \mathcal{N}$, denoted as $|\mathcal{N}|$. We use folded Gaussian to ensure layer 1 weights are nonnegative.




Figure 1: **Output errors by SGD and our method for different dimensions and sample sizes.** For the same d , we fix the ground-truth network as sample size grows.

Table 1: **Means / Standard deviations of network estimate errors for different network weights.** Values are computed from the process of learning 128 different ground-truth networks with $d = 16$. 512 training samples are drawn for each learning trial.

	Layer 1		Layer 2		Output	
	Mean	Std	Mean	Std	Mean	Std
SGD	0.715	0.090	1.203	0.134	0.431	0.038
Ours	0.039	0.008	≈ 0	≈ 0	0.055	0.008




Figure 2: **Respective errors of layers 1 and 2 and outputs of different label noise strengths for SGD, LP and QP.** We fix the ground-truth weights with $d = 10$ and only the noise strength varies. 512 training samples are drawn for each learning trial.

Noise Robustness: Fig. 2 confirms the robustness of our methods when output noise exists. Samples are generated by a ground-truth residual unit with output noise being i.i.d. zero-mean Gaussian in different strengths (i.e. standard deviations). We try both QP and LP because in noisy setting the two approaches are not equivalent w.r.t. the solution space⁷. First, SGD always gives larger errors than our methods, even though it is hardly affected by tuning the noise strength. For QP/LP, all the errors for layer 1, 2 and output grow almost linearly as noise strength increases, indicating that both QP/LP learn the optima robustly when output noise is present, where QP slightly outperforms LP.

8 CONCLUSION

In this paper, we address the problem of learning a general class of two-layer residual units and propose an algorithm based on landscape design and convex optimization: We demonstrate firstly that minimizers of our objective functionals can express the *exact* ground-truth network. Then, we show that the corresponding ERM with nonparametric function estimation can be solved using *convex* QP/LP, which indicates *polynomial-time solvability* w.r.t. sample size and dimension. Moreover, our algorithms that are used to estimate both layers as well as the whole networks are *strongly consistent*, with very weak conditions on input distributions.

Our work opens the door to a variety of open problems to explore. We provide a strong consistency result without assuming a particular input distribution. It would be interesting to explore the sample complexity of our algorithm when stronger assumptions are placed on input distribution. Another extension could be to dispose of the limitations on ground-truth weights, e.g. $m \geq d$, non-degeneration, nonnegativity of \mathbf{A}^* . In addition, given that our algorithm solves ReLU by using nonparametric estimation through convex optimization, this might provide an inspiration to solving learning problems with other nonlinearities using nonparametric methods that improve or change the optimization landscape.

⁷See Appendix B for noisy model discussion.

REPRODUCIBILITY STATEMENT

The full set of assumptions of all theoretical results in this paper are stated - section 2, Appendix F and G. The complete proofs of all theoretical results are all included - Appendix F, G, H, I and J. Code in MATLAB is zipped as supplementary file. The numerical simulation details, including data sizes and splits, hyperparameters and how they were chosen, clear definitions of the specific measures or statistics used to report results, are discussed in detail - the **Setup** and **Evaluation** paragraphs in section 7. Results with central tendency (e.g. mean) and variation (e.g. stddev) are reported - Tab. 1. The use of external computing infrastructure, CVX, is clarified, and the consent obtaining and licensing of CVX and SDPT3 solver used in this work exactly follows the original author instructions: <http://cvxr.com/cvx/doc/citing.html> - footnote 5.

REFERENCES

- Zeyuan Allen-Zhu and Yuanzhi Li. What can ResNet learn efficiently, going beyond kernels? In *Advances in Neural Information Processing Systems*, pp. 9015–9025, 2019.
- Zeyuan Allen-Zhu, Yuanzhi Li, and Yingyu Liang. Learning and generalization in overparameterized neural networks, going beyond two layers. In *Advances in neural information processing systems*, pp. 6155–6166, 2019.
- Sanjeev Arora, Aditya Bhaskara, Rong Ge, and Tengyu Ma. Provable bounds for learning some deep representations. In *International Conference on Machine Learning*, pp. 584–592, 2014.
- Miriam Ayer, H Daniel Brunk, George M Ewing, William T Reid, and Edward Silverman. An empirical distribution function for sampling with incomplete information. *The annals of mathematical statistics*, pp. 641–647, 1955.
- Eugene Belilovsky, Michael Eickenberg, and Edouard Oyallon. Greedy layerwise learning can scale to imagenet. In *International conference on machine learning*, pp. 583–593. PMLR, 2019.
- Yoshua Bengio and Olivier Delalleau. On the expressive power of deep architectures. In *International Conference on Algorithmic Learning Theory*, pp. 18–36. Springer, 2011.
- M Émile Borel. Les probabilités dénombrables et leurs applications arithmétiques. *Rendiconti del Circolo Matematico di Palermo (1884-1940)*, 27(1):247–271, 1909.
- Hugh D Brunk. Maximum likelihood estimates of monotone parameters. *The Annals of Mathematical Statistics*, pp. 607–616, 1955.
- Alon Brutzkus and Amir Globerson. Globally optimal gradient descent for a ConvNet with Gaussian inputs. *ArXiv*, abs/1702.07966, 2017.
- Yining Chen and Richard J Samworth. Generalized additive and index models with shape constraints. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 78(4):729–754, 2016.
- Simon S. Du and Surbhi Goel. Improved learning of one-hidden-layer convolutional neural networks with overlaps. *ArXiv*, abs/1805.07798, 2018.
- Simon S. Du, Jason D. Lee, Yuandong Tian, Barnabás Póczos, and Amrendra Kumar Singh. Gradient descent learns one-hidden-layer cnn: Don’t be afraid of spurious local minima. *ArXiv*, abs/1712.00779, 2017.
- Tolga Ergen and Mert Pilanci. Revealing the structure of deep neural networks via convex duality. *arXiv preprint arXiv:2002.09773*, 2020.
- Tolga Ergen and Mert Pilanci. Implicit convex regularizers of cnn architectures: Convex optimization of two-and three-layer networks in polynomial time. In *International Conference on Learning Representations (ICLR)*, 2021.
- J Farkas. Ober die theorie der einfachen ungleichungen. *J. Reine Angew. Math*, 124:1–24, 1902.
- Rong Ge, Jason D Lee, and Tengyu Ma. Learning one-hidden-layer neural networks with landscape design. *arXiv preprint arXiv:1711.00501*, 2017.

- Rong Ge, Rohith Kuditipudi, Zhize Li, and Xiang Wang. Learning two-layer neural networks with symmetric inputs. *arXiv preprint arXiv:1810.06793*, 2018.
- Surbhi Goel and Adam R. Klivans. Learning neural networks with two nonlinear layers in polynomial time. In *COLT*, 2017.
- Surbhi Goel, Adam R. Klivans, and Raghu Meka. Learning one convolutional layer with overlapping patches. *ArXiv*, abs/1802.02547, 2018.
- Michael Grant and Stephen Boyd. Graph implementations for nonsmooth convex programs. In V. Blondel, S. Boyd, and H. Kimura (eds.), *Recent Advances in Learning and Control*, Lecture Notes in Control and Information Sciences, pp. 95–110. Springer-Verlag Limited, 2008. http://stanford.edu/~boyd/graph_dcp.html.
- Michael Grant and Stephen Boyd. CVX: Matlab software for disciplined convex programming, version 2.1. <http://cvxr.com/cvx>, March 2014.
- Adityanand Guntuboyina, Bodhisattva Sen, et al. Nonparametric shape-restricted regression. *Statistical Science*, 33(4):568–594, 2018.
- James Douglas Hamilton. *Time series analysis*, volume 2. Princeton university press Princeton, NJ, 1994.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778, 2016a.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Identity mappings in deep residual networks. In *European conference on computer vision*, pp. 630–645. Springer, 2016b.
- Majid Janzamin, Hanie Sedghi, and Anima Anandkumar. Beating the perils of non-convexity: Guaranteed training of neural networks using tensor methods. *arXiv preprint arXiv:1506.08473*, 2015.
- Florian Jarre. On the convergence of the method of analytic centers when applied to convex quadratic programs. *Mathematical Programming*, 49(1-3):341–358, 1990.
- Robert I Jennrich. Asymptotic properties of non-linear least squares estimators. *The Annals of Mathematical Statistics*, 40(2):633–643, 1969.
- Iain M Johnstone. High dimensional statistical inference and random matrices. *arXiv preprint math/0611589*, 2006.
- Sham M. Kakade, Adam Tauman Kalai, Varun Kanade, and Ohad Shamir. Efficient learning of generalized linear and single index models with isotonic regression. In *NIPS*, 2011.
- Narendra Karmarkar. A new polynomial-time algorithm for linear programming. In *Proceedings of the sixteenth annual ACM symposium on Theory of computing*, pp. 302–311, 1984.
- Yoon Kim. Convolutional neural networks for sentence classification. *arXiv preprint arXiv:1408.5882*, 2014.
- Mikhail K Kozlov, Sergei P Tarasov, and Leonid G Khachiyan. The polynomial solvability of convex quadratic programming. *USSR Computational Mathematics and Mathematical Physics*, 20(5): 223–228, 1980.
- Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pp. 1097–1105, 2012.
- Arun K Kuchibhotla, Rohit K Patra, and Bodhisattva Sen. Efficient estimation in convex single index models. *Preprint. Available at*, 2017.
- Yann LeCun, Léon Bottou, Yoshua Bengio, Patrick Haffner, et al. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.

- Yuanzhi Li and Yang Yuan. Convergence analysis of two-layer neural networks with ReLU activation. *ArXiv*, abs/1705.09886, 2017.
- Roi Livni, Shai Shalev-Shwartz, and Ohad Shamir. On the computational efficiency of training neural networks. In *Advances in neural information processing systems*, pp. 855–863, 2014.
- David G Luenberger. *Optimization by vector space methods*. John Wiley & Sons, 1997.
- Andrew L Maas, Awni Y Hannun, and Andrew Y Ng. Rectifier nonlinearities improve neural network acoustic models. In *ICML Workshop on Deep Learning for Audio, Speech and Language Processing*, 2013.
- Henry B Mann and Abraham Wald. On stochastic limit and order relationships. *The Annals of Mathematical Statistics*, 14(3):217–226, 1943.
- Mary C. Meyer. A simple new algorithm for quadratic programming with applications in statistics. *Communications in Statistics - Simulation and Computation*, 42:1126–1139, 2013.
- Renato DC Monteiro and Ilan Adler. Interior path following primal-dual algorithms. part ii: Convex quadratic programming. *Mathematical Programming*, 44(1-3):43–66, 1989.
- Katta G Murty. Computational complexity of parametric linear programming. *Mathematical programming*, 19(1):213–219, 1980.
- Vinod Nair and Geoffrey E Hinton. Rectified linear units improve restricted boltzmann machines. In *Proceedings of the 27th international conference on machine learning (ICML-10)*, pp. 807–814, 2010.
- Mert Pilanci and Tolga Ergen. Neural networks are convex regularizers: Exact polynomial-time convex optimization formulations for two-layer networks. In *International Conference on Machine Learning*, pp. 7695–7705. PMLR, 2020.
- Guillaume Rabusseau, Tianyu Li, and Doina Precup. Connecting weighted automata and recurrent neural networks through spectral learning. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pp. 1630–1639. PMLR, 2019.
- R Tyrrell Rockafellar and Roger J-B Wets. *Variational analysis*, volume 317. Springer Science & Business Media, 2009.
- Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. Imagenet large scale visual recognition challenge. *International journal of computer vision*, 115(3):211–252, 2015.
- Arda Sahiner, Tolga Ergen, John Pauly, and Mert Pilanci. Vector-output ReLU neural network problems are copositive programs: Convex analysis of two layer networks and polynomial-time algorithms. *arXiv preprint arXiv:2012.13329*, 2020.
- Emilio Seijo, Bodhisattva Sen, et al. Nonparametric least squares estimation of a multivariate convex regression function. *The Annals of Statistics*, 39(3):1633–1657, 2011.
- Pranab K Sen and Julio M Singer. *Large sample methods in statistics: an introduction with applications*, volume 25. CRC press, 1994.
- Alexander Shapiro, Darinka Dentcheva, and Andrzej Ruszczyński. *Lectures on stochastic programming: modeling and theory*. SIAM, 2014.
- Mahdi Soltanolkotabi. Learning ReLUs via gradient descent. In *NIPS*, 2017.
- Ilya Sutskever, Oriol Vinyals, and Quoc V Le. Sequence to sequence learning with neural networks. In *Advances in neural information processing systems*, pp. 3104–3112, 2014.
- Yuandong Tian. An analytical formula of population gradient for two-layered ReLU network and its applications in convergence and critical point analysis. In *ICML*, 2017.

- Kim-Chuan Toh, Michael J Todd, and Reha H Tütüncü. Sdpt3—a matlab software package for semidefinite programming, version 1.3. *Optimization methods and software*, 11(1-4):545–581, 1999.
- Leslie G. Valiant. A theory of the learnable. *Commun. ACM*, 27:1134–1142, 1984.
- Constance van Eeden. Maximum likelihood estimation of ordered probabilities, 2. *Stichting Mathematisch Centrum. Statistische Afdeling*, (S 196/56), 1956.
- Vladimir Vapnik. Principles of risk minimization for learning theory. In *Advances in neural information processing systems*, pp. 831–838, 1992.
- Shanshan Wu, Alexandros G. Dimakis, and Sujay Sanghavi. Learning distributions generated by one-layer ReLU networks. In *NeurIPS*, 2019.
- Yinyu Ye and Edison Tse. An extension of karmarkar’s projective algorithm for convex quadratic programming. *Mathematical programming*, 44(1-3):157–179, 1989.
- Xiao Zhang, Yaodong Yu, Lingxiao Wang, and Quanquan Gu. Learning one-hidden-layer ReLU networks via gradient descent. *ArXiv*, abs/1806.07808, 2018.
- Yin Zhang, Richard A Tapia, and John E Dennis, Jr. On the superlinear and quadratic convergence of primal-dual interior point linear programming algorithms. *SIAM Journal on Optimization*, 2(2): 304–324, 1992.
- Kai Zhong, Zhao Song, Prateek Jain, Peter L. Bartlett, and Inderjit S. Dhillon. Recovery guarantees for one-hidden-layer neural networks. In *ICML*, 2017.

APPENDICES

We include an outline of the appendices:

- Appendix A: An empirical analysis of the running time of our algorithm and an analysis of the effect of condition number on the estimation accuracy.
- Appendix B: A brief discussion of our learning algorithms in a noisy context, referring to the LP slack variable technique that is used in our experiments with noise in the main paper. Also, this section provides an insight to potential future direction of this work.
- Appendix C: A detailed explanation of the computational complexity and the methods through which the QPs/LPs in this paper are solved.
- Appendix D: A preliminarily proposed general optimization formula for the extension of learning multi-layer ResNets, to show the scalability of our algorithm.
- Appendix E: A justification of how our QP/LP approaches get rid of the exponentiality that vanilla LR approach (section 3) has.
- Appendix F: A formal result and an empirical validation of the example given in the main paper in the context of the vanilla linear regression method (section 3).
- Appendix G: A generalization of layer 2 learning to the case where Cond. 2.1 is not satisfied.
- Appendix H: Proofs regarding the minimizers of the objective functions we use.
- Appendix I: Proofs that justify the convexity of our QPs.
- Appendix J: Proofs that show our estimation algorithm is strongly consistent.

A RUNNING TIME AND SENSITIVITY TO CONDITION NUMBER

Following the setup in section 7, we report additional experiments testing the running time of our algorithm in comparison to SGD and the feasibility of estimating the inverse of \mathbf{B}^* when ill-conditioned.




Figure 3: **Output errors against running time by SGD and our LP algorithm for different dimensions of networks.** For each dimension size of $d = 8$, $d = 16$ and $d = 32$, we report the learning curves of SGD and the final output error of our LP algorithm against time used on 5 different ground-truth networks. We used 512 training samples, drawn for learning each ground-truth network. (CPU specification: 2.8 GHz Quad-Core Intel Core i7.)

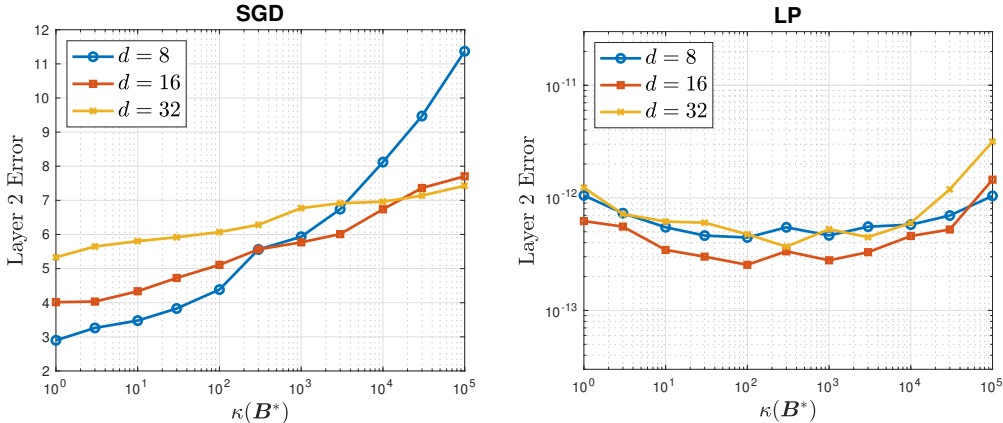


Figure 4: **Layer 2 errors against layer 2 ground-truth weight condition numbers for different dimensions of networks by SGD and our LP algorithm.** Each data point is the mean across learning 32 different ground-truth networks with the same pair of d and a condition number for B^* . The condition number is denoted by $\kappa(B^*)$. We used 512 training samples, drawn for learning each ground-truth network.

A.1 RUNNING TIME EFFICIENCY VERSUS SGD

In Fig. 3, we compare the running time of our algorithm against the running time of SGD for ground-truth networks sampled at random. The running time of our algorithm is significantly lower, with an error level which is also significantly lower. Furthermore, this is demonstrated across input dimensions, where the difference between the running time of SGD until convergence and our algorithm’s running time increases as d increases. In addition, the running time of our algorithm is fixed, and does not depend on the ground-truth network, while the running time of SGD varies based on the network sampled.

A.2 LAYER 2 WEIGHTS CONDITION NUMBER ROBUSTNESS

Our algorithm depends on a matrix-inversion step to obtain the estimated parameters. This affects the estimation of B^* . Due to this, we further investigate the effect of the condition number of B^* in the ground-truth parameters on the accuracy of estimation. For this experiment, to obtain ground-truth network parameters with different condition number, we follow the procedure of Ge et al. (2018), and multiply a diagonal matrix with exponentially-dropping values on the diagonal (λ^{-i} for the i -th element) by two random orthonormal matrices, U and V^T , on the left and on the right respectively: $B = U \text{diag}(\lambda^{-1}, \dots, \lambda^{-d}) V^T$.

In Fig. 4, we compare the effect of the condition number of \mathbf{B}^* on SGD and on our algorithm. SGD turns out to be quite sensitive to the condition number, especially for lower dimensions. Our algorithm, on the other hand, is almost not affected by the condition number, with a final estimation error consistently close to 0, recovering the matrix \mathbf{B}^* .

B DISCUSSION: NOISY MODEL

Here, we discuss our learning methods in the noisy case. We introduce output noise to our model, namely

$$\mathbf{y} = \mathbf{B}^* \left[(\mathbf{A}^* \mathbf{x})^+ + \mathbf{x} \right] + \mathbf{z}, \quad (9)$$

where the label noise $\mathbf{z} \in \mathbb{R}^d$ is an i.i.d. random vector with respect to each component, satisfying $\mathbb{E}[\mathbf{z}] = 0$. In addition \mathbf{z} and \mathbf{x} are statistically independent.

Taking layer 2 as an example, our original objective functional does not reach zero by substituting the ground-truth: $G_2(\mathbf{C}^*, \mathbf{x} \mapsto (\mathbf{A}^* \mathbf{x})^+) = \mathbb{E}_{\mathbf{z}}[\|\mathbf{C}^* \mathbf{z}\|^2] = \sigma^2 \text{Tr } \mathbf{C}^* \mathbf{C}^{*\top}$ where σ is the noise strength. However, if layer 2 is well conditioned, the ground-truth will still assign a value close to zero to the objective. In this sense, with G_2 's continuity, the ground-truth can approximately minimize G_2 , which validates our QP approach in terms of learning noisy residual units.

Our original LP (Eq. 5) fails to give feasible solutions due to possible violation of the inequality $\mathbf{C}^* \mathbf{y} - \mathbf{x} \geq 0$, since $\mathbf{C}^* \mathbf{y} - \mathbf{x} = (\mathbf{A}^* \mathbf{x})^+ + \mathbf{C}^* \mathbf{z}$ is not necessarily an entrywise nonnegative vector because the term $\mathbf{C}^* \mathbf{z}$ might have negative entries. So we introduce slack variables $\zeta^{(i)}$ to soften the constraints:

$$\min_{\mathbf{C}, \mathbf{z}} \frac{1}{n} \sum_{i \in [n]} \mathbf{1}^\top \cdot \zeta^{(i)}, \quad (10)$$

$$\text{s.t. } \mathbf{C} \mathbf{y}^{(i)} - \mathbf{x}^{(i)} \geq -\zeta^{(i)}, \zeta^{(i)} \geq \mathbf{0}. \quad (11)$$

For a sample $(\mathbf{x}^{(i)}, \mathbf{y}^{(i)})$ with noise $\mathbf{z}^{(i)}$, if $(\mathbf{A}^* \mathbf{x}^{(i)})^+ + \mathbf{C}^* \mathbf{z}^{(i)} < 0$, then its L^1 norm would be added to the objective. This method remedies violations to the inequality $\mathbf{C}^* \mathbf{y} - \mathbf{x} \geq 0$. With access to a sufficiently large sample and with the ‘‘stability’’ assumption, our solution $\hat{\mathbf{C}}$ would be close to \mathbf{C}^* since large deviations seldom rise and their penalties are diluted in the objective.

C DISCUSSION: SOLVING CONVEX PROGRAMS

The theoretical foundation of solving convex QP and LP has been driven to maturity in terms of computational complexity (Kozlov et al., 1980; Murty, 1980) and convergence analysis (Jarre, 1990; Zhang et al., 1992). The time complexity for a convex QP/LP is analyzed in terms of the number of scalar variables N and the number of bits L in the input (Ye & Tse, 1989). For example, under Cond. 2.1, $N = d^2 + nd$ for the QP because the matrix variables have d^2 scalars and the function estimators have nd scalars. Similarly we have $N = d^2$ for the noiseless LP and $N = d^2 + nd$ for the noisy LP. It is guaranteed that primal-dual interior point methods can solve the convex QP/LP in a polynomial number of iterations $\mathcal{O}(\sqrt{N}L)$, where each iteration costs at worst $\mathcal{O}(N^{2.5})$ arithmetic operations from the Cholesky decomposition for the needed matrix inversion (Monteiro & Adler, 1989; Karmarkar, 1984). This indicates that our QPs/LPs are guaranteed to be solvable in $\mathcal{O}(N^3L)$ arithmetic operations, which is generally an $\mathcal{O}(\text{poly}(n, m, d, L))$ complexity.

To solve convex QPs/LPs in this paper, we use CVX, a commonly used package for specifying and solving convex programs (Grant & Boyd, 2014; 2008). For both QPs and LPs, CVX calls a solver, SDPT3 (Toh et al., 1999), which is specified for semidefinite-quadratic-linear programming and applies interior-point methods with the computational complexity mentioned above. Experimentally, SDPT3 indeed specifies and solves our programs fast and makes our numerical results robust and stable.

D DISCUSSION: EXTENSION TO MULTI-LAYER

In this appendix, we propose a generalized optimization formula of nonparametric learning of multi-layer residual networks to show that our algorithm is scalable. Assume we have an L -layer network of the form

$$\begin{aligned} \mathbf{r}_0 &= \mathbf{x} \\ \mathbf{r}_l &= (\mathbf{W}_{l-1}^* \mathbf{r}_{l-1})^+ + \mathbf{r}_{l-1}, l = 1, \dots, L-1 \\ \mathbf{y} &= \mathbf{W}_{L-1}^* \mathbf{r}_{L-1}. \end{aligned}$$

We can scale the objective to use nonparameter variables $\{\xi_l^{(i)}\}_{l \in [L-1]}$ to estimate every ReLU activation $(\mathbf{W}_l^* \mathbf{r}_l)^+$ in this network. Generally, for the l -th ReLU we can have the outer layer objective

$$\min_{\xi_{[l]}^{(i)}, \mathbf{V}_l} \left\| \xi_l^{(i)} + \dots + \xi_1^{(i)} + \mathbf{x}^{(i)} - \mathbf{V}_l \hat{\mathbf{h}}_l^{(i)} \right\|, \text{ s.t. } \xi_{[l]}^{(i)} \geq 0$$

where hat denotes estimated values and $\hat{\mathbf{h}}_{L-1}^{(i)} = \mathbf{y}^{(i)}$, and its corresponding inner layer objective

$$\min_{\phi_l^{(i)}, \mathbf{W}_{l-1}, \xi_l^{(i)}} \left\| \phi_l^{(i)} + \mathbf{W}_{l-1} \left(\hat{\xi}_{l-1}^{(i)} + \dots + \hat{\xi}_1^{(i)} + \mathbf{x}^{(i)} \right) - \xi_l^{(i)} \right\|, \text{ s.t. } \phi_l^{(i)} \geq 0, \xi_l^{(i)} \geq 0.$$

We can apply techniques, e.g., tune the combinations of those objectives with respect to layers or whether to have a hat or not, or even get them weighted, to keep its convex quadratic form and at the same time improve its optimization landscape. Specifically when $L = 2$, this becomes the two-layer ResNet learning discussed in the main paper.

E SAMPLE EFFICIENCY: VANILLA LR VERSUS OURS

In this appendix, we describe an example to justify how our core QP/LP approaches eliminate the exponential sample efficiency compared to the vanilla LR approach. In layer 2 learning, \mathbf{C} 's optimization is to find a feasible point in the space determined by n inequalities each of which corresponds to a sample. Taking the j -th row of the inequality, i.e. $C_{j,:} \mathbf{y} \geq x_j$. Every time we get a new sample $\mathbf{x}^{(i)}, \mathbf{y}^{(i)}$ in,

- The inequality $C_{j,:} \mathbf{y}^{(i)} \geq x_j^{(i)}$ eliminates the solution space of $C_{j,:}$ by one of the spaces divided by plane $C_{j,:} \mathbf{y}^{(i)} = x_j^{(i)}$. This property guarantees fast convergence speed at early phase since $C_{j,:}$'s feasibility starts from \mathbb{R}^d .
- In fact, when ReLU does not activate $\mathbf{A}^* \mathbf{x}^{(i)}$ at the j -th row, i.e. $\mathbf{A}^* \mathbf{x}^{(i)} \leq 0$, the theoretical solution of $C_{j,:}$ that $C_{j,:} \mathbf{B}^* = (0, \dots, 0, 1, 0, \dots, 0)$ (a 1 at index j) directly lies on the divisive plane $C_{j,:} \mathbf{y}^{(i)} = x_j^{(i)}$:

$$C_{j,:} \mathbf{y}^{(i)} = C_{j,:} \mathbf{B}^* \left[\left(\mathbf{A}^* \mathbf{x}^{(i)} \right)^+ + \mathbf{x}^{(i)} \right] = x_j^{(i)}.$$

This property remarkably speeds up further constraints on the solution space of $C_{j,:}$ to the correct estimate and is with high probability, since it directly depends on the sign of $\mathbf{A}^* \mathbf{x}^{(i)}$.

With the vanilla LR approach mentioned in section 3 all the dimensions need to have the correct sign at the same time. This is the reason for the exponential complexity of this approach. Our main algorithm only requires one dimension to get the correct sign for a single sample, which avoids the exponential sample size expectation. The convergence speed is determined by the actual probability of each dimension not getting activated by ReLU, and such probability is determined by specified input distribution.

Algorithm 3 Learn a ReLU residual unit by LR.

-
- 1: **Input:** $\{(\mathbf{x}^{(i)}, \mathbf{y}^{(i)})\}_{i=1}^n$, n samples drawn by Eq. 1.
 - 2: **Output:** $\hat{\mathbf{A}}, \hat{\mathbf{B}}$, estimated weight matrices.
 - 3: $\hat{\mathbf{B}} \leftarrow \text{LR}\{(\mathbf{x}, \mathbf{y}) \in \{(\mathbf{x}^{(i)}, \mathbf{y}^{(i)})\}_{i=1}^{\lfloor n/2 \rfloor} \mid \mathbf{x} < 0\}$.
 - 4: $\hat{\mathbf{D}} \leftarrow \text{LR}\{(\mathbf{x}, \mathbf{y}) \in \{(\mathbf{x}^{(i)}, \mathbf{y}^{(i)})\}_{i=\lfloor n/2 \rfloor+1}^n \mid \mathbf{x} > 0\}$. {the estimation of $\mathbf{B}^* (\mathbf{A}^* + \mathbf{I}_d)$ }
 - 5: Solve full-rank linear equation system $\hat{\mathbf{B}} \cdot \tilde{\mathbf{A}} = \hat{\mathbf{D}}$
 - 6: **return** $\tilde{\mathbf{A}} - \mathbf{I}_d, \hat{\mathbf{B}}$.
-

F VANILLA LINEAR REGRESSION: MORE DETAILS

Alg. 3 gives the full list of steps of learning two-layer residual units by LR, where we first split the drawn samples into two halves for the respective two key steps, then for both halves we filter negative and positive vectors, respectively. By running LR on both filtered sets we obtain an estimation of the ground-truth network.

In the following sections, we formalize this intuition and describe the exponential sample complexity of this approach under entry-wise i.i.d. setting as an example of the inefficiency of this method.

Theorem F.1 (exponential sample complexity, vanilla LR). *Assume the input vectors are i.i.d. with respect to each component. Then Alg. 3 learns a neural network from a ground-truth residual unit with at least $\mathcal{O}(d \cdot 2^{d+1})$ expected number of samples.*

Proof. For each $j \in [d]$, let P_j be the marginal probability of x_j being positive, i.e. $P_j = P(x_j > 0) > 0$. Thus, the probability that a sample is positive $P(\mathbf{x} > 0) = \prod_{j \in [d]} P_j$. Then the expected number of sampling trials to obtain d positive samples is $d / \prod_{j \in [d]} P_j$. Similarly, the expected number of sampling trials to obtain d negative samples is $d / \prod_{j \in [d]} [1 - P_j - P(\mathbf{x}_j = 0)]$. Let n be a random natural number s.t. with n samples the network has learned. The expected number of samples that guarantees successful learning is the sum of the two expectations

$$\begin{aligned}
 \mathbb{E}[n] &= d \left\{ \frac{1}{\prod_{j \in [d]} P_j} + \frac{1}{\prod_{j \in [d]} [1 - P_j - P(\mathbf{x}_j = 0)]} \right\} \\
 &\geq d \left[\frac{1}{\prod_{j \in [d]} P_j} + \frac{1}{\prod_{j \in [d]} (1 - P_j)} \right] \\
 &\geq d \cdot 2 \prod_{j \in [d]} \frac{1}{\sqrt{P_j (1 - P_j)}} \\
 &\geq d \cdot 2 \prod_{j \in [d]} \frac{1}{(P_j + 1 - P_j) / 2} \\
 &= d \cdot 2^{d+1}.
 \end{aligned}$$

Thus $\mathbb{E}[n] \geq \mathcal{O}(d \cdot 2^{d+1})$. □

With an exponential sample complexity, the time complexity of the LR approach is thereby also exponential because filtering through all the samples costs time with the same complexity as the number of samples, even though the LR itself only costs polynomial time in the number of samples. To summarize, this vanilla LR approach learns exact ground-truth residual units but with exponential complexity in terms of both computational cost and sample size. While this approach to learning a residual unit such as ours is simple and intuitive, we aim to improve on this approach, making full use of all samples available.

F.1 EXPERIMENTS: SAMPLE EFFICIENCY

We present the results of the vanilla LR in learning the residual units. As discussed in section 3, LR requires full-rank linear systems parameterized by samples to learn the exact ground-truth parameters.

Table 2: **Learning success rate of vanilla LR on residual units with input $\mathbf{x} \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, 1)$.** The success rates shows an exponentially decreasing trend with the same number of samples. The sample sizes to achieve close to the same rate grows exponentially as the number of dimensions grows linearly.

$d \backslash n$	1e1	1e2	5e2	1e3	5e3	1e4	5e4	1e5
4	0.002	0.137	0.999	1	1	1	1	1
6	0	0	0.041	0.614	1	1	1	1
8	0	0	0	0	0.579	0.998	1	1
10	0	0	0	0	0	0.002	1	1
12	0	0	0	0	0	0	0.001	0.322

Table 3: **Learning success rate of vanilla LR on residual units with input $\mathbf{x} \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0.1, 1)$.** With the non-zero Gaussian mean, the success rates show an overall decline compared with the rates shown in Tab. 2 for zero-mean Gaussian inputs.

$d \backslash n$	1e1	1e2	5e2	1e3	5e3	1e4	5e4	1e5
4	0.001	0.122	0.996	1	1	1	1	1
6	0	0	0.030	0.346	1	1	1	1
8	0	0	0	0	0.116	0.792	1	1
10	0	0	0	0	0	0	0.619	1
12	0	0	0	0	0	0	0	0.001

However, with degenerate linear systems, LR is completely incapable of learning the parameters. Therefore, there exists a hard threshold for the number of samples required by LR that makes the linear equation system full-rank. With such a dichotomous constraint on LR learning, we take the learning success rate among 1000 trials as the metric to evaluate the performance of LR with different sample sizes and number of dimensions.

Tab. 2 shows the learning success rates with zero-mean Gaussian inputs. For each fixed number of dimensions, it appears there is a hard threshold that switches the learnability of LR. The exponential sample complexity is also reflected there as the number of dimensions grows linearly. Tab. 3 shows the rates with input mean non-zero, where an overall decline in success rates happens. This observation is explainable because the bottleneck of LR learning is the lower value found between the probability of sampling a positive and a negative vector. A positive mean reduces the probability of the latter, and thereby increases the sample size required.

G GENERALIZATION OF LAYER 2 LEARNING

In this appendix, we discuss how layer 2 is learned without satisfying Cond. 2.1. It is a fact that \mathbf{A}^* being a scale transformation w.r.t. some components causes scaling equivalence to our layer 2 objective functional minimizers just as in layer 1. To be more specific, we give a general version of Thm. 4.1 that handles the case without Cond. 2.1 and describes the scaling equivalence of G_2 minimizers. See Thm. G.1 for a formal description.

Theorem G.1 (objective minimizer space, layer 2, general). *Let G_2 be $\frac{1}{2}\mathbb{E}_{\mathbf{x}} [\|h(\mathbf{x}) + \mathbf{x} - \mathbf{C}\mathbf{y}\|^2]$ as a functional, where $\mathbf{C} \in \mathbb{R}^{d \times m}$, $h \in \mathbb{C}_{\geq 0}^0$. Then $G_2(\mathbf{C}, h)$ reaches its zero minimum iff $\mathbf{C}\mathbf{B}^* = \text{diag}(\mathbf{k})$ where for each $j \in [d]$,*

- $\frac{1}{1+A_{j,j}^*} \leq k_j \leq 1$, if \mathbf{A}^* is a scale transformation w.r.t. the j -th component.
- $k_j = 1$, if \mathbf{A}^* is not a scale transformation w.r.t. the j -th component.

As in layer 1, the scaling equivalence in layer 2 can also be obtained by assigning \mathbf{A}^* as a scale transformation w.r.t. some component and using the properties of the ReLU nonlinearity. See subsection H.1 for a detailed explanation of Thm. G.1. To obtain the scale factors \mathbf{k} and correct a scale-equivalent G_2 minimizer \mathbf{C} to the ground-truth, we observe linear models parameterized by the scale factors, where for each $j \in [d]$, $[\mathbf{C}\mathbf{y}]_j = k_j x_j$ given that $x_j < 0$. See Thm. G.2 and its proof for justifications of the linear models that are used to compute \mathbf{k} .

Theorem G.2 (scale factor, layer 2, general). *Assume \mathbf{C} is a minimizer of G_1 in the context of Thm. G.1. Then for any $j \in [d]$, the following three propositions are equivalent:*

- a) \mathbf{A}^* is a scale transformation w.r.t. the j -th component.
- b) $[\mathbf{C}\mathbf{y}]_j / x_j$ is a constant c_j^n given that $x_j < 0$.
- c) $[\mathbf{C}\mathbf{y}]_j / x_j$ is a constant c_j^p given that $x_j > 0$.

Additionally, if one of the above three propositions is true, then $k_j = c_j^n$.

Proof. For $j \in [d]$, we first prove a) \implies b), c): From $(\mathbf{A}_{j,:}^*)^\top = A_{j,j}^* \cdot \mathbf{e}^{(j)}$ we have

$$\begin{aligned} \frac{[\mathbf{C}\mathbf{y}]_j}{x_j} &= \frac{[k_j \mathbf{e}^{(j)}]^\top [(\mathbf{A}^* \mathbf{x})^+ + \mathbf{x}]}{x_j} \\ &= \frac{k_j [(\mathbf{A}_{j,:}^* \mathbf{x})^+ + x_j]}{x_j} = \frac{k_j [(A_{j,j}^* x_j)^+ + x_j]}{x_j} \\ &= \begin{cases} k_j, & x_j < 0 \\ k_j (A_{j,j}^* + 1), & x_j > 0 \end{cases} \end{aligned}$$

Then we prove \neg a) $\implies \neg$ b), \neg c): \neg a) $\implies \exists j' \in [d]$ and $j' \neq j$ s.t. $A_{j,j'}^* \neq 0$. Recall

$$\frac{[\mathbf{C}\mathbf{y}]_j}{x_j} = \frac{k_j [(\mathbf{A}_{j,:}^* \mathbf{x})^+ + x_j]}{x_j}.$$

The value of $x_{j'}$ affects the value of $[\mathbf{C}\mathbf{y}]_j / x_j$ because $A_{j,j'}^* \neq 0$. In fact, from $\text{supp } p(\mathbf{x}) = \mathbb{R}^d$ we have $\text{supp } p(x_{j'} \mid x_j) = \mathbb{R}$, indicating that the value of $[\mathbf{C}\mathbf{y}]_j / x_j$ can never be kept as a constant when given both $x_j < 0$ and $x_j > 0$ because $x_{j'}$ can be any real number. \square

With the exact derivations above, we are able to obtain a left inverse of \mathbf{B}^* , namely \mathbf{C} , that satisfies $\mathbf{C}\mathbf{B}^* = \mathbf{I}_d$. Consider Eq. 1 left multiplied by $\mathbf{B}^*\mathbf{C}$

$$\mathbf{B}^* \mathbf{C} \mathbf{y} = \mathbf{y}. \quad (12)$$

Eq. 12 is also a noiseless unbiased linear model where $\mathbf{C}\mathbf{y}$ and \mathbf{y} are observable, and thereby \mathbf{B}^* is computable due to the easy solvability of its linearity.

Alg. 5 describes how a residual unit second layer is learned without satisfying Cond. 2.1: We first solve the QP/LP and obtain a scaling equivalence to an estimated left inverse of \mathbf{B}^* and the function estimate, namely $\hat{\mathbf{C}}$ and $\hat{\mathbf{\Xi}}$. Then we compute the scale factor estimate $\hat{\mathbf{k}}$ by running Alg. 4, where for each component index $j \in [d]$, we first use a tolerance parameter as a threshold to determine whether ground-truth layer 1 is a scale transformation, and if so, we run LR to estimate the model $[\mathbf{C}\mathbf{y}]_j = k_j x_j$, otherwise the scale factor is directly assigned by 1. Upon correcting $\hat{\mathbf{C}}$ and $\hat{\mathbf{\Xi}}$ by $\hat{\mathbf{k}}$, we obtain a layer 2 estimate $\hat{\mathbf{B}}$ by running LR to estimate the linear model Eq. 12. The strong consistency of our results in this appendix is justified in Appendix J.

H EXACT DERIVATION OF OBJECTIVE FUNCTIONAL MINIMIZERS

H.1 LAYER 2

Lem. 4.2 is proved as follows.

Algorithm 4 Rescale a \hat{G}_2^{NPE} minimizer.

- 1: **Parameters:** $\varepsilon_{\text{tol}} > 0$, LR objective tolerance.
 - 2: **Input:** $\{(\mathbf{x}^{(i)}, \mathbf{y}^{(i)})\}_{i=1}^n$, samples drawn by Eq. 1; $\hat{\mathbf{C}}$, a \hat{G}_2^{NPE} minimizer.
 - 3: **Output:** $\hat{\mathbf{k}}$, a layer 2 scale factor estimate.
 - 4: **for** each $j \in [d]$ **do**
 - 5: $\hat{k}_j \leftarrow \text{LR} \left\{ x_j^{(i)}, [\hat{\mathbf{C}}\mathbf{y}^{(i)}]_j \right\}_{x_j^{(i)} < 0}$.
 - 6: **if** the LR objective optimal $\hat{R}_j(\hat{k}_j) > \varepsilon_{\text{tol}}$ **then**
 - 7: $\hat{k}_j \leftarrow 1$.
 - 8: **end if**
 - 9: **end for**
 - 10: **return** $\hat{\mathbf{k}}$.
-

Algorithm 5 Learn a ReLU residual unit layer 2.

- 1: **Input:** $\{(\mathbf{x}^{(i)}, \mathbf{y}^{(i)})\}_{i=1}^n$, samples drawn by Eq. 1.
 - 2: **Output:** $\hat{\mathbf{B}}, \hat{\Xi}$, a layer 2 and $\mathbf{x} \mapsto (\mathbf{A}^*\mathbf{x})^+$ estimate.
 - 3: Go to line 4 if QP, line 5 if LP.
 - 4: Solve QP: Eq. 4 and obtain a \hat{G}_2^{NPE} minimizer, denoted by $\hat{\mathbf{C}}, \hat{\Xi}$. Go to line 6.
 - 5: Solve LP: Eq. 5 and obtain a minimizer $\hat{\mathbf{C}}$, then assign $\hat{\xi}^{(i)} \leftarrow \hat{\mathbf{C}}\mathbf{y}^{(i)} - \mathbf{x}^{(i)}$ for each $i \in [n]$.
 - 6: Run Alg. 4 on $\{(\mathbf{x}^{(i)}, \mathbf{y}^{(i)})\}_{i \in [n]}$, $\hat{\mathbf{C}}$ and obtain $\hat{\mathbf{k}}$.
 - 7: $\hat{\mathbf{B}} \leftarrow \text{LR} \left\{ \text{diag}^{-1}(\hat{\mathbf{k}}) \cdot \hat{\mathbf{C}}\mathbf{y}^{(i)}, \mathbf{y}^{(i)} \right\}_{i \in [n]}$.
 - 8: $\hat{\xi}^{(i)} \leftarrow \text{diag}^{-1}(\hat{\mathbf{k}}) [\hat{\xi}^{(i)} + \mathbf{x}^{(i)}] - \mathbf{x}^{(i)}$ for each $i \in [n]$. {Correct $\hat{\Xi}$ to the function estimation of $(\mathbf{A}^*\mathbf{x})^+$.}
 - 9: **return** $\hat{\mathbf{B}}, \hat{\Xi}$.
-

Proof. “ \Leftarrow ”: Since $h(\mathbf{x}) = \mathbf{C}\mathbf{y} - \mathbf{x} \geq 0$, random vector $h(\mathbf{x}) + \mathbf{x} - \mathbf{C}\mathbf{y}$ is always a zero vector, which implies $G_2(\mathbf{C}, h) = 0$. Hence “ \Leftarrow ” holds.

“ \Rightarrow ”: Since r.v. $\|h(\mathbf{x}) + \mathbf{x} - \mathbf{C}\mathbf{y}\|^2 \geq 0$ we have

$$G_2(\mathbf{C}, h) = 0 \implies \lambda \left(\left\{ \mathbf{x} \in \mathbb{R}^d \mid \|h(\mathbf{x}) + \mathbf{x} - \mathbf{C}\mathbf{y}\|^2 > 0 \right\} \right) = 0 \quad (13)$$

where λ is the Lebesgue measure on \mathbb{R}^d .

Proof by contradiction: Assume that $\exists \mathbf{x}' \in \mathbb{R}^d, \exists i \in [d]$, s.t. $(\mathbf{C}\mathbf{y}' - \mathbf{x}')_i < 0$, where \mathbf{y}' is the corresponding network output. Let $f(\mathbf{x}) = (\mathbf{C}\mathbf{y} - \mathbf{x})_i$ which is continuous on \mathbb{R}^d . Therefore for $\epsilon = -f(\mathbf{x}') > 0, \exists \delta > 0, \forall \mathbf{x} \in B(\mathbf{x}; \delta)$ i.e. $\|\mathbf{x} - \mathbf{x}'\| < \delta$,

$$\begin{aligned} |f(\mathbf{x}) - f(\mathbf{x}')| < \epsilon &\implies 2f(\mathbf{x}') < f(\mathbf{x}) < 0 \implies [\mathbf{x} - \mathbf{C}\mathbf{y}]_i > 0 \\ \implies [h(\mathbf{x}) + \mathbf{x} - \mathbf{C}\mathbf{y}]_i > 0 &\implies \|h(\mathbf{x}) + \mathbf{x} - \mathbf{C}\mathbf{y}\|^2 > 0. \end{aligned}$$

Since $\lambda(B(\mathbf{x}; \delta)) = \frac{\pi^{d/2}}{\Gamma(d/2 + 1)} \delta^d > 0$ and $B(\mathbf{x}; \delta)$ is a subset of the measured set in Eq. 13,

we have a contradiction. Thus $\mathbf{C}\mathbf{y} - \mathbf{x} \geq 0$ holds in a pointwise manner in \mathbb{R}^d , indicating $\{\mathbf{x} \in \mathbb{R}^d \mid h(\mathbf{x}) \neq \mathbf{C}\mathbf{y} - \mathbf{x}\}$ must be a null set. With h 's continuity, h must be $\mathbf{x} \mapsto \mathbf{C}\mathbf{y} - \mathbf{x}$. Therefore “ \Rightarrow ” holds. \square

Lemma H.1. $\mathbf{C}\mathbf{y} - \mathbf{x} \geq 0$ holds for any $\mathbf{x} \in \mathbb{R}^d$ and its corresponding output \mathbf{y} only if $\mathbf{C}\mathbf{B}^*$ is a diagonal matrix.

Proof. Let $\mathbf{D} = \mathbf{C}\mathbf{B}^*$. We rewrite $\mathbf{C}\mathbf{y} - \mathbf{x} \geq 0$ as

$$\mathbf{D} \left[(\mathbf{A}^*\mathbf{x})^+ + \mathbf{x} \right] - \mathbf{x} \geq 0. \quad (14)$$

For further use, we substitute \mathbf{x} by $-\mathbf{x}$ and the resulting inequality $\mathbf{D} \left[(-\mathbf{A}^* \mathbf{x})^+ - \mathbf{x} \right] + \mathbf{x} \geq 0$ still holds. Added by Eq. 14 we have

$$\mathbf{D} \left(|\mathbf{A}_{k,:}^* \cdot \mathbf{x}| \right)_{d \times 1} \geq 0. \quad (15)$$

Proof by contradiction: Assume \mathbf{D} is not diagonal, then $\exists i \neq j$, such that $D_{i,j} \neq 0$. Consider the following two cases:

a) $D_{i,j} > 0$. We take the i -th row of Eq. 14 as follows

$$\sum_{k=1}^d D_{i,k} \left[(\mathbf{A}_{k,:}^* \cdot \mathbf{x})^+ + x_k \right] \geq x_i. \quad (16)$$

Let $x_{-j} = 0$ and $x_j < 0$. Then $\sum_{k=1}^d D_{i,k} \left[(\mathbf{A}_{k,:}^* \cdot \mathbf{x})^+ + x_k \right] = D_{i,j} \cdot x_j < 0 \implies \perp$.

b) $D_{i,j} < 0$. We take the i -th row of Eq. 15 as follows

$$\sum_{k=1}^d D_{i,k} |\mathbf{A}_{k,:}^* \cdot \mathbf{x}| \geq 0. \quad (17)$$

Let $\mathbf{x} = (\mathbf{A}^*)^{-1} \mathbf{v}$, where $\mathbf{v} = \mathbf{e}^{(j)}$. Then $\sum_{k=1}^d D_{i,k} |\mathbf{A}_{k,:}^* \cdot \mathbf{x}| = D_{i,j} < 0 \implies \perp$.

□

With Lem. 4.2 and Lem. H.1, we prove Thm. G.1 as follows.

Proof. With Lem. 4.2, we only need to prove $\forall \mathbf{x} \in \mathbb{R}^d, \mathbf{C}\mathbf{y} - \mathbf{x} \geq 0 \iff \mathbf{C}\mathbf{B}^* = \text{diag}(\mathbf{k})$.

“ \Leftarrow ”: We have

$$\mathbf{C}\mathbf{y} - \mathbf{x} = \mathbf{C}\mathbf{B}^* \left[(\mathbf{A}^* \mathbf{x})^+ + \mathbf{x} \right] - \mathbf{x} = \text{diag}(\mathbf{k}) \left[(\mathbf{A}^* \mathbf{x})^+ + \mathbf{x} \right] - \mathbf{x}. \quad (18)$$

For the i -th row of Eq. 18, consider the following two cases:

a) $\mathbf{A}_{i,:}^* \cdot \mathbf{x} = A_{i,i} x_i$, i.e. \mathbf{A}^* is a scale transformation w.r.t. the i -th row. With $\frac{1}{1+A_{i,i}^*} \leq k_i \leq 1$ we have

$$[\mathbf{C}\mathbf{y} - \mathbf{x}]_i = k_i \left[(A_{i,i}^* x_i)^+ + x_i \right] - x_i = k_i (A_{i,i}^* x_i)^+ + (k_i - 1) x_i.$$

For $x_i \geq 0$, $[\mathbf{C}\mathbf{y} - \mathbf{x}]_i = (k_i A_{i,i}^* + k_i - 1) x_i \geq 0$.

For $x_i < 0$, $[\mathbf{C}\mathbf{y} - \mathbf{x}]_i = (k_i - 1) x_i \geq 0$.

b) $\mathbf{A}_{i,:}^* \cdot \mathbf{x} \neq A_{i,i} x_i$. With $k_i = 1$ we have

$$[\mathbf{C}\mathbf{y} - \mathbf{x}]_i = k_i \left[(\mathbf{A}_{i,:}^* \cdot \mathbf{x})^+ + x_i \right] - x_i = (\mathbf{A}_{i,:}^* \cdot \mathbf{x})^+ \geq 0.$$

Hence “ \Leftarrow ” holds.

“ \Rightarrow ”: Let $\mathbf{D} = \mathbf{C}\mathbf{B}^*$. With Lem. H.1, \mathbf{D} is diagonal. Consider the following two cases:

a) $\mathbf{A}_{i,:}^* \cdot \mathbf{x} = A_{i,i} x_i$. The i -th inequality can be written as

$$\begin{aligned} [\mathbf{C}\mathbf{y} - \mathbf{x}]_i &= D_{i,i} \left[(A_{i,i}^* x_i)^+ + x_i \right] - x_i \\ &= D_{i,i} (A_{i,i}^* x_i)^+ + (D_{i,i} - 1) x_i \geq 0. \end{aligned} \quad (19)$$

Proof by contradiction: we need to find $\mathbf{x} \in \mathbb{R}^d$ which contradicts with Eq. 19 in the following three cases:

- If $D_{i,i} \leq 0$, then, $x_i > 0 \implies D_{i,i} (A_{i,i}^* x_i)^+ + (D_{i,i} - 1) x_i < 0 \implies \perp$.
- If $D_{i,i} > 1$, then, $x_i < 0 \implies \perp$.

c) If $0 < D_{i,i} < \frac{1}{1+A_{i,i}^*}$, then $\exists a > 0$, s.t. $D_{i,i} = \frac{1}{1+A_{i,i}^*+a}$. Letting $x_i > 0$, we have

$$D_{i,i} (A_{i,i}^* x_i)^+ + (D_{i,i} - 1) x_i = \frac{(A_{i,i}^* x_i)^+}{1 + A_{i,i}^* + a} - \frac{(A_{i,i}^* + a) x_i}{1 + A_{i,i}^* + a} < 0 \implies \perp.$$

Hence $\frac{1}{1+A_{i,i}^*} \leq D_{i,i} \leq 1$.

b) $\mathbf{A}_{i,:}^* \cdot \mathbf{x} \neq A_{i,i}^* x_i$, i.e. $\exists j \neq i$, s.t. $\mathbf{A}_{i,j}^* > 0$. The i -th inequality can be written as

$$\begin{aligned} [\mathbf{C}\mathbf{y} - \mathbf{x}]_i &= D_{i,i} \left[(\mathbf{A}_{i,:}^* \mathbf{x})^+ + x_i \right] - x_i \\ &= D_{i,i} (\mathbf{A}_{i,:}^* \mathbf{x})^+ + (D_{i,i} - 1) x_i \geq 0. \end{aligned} \quad (20)$$

Proof by contradiction: we need to find $\mathbf{x} \in \mathbb{R}^d$ which contradicts with Eq. 20 in the following three cases:

- a) If $D_{i,i} \leq 0$, then, $x_i > 0 \implies D_{i,i} (\mathbf{A}_{i,:}^* \mathbf{x})^+ + (D_{i,i} - 1) x_i < 0 \implies \perp$.
- b) If $D_{i,i} > 1$, then, $x_i < 0 \wedge x_{-i} \leq 0 \implies \perp$.
- c) If $0 < D_{i,i} < 1$, then, $x_i > 0 \wedge x_j \leq -\frac{A_{i,i}^*}{A_{i,j}^*} x_i \wedge x_k \leq 0 \implies \perp$, where $k \neq i, j$.

Hence $D_{i,i} = 1$.

Hence $\mathbf{D} = \mathbf{C}\mathbf{B}^* = \text{diag}(\mathbf{k})$, and thereby “ \implies ” holds. \square

H.2 LAYER 1

The proof of Lem. 5.3 is similar to that of Lem. 4.2 because the two lemmas follow the same idea, which is to link objective functional minimization with always-hold inequalities. With Lem. 5.3, we prove Thm. 5.1 as follows.

Proof. With Lem. 5.3, we only need to prove $\forall \mathbf{x} \in \mathbb{R}^d, (\mathbf{A}^* \mathbf{x})^+ - \mathbf{A} \mathbf{x} \geq 0 \iff \forall i \in [d], \mathbf{A}_{i,:} = k_i \mathbf{A}_{i,:}^*$, where $0 \leq k_i \leq 1$. This is equivalent to what it is for a single row, i.e. $\forall \mathbf{x} \in \mathbb{R}^d, (\mathbf{a}^{*\top} \mathbf{x})^+ - \mathbf{a}^\top \mathbf{x} \geq 0 \iff \mathbf{a} = k \mathbf{a}^*$, where $0 \leq k \leq 1$.

“ \Leftarrow ”: The case where $\mathbf{a}^* = 0$ is clear. If \mathbf{a}^* is not a zero vector and $\mathbf{a} = k \mathbf{a}^*$, we have

- a) If $\mathbf{a}^{*\top} \mathbf{x} \geq 0$, $(\mathbf{a}^{*\top} \mathbf{x})^+ - \mathbf{a}^\top \mathbf{x} = \mathbf{a}^{*\top} \mathbf{x} - k \mathbf{a}^{*\top} \mathbf{x} = (1 - k) \mathbf{a}^{*\top} \mathbf{x} \geq 0$.
- b) If $\mathbf{a}^{*\top} \mathbf{x} < 0$, $(\mathbf{a}^{*\top} \mathbf{x})^+ - \mathbf{a}^\top \mathbf{x} = -k \mathbf{a}^{*\top} \mathbf{x} \geq 0$.

Hence “ \Leftarrow ” holds.

“ \Rightarrow ”: If $\mathbf{a}^* = 0$, \mathbf{a} must be a zero vector, otherwise let $\mathbf{x} = \mathbf{a}$, then $-\mathbf{a}^\top \mathbf{x} < 0 \implies \perp$. If \mathbf{a}^* is not a zero vector, consider two cases below:

- a) If $\mathbf{a} \neq k \mathbf{a}^*$ where $k \geq 0$. Let $\mathbf{x} = \frac{\|\mathbf{a}^*\|}{\|\mathbf{a}\|} \cdot \mathbf{a} - \mathbf{a}^*$, then \mathbf{x} is not a zero vector, and

$$\left. \begin{aligned} \mathbf{a}^{*\top} \mathbf{x} &= \|\mathbf{a}^*\|^2 (\cos \theta - 1) < 0, \\ \mathbf{a}^\top \mathbf{x} &= \|\mathbf{a}\| \|\mathbf{a}^*\| (1 - \cos \theta) > 0 \end{aligned} \right\} \implies (\mathbf{a}^{*\top} \mathbf{x})^+ - \mathbf{a}^\top \mathbf{x} < 0 \implies \perp$$

where θ denotes the angle between \mathbf{a} and \mathbf{a}^* .

- b) If $\mathbf{a} = k \mathbf{a}^*$ where $k > 1$, then let $\mathbf{x} = \mathbf{a}^*$ we have

$$(\mathbf{a}^{*\top} \mathbf{x})^+ - \mathbf{a}^\top \mathbf{x} = (1 - k) \|\mathbf{a}^*\|^2 < 0 \implies \perp.$$

Hence $\mathbf{a} = k \mathbf{a}^*$ where $0 \leq k \leq 1$ if \mathbf{a}^* is not zero, and thereby “ \Rightarrow ” holds. \square

I QP CONVEXITY

The LPs in the main paper are trivially convex. So in this appendix, we only justify the convexity of our QPs: We first prove the convexity of single-sample objectives, then the convexity of the empirical objectives with nonparametric estimation, i.e. \hat{G}_1^{NPE} and \hat{G}_2^{NPE} is obtained by the convexity of convex function summations.

J.1 LAYER 2

In this subsection, we prove the strong consistency of the layer 2 estimator.

J.1.1 OBJECTIVE MINIMIZER SPACE ESTIMATOR

For simplicity of notation, we use $\hat{\mathbb{S}}_n$ to denote our layer 2 QP/LP solution space by n random samples⁸ as a random set

$$\hat{\mathbb{S}}_n := \{\mathbf{C} \in \mathbb{R}^{d \times m} : \mathbf{C}\mathbf{y}^{(i)} - \mathbf{x}^{(i)} \geq 0, \forall i \in [n]\} \quad (26)$$

and \mathbb{S}^* to denote the value space of \mathbf{C} in Lem. 4.2 which minimizes layer 2 objective functional

$$\mathbb{S}^* := \{\mathbf{C} \in \mathbb{R}^{d \times m} : \mathbf{C}\mathbf{y} - \mathbf{x} \geq 0, \forall \mathbf{x} \in \mathbb{R}^d \text{ and its corresponding output } \mathbf{y}\}. \quad (27)$$

For further use, we name $\hat{\mathbb{P}}_n$ as the set of n sampled inputs which define $\hat{\mathbb{S}}_n$, i.e. $\hat{\mathbb{P}}_n := \{\mathbf{x}^{(i)}\}_{i \in [n]}$ where each $\mathbf{x}^{(i)}$ is the same random variable in Eq. 26. By the definitions above, we describe the strong consistency of our QP/LP as Lem. J.1.

Lemma J.1 (QP/LP strong consistency, layer 2). $D_{\text{H}}(\hat{\mathbb{S}}_n, \mathbb{S}^*) \xrightarrow{a.s.} 0$ as $n \rightarrow \infty$.

Proof. First we prove that $D_{\text{H}}(\hat{\mathbb{S}}_n, \mathbb{S}^*) \xrightarrow{p} 0$ as $n \rightarrow \infty$. Recall Thm. G.1, $\mathbf{C}\mathbf{y} - \mathbf{x} \geq 0$ only if $\mathbf{C}\mathbf{B}^* = \text{diag}(\mathbf{k})$. We inherit the notation as defining $\mathbf{D} = \mathbf{C}\mathbf{B}^*$. Theorem G.1 is based on Lem. H.1, and we prove them by raising points that show contradiction, i.e. violate the inequality that holds in a pointwise fashion:

- a) In the proof of Lem. H.1, we use d points: $-e^{(i)}$, for $i \in [d]$, to make $\sum_{k=1}^d D_{i,k} \left[\left(\mathbf{A}_{k,:}^* \cdot \mathbf{x} \right)^+ + x_k \right] < x_i$, so that $D_{i,j}$ ($i \neq j$) cannot be positive; and another d points: $(\mathbf{A}^*)^{-1}e^{(i)}$, to show $\sum_{k=1}^d D_{i,k} \left| \mathbf{A}_{k,:}^* \cdot \mathbf{x} \right| < 0$, so that $D_{i,j}$ ($i \neq j$) cannot be negative. For each $-e^{(i)}$ we use here, since the violations follow strict inequalities, we know there exists a neighborhood of $-e^{(i)}$, $\mathbb{N}_i = \mathbb{N}(-e^{(i)})$, such that $\forall \mathbf{z} \in \mathbb{N}_i$, $\sum_{k=1}^d D_{i,k} \left[\left(\mathbf{A}_{k,:}^* \cdot \mathbf{z} \right)^+ + z_k \right] < z_i$. We can similarly find such neighborhood of each $(\mathbf{A}^*)^{-1}e^{(i)}$ that the strict inequality holds within the neighborhood respectively. We index them as \mathbb{N}_{d+1} to \mathbb{N}_{2d} .
- b) In the proof of Thm. G.1, we further construct d points: for each $i \in [d]$, we take a point x such that $x_i > 0 \wedge x_j \leq -\frac{A_{i,i}^*}{A_{i,j}^*}x_i \wedge x_k \leq 0$, where $k \neq i, j$. This counterexample shows $[\mathbf{C}\mathbf{y} - \mathbf{x}]_i < 0$, and eliminates the possibility of $0 < D_{i,i} < 1$ when $\mathbf{A}_{i,:}^*$ is not a scale transformation. We can similarly find neighborhood of each point and index them as \mathbb{N}_{2d+1} to \mathbb{N}_{3d} . Note that we omit some cases in the proof of Thm. G.1, because the first $2d$ points are sufficient to use in those cases to show contradiction.

In the sampling procedure, if we sample at least one point in each neighborhood \mathbb{N}_i , Thm. G.1 assures the solution we get $\hat{\mathbb{C}}_n$ would lie in the true optimal set \mathbb{S}^* . The probability that the sampling procedure ‘‘omits’’ any of the neighborhoods is

$$\begin{aligned} P\left(\hat{\mathbb{P}}_n \cap \mathbb{N}_1 = \emptyset \text{ or } \hat{\mathbb{P}}_n \cap \mathbb{N}_2 = \emptyset \text{ or } \dots \text{ or } \hat{\mathbb{P}}_n \cap \mathbb{N}_{3d} = \emptyset\right) &\leq \sum_{i=1}^{3d} P\left(\hat{\mathbb{P}}_n \cap \mathbb{N}_i = \emptyset\right) \\ &\leq 3d \left[1 - \min_{i \in [3d]} P(\mathbb{N}_i)\right]^n \end{aligned} \quad (28)$$

Since the measure on each neighborhood $P(\mathbb{N}_i) = \int_{\mathbf{x} \in \mathbb{N}_i} p(\mathbf{x}) > 0$,

$$P\left(\hat{\mathbb{P}}_n \cap \mathbb{N}_1 = \emptyset \text{ or } \hat{\mathbb{P}}_n \cap \mathbb{N}_2 = \emptyset \text{ or } \dots \text{ or } \hat{\mathbb{P}}_n \cap \mathbb{N}_{3d} = \emptyset\right) \rightarrow 0, \text{ as } n \rightarrow \infty. \quad (29)$$

⁸Here, we take samples as random variables for empirical analysis.

Here we obtain $D_H(\hat{\mathbb{S}}_n, \mathbb{S}^*) \xrightarrow{P} 0$ as $n \rightarrow \infty$. Now we take the infinite sum over the both sides of Eq. 28

$$\begin{aligned} & \sum_{n \in [\infty]} P \left(\hat{\mathbb{P}}_n \cap \mathbb{N}_1 = \emptyset \text{ or } \hat{\mathbb{P}}_n \cap \mathbb{N}_2 = \emptyset \text{ or } \dots \text{ or } \hat{\mathbb{P}}_n \cap \mathbb{N}_{3d} = \emptyset \right) \\ & \leq 3d \sum_{n \in [\infty]} \left[1 - \min_{i \in [3d]} P(\mathbb{N}_i) \right]^n = 3d \left[\frac{1}{\min_{i \in [3d]} P(\mathbb{N}_i)} - 1 \right] < +\infty. \end{aligned}$$

By the Borel-Cantelli lemma (Borel, 1909), $D_H(\hat{\mathbb{S}}_n, \mathbb{S}^*) \xrightarrow{\text{a.s.}} 0$ as $n \rightarrow \infty$. \square

J.1.2 SCALE FACTOR ESTIMATOR

To avoid ambiguity, we use n_{sf} to denote the number of samples used in Alg. 4. The samples pairs are $\{(\mathbf{x}^{(i)}, \mathbf{y}^{(i)})\}_{i=1}^{n_{\text{sf}}}$. Without loss of generality, the following discussion focuses on some fixed index $j \in [d]$. In Alg. 4, we plug in our estimator $\hat{\mathbf{C}}_n$ and use LR to estimate k_j given that $\mathbf{x}_j^{(i)} < 0$

$$\mathbf{k}_{n_{\text{sf}}}(\hat{\mathbf{C}}_n) = \arg \min_k \frac{1}{2n_{\text{sf}}} \sum_{i \in [n_{\text{sf}}]} \left\| [\hat{\mathbf{C}}_n \mathbf{y}^{(i)}]_j - k \mathbf{x}_j^{(i)} \right\|^2 = \frac{\sum_{i \in [n_{\text{sf}}]} \mathbf{x}_j^{(i)} [\hat{\mathbf{C}}_n \mathbf{y}^{(i)}]_j}{\sum_{i \in [n_{\text{sf}}]} \left(\mathbf{x}_j^{(i)} \right)^2}. \quad (30)$$

We first give the strong consistency of layer 2 scale factor estimator for as $n_{\text{sf}} \rightarrow \infty$, as described in Lem. J.2.

Lemma J.2 (scale factor estimator strong consistency, layer 2). *Suppose \mathbf{A}^* is a scale transformation w.r.t. the j -th component, and $\mathbf{k}_{n_{\text{sf}}}(\hat{\mathbf{C}}_n)$ is the n_{sf} -sample estimator of k_j via LR: Eq. 30 given that $\mathbf{x}_j^{(i)} < 0$. Define sets*

$$\mathbb{U}_{n_{\text{sf}}, n} := \{ \mathbf{k}_{n_{\text{sf}}}(\mathbf{C}) : \mathbf{C} \in \hat{\mathbb{S}}_n \}, \text{ and } \mathbb{U}^* := \left[\frac{1}{1 + A_{j,j}^*}, 1 \right]. \quad (31)$$

Then $\lim_{n_{\text{sf}} \rightarrow \infty} \lim_{n \rightarrow \infty} D_H(\mathbb{U}_{n_{\text{sf}}, n}, \mathbb{U}^*) \stackrel{\text{a.s.}}{=} 0$.

Proof. Following our notation, Thm. G.1 and Thm. G.2 ensure that if \mathbf{A}^* is a scale transformation w.r.t. the j -th component, for any \mathbf{C} belonging to the true optimal set \mathbb{S}^* , $\mathbf{k}_{n_{\text{sf}}}(\mathbf{C}) \in \left[\frac{1}{1 + A_{j,j}^*}, 1 \right]$. And the ‘‘iff’’ statement strengthens that $\mathbb{U}^* = \{ \mathbf{k}_{n_{\text{sf}}}(\mathbf{C}) : \mathbf{C} \in \mathbb{S}^* \}$ for any $n_{\text{sf}} \in \mathbb{Z}^+$. Note that since $\mathbb{S}^* \subset \hat{\mathbb{S}}_n$, we have $\mathbb{U}^* \subset \mathbb{U}_{n_{\text{sf}}, n}$. We only need to prove $D(\mathbb{U}_{n_{\text{sf}}, n}, \mathbb{U}^*) \xrightarrow{\text{a.s.}} 0$ as $n \rightarrow \infty$.

$\forall \hat{\mathbf{C}}_n \in \hat{\mathbb{S}}_n$ and $\forall \mathbf{C} \in \mathbb{S}^*$,

$$\begin{aligned} & \left| \mathbf{k}_{n_{\text{sf}}}(\hat{\mathbf{C}}_n) - \mathbf{k}_{n_{\text{sf}}}(\mathbf{C}) \right| \\ & = \frac{1}{\sum_{i \in [n_{\text{sf}}]} \left(\mathbf{x}_j^{(i)} \right)^2} \left| \sum_{i \in [n_{\text{sf}}]} \mathbf{x}_j^{(i)} \left[(\mathbf{e}^{(j)})^\top (\hat{\mathbf{C}}_n - \mathbf{C}) \mathbf{B}^* \left[(\mathbf{A}^* \mathbf{x}^{(i)})^+ + \mathbf{x}^{(i)} \right] \right] \right| \\ & \leq \frac{1}{\sum_{i \in [n_{\text{sf}}]} \left(\mathbf{x}_j^{(i)} \right)^2} \left[\sum_{i \in [n_{\text{sf}}]} \left| \mathbf{x}_j^{(i)} \right| \left\| (\mathbf{e}^{(j)})^\top (\hat{\mathbf{C}}_n - \mathbf{C}) \mathbf{B}^* \right\|_2 \left(\left\| \mathbf{A}^* \mathbf{x}^{(i)} \right\|_2 + \left\| \mathbf{x}^{(i)} \right\|_2 \right) \right] \\ & \leq \frac{1}{\sum_{i \in [n_{\text{sf}}]} \left(\mathbf{x}_j^{(i)} \right)^2} \left[\sum_{i \in [n_{\text{sf}}]} \left| \mathbf{x}_j^{(i)} \right| \left\| \hat{\mathbf{C}}_n - \mathbf{C} \right\|_{\text{F}} \left\| \mathbf{B}^* \right\|_{\text{F}} \left(\left\| \mathbf{A}^* \right\|_{\text{F}} + 1 \right) \left\| \mathbf{x}^{(i)} \right\|_2 \right] \\ & = \frac{\sum_{i \in [n_{\text{sf}}]} \left[\left| \mathbf{x}_j^{(i)} \right| \left\| \mathbf{B}^* \right\|_{\text{F}} \left(\left\| \mathbf{A}^* \right\|_{\text{F}} + 1 \right) \left\| \mathbf{x}^{(i)} \right\|_2 \right]}{\sum_{i \in [n_{\text{sf}}]} \left(\mathbf{x}_j^{(i)} \right)^2} \left\| \hat{\mathbf{C}}_n - \mathbf{C} \right\|_{\text{F}}. \end{aligned}$$

Then

$$\sup_{\hat{\mathbf{C}}_n \in \hat{\mathbb{S}}_n} \inf_{\mathbf{C} \in \mathbb{S}^*} \left| \mathbf{k}_{n_{sf}}(\hat{\mathbf{C}}_n) - \mathbf{k}_{n_{sf}}(\mathbf{C}) \right| \quad (32)$$

$$\leq \frac{\sum_{i \in [n_{sf}]} \left[\left| \mathbf{x}_j^{(i)} \right| \|\mathbf{B}^*\|_F (\|\mathbf{A}^*\|_F + 1) \|\mathbf{x}^{(i)}\|_2 \right]}{\sum_{i \in [n_{sf}]} \left(\mathbf{x}_j^{(i)} \right)^2} \sup_{\hat{\mathbf{C}}_n \in \hat{\mathbb{S}}_n} \inf_{\mathbf{C} \in \mathbb{S}^*} \left\| \hat{\mathbf{C}}_n - \mathbf{C} \right\|_F \quad (33)$$

which implies

$$D(\mathbb{U}_{n_{sf}, n}, \mathbb{U}) \leq \frac{\sum_{i \in [n_{sf}]} \left[\left| \mathbf{x}_j^{(i)} \right| \|\mathbf{B}^*\|_F (\|\mathbf{A}^*\|_F + 1) \|\mathbf{x}^{(i)}\|_2 \right]}{\sum_{i \in [n_{sf}]} \left(\mathbf{x}_j^{(i)} \right)^2} D(\hat{\mathbb{S}}_n, \mathbb{S}^*). \quad (34)$$

Take the $n_{sf} \rightarrow \infty$ limit over both sides of Eq. 34. With the strong law of large numbers⁹ we have

$$\lim_{n_{sf} \rightarrow \infty} D(\mathbb{U}_{n_{sf}, n}, \mathbb{U}) \leq \frac{\mathbb{E} \left[\left| \mathbf{x}_j^{(i)} \right| \|\mathbf{B}^*\|_F (\|\mathbf{A}^*\|_F + 1) \|\mathbf{x}^{(i)}\|_2 \right]}{\mathbb{E} \left[\mathbf{x}_j^{(i)} \right]^2} D(\hat{\mathbb{S}}_n, \mathbb{S}^*), \text{ w.p. 1.} \quad (35)$$

Since $D(\hat{\mathbb{S}}_n, \mathbb{S}^*) \xrightarrow{\text{a.s.}} 0$ as $n \rightarrow \infty$, we have $D(\mathbb{U}_{n_{sf}, n}, \mathbb{U}^*) \xrightarrow{\text{a.s.}} 0$ as $n \rightarrow \infty$ then $n_{sf} \rightarrow \infty$. \square

J.1.3 LAYER 2 WEIGHTS ESTIMATOR

In Alg. 5, we solve \mathbf{B} via LR. Let $\hat{\mathbf{z}}^{(i)} = \text{diag}^{-1}(\hat{\mathbf{k}}) \cdot \hat{\mathbf{C}}_n \mathbf{y}^{(i)} \in \mathbb{R}^d$, where $\hat{\mathbf{k}}$ is obtained through Alg. 4 with input $\hat{\mathbf{C}}_n$. Assume we are using sample size of n_w to do the LR. The optimization problem is

$$\min_{\mathbf{B}} \sum_{i \in [n_w]} \left\| \mathbf{y}^{(i)} - \mathbf{B} \hat{\mathbf{z}}^{(i)} \right\|^2 \quad (36)$$

Now we present the strong consistency of layer 2 estimator, as described in Thm. J.3.

Theorem J.3 (strong consistency, layer 2). *Suppose $\hat{\mathbf{B}}_{n_{sf}}$ is the solution to Eq. 36. Then $\hat{\mathbf{B}}_{n_w} \xrightarrow{\text{a.s.}} \mathbf{B}^*$ as $n, n_{sf}, n_w \rightarrow \infty$.*

Proof. Let $\boldsymbol{\beta}$ denote $\text{vec } \mathbf{B}$ (flattening \mathbf{B} into a vector), then $\mathbf{B} \hat{\mathbf{z}}^{(i)} = \left(\hat{\mathbf{z}}^{(i)} \right)^\top \otimes \mathbf{I}_m$. Here the operation \otimes denotes the Kronecker product. Then we can define an equivalent optimization problem

$$\min_{\boldsymbol{\beta}} \frac{1}{2n_w} \sum_{i \in [n_w]} \left\| \mathbf{y}^{(i)} - \left[\left(\hat{\mathbf{z}}^{(i)} \right)^\top \otimes \mathbf{I}_m \right] \boldsymbol{\beta} \right\|^2. \quad (37)$$

Take the derivatives of $\boldsymbol{\beta}$, we obtain

$$-2 \sum_{i \in [n_w]} \left[\left[\left(\hat{\mathbf{z}}^{(i)} \right)^\top \otimes \mathbf{I}_m \right]^\top \left(\mathbf{y}^{(i)} - \left[\left(\hat{\mathbf{z}}^{(i)} \right)^\top \otimes \mathbf{I}_m \right] \boldsymbol{\beta} \right) \right] = 0 \quad (38)$$

⁹Here we assume that the Kolmogorov's strong law assumption on moments (Sen & Singer, 1994) is met as is commonly done in empirical analysis.

Then the optimal solution $\hat{\beta}_{n_w}$ of this optimization can be written in closed form:

$$\begin{aligned}
\hat{\beta}_{n_w} &= \left[\sum_{i \in [n_w]} \left[\left(\hat{\mathbf{z}}^{(i)} \right)^\top \otimes \mathbf{I}_m \right]^\top \left[\left(\hat{\mathbf{z}}^{(i)} \right)^\top \otimes \mathbf{I}_m \right] \right]^{-1} \left(\sum_{i \in [n_w]} \left[\left(\hat{\mathbf{z}}^{(i)} \right)^\top \otimes \mathbf{I}_m \right]^\top \mathbf{y}^{(i)} \right) \\
&= \left[\sum_{i \in [n_w]} \left[\hat{\mathbf{z}}^{(i)} \otimes \mathbf{I}_m \right] \left[\left(\hat{\mathbf{z}}^{(i)} \right)^\top \otimes \mathbf{I}_m \right] \right]^{-1} \left(\sum_{i \in [n_w]} \left[\hat{\mathbf{z}}^{(i)} \otimes \mathbf{I}_m \right] \mathbf{y}^{(i)} \right) \\
&= \left[\sum_{i \in [n_w]} \left[\hat{\mathbf{z}}^{(i)} \left(\hat{\mathbf{z}}^{(i)} \right)^\top \right] \otimes \mathbf{I}_m \right]^{-1} \left(\left[\sum_{i \in [n_w]} \hat{\mathbf{z}}^{(i)} \otimes \mathbf{I}_m \right] \mathbf{y}^{(i)} \right) \\
&= \left[\left[\sum_{i \in [n_w]} \hat{\mathbf{z}}^{(i)} \left(\hat{\mathbf{z}}^{(i)} \right)^\top \right] \otimes \mathbf{I}_m \right]^{-1} \left(\left[\sum_{i \in [n_w]} \hat{\mathbf{z}}^{(i)} \otimes \mathbf{I}_m \right] \mathbf{y}^{(i)} \right)
\end{aligned}$$

We inherit the notation from the last two subsections. By Lem. J.1, $d(\hat{\mathbf{C}}_n, \mathbb{S}^*) \xrightarrow{\text{a.s.}} 0$ as $n \rightarrow \infty$. Thus $\forall \varepsilon > 0, \exists N$ such that $\forall n \geq N, d(\hat{\mathbf{C}}_n, \mathbb{S}^*) \leq \varepsilon$ w.p. 1, i.e. $\exists \mathbf{C}_n \in \mathbb{S}^*$ s.t. $d(\hat{\mathbf{C}}_n, \mathbf{C}_n) \leq \varepsilon$ w.p. 1. Then by Lem. J.2, $\exists K > 0$, for ε that is small enough, $\exists N_{\text{sf}}$ such that $\forall n_{\text{sf}} \geq N_{\text{sf}}$, we have $|\mathbf{k}_{n_{\text{sf}}}(\hat{\mathbf{C}}_n) - \mathbf{k}_{n_{\text{sf}}}(\mathbf{C}_n)| \leq K\varepsilon$ w.p.1. For simplicity, we omit the under-script n_{sf} of $\mathbf{k}_{n_{\text{sf}}}$ in the following discussion. In fact,

$$\begin{aligned}
\left| \hat{\mathbf{z}}_j^{(i)} - \mathbf{z}_j^{(i)} \right| &= \left| \frac{1}{\mathbf{k}(\hat{\mathbf{C}}_n)} \left(\mathbf{e}^{(i)} \right)^\top \hat{\mathbf{C}}_n \mathbf{y}^{(i)} - \frac{1}{\mathbf{k}(\mathbf{C}_n)} \left(\mathbf{e}^{(i)} \right)^\top \mathbf{C}_n \mathbf{y}^{(i)} \right| \\
&= \left| \frac{1}{\mathbf{k}(\hat{\mathbf{C}}_n) \mathbf{k}(\mathbf{C}_n)} \left[\mathbf{k}(\mathbf{C}_n) \left(\mathbf{e}^{(i)} \right)^\top \hat{\mathbf{C}}_n - \mathbf{k}(\hat{\mathbf{C}}_n) \left(\mathbf{e}^{(i)} \right)^\top \mathbf{C}_n \right] \mathbf{y}^{(i)} \right| \\
&\leq \left| \frac{1}{\mathbf{k}(\hat{\mathbf{C}}_n) \mathbf{k}(\mathbf{C}_n)} \right| \left| \mathbf{k}(\mathbf{C}_n) \left(\mathbf{e}^{(i)} \right)^\top \hat{\mathbf{C}}_n \mathbf{y}^{(i)} - \mathbf{k}(\mathbf{C}_n) \left(\mathbf{e}^{(i)} \right)^\top \mathbf{C}_n \mathbf{y}^{(i)} \right| \\
&\quad + \left| \frac{1}{\mathbf{k}(\hat{\mathbf{C}}_n) \mathbf{k}(\mathbf{C}_n)} \right| \left| \mathbf{k}(\mathbf{C}_n) \left(\mathbf{e}^{(i)} \right)^\top \mathbf{C}_n \mathbf{y}^{(i)} - \mathbf{k}(\hat{\mathbf{C}}_n) \left(\mathbf{e}^{(i)} \right)^\top \mathbf{C}_n \mathbf{y}^{(i)} \right| \\
&\leq \frac{(1 + \mathbf{A}_{j,j}^*)^2}{1 - K\varepsilon (1 + \mathbf{A}_{j,j}^*)} \left(\left\| \mathbf{y}^{(i)} \right\| \left\| \hat{\mathbf{C}}_n - \mathbf{C}_n \right\|_{\text{F}} + \left| \mathbf{k}(\hat{\mathbf{C}}_n) - \mathbf{k}(\mathbf{C}_n) \right| \left\| \mathbf{C}_n \mathbf{B}^* \right\|_{\text{F}} (\left\| \mathbf{A}^* \right\|_{\text{F}} + 1) \left\| \mathbf{x}^{(i)} \right\| \right) \\
&\leq \frac{(1 + \mathbf{A}_{j,j}^*)^2}{1 - K\varepsilon (1 + \mathbf{A}_{j,j}^*)} \left(\left\| \mathbf{y}^{(i)} \right\| + dK (\left\| \mathbf{A}^* \right\|_{\text{F}} + 1) \left\| \mathbf{x}^{(i)} \right\| \right) \varepsilon \tag{39}
\end{aligned}$$

Then

$$\begin{aligned}
&\left\| \hat{\mathbf{z}}^{(i)} \left(\hat{\mathbf{z}}^{(i)} \right)^\top - \mathbf{z}^{(i)} \left(\mathbf{z}^{(i)} \right)^\top \right\|_{\text{F}} \\
&= \left\| \hat{\mathbf{z}}^{(i)} \left(\hat{\mathbf{z}}^{(i)} \right)^\top - \mathbf{z}^{(i)} \left(\hat{\mathbf{z}}^{(i)} \right)^\top + \mathbf{z}^{(i)} \left(\hat{\mathbf{z}}^{(i)} \right)^\top - \mathbf{z}^{(i)} \left(\mathbf{z}^{(i)} \right)^\top \right\|_{\text{F}} \\
&\leq \left\| \left[\hat{\mathbf{z}}^{(i)} - \mathbf{z}^{(i)} \right] \left(\hat{\mathbf{z}}^{(i)} \right)^\top \right\|_{\text{F}} + \left\| \mathbf{z}^{(i)} \left[\left(\hat{\mathbf{z}}^{(i)} \right)^\top - \left(\mathbf{z}^{(i)} \right)^\top \right] \right\|_{\text{F}} \\
&\leq \left\| \left[\hat{\mathbf{z}}^{(i)} - \mathbf{z}^{(i)} \right] \left[\hat{\mathbf{z}}^{(i)} - \mathbf{z}^{(i)} \right]^\top \right\|_{\text{F}} + 2 \left\| \mathbf{z}^{(i)} \left[\left(\hat{\mathbf{z}}^{(i)} \right)^\top - \left(\mathbf{z}^{(i)} \right)^\top \right] \right\|_{\text{F}} \\
&\leq \sum_{j \in [d]} \left(\hat{\mathbf{z}}_j^{(i)} - \mathbf{z}_j^{(i)} \right)^2 + 2 \left\| \mathbf{z}^{(i)} \right\| \left\| \hat{\mathbf{z}}^{(i)} - \mathbf{z}^{(i)} \right\|
\end{aligned}$$

It follows that $\left\| \hat{\mathbf{z}}^{(i)} (\hat{\mathbf{z}}^{(i)})^\top - \mathbf{z}^{(i)} (\mathbf{z}^{(i)})^\top \right\|_{\text{F}}$ is also bounded by $\mathcal{O}(\varepsilon)$. With similar techniques we can prove

$$\left\| \sum_{i \in [n_w]} \hat{\mathbf{z}}^{(i)} (\hat{\mathbf{z}}^{(i)})^\top - \sum_{i \in [n_w]} \mathbf{z}^{(i)} (\mathbf{z}^{(i)})^\top \right\|_{\text{F}} \leq \mathcal{O}(\varepsilon) \quad (40)$$

and

$$\left\| \hat{\mathbf{z}}^{(i)} \otimes \mathbf{I}_m - \mathbf{z}^{(i)} \otimes \mathbf{I}_m \right\|_{\text{F}} \leq \mathcal{O}(\varepsilon) \quad (41)$$

Denote $\left[\sum_{i \in [n_w]} \mathbf{z}^{(i)} (\mathbf{z}^{(i)})^\top \right]$ as \mathbf{P} and $\sum_{i \in [n_w]} \mathbf{z}^{(i)}$ as \mathbf{Q} . Substitute $\mathbf{z}^{(i)}$ with $\hat{\mathbf{z}}^{(i)}$ in the above expression we have $\hat{\mathbf{P}}$ and $\hat{\mathbf{Q}}$. Hence,

$$\begin{aligned} \left\| \hat{\beta}_{n_w} - \beta^* \right\|_{\text{F}} &= \left\| \left[\hat{\mathbf{P}} \otimes \mathbf{I}_m \right]^{-1} \left(\left[\hat{\mathbf{Q}} \otimes \mathbf{I}_m \right] \mathbf{y}^{(i)} \right) - \left[\mathbf{P} \otimes \mathbf{I}_m \right]^{-1} \left(\left[\mathbf{Q} \otimes \mathbf{I}_m \right] \mathbf{y}^{(i)} \right) \right\|_{\text{F}} \\ &\leq \left\| \left[\hat{\mathbf{P}} \otimes \mathbf{I}_m \right]^{-1} \left(\left[\hat{\mathbf{Q}} \otimes \mathbf{I}_m \right] \mathbf{y}^{(i)} \right) - \left[\hat{\mathbf{P}} \otimes \mathbf{I}_m \right]^{-1} \left(\left[\mathbf{Q} \otimes \mathbf{I}_m \right] \mathbf{y}^{(i)} \right) \right\|_{\text{F}} \\ &\quad + \left\| \left[\hat{\mathbf{P}} \otimes \mathbf{I}_m \right]^{-1} \left(\left[\mathbf{Q} \otimes \mathbf{I}_m \right] \mathbf{y}^{(i)} \right) - \left[\mathbf{P} \otimes \mathbf{I}_m \right]^{-1} \left(\left[\mathbf{Q} \otimes \mathbf{I}_m \right] \mathbf{y}^{(i)} \right) \right\|_{\text{F}} \\ &= \left\| \left[\hat{\mathbf{P}} \otimes \mathbf{I}_m \right]^{-1} \right\|_{\text{F}} \left\| \left(\left[\hat{\mathbf{Q}} - \mathbf{Q} \right] \otimes \mathbf{I}_m \right) \mathbf{y}^{(i)} \right\|_{\text{F}} + \left\| \left[\hat{\mathbf{P}} \otimes \mathbf{I}_m \right]^{-1} - \left[\mathbf{P} \otimes \mathbf{I}_m \right]^{-1} \right\|_{\text{F}} \\ &\quad \left\| \left[\mathbf{Q} \otimes \mathbf{I}_m \right] \mathbf{y}^{(i)} \right\|_{\text{F}} \end{aligned}$$

In the first part, by triangle inequality,

$$\left\| \left[\hat{\mathbf{P}} \otimes \mathbf{I}_m \right]^{-1} \right\|_{\text{F}} \leq \left\| \left[\hat{\mathbf{P}} \otimes \mathbf{I}_m \right]^{-1} - \left[\mathbf{P} \otimes \mathbf{I}_m \right]^{-1} \right\|_{\text{F}} + \left\| \left[\mathbf{P} \otimes \mathbf{I}_m \right]^{-1} \right\|_{\text{F}} \quad (42)$$

So we only need to prove $\left\| \left[\hat{\mathbf{P}} \otimes \mathbf{I}_m \right]^{-1} - \left[\mathbf{P} \otimes \mathbf{I}_m \right]^{-1} \right\|_{\text{F}} \leq \mathcal{O}(\varepsilon)$ to claim $\left\| \hat{\beta}_{n_w} - \beta^* \right\|_{\text{F}} \leq \mathcal{O}(\varepsilon)$.

Denote $\hat{\mathbf{P}} - \mathbf{P} = \Delta \mathbf{P}$. From Eq. 39, we know every entry of $\Delta \mathbf{P} \otimes \mathbf{I}_m$ can be bounded by $\mathcal{O}(\varepsilon)$. By simple calculation we have

$$\left(\mathbf{P} + \Delta \mathbf{P} \right)^{-1} = \mathbf{P}^{-1} - \mathbf{P}^{-1} \Delta \mathbf{P} \mathbf{P}^{-1} + \mathcal{O}(\varepsilon^2). \quad (43)$$

Then we have

$$\mathbf{P}^{-1} - \hat{\mathbf{P}}^{-1} = \mathbf{P}^{-1} \Delta \mathbf{P} \mathbf{P}^{-1} + \mathcal{O}(\varepsilon^2) = \mathcal{O}(\varepsilon). \quad (44)$$

□

J.2 LAYER 1

In this subsection, we justify the strong consistency of layer 1 objective functional minimizer estimator in detail, i.e. the layer 1 QP/LP solution space. We will omit the detailed proof of Alg. 2 line 4 to 7 strong consistency since it is similar to the proof of Lem. J.2. Additionally, we also omit the strong consistency of the $\mathbf{x} \mapsto (\mathbf{A}^* \mathbf{x})^+$ function estimator because it can be directly obtained by Lem. J.1 and J.2 and the continuous mapping theorem.

We use a new optimization problem equivalent to the optimization of G_1 . Before that, we first define the equivalence between two optimization problems as follows.

Definition J.3. Let opt1 and opt2 be two optimization problems, with f_1, f_2 as the respective objective functions. Then opt1 and opt2 are said to be equivalent if given a feasible solution to opt1 , namely \mathbf{x}_1 , a feasible solution to opt2 is uniquely corresponded, namely \mathbf{x}_2 , such that $f_1(\mathbf{x}_1) = f_2(\mathbf{x}_2)$, and vice versa.

The new optimization problem and its equivalence to the optimization of G_1 is described in Lem. J.4.

Lemma J.4. *The optimization of G_1 (Eq. 6) is equivalent to*

$$\min_{\mathbf{A}} f(\mathbf{A}) = \frac{1}{2} \mathbb{E}_{\mathbf{x}} \left[\left\| (\mathbf{A} \mathbf{x} - \mathbf{h})^+ \right\|^2 \right]. \quad (45)$$

Proof. To see this, suppose \mathbf{A}_1 is one optimal solution to Eq. 45, then we can construct $r_1(\mathbf{x}) = (\mathbf{A}_1\mathbf{x} - \mathbf{h})^+$ so that $G_1(\mathbf{A}_1, r_1) = f(\mathbf{A}_1)$ and the optimality implies

$$\min_{\mathbf{A}, r} G_1(\mathbf{A}, r) \leq \min_{\mathbf{A}} f(\mathbf{A}).$$

On the other hand, suppose (\mathbf{A}_2, r_2) is an optimum of G_1 . Let $r_3(\mathbf{x}) = (\mathbf{h} - \mathbf{A}_2\mathbf{x})^+$, then $\forall \mathbf{x} \in \mathbb{R}^d$ and \mathbf{h} be the corresponding hidden output, if $[\mathbf{h} - \mathbf{A}_2\mathbf{x}]_j \geq 0$, then $[r_3(\mathbf{x}) + \mathbf{A}_2\mathbf{x} - \mathbf{h}]_j = 0$, otherwise $[r_3(\mathbf{x}) + \mathbf{A}_2\mathbf{x} - \mathbf{h}]_j^2 = [\mathbf{A}_2\mathbf{x} - \mathbf{h}]_j^2 \leq [h_2(\mathbf{x}) + \mathbf{A}_2\mathbf{x} - \mathbf{h}]_j^2$ since h_2 is nonnegative. So that we know that

$$\min_{\mathbf{A}, r} G_1(\mathbf{A}, r) = G_1(\mathbf{A}_2, r_2) = G_1(\mathbf{A}_2, r_3) = f(\mathbf{A}_2)$$

From the optimality, we further have

$$\min_{\mathbf{A}, r} G_1(\mathbf{A}, r) \geq \min_{\mathbf{A}} f(\mathbf{A})$$

From the simple calculation above, we can see that one optimal solution to Eq. 45 has a one-to-one correspondence to an optimal solution to G_1 . \square

Similarly, the empirical version of the two problems are equivalent, which indicates their consistency in the empirical estimation being equivalent. In the following, we justify the strong consistency of empirical Eq. 45 instead of G_1

$$\min_{\mathbf{A}} \hat{f}_n(\mathbf{A}) = \frac{1}{2n} \sum_{i \in [n]} \left\| \left(\mathbf{A}\mathbf{x}^{(i)} - \mathbf{h}^{(i)} \right)^+ \right\|^2. \quad (46)$$

Denote $\mathbb{T}^* := \{\text{diag}(\mathbf{k}) \cdot \mathbf{A}^* \mid 0 \leq k_j \leq 1, j \in [d]\}$ as the true optimal solution set, and $\hat{\mathbb{T}}_n$ as the optimal solution set corresponding to the n -sample problem. In the following, we justify four conditions in a row that hold for f to derive the strong consistency of its optimal solution estimator.

Lemma J.5. *Let $\hat{\mathbb{T}}_{n'}$ be the layer 1 QP/LP solution space by n' samples. Then there exists a compact set \mathbb{C} determined by \mathbf{A}^* , namely $\mathbb{C}(\mathbf{A}^*)$, s.t. $\hat{\mathbb{T}}_{n'} \subset \mathbb{C}(\mathbf{A}^*)$ w.p. 1 as $n \rightarrow \infty$.*

Proof. $\forall l \in [d]$, let \mathbf{a}^* be the l -th row of \mathbf{A}^* , and $\hat{\mathbf{a}}_l$ be the l -th row of $\hat{\mathbf{A}}_{n'}$. We'd like first to prove that the set

$$\hat{\mathbb{T}}_{n'}^l = \{\mathbf{a}_l : \mathbf{a}_l \text{ is the } l\text{-th row of } \mathbf{A}, \text{ where } \mathbf{A} \in \hat{\mathbb{T}}_{n'}\} \quad (47)$$

is compact w.p.1.

Suppose $n' > d$, and among the n' samples, we classify them into two folds. To avoid ambiguity, let $\mathbf{u}^{(i)}$ be the points such that $(\mathbf{a}^*)^\top \mathbf{u}^{(i)} > 0$, $i \in [q]$; and $\mathbf{v}^{(j)}$ be the points such that $(\mathbf{a}^*)^\top \mathbf{v}^{(j)} < 0$, $j \in [n - q]$. From the analysis of Thm. 5.1, we have $(\mathbf{a}^*)^\top \mathbf{u}^{(i)} \geq \hat{\mathbf{a}}^\top \mathbf{u}^{(i)}$, $\forall i \in [q]$ and $\hat{\mathbf{a}}^\top \mathbf{v}^{(j)} \leq 0$, $\forall j \in [n - q]$. It follows that we can rewrite $\hat{\mathbb{T}}_{n'}^l$ as a polyhedron

$$\hat{\mathbb{T}}_{n'}^l = \{\mathbf{a} \in \mathbb{R}^d : \mathbf{a}^\top \mathbf{u}^{(i)} \leq (\mathbf{a}^*)^\top \mathbf{u}^{(i)}, \mathbf{a}^\top \mathbf{v}^{(j)} \leq 0\} \quad (48)$$

We are going to show the polyhedron $\hat{\mathbb{T}}_{n'}^l$ is bounded by contradiction. If it is not bounded, then $\exists \mathbf{d} \in \mathbb{R}^d$, $\mathbf{d} \neq \mathbf{0}$ and $\tilde{\mathbf{a}} \in \hat{\mathbb{T}}_{n'}^l$, such that $\forall \lambda > 0$, $\tilde{\mathbf{a}} + \lambda \mathbf{d} \in \hat{\mathbb{T}}_{n'}^l$. Then

$$(\tilde{\mathbf{a}} + \lambda \mathbf{d})^\top \mathbf{u}^{(i)} = \tilde{\mathbf{a}}^\top \mathbf{u}^{(i)} + \lambda \mathbf{d}^\top \mathbf{u}^{(i)} \leq (\mathbf{a}^*)^\top \mathbf{u}^{(i)} \iff \lambda \mathbf{d}^\top \mathbf{u}^{(i)} \leq (\mathbf{a}^*)^\top \mathbf{u}^{(i)} - \tilde{\mathbf{a}}^\top \mathbf{u}^{(i)} \quad (49)$$

similarly,

$$(\tilde{\mathbf{a}} + \lambda \mathbf{d})^\top \mathbf{v}^{(j)} = \tilde{\mathbf{a}}^\top \mathbf{v}^{(j)} + \lambda \mathbf{d}^\top \mathbf{v}^{(j)} \leq 0 \quad (50)$$

From the definition, λ can be arbitrarily big, then $\mathbf{d}^\top \mathbf{u}^{(i)} \leq 0$, $\forall i \in [q]$, and $\mathbf{d}^\top \mathbf{v}^{(j)} \leq 0$, $\forall j \in [n - q]$.

Since we know $\text{span}\{\mathbf{u}^{(i)}\} = \mathbb{R}^d$ w.p. 1, then there \exists some i^* such that $\mathbf{d}^\top \mathbf{u}^{(i^*)} < 0$, w.p. 1. (Otherwise if $\mathbf{d}^\top \mathbf{u}^{(i)} = 0$ for all $i \in [q]$, then either $\text{span}\{\mathbf{u}^{(i)}\} \neq \mathbb{R}^d$ or $\mathbf{d} = \mathbf{0}$.) Under our assumption that, $\hat{\mathbb{T}}_{n'}^l$ is not bounded, we know the following system (w.r.t \mathbf{x}) has a feasible solution w.p. 1

$$\begin{bmatrix} -\mathbf{U} \\ -\mathbf{V} \end{bmatrix} \mathbf{x} \geq 0, \mathbf{x}^\top \mathbf{u}^{(i^*)} < 0 \quad (\text{I})$$

where every row of \mathbf{U} and \mathbf{V} is $(\mathbf{u}^{(i)})^\top$ and $(\mathbf{v}^{(j)})^\top$ respectively. By Farkas' lemma (Farkas, 1902), the system

$$[-\mathbf{U}^\top, -\mathbf{V}^\top] \cdot \mathbf{x} = \mathbf{u}^{(i^*)}, \mathbf{x} \geq 0 \quad (\text{II})$$

is not feasible (w.p. 1). We claim that $\mathbf{u}^{(i^*)}$ lies in the conic hull of $-\mathbf{v}^{(j)}$'s w.p. 1. So that the second system actually has a feasible solution and thus it raises the contradiction.

Denote the conic hull as

$$\mathbb{H} = \left\{ \mathbf{t} \in \mathbb{R}^d : \mathbf{t} = \sum_{j \in [n-q]} \lambda_j (-\mathbf{v}^{(j)}), \lambda_j \geq 0 \text{ for } \forall j \in [n-q] \right\}$$

Now suppose $\mathbf{u}^{(i^*)} \notin \mathbb{H}$, by the supporting hyperplane theorem (Luenberger, 1997), $\exists \mathbf{b} \in \mathbb{R}^d, \mathbf{b} \neq \mathbf{0}$, such that $\mathbf{b}^\top \mathbf{u}^{(i^*)} \leq \mathbf{b}^\top \mathbf{t}$ for $\forall \mathbf{t} \in \mathbb{H}$. Then by definition,

$$-\mathbf{v}^{(j)} \in \{ \mathbf{t} : \mathbf{b}^\top \mathbf{t} \leq \mathbf{b}^\top \mathbf{u}^{(i^*)} \}, \text{ for } \forall j \in [n-q]$$

Denote the hyperplane $\mathbb{J} = \{ \mathbf{t} : \mathbf{b}^\top \mathbf{t} \leq \mathbf{b}^\top \mathbf{u}^{(i^*)} \}$, then

$$P(\mathbf{u}^{(i^*)} \notin \mathbb{H}) \leq P(-\mathbf{v}^{(j)} \in \mathbb{J}, \text{ for } \forall j \in [n-q]) = P(-\mathbf{v}^{(j)} \in \mathbb{J})^{n-q}$$

Since we have in this case a geometric sequence, we know its infinite sum is bounded. By Borel-Cantelli lemma (Borel, 1909), we conclude that $\mathbf{u}^{(i^*)} \in \mathbb{H}$ w.p. 1. Then system II is feasible w.p. 1. So that $\hat{\mathbb{T}}_{n'}^l$ is compact w.p. 1.

Now we prove that there exists a compact set $\mathbb{C}(\mathbf{A}^*)$, s.t. $\hat{\mathbb{T}}_{n'} \subset \mathbb{C}(\mathbf{A}^*)$ w.p. 1 as $n \rightarrow \infty$. Similarly, we focus on the analysis of one row. As discussed above, $\hat{\mathbb{T}}_{n'}^l$ is compact w.p. 1. Let $\hat{\mathbb{W}}_{n'}^1$ be the set of all $\mathbf{u}^{(i)}$ sampled in estimating $\hat{\mathbb{T}}_{n'}^l$, and $\hat{\mathbb{W}}_{n'}^2$ be the set of all $\mathbf{v}^{(j)}$ sampled. Now for another set $\hat{\mathbb{T}}_{n''}$, similarly define sample point sets $\hat{\mathbb{W}}_{n''}^1$ and $\hat{\mathbb{W}}_{n''}^2$. We claim that $\forall \mathbf{u}^{(i)} \in \hat{\mathbb{W}}_{n''}^1, \mathbf{u}^{(i)}$ lies in the conic hull of $\hat{\mathbb{W}}_{n''}^2$. Actually, this part of the proof is very similar to the way we prove $\mathbf{u}^{(i^*)} \in \mathbb{H}$ w.p. 1, so we will omit it here.

$\forall \mathbf{a} \in \hat{\mathbb{T}}_{n''}$, let $\mathbf{x}^{(1)}$ and $\mathbf{x}^{(2)}$ be two different points in $\hat{\mathbb{W}}_{n''}^2$. Then $\mathbf{a}^\top (\lambda_1 \mathbf{x}^{(1)} + \lambda_2 \mathbf{x}^{(2)}) \leq (\mathbf{a}^*)^\top (\lambda_1 \mathbf{x}^{(1)} + \lambda_2 \mathbf{x}^{(2)})$ for $\forall \lambda_1 \geq 0$ and $\lambda_2 \geq 0$. This simple calculation reveals $\mathbf{a}^\top \mathbf{u}^{(i)} \leq (\mathbf{a}^*)^\top \mathbf{u}^{(i)}$ for $\forall \mathbf{u}^{(i)} \in \hat{\mathbb{W}}_{n''}^1$ (from the claim we made). This implies that $\hat{\mathbb{T}}_{n''} \subset \mathbb{C}(\mathbf{A}^*)$ w.p. 1 as $n \rightarrow \infty$, too. \square

Lemma J.6. *The minimizer space of $f(\mathbf{A})$, i.e. $\mathbb{T}^* = \{ \text{diag}(\mathbf{k}) \cdot \mathbf{A}^* \mid 0 \leq k_j \leq 1, j \in [d] \}$, is contained in $\mathbb{C}(\mathbf{A}^*)$.*

Lemma J.7. *$f(\mathbf{A})$ is finite valued and continuous on $\mathbb{C}(\mathbf{A}^*)$.*

Lem. J.6 and J.7 are easily obtained by the formulation of f (see Eq. 45) and Lem. J.5.

Lemma J.8 (uniform a.s. convergence). *$\hat{f}_n(\mathbf{A}) \xrightarrow{\text{a.s.}} f(\mathbf{A})$ as $n \rightarrow \infty$, uniformly in $\mathbf{A} \in \mathbb{C}(\mathbf{A}^*)$.*

Proof. Name single-sample objective $g(\mathbf{x}, \mathbf{A}) = \frac{1}{2} \left\| (\mathbf{A}\mathbf{x} - \mathbf{h})^+ \right\|^2$. The uniform a.s. convergence is guaranteed by the uniform law of large numbers (Jennrich, 1969):

- By Lem. J.5, $\mathbb{C}(\mathbf{A}^*)$ is a compact set.
- g is continuous w.r.t. \mathbf{A} by its formulation and measurable over \mathbf{x} at each $\mathbf{A} \in \mathbb{C}(\mathbf{A}^*)$.
- In fact,

$$\begin{aligned} g(\mathbf{x}, \mathbf{A}) &= \frac{1}{2} \left\| (\mathbf{A}\mathbf{x} - \mathbf{h})^+ \right\|^2 \\ &\leq \|\mathbf{A}\mathbf{x}\|^2 + \|\mathbf{A}^*\mathbf{x}\|^2 \\ &\leq (\|\mathbf{A}\|_F + \|\mathbf{A}^*\|_F) \|\mathbf{x}\|^2. \end{aligned}$$

Since $\mathbf{A} \in \mathbb{C}(\mathbf{A}^*)$ is in a compact set,

$$g(\mathbf{x}, \mathbf{A}) \leq \left[\sup_{\mathbf{A} \in \mathbb{C}(\mathbf{A}^*)} \|\mathbf{A}\|_F + \|\mathbf{A}^*\|_F \right] \|\mathbf{x}\|^2. \quad (51)$$

Thus the dominating function exists¹⁰. □

By Lem. J.5, J.6, J.7 and J.8, all of the conditions are satisfied in (Shapiro et al., 2014, Thm. 5.3). Thus, we have the strong consistency of layer 1 objective optima estimator as described in Lem. J.9.

Lemma J.9 (QP/LP strong consistency, layer 1). $D_{\mathbb{H}}(\hat{\mathbb{T}}_{n'}, \mathbb{T}^*) \xrightarrow{a.s.} 0$ as $n \rightarrow \infty$.

Similar to Lem. J.2, we have the strong consistency of the layer 1 scale factor estimator as described in Lem. J.10.

Lemma J.10 (scale factor estimator strong consistency, layer 1). Let $\mathbf{k}_{n'_{sf}}(\hat{\mathbf{C}}_n)$ be the n'_{sf} -sample estimator of k_j via LR: Alg. 2 line 5. given that $h_j^{(i)} > 0$. Define sets

$$\mathbb{V}_{n'_{sf}, n} := \{\mathbf{k}_{n'_{sf}}(\mathbf{A}) : \mathbf{A} \in \hat{\mathbb{T}}_{n'}\}, \text{ and } \mathbb{V}^* := [0, 1]. \quad (52)$$

Then $\lim_{n'_{sf} \rightarrow \infty} \lim_{n' \rightarrow \infty} D_{\mathbb{H}}(\mathbb{V}_{n'_{sf}, n}, \mathbb{V}^*) \xrightarrow{a.s.} 0$.

Remark. In case $k_j = 0$, suppose the algorithm finds a solution over a continuous distribution with $[0, 1]$ as support and the probability that it finds a solution with scale factor 0 is 0.

With Thm. 5.2 and the continuous mapping theorem, the strong consistency of layer 1 estimation is guaranteed.

Theorem J.11 (strong consistency, layer 1). Suppose $\hat{\mathbf{A}}_{n'}$ is scaled by $\hat{\mathbf{k}}_{n'_{sf}}$. Then $\hat{\mathbf{A}}_{n'} \xrightarrow{a.s.} \mathbf{A}^*$ as $n', n'_{sf} \rightarrow \infty$.

By Thm. J.3 and J.11, Thm. 6.3 is guaranteed by the continuous mapping theorem.

¹⁰Here, we assume $\mathbb{E}[\|\mathbf{x}\|^2] < +\infty$ as is commonly done in empirical analysis.