

# RedEx: Beyond Fixed Representation Methods via Convex Optimization

**Amit Daniely**

*Hebrew University of Jerusalem and Google Research Tel Aviv*

**Mariano Schain**

*Google Research Tel Aviv*

**Gilad Yehudai**

*Weizmann Institute of Science and Google Research Tel Aviv*

**Editors:** Claire Vernade and Daniel Hsu

## Abstract

Optimizing Neural networks is a difficult task which is still not well understood. On the other hand, fixed representation methods such as kernels and random features have provable optimization guarantees but inferior performance due to their inherent inability to learn the representations. In this paper, we aim at bridging this gap by presenting a novel architecture called RedEx (Reduced Expander Extractor) that is as expressive as neural networks and can also be trained in a layer-wise fashion via a convex program with semi-definite constraints and optimization guarantees. We also show that RedEx provably surpasses fixed representation methods, in the sense that it can efficiently learn a family of target functions which fixed representation methods cannot.

## 1. Introduction

Neural networks have demonstrated unparalleled performance in various tasks, including Computer Vision and Natural Language Processing (NLP). However, training them remains a challenging task that is not yet fully understood. On the theoretical side, the optimization landscape of neural networks is highly non-convex, characterized by numerous spurious local minima (Safran and Shamir, 2018; Yun et al., 2018) and are often also non-smooth. Consequently, proving optimization results for non-convex and non-smooth functions is generally deemed unfeasible (Kornowski and Shamir (2021)). On the practical side, the optimization process for neural networks primarily employs gradient-based methods like Stochastic Gradient Descent (SGD) or ADAM (Kingma and Ba (2014)), necessitating a meticulous search for hyperparameters. This process often relies on trial and error rather than being firmly grounded in theory.

On the contrary, fixed representation methods, such as kernels and random features, can be efficiently learned with provable guarantees using convex optimization techniques. However, recent research has highlighted a limitation: as they do not learn a representation, these methods are inherently less powerful than neural networks. There are learning scenarios where neural networks demonstrate efficient learning, while fixed representation methods falter (e.g. Yehudai and Shamir (2019); Kamath et al. (2020); Malach et al. (2021b); Ghorbani et al. (2019); Daniely and Malach (2020)). Different approaches that do facilitate efficient and provable representation learning often rely on overly simplistic models (e.g. Yehudai and Shamir (2020); Vardi et al. (2021); Bietti et al. (2022)) or necessitate stringent assumptions about the data and employ specialized algorithms

tailored to specific learning contexts (e.g. [Ge et al. \(2017\)](#); [Allen-Zhu and Li \(2020\)](#); [Abbe et al. \(2021\)](#)).

A natural question that arises is whether there exists a model-class which benefits from the "best of all worlds", namely:

*Is there a model class that can be learned efficiently without making assumptions about the input distribution, matches the expressiveness of neural networks, and is capable of learning meaningful representations rather than relying on fixed ones?*

In this paper, we provide an affirmative answer to this question by introducing the **Reduced Extractor Expander (RedEx)** architecture. We demonstrate that RedEx is as expressive as neural networks and can be learned using a convex program without any assumptions on the input data. Moreover, we establish that RedEx learns non-trivial representations, as evidenced by a novel learning problem we introduce that RedEx can efficiently learn, while fixed representation methods cannot. In more details, our main contributions are:

1. We introduce the RedEx architecture and show that it can efficiently express any Boolean circuit (Thm. 2).
2. We present an efficient polynomial-time algorithm for training RedEx, based on convex Semidefinite Programming (SDP) (Thm. 3 and Algorithm 2).
3. We introduce a learning problem, based on a variation of the sparse-parity task which RedEx can learn efficiently, while any fixed-representation methods cannot (Thm. 7 and Thm. 8).

Furthermore, we demonstrate that if the output is one-dimensional, RedEx can be trained using standard gradient-based methods like gradient descent, without the need for SDP. Finally we extend the RedEx architecture to the convolutional setting.

## 1.1. Related Works

**Fixed representation methods and NTK.** Fixed representation methods are models which can be viewed as a feature mapping which is non-linear and fixed followed by a learned linear mapping. This includes kernel methods, random features ([Rahimi and Recht, 2007](#)), and others. In recent years, neural networks under certain assumptions were analyzed in the so called "kernel regime" ([Woodworth et al., 2020](#)). In this approach, it is assumed that the training takes place near the initial weights. This allows to analyse neural networks as if it is a fixed representation method. This approach was popularized through the Neural Tangent Kernel (NTK) model ([Jacot et al., 2018](#)). Many similar works have shown positive results where neural networks can provably learn under different assumptions, e.g. [Daniely \(2020\)](#); [Andoni et al. \(2014\)](#); [Du et al. \(2017\)](#); [Daniely \(2017\)](#); [Allen-Zhu et al. \(2019\)](#); [Li and Liang \(2018\)](#); [Cao and Gu \(2019\)](#).

**Limitations of fixed representation methods.** Several works in recent years have focused on the limitations of fixed representation methods, NTK and learning under the "kernel regime". [Yehudai and Shamir \(2019\)](#) and [Kamath et al. \(2020\)](#) have shown that fixed representation methods cannot learn even a single ReLU neuron under Gaussian distribution, unless the number of features is exponential in the input dimension. On the other hand, neural networks were shown to be able to

efficiently learn single neurons (Yehudai and Shamir, 2020; Vardi et al., 2021). Several other works have shown that under certain distributional assumptions fixed representation methods cannot learn parity functions while neural networks can (see e.g. Malach et al. (2021b); Daniely and Malach (2020); Malach et al. (2021a)). Finally, Ghorbani et al. (2019, 2021) have shown that the NTK and random features methods can essentially learn efficiently only low degree polynomials.

**Provable optimization beyond fixed representations.** Several works consider model-classes which go beyond fixed representations, but can be efficiently and provably learned. These works usually consider either overly-simplistic models, or have strong assumptions on the input data. Yehudai and Shamir (2020); Vardi et al. (2021); Bietti et al. (2022); Bruna et al. (2023); Frei et al. (2020) consider learning single neurons or single index neurons with provable optimization guarantees. However these models are overly-simplistic and have very limited expressiveness. Abbe et al. (2021); Allen-Zhu and Li (2020) consider a certain hierarchical model resembling RedEx and show a family of functions that these models can learn. However the guarantees are for a specific family of input distributions, with a training algorithm that is tailored for these specific learning problems. Ge et al. (2017); Tian (2017) consider learning a one-hidden layer neural network with gradient descent for Gaussian inputs using a specific analytic formula relying on the distribution of the data. Our model also share similarity to phase retrieval methods (e.g. Candes et al. (2013, 2015)), although these works mostly consider Gaussian data or data distributed uniformly on a sphere.

Livni et al. (2014) consider the problem of learning a one-hidden layer network with square activation under trace norm constraints. They prove learnability using a reduction to a convex program relying on the GECO algorithm (Shalev-Shwartz et al., 2011). Our work is similar to that in nature, however we provide several extensions: (1) A separation result between fixed representation methods and RedEx, which do not appear in Livni et al. (2014); (2) Extensions to multivariate output and a convolutional structure, which is not possible using the convex reduction in Livni et al. (2014); and (3) A multilayer version of RedEx which enables to express any Boolean circuit, and thus match the expressive power of neural networks.

## 2. Notations and Settings

We denote vectors in bold-face:  $\mathbf{x}$ . We will assume the input space is  $\mathbb{R}^d$ . The output space will be denoted  $\mathcal{Y}$ . We will consider algorithms that learn functions from  $\mathbb{R}^d$  to  $\mathbb{R}^k$  and are evaluated by a convex loss function  $\ell : \mathbb{R}^k \times \mathcal{Y} \rightarrow [0, \infty)$ . Given a distribution  $\mathcal{D}$  on  $\mathbb{R}^d \times \mathcal{Y}$  and  $h : \mathbb{R}^d \rightarrow \mathbb{R}^k$  we denote  $\ell_{\mathcal{D}}(h) = \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} \ell(h(\mathbf{x}), y)$ . Likewise, for a dataset  $S = \{(\mathbf{x}_1, y_1), \dots, (\mathbf{x}_m, y_m)\}$  we denote  $\ell_S(h) = \frac{1}{m} \sum_{i=1}^m \ell(h(\mathbf{x}_i), y_i)$ .

For  $\mathbf{x} \in \mathbb{R}^d$  we denote by  $\mathbf{x}^{\otimes 2} := \mathbf{x}\mathbf{x}^\top \in \mathbb{R}^{d \times d}$  the outer product. For a vector,  $\mathbf{x} \in \mathbb{R}^d$  we will use  $\|\mathbf{x}\| = \sqrt{\sum_i x_i^2}$  to denote the Euclidean norm. Given  $I \subset [d]$  we denote  $\mathbf{x}(I) = \sum_{i \in I} x_i$ . For a matrix  $A$ , we will use  $\|A\|_{\text{fr}} = \sqrt{\sum_{ij} A_{ij}^2}$  to denote the Frobenius norm,  $\|A\|_{\text{sp}} = \max_{\|\mathbf{x}\|=1} \|A\mathbf{x}\|$  to denote the spectral norm, and  $\|A\|_{\text{Tr}}$  to denote the trace norm which is the sum of  $A$ 's singular values. For a diagonal matrix  $D$  we denote by  $|D|$  the diagonal matrix whose  $i$ -th diagonal coordinate is equal to  $|D_{i,i}|$ . If  $D$  is also Positive Semi-Definite (PSD) we define by  $\sqrt{D}$  the matrix whose  $i$ -th diagonal coordinate is equal to  $\sqrt{D_{i,i}}$ . We will use  $\mathbf{A}$  to denote a tuple of matrices  $\mathbf{A} = (A_1, \dots, A_k) \in (\mathbb{R}^{d \times n})^k$ . We will let  $\|\mathbf{A}\|_{\text{fr}}^2 = \sum_{i=1}^k \|A_i\|_{\text{fr}}^2$ . For  $B \in \mathbb{R}^{m \times d}$

and  $C \in \mathbb{R}^{n \times m}$  we denote  $BA = (BA_1, \dots, BA_k)$  and  $AC = (A_1C, \dots, A_kC)$ . We denote by  $B_M^d \subset \mathbb{R}^d$  the Euclidean ball of radius  $M$  centered at 0.

For a symmetric matrix  $A \in \mathbb{R}^{d \times d}$ , we say that  $A = U^\top DU$  is a *compact orthogonal diagonalization* if  $U \in \mathbb{R}^{d' \times d}$  is a matrix with  $d' \leq d$  orthonormal rows and  $D \in \mathbb{R}^{d' \times d'}$  is a diagonal matrix with non-zero diagonal entries, or the 0 matrix in  $\mathbb{R}^{1 \times 1}$ . Note that any symmetric matrix has a compact orthogonal diagonalization. For a linear subspace  $V \subset \mathbb{R}^m$ , we denote by  $P_V$  the projection on  $V$  and say that  $A$  is  $V$ -supported if  $A = P_V^\top AP_V$ .

For an embedding  $\Psi : \mathcal{X} \rightarrow \mathbb{R}^n$  we denote by  $\mathcal{H}_\Psi$  the space of all function for which there is  $\mathbf{v} \in \mathbb{R}^n$  such that  $\forall \mathbf{x} \in \mathcal{X}, h(\mathbf{x}) = \langle \mathbf{v}, \Psi(\mathbf{x}) \rangle$ . We also define a norm on  $\mathcal{H}_\Psi$  by  $\|h\|_\Psi = \min\{\|\mathbf{v}\| : \forall \mathbf{x} \in \mathcal{X}, h(\mathbf{x}) = \langle \mathbf{v}, \Psi(\mathbf{x}) \rangle\}$ . We note that  $\|\cdot\|_\Psi$  turn  $\mathcal{H}_\Psi$  into a Hilbert space. We also define  $k_\Psi(\mathbf{x}, \mathbf{y}) = \langle \Psi(\mathbf{x}), \Psi(\mathbf{y}) \rangle$ . For two matrices  $A, B \in \mathbb{R}^{n \times m}$  we define  $\langle A, B \rangle := \text{Tr}(B^\top A)$ .

### 3. Reduced Extractor-Expanders (RedEx)

In this section we present the main architecture that we analyse throughout the paper. This architecture is aimed at being analogous to 2-layer neural networks with a quadratic activation. Its base component consists of two layers: The first layer "extracts" the most informative directions of the data using a matrix with orthogonal rows, and bounded Frobenius norm, and then "expands" it quadratically by increasing the dimension. The second layer is a linear transformation with bounded norm over the extracted features.

**Definition 1 (RedEx - Reduced Extractor-Expander)** A function  $\Lambda : \mathbb{R}^d \rightarrow \mathbb{R}^k$  is called a **RedEx** (Reduced Expander Extractor) of width  $M$  if it is of the form:  $\Lambda = \Psi_{\mathbf{P}} \circ \Psi_V$  where:

1. The function  $\Psi_V : \mathbb{R}^d \rightarrow \mathbb{R}^{d' \times d'}$  is of the form:

$$\Psi_V(\mathbf{x}) = (V\mathbf{x})^{\otimes 2}$$

for  $V \in \mathbb{R}^{d' \times d}$  with  $d' \leq d$ , orthogonal rows, and  $\|V\|_{\text{fr}}^2 \leq M$ . We call the matrix  $V$  extractor.

2. The function  $\Psi_{\mathbf{P}} : \mathbb{R}^{d' \times d'} \rightarrow \mathbb{R}^k$  is of the form:

$$\Psi_{\mathbf{P}}(X) = (\langle P_1, X \rangle, \dots, \langle P_k, X \rangle)$$

for  $P_i \in \mathbb{R}^{d' \times d'}$  with  $\|P_i\|_{\text{sp}} \leq 1$ .

The function  $\Psi_V$  is called an **extractor-expander**.

To have some intuition, the width of an extractor  $V$  can be thought of as a continuous surrogate to the number of orthogonal directions  $V$  extracts. Indeed, in order to extract  $d'$  orthogonal dimensions defined by unit vectors  $\mathbf{e}_1, \dots, \mathbf{e}_{d'}$  we can use the extractor  $V \in \mathbb{R}^{d' \times d}$  whose  $i$ 'th row is  $\mathbf{e}_i^\top$ . In this case, the width of  $V$  is  $d'$ . Definition 1 generalizes such extractor matrices, and allows to give larger weights to directions which are "more important". The width  $M$  controls the expressivity of the architecture. Allowing large width will result with more functions that can be expressed, but on the other hand will require more examples to learn them. Alternatively, small width will result with a less expressive class of function, but with better generalization capabilities.

Given a data set  $(\mathbf{x}_1, \mathbf{y}_1), \dots, (\mathbf{x}_m, \mathbf{y}_m) \in \mathbb{R}^d \times \mathbb{R}^k$  our algorithm will seek a RedEx  $\Psi_{\mathbf{P}} \circ \Psi_V$  that minimizes the loss subject to a width constraint. We will also allow for an additional small regularization term, that will be used to guarantee generalization. Specifically, given a loss function  $\ell$ , and for the regularization function

$$\mathcal{R}(\mathbf{P}, V) = \|V^\top V\|_{\text{fr}}^2 + \sum_{i=1}^k \|V^\top P_i V\|_{\text{fr}}^2 \quad (1)$$

our algorithm will minimize:  $\ell_S(\Psi_{\mathbf{P}} \circ \Psi_V) + \lambda \mathcal{R}(\mathbf{P}, V)$ , subject to the constraint that the width of  $V$  is at most  $M$ . We will later explain how this can be done in polynomial time, and will provide guaranties on its performance. We first extend the above architecture to multi-layer RedEx in the same manner that a 2-layer neural network is extended to multi-layer. The basic idea is to use several extractor-expanders  $\Psi_{V^t}$  in a sequential manner, and at the last layer use a linear transformation  $\Psi_{\mathbf{P}}$ . As with the basic depth-two RedEx architecture, the width of the extractors  $V^t$  will control the complexity of the functions computed by the architecture, and will be used to trade-off expressive power and sample complexity.

There are two issues that arise when doing such a generalization: (1) The representation dimension grows exponentially with the number of layers. This is because the function  $\Psi_V(\mathbf{x}) = (V\mathbf{x})^{\otimes 2}$  expands the dimension quadratically. (2) Training all the layers simultaneously is computationally hard (as we show later, poly-sized deep RedEx architectures can express any poly-sized Boolean circuit, which implies that they are hard to learn [Kearns and Valiant \(1994\)](#)), whereas one of our main goals is to obtain provable guarantees for the optimization process.

To deal with the first issue, we allow extractors whose output dimension is at most the number of examples  $m$ . This limits the representation dimension to  $m^2$ . This is a convenient way to deal with this issue theoretically, however in practice alternative approaches might be favourable. For instance, we can simply delete all the rows in the extractors  $V$  whose norm is  $\leq \epsilon$  for some tunable parameter  $\epsilon$ . For a sufficiently small  $\epsilon$ , this will not alter the solution by much. We note that the number of rows with norm that is larger than  $\epsilon$  is at most the width of the extractor, divided by  $\epsilon^2$ . It is also possible to use the kernel trick. Lastly, a more practically oriented way, is to apply a dimension reduction method, such as PCA, after every extractor-expander layer.

To address the second issue, instead of training all the layers simultaneously, we train them sequentially. This is analogous to layer-wise training in neural networks [Bengio et al. \(2006\)](#). The idea is that for  $t$ -th layer, we find functions  $\Psi_{V^t}$  and  $\Psi_{\mathbf{P}^t}$  which minimizes the target loss. After training is finished we only keep the representation function  $\Psi_{V^t}$ , while discarding the linear transformations. The input for the next layer  $t + 1$  is the output of the extractor-expander  $\Psi_{V^t}$ .

Another small issue that arises for extractor-expanders is that they only allow to compute degree-2 polynomials which are even (i.e. satisfy  $p(\mathbf{x}) = p(-\mathbf{x})$ ). More generally, multi-layer RedEx would only allow to compute even high degree polynomials. This issue can be easily fixed by adding an extra fixed coordinate to the inputs. We summarize all the above in [Algorithm 1](#).

The heart of the above algorithm is to minimize [Eq. \(2\)](#) in [step 5](#). On one hand, it can be done using standard gradient methods such as GD or SGD. The problem with this approach is that the function being optimized is not convex, even without the norm and orthogonality constraints. Thus, it is not clear that it converges to a global optimum. In the next section we provide an efficient algorithm for finding this optimum, thus providing an efficient and provable algorithm for layer-wise learning of multi-layer RedEx.

**Algorithm 1** Training multi-layer RedEx

- 
- 1: **Parameters:** A loss  $\ell : \mathbb{R}^k \times \mathcal{Y} \rightarrow [0, \infty)$ , number of layers  $L$ , width parameters  $M_1, \dots, M_L$ , regularization parameters  $\lambda_1, \dots, \lambda_L$ , and constant parameter  $c$ .
  - 2: **Input:** A dataset  $(\mathbf{x}_1, y_1), \dots, (\mathbf{x}_m, y_m) \in \mathbb{R}^d \times \mathcal{Y}$ .
  - 3: Define  $\mathbf{x}_1^0 = \begin{pmatrix} c \\ \mathbf{x}_1 \end{pmatrix}, \dots, \mathbf{x}_m^0 = \begin{pmatrix} c \\ \mathbf{x}_m \end{pmatrix}$  and  $d_0 := d + 1$
  - 4: **for**  $t=1, \dots, L$  **do**
  - 5: Find  $V^t$  and  $\mathbf{P}^t = (P_1^t, \dots, P_k^t)$  that minimizes:

$$\frac{1}{m} \sum_{i=1}^m \ell(\Psi_{\mathbf{P}^t} \circ \Psi_{V^t}(\mathbf{x}_i^{t-1}), y_i) + \lambda_t \mathcal{R}(\mathbf{P}^t, V^t) \quad (2)$$

where:

1.  $\Psi_{\mathbf{P}^t} \circ \Psi_{V^t}(\mathbf{x}) := (\langle P_1^t, (V^t \mathbf{x})^{\otimes 2} \rangle, \dots, \langle P_k^t, (V^t \mathbf{x})^{\otimes 2} \rangle)$ .
  2.  $V^t$  has  $n_t \leq m$  orthogonal rows. Each row is a vector in  $\mathbb{R}^{d_{t-1}}$ , and  $\|V^t\|_F^2 \leq M_t$
  3. For each  $j \in [k]$ ,  $P_j \in \mathbb{R}^{n_t \times n_t}$  and  $\|P_j^t\|_{\text{sp}} \leq 1$ .
- 6: Define  $d_t = n_t^2$  and  $\mathbf{x}_i^t = \Psi_{V^t}(\mathbf{x}_i^{t-1})$  for  $i \in [m]$ . We will view  $\mathbf{x}_i^t$  as a vector in  $\mathbb{R}^{d_t}$
  - 7: **end for**
  - 8: **Output:** Output the hypothesis  $\Psi_{\mathbf{P}^L} \circ \Psi_{V^L} \circ \Psi_{V^{L-1}} \circ \dots \circ \Psi_{V^1}$
- 

We emphasize that one caveat of the layer-wise optimization approach is that a global minimizer for all the layers simultaneously might achieve better performance than a global minimizer for each layer separately. It can be viewed as a greedy algorithm, where at each step we optimize the current layer which is locally the best possible step, but it may not be the best step globally (i.e. for all layers simultaneously). The main advantage of this sequential approach is that it will allow us to use convex optimization and give provable guarantees for optimizing our model, while also providing separation between this model and fixed representation methods. Note that there are no layer-wise training guarantees for neural networks in general.

We now show an expressivity result, namely that the RedEx architecture can approximate any Boolean circuit with only a quadratic increase of the size of the circuit. This shows that like neural networks, our architecture can express a very large class of functions – virtually any function of interest. Indeed, any function that can be computed efficiently has a small circuit that computes it [Vollmer \(1999\)](#).

**Theorem 2** *Let  $B : \{0, 1\}^d \rightarrow \{0, 1\}$  be a function computed by a Boolean circuit of size  $T$ . Then we can define a RedEx with depth  $O(T)$  and intermediate feature dimension at most  $O(T^2)$  that computes  $B$ .*

The proof can be found in [Appendix A](#). [Thm. 2](#) implies that the RedEx architecture is as expressive as neural networks. This can be seen using the following simple argument: Any neural network (with inputs in  $\{0, 1\}^d$  and output in  $\{0, 1\}$ ) can be simulated by a boolean circuit, where the number of nodes in the circuit is at most polynomial in the number of parameters (see [Maass \(1997\)](#)). [Thm. 2](#)

shows that any Boolean circuit can be simulated by a multilayer RedEx architecture with at most polynomial blow-up in the size of the circuit. Thus, given a neural network, it can be simulated by a multilayer RedEx architecture with at most polynomial blow-up in the number of parameters.

#### 4. Efficient and Provable Learnability of RedEx

In this section we present an efficient algorithm for learning a single layer RedEx, and present generalization guaranties for it. This algorithm can be used to minimize objective Eq. (2), thus leading to an efficient implementation of algorithm 1.

Our approach is to reduce the problem of minimizing Eq. (2) under the norm and orthogonality constraints to a convex semi-definite program, which contains PSD constraints. Such a problem can be solved using convex SDP algorithms in polynomial time.

In order to do so, we present a different parametrization of RedEx functions. Let

$$\mathcal{P}_M = \{(\mathbf{P}, V) : \|V\|_{\text{fr}}^2 \leq M \text{ and } \forall i, \|P_i\|_{\text{sp}} \leq 1\}$$

A RedEx function of width  $M$  is defined by  $(\mathbf{P}, V) \in \mathcal{P}_M$ . Now, let

$$\mathcal{P}^M = \{(\mathbf{A}, R) : \text{Tr}(R) \leq M \text{ and } \forall i, -R \preceq A_i \preceq R\}$$

As the following lemma shows, we can alternatively define width  $M$  RedEx functions via  $(\mathbf{A}, R) \in \mathcal{P}^M$ . We then show that under this alternative parameterization of RedEx functions, objective Eq. (2) becomes convex. Furthermore, we can efficiently convert the alternative parameterization to the original. These two facts enable us to efficiently implement algorithm 2.

**Theorem 3** *Let  $\mathcal{H}_M$  be the class of functions of the form:*

$$h_{V, \mathbf{P}}(\mathbf{x}) = (\langle P_1, (V\mathbf{x})^{\otimes 2} \rangle, \dots, \langle P_k, (V\mathbf{x})^{\otimes 2} \rangle)$$

*For  $(\mathbf{P}, V) \in \mathcal{P}_M$ . Let  $\mathcal{H}^M$  be the class of functions of the form:*

$$h^{R, \mathbf{A}}(\mathbf{x}) = (\langle \mathbf{x}, A_1 \mathbf{x} \rangle, \dots, \langle \mathbf{x}, A_k \mathbf{x} \rangle)$$

*For  $(\mathbf{A}, R) \in \mathcal{P}^M$ . We have:*

1.  $\mathcal{H}_M = \mathcal{H}^M$ .
2. Fix  $(\mathbf{A}, R) \in \mathcal{P}^M$ . Diagonalize  $R = U^\top D U$  for unitary  $U$  and diagonal  $D$  and let

$$V = \sqrt{D} U, \quad P_i = (V^\dagger)^\top A_i (V^\dagger)$$

*then  $(\mathbf{P}, V) \in \mathcal{P}_M$  and  $h_{V, \mathbf{P}} = h^{R, \mathbf{A}}$ . Furthermore,  $\text{Tr}(R) = \|V\|_{\text{fr}}^2$  and  $\mathcal{R}(\mathbf{P}, V) = \|R\|_{\text{fr}}^2 + \|\mathbf{A}\|_{\text{fr}}^2$*

The proof can be found in Appendix B. Theorem 3 suggests the following algorithm for training a single RedEx layer, and to optimize objective Eq. (2) in algorithm 1.



---

**Algorithm 2** Training 1-layer RedEx

---

- 1: **Parameters:** Loss  $\ell : \mathbb{R}^k \times \mathcal{Y} \rightarrow [0, \infty)$ , width parameter  $M$  and regularization parameter  $\lambda$
- 2: **Input:** A dataset  $(\mathbf{x}_1, y_1), \dots, (\mathbf{x}_m, y_m) \in \mathbb{R}^d \times \mathcal{Y}$
- 3: Find symmetric  $d \times d$  matrices  $A_1, \dots, A_k, R$  by solving the semi-definite program:

$$\min \quad \frac{1}{m} \sum_{i=1}^m \ell_{y_i} \left( \mathbf{x}_i^\top A_1 \mathbf{x}_i, \dots, \mathbf{x}_i^\top A_k \mathbf{x}_i \right) + \lambda (\|R\|_{\text{fr}}^2 + \|\mathbf{A}\|_{\text{fr}}^2) \quad (3)$$

$$\begin{aligned} \text{s.t.} \quad & -R \preceq A_i \preceq R \\ & R \succeq 0 \\ & \text{Tr}(R) \leq M \end{aligned} \quad (4)$$

- 4: Compute an orthogonal diagonalization  $R = U^\top D U$
  - 5: Output  $V = \sqrt{D} U$  and  $P_i = (V^\dagger)^\top A_i V^\dagger$ .
- 

**Remark 4** We note that Algorithm 2 can be performed in polynomial time, although we don't specify the exact training time for the algorithm. The reason is that there are different convex program solvers with different pros and cons, and the training time depends on which solver is chosen. The minimization objective in Eq. (3) is a general strongly convex function under semi-definite constraints. It can be solved using general interior point method (Potra and Wright, 2000), the ellipsoid algorithm (see e.g. Ch.2 in Bubeck et al. (2015)), conic optimization (e.g. Auslender and Teboulle (2006); Dahl and Andersen (2022)) or any other method which solves convex SDP problems.

**Remark 5** We have that  $\text{rank}(R) \leq m$ , and hence  $V$  has at most  $m$  non-zero rows, that is the number of "important features" is bounded by the size of the dataset. Indeed, let  $P : \mathbb{R}^d \rightarrow \mathbb{R}^d$  be the projection on  $\text{span}\{\mathbf{x}_1, \dots, \mathbf{x}_m\}$ , and let  $(\mathbf{A}, R)$  be an optimal solution to program Eq. (3). Note that there is a single optimal solution, as program Eq. (3) is strongly convex. It is not hard to see that the objective value of  $(P^\top \mathbf{A} P, P^\top R P)$  is as good as the objective value of  $(\mathbf{A}, R)$ , this is because projection on the data samples produces the same class of functions, while it does not increase both the trace and Frobenius norms. As the optimal solution is unique, we conclude that  $R = P^\top R P$ .

We next state a generalization result for algorithm 2. The result follows directly from Corollary 13.6 in Shalev-Shwartz and Ben-David (2014), by noticing that the objective is convex with an appropriate regularization term. To this end, we define

$$\text{Val}_{\mathcal{D}}(\mathbf{A}, R) := \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} \ell_y \left( \mathbf{x}^\top \mathbf{A} \mathbf{x} \right) \quad (5)$$

and

$$\text{Val}_{\mathcal{D}, M} := \inf_{-R \preceq A_i \preceq R \text{ and } \text{Tr}(R) \leq M} \text{Val}_{\mathcal{D}}(\mathbf{A}, R)$$

**Theorem 6** Assume that the dataset is an i.i.d. sample from a distribution  $\mathcal{D}$  on  $B_{M_1}^d \times \mathcal{Y}$  and that the loss is  $L$ -Lipschitz. Let  $(\mathbf{A}, R)$  be the output of algorithm 2. Then

$$\mathbb{E}_S \text{Val}_{\mathcal{D}}(\mathbf{A}, R) \leq \text{Val}_{\mathcal{D}, M} + \lambda(k+1)M^2 + \frac{M_1^4 L^2}{\lambda m}$$



Thm. 6 shows a trade-off in generalization capabilities by choosing the parameter  $M$ , similar to the well known bias-variance trade-off. Namely, larger value of  $M$  allows for better expressive power but requires more samples to achieve good generalization capabilities. We note that our generalization result scales at a rate of  $O\left(\frac{1}{\sqrt{m}}\right)$  by choosing an appropriate  $\lambda$ . It is an interesting question whether this rate can be improved to  $O\left(\frac{1}{m}\right)$ , similarly to what is done in Wang and Lin (2021), but for non-smooth regularizers.

## 5. Layerwise RedEx surpasses Kernel Methods

In this section we will provide a learning problem which demonstrates a separation between RedEx and fixed representation methods. The problem we choose is inspired by Daniely and Malach (2020) where they show that neural networks can learn the sparse parity function under a certain distribution which "leaks" the coordinates of the parity. In more details, given an input space  $\{\pm 1\}^d$ , the sparse parity function on the  $k$  coordinates  $\mathbf{x}_{i_1}, \dots, \mathbf{x}_{i_k}$  is defined as:  $\prod_{j=1}^k x_{i_j}$ . Since RedEx is learned in a layer-wise fashion, we consider a slightly different learning problem which better aligns with the RedEx architecture and still cannot be learned by fixed representation methods.

Namely, we consider the problem of learning the following family of models: The input space is  $\{\pm 1\}^d$ , the output space is  $\mathbb{R}^{(1+k/2)}$  for even  $k = \Theta\left(\frac{\sqrt{\log(d)}}{(\log \log(d))^2}\right)$ , and the input distribution is uniform on  $\{\pm 1\}^d$ . Denote by  $p_0, \dots, p_k$ , the set of orthogonal polynomials w.r.t. the distribution of  $\sum_{i=1}^k X_i$  for i.i.d. Radamacher r.v.  $X_1, \dots, X_k \in \{\pm 1\}$ . These polynomials are called *Kravchuk Polynomials* Nikiforov et al. (1991) and are given by the recursion formula

$$p_0(x) = 1, \quad p_1(x) = \frac{x}{\sqrt{k}}, \quad xp_i(x) = \sqrt{(i+1)(k-i)}p_{i+1}(x) + \sqrt{i(k-i+1)}p_{i-1}(x) \quad (6)$$

We consider the problem of learning a function of the form

$$h_{\mathcal{I}}(\mathbf{x}) = \left( p_0 \left( \sum_{i \in \mathcal{I}} x_i \right), p_2 \left( \sum_{i \in \mathcal{I}} x_i \right), p_4 \left( \sum_{i \in \mathcal{I}} x_i \right), \dots, p_k \left( \sum_{i \in \mathcal{I}} x_i \right) \right)$$

for an unknown set of coordinates  $\mathcal{I} \subset [d]$  with  $|\mathcal{I}| = k$  and w.r.t. the square loss  $\ell(\hat{\mathbf{y}}, \mathbf{y}) = \|\hat{\mathbf{y}} - \mathbf{y}\|^2$ . Our first result shows that algorithm 1 learns a function with loss of  $o(1)$ . Note that the coordinates of  $h_{\mathcal{I}}$  are polynomials of increasing degree, while its last coordinate is the sparse parity function  $\mathbf{x} \mapsto \prod_{i \in \mathcal{I}} x_i$ . To see that this is indeed the sparse parity function, note that by definition it is orthogonal to any Kravchuk polynomial of degree  $i \leq k$ , and the  $k$ -th Kravchuk polynomial is the unique polynomial with this property, hence it must be the sparse parity function. Thus, our function can be seen as learning the parity function, but using a kind of "staircase property" (Abbe et al., 2021) which the RedEx architecture exploits due to its layer-wise training.

**Theorem 7** *Assume we run algorithm 1 on  $m$  i.i.d. examples,  $L = \lceil \log_2(k) \rceil$  layers, regularization parameters  $\lambda_1 = \lambda_2 = \dots = \lambda_L = \frac{1}{\sqrt{m}}$ , width parameters  $M_1 = \frac{1}{2} + \frac{1}{\sqrt{2-2/k}}$ ,  $M_2 = \dots = M_L = M = 2^{3k}(3k)^k$ , and constant parameter  $c = \sqrt{2}$ . Assume furthermore that each layer is trained using a fresh sample. Then, w.p.  $1 - \delta$ , for the output hypothesis  $h$ ,  $\ell_{\mathcal{D}}(h) = \frac{d^4 2^{O(k^2 \log^2(k))}}{\delta m^{1/4}} = \frac{O(d^5)}{\delta m^{1/4}}$*

The reason for sampling a batch of fresh samples when training each layer is a technical artifact of the proof, aimed at eliminating the dependence between the training of each layer. It can be seen alternatively as if the original dataset is larger by a factor of  $\log(k) := O(\log \log(d))$ , and we only use a part of it for training each layer.

We compliment the above result by showing that polynomial-time fixed-representation methods, such as kernels and random features, cannot achieve the guarantee in Theorem 7. The reason is that the last coordinate of  $h_{\mathcal{I}}$  is the parity function  $\mathbf{x} \mapsto \prod_{i \in \mathcal{I}} x_i$ . This implies that any fixed-representation method that is guaranteed to find a function  $h$  with  $\mathbb{E}_{\mathbf{x}} \|h_{\mathcal{I}}(\mathbf{x}) - h(\mathbf{x})\|_2^2 = o(1)$  has super-polynomial complexity of  $d^{\Omega(k)}$ . Specifically, Corollary 13 from Ben-David et al. (2002) implies:

**Theorem 8** *Let  $\Psi : \{\pm 1\}^d \rightarrow B_{M_1}^m$  be any, possibly random, embedding. Assume that for any  $\mathcal{I} \subset [d]$  with  $|\mathcal{I}| = k$ , w.p.  $\geq 1/2$  over the choice of  $\Psi$ , there are vectors  $\mathbf{w}_0, \dots, \mathbf{w}_k \in B_{M_2}^m$  such that*

$$\mathbb{E}_{\mathbf{x}} \sum_{j=0}^{k/2} \left( p_{2j} \left( \sum_{i \in \mathcal{I}} x_i \right) - \langle \mathbf{w}_j, \Psi(\mathbf{x}) \rangle \right)^2 \leq 0.99$$

Then  $M_1 M_2 \geq d^{\Omega(k)}$

### 5.1. On the proof of theorem 7

Theorem 7 is proved in Appendix C. In section C.1 it is shown that  $V^1$ , the first layer’s extractor, ”reveal” the important coordinates, in the sense that  $V^1 \approx V^1 P_{\mathcal{I}}$ , where  $P_{\mathcal{I}}$  is the projection on the coordinates in  $\mathcal{I}$ . It is also shown that the representation  $\Psi_1$  computed by the first layer, is expressive enough so that  $p_0(\sum_{i \in \mathcal{I}} x_i)$  and  $p_2(\sum_{i \in \mathcal{I}} x_i)$  can be well approximated by functions  $\mathbf{x} \mapsto \langle \mathbf{v}_0, \Psi_1(\mathbf{x}) \rangle$  and  $\mathbf{x} \mapsto \langle \mathbf{v}_2, \Psi_1(\mathbf{x}) \rangle$  for vectors  $\mathbf{v}_0, \mathbf{v}_2$  with a norm bound that do not depend on  $d$ , but only on  $|\mathcal{I}|$ . In section C.2 it is then shown by induction that  $\Psi_t$  is expressive enough so that  $p_0(\sum_{i \in \mathcal{I}} x_i), \dots$ , and  $p_{2^t}(\sum_{i \in \mathcal{I}} x_i)$  can be well approximated by functions  $\mathbf{x} \mapsto \langle \mathbf{v}_i, \Psi_t(\mathbf{x}) \rangle$  for vectors  $\mathbf{v}_i$  with a norm bound that depend only on  $|\mathcal{I}|$ . The reason is that each  $p_i$  can be represented as a quadratic polynomial of  $p_j$  for  $j \leq \lceil \frac{i}{2} \rceil$ , with bounded coefficients, together with the fact that by the induction hypothesis these  $p_j$ ’s can be expressed as a linear function on top of  $\Psi_{t-1}$

## 6. Extensions and Discussion

In the following section we will show two extensions of the RedEx architecture – using a norm formulation of the objective and extension to a convolutional structure. Importantly, the norm formulation of RedEx for a one-dimensional output can be trained without the semi-definite constraints, and thus be trained using standard gradient descent or any other non-constrained convex optimization methods.

### 6.1. Norm Formulation of RedEx and Relation to Trace norm

In Algorithm 2 we gave a constrained optimization problem which can be solved using SDPs. In this section we show how to reformulate this problem as an unconstrained optimization via a new norm we define:

**Definition 9** For  $\mathbf{A}_{1:k} = (A_1, \dots, A_k) \in \mathbb{R}^{d \times d}$  where each  $A_i$  is symmetric we define the **RedEx norm** as:

$$\|\mathbf{A}_{1:k}\|_{R_x} := \min \{ \text{Tr}(R) : R \succeq 0, \text{ and } \forall i, -R \preceq A_i \preceq R \} \quad (7)$$

We first show that the above defined norm satisfies several properties:

**Lemma 10** *Properties of  $\|\cdot\|_{R_x}$*

1.  $\|\cdot\|_{R_x}$  is a norm on  $k$ -tuples of symmetric matrices.
2. If  $k = 1$ , then  $\|\cdot\|_{R_x}$  is equivalent to the trace norm. Additionally, if we write  $A = U^\top D U$  for an orthogonal  $U$  and diagonal  $D$ , then  $R = U^\top |D| U$ .

The proof can be found in Appendix D. Item (3) gives a very simple expression for the RedEx norm in the case for  $k = 1$ , however we are not aware of a simple expression for  $\|\cdot\|_{R_x}$  where  $k \geq 2$ . We can now optimize Eq. (2) using the RedEx norm. For that we replace the minimization problem in Algorithm 2 by:

$$\min \quad \frac{1}{m} \sum_{i=1}^m \ell_{y_i} (\mathbf{x}_i^\top A_1 \mathbf{x}_i, \dots, \mathbf{x}_i^\top A_k \mathbf{x}_i) + \lambda_1 \|\mathbf{A}\|_{R_x} + \lambda_2 \|\mathbf{A}\|_{\text{fr}}^2 \quad (8)$$

We give the full algorithm in Appendix D. Note that we don't need to minimize over the Frobenius norm of  $R$ , as it is already done by minimizing the RedEx norm. The caveat of Eq. (8) is that we currently don't know how to calculate the gradient of the RedEx norm directly (i.e. without calculating  $R$ ), or the projection on norm-induced balls unless we resort to general convex SDP solvers. Hence, at the moment we don't know how to utilize the norm formulation in order to design faster algorithms.

One major practical improvement on the training of RedEx that we can make is in the case where our goal is to learn a function  $f : \mathbb{R}^d \rightarrow \mathbb{R}$  (i.e. the output dimension  $k = 1$ ). For this case, we can use the characterization in Lemma 10 (3), where for  $k = 1$  the RedEx norm is equivalent to the trace norm. In this case, we can replace the minimization problem in Algorithm 2 by:

$$\min \quad \frac{1}{m} \sum_{i=1}^m \ell_{y_i} (\mathbf{x}_i^\top A \mathbf{x}_i) + \lambda_1 \|A\|_{\text{Tr}} + 2\lambda_2 \|A\|_{\text{fr}}^2$$

This problem is substantially easier than minimizing Eq. (3), since it is an unconstrained convex optimization problem that can be solved by standard GD or SGD. Note that to find  $V$  we don't need to find  $R$ , since by Lemma 10 (3) we can compute a diagonalization  $A = U^\top D U$ , and then output  $V = \sqrt{|D|} U$  and  $P = (V^\dagger)^\top A V^\dagger$ . We give the full algorithm in Appendix D.

## 6.2. Convolutions

One of the advantages of neural networks is that it allows to choose an architecture according to the structure of the data. A central example is convolutional networks for data which is translation invariant such as images. The input vector for a convolutional layer is divided into patches. In

other words, it is a vector  $(\mathbf{x}_1, \dots, \mathbf{x}_p) \in (\mathbb{R}^d)^p$ . A convolutional layer applies on each patch the same linear function followed by a non-linearity. That is, it computes a mapping of the form  $(\mathbf{x}_1, \dots, \mathbf{x}_p) \mapsto (\sigma(W\mathbf{x}_1), \dots, \sigma(W\mathbf{x}_p))$

For a matrix  $W \in \mathbb{R}^{d' \times d}$  and some non-linearity  $\sigma$ . A convolutional extractor-expander works in a similar fashion. It applies the same extractor  $V$  to all patches, and then expand each patch quadratically. This is detailed in the following definition.

**Definition 11 (Convolutional RedEx)** A function  $\Lambda : (\mathbb{R}^d)^p \rightarrow \mathbb{R}^k$  is called a **Convolutional RedEx** of width  $M$  if it is of the form:  $\Lambda = \Psi_P \circ \Psi_V$  where:

1. The function  $\Psi_V : (\mathbb{R}^d)^p \rightarrow (\mathbb{R}^{d' \times d'})^p$  is of the form:

$$\Psi_V(\mathbf{x}_1, \dots, \mathbf{x}_p) = ((V\mathbf{x}_1)^{\otimes 2}, \dots, (V\mathbf{x}_p)^{\otimes 2})$$

for  $V \in \mathbb{R}^{d' \times d}$  with  $d' \leq d$ , orthogonal rows, and  $\|V\|_{\text{fr}}^2 \leq M$ .

2. The function  $\Psi_P : (\mathbb{R}^{d' \times d'})^p \rightarrow \mathbb{R}^k$  is of the form:

$$\Psi_P(X_1, \dots, X_p) = \left( \sum_{j=1}^p \langle P_{1,j}, X_j \rangle, \dots, \sum_{j=1}^p \langle P_{k,j}, X_j \rangle \right)$$

for  $P_{i,j} \in \mathbb{R}^{d' \times d'}$  with  $\|P_{i,j}\| \leq 1$ .

The function  $\Psi_V$  is called a **convolutional extractor-expander**.

As with the basic version of RedEx, we can extend the basic convolutional RedEx architecture to a multilayer architecture. Likewise, a single layer of convolutional RedEx s can be trained efficiently, similarly to a single layer of RedEx. A multilayer convolutional RedEx can be trained efficiently in a layerwise manner, as basic RedEx. We outline next the algorithm for learning a single convolutional RedEx layer. The extension to multilayer is straight forward.

---

### Algorithm 3 Training 1-layer convolutional RedEx

---

- 1: **Parameters:** A loss  $\ell : \mathbb{R}^k \times \mathcal{Y} \rightarrow [0, \infty)$ , width parameter  $\lambda_1$  and a regularization parameter  $\lambda_2$
- 2: **Input:** A dataset  $(\mathbf{x}_1, y_1), \dots, (\mathbf{x}_m, y_m) \in (\mathbb{R}^d)^p \times \mathcal{Y}$
- 3: Find symmetric  $d \times d$  matrices  $A_{i,j}, R$  for  $1 \leq i \leq k$  and  $1 \leq j \leq p$  by solving the semi-definite program:

$$\begin{aligned} \min \quad & \frac{1}{m} \sum_{i=1}^m \ell_{y_i} \left( \sum_{j=1}^p \mathbf{x}_{i,j}^\top A_{1,j} \mathbf{x}_{i,j}, \dots, \sum_{j=1}^p \mathbf{x}_{i,j}^\top A_{k,j} \mathbf{x}_{i,j} \right) + \lambda_1 \text{Tr}(R) + \lambda_2 (\|R\|_{\text{fr}}^2 + \|A\|_{\text{fr}}^2) \\ \text{s.t.} \quad & -R \preceq A_{i,j} \preceq R \\ & R \succeq 0 \end{aligned}$$

- 4: Compute an orthogonal diagonalization  $R = U^\top D U$
  - 5: Output  $V = \sqrt{D} U$  and  $P_{i,j} = (V^\dagger)^\top A_{i,j} V^\dagger$ .
-

### 6.3. Conclusions and Future Work

In this work we presented the novel RedEx architecture. This architecture is as expressive as neural networks, and can be trained in a layer-wise fashion using convex programs with semi-definite constraints. We also provided a separation result between RedEx and fixed representation methods based on a variation of the sparse-parity problem. Finally, we have shown several extensions of RedEx to the convolutional setting and replacing the semi-definite constraints to adding norm regularizers based on the newly introduced RedEx norm. Notably, for a one-dimensional input, it allows training of RedEx using non-constrained convex optimization algorithms such as gradient descent.

We believe our work can lead to more efficient representation learning methods based on convex optimization. This can include better and richer architectures, which may allow more efficient implementations that can be provably learned without the use of heavy convex SDP algorithms. Finally, it is interesting to provide stronger separation results between RedEx and fixed representation methods under milder assumptions, e.g. in the case where the output is one-dimensional.

### References

- Emmanuel Abbe, Enric Boix-Adsera, Matthew S Brennan, Guy Bresler, and Dheeraj Nagaraj. The staircase property: How hierarchical structure can guide deep learning. *Advances in Neural Information Processing Systems*, 34:26989–27002, 2021.
- Zeyuan Allen-Zhu and Yuanzhi Li. Backward feature correction: How deep learning performs deep learning. *arXiv preprint arXiv:2001.04413*, 2020.
- Zeyuan Allen-Zhu, Yuanzhi Li, and Yingyu Liang. Learning and generalization in overparameterized neural networks, going beyond two layers. *Advances in neural information processing systems*, 32, 2019.
- A. Andoni, R. Panigrahy, G. Valiant, and L. Zhang. Learning polynomials with neural networks. In *Proceedings of the 31st International Conference on Machine Learning*, pages 1908–1916, 2014.
- Alfred Auslender and Marc Teboulle. Interior gradient and proximal methods for convex and conic optimization. *SIAM Journal on Optimization*, 16(3):697–725, 2006.
- Shai Ben-David, Nadav Eiron, and Hans Ulrich Simon. Limitations of learning via embeddings in euclidean half spaces. *Journal of Machine Learning Research*, 3(Nov):441–461, 2002.
- Yoshua Bengio, Pascal Lamblin, Dan Popovici, and Hugo Larochelle. Greedy layer-wise training of deep networks. *Advances in neural information processing systems*, 19, 2006.
- Alberto Bietti, Joan Bruna, Clayton Sanford, and Min Jae Song. Learning single-index models with shallow neural networks. *Advances in Neural Information Processing Systems*, 35:9768–9783, 2022.
- Joan Bruna, Loucas Pillaud-Vivien, and Aaron Zweig. On single index models beyond gaussian data. *arXiv preprint arXiv:2307.15804*, 2023.
- Sébastien Bubeck et al. Convex optimization: Algorithms and complexity. *Foundations and Trends® in Machine Learning*, 8(3-4):231–357, 2015.

- Emmanuel J Candes, Thomas Strohmer, and Vladislav Voroninski. Phaselift: Exact and stable signal recovery from magnitude measurements via convex programming. *Communications on Pure and Applied Mathematics*, 66(8):1241–1274, 2013.
- Emmanuel J Candes, Xiaodong Li, and Mahdi Soltanolkotabi. Phase retrieval via wirtinger flow: Theory and algorithms. *IEEE Transactions on Information Theory*, 61(4):1985–2007, 2015.
- Yuan Cao and Quanquan Gu. Generalization bounds of stochastic gradient descent for wide and deep neural networks. *arXiv preprint arXiv:1905.13210*, 2019.
- Joachim Dahl and Erling D Andersen. A primal-dual interior-point algorithm for nonsymmetric exponential-cone optimization. *Mathematical Programming*, 194(1-2):341–370, 2022.
- Amit Daniely. Sgd learns the conjugate kernel class of the network. In *Advances in Neural Information Processing Systems*, pages 2422–2430, 2017.
- Amit Daniely. Neural networks learning and memorization with (almost) no over-parameterization. In *NeurIPS*, 2020.
- Amit Daniely and Eran Malach. Learning parities with neural networks. *Advances in Neural Information Processing Systems*, 33:20356–20365, 2020.
- Simon S Du, Jason D Lee, Yuandong Tian, Barnabas Poczos, and Aarti Singh. Gradient descent learns one-hidden-layer cnn: Don’t be afraid of spurious local minima. *arXiv preprint arXiv:1712.00779*, 2017.
- Spencer Frei, Yuan Cao, and Quanquan Gu. Agnostic learning of a single neuron with gradient descent. *Advances in Neural Information Processing Systems*, 33:5417–5428, 2020.
- Rong Ge, Jason D Lee, and Tengyu Ma. Learning one-hidden-layer neural networks with landscape design. *arXiv preprint arXiv:1711.00501*, 2017.
- Behrooz Ghorbani, Song Mei, Theodor Misiakiewicz, and Andrea Montanari. Limitations of lazy training of two-layers neural network. In *Advances in Neural Information Processing Systems*, pages 9108–9118, 2019.
- Behrooz Ghorbani, Song Mei, Theodor Misiakiewicz, and Andrea Montanari. Linearized two-layers neural networks in high dimension. *The Annals of Statistics*, 49(2):1029–1054, 2021.
- Arthur Jacot, Franck Gabriel, and Clément Hongler. Neural tangent kernel: Convergence and generalization in neural networks. In *Advances in neural information processing systems*, pages 8571–8580, 2018.
- Haotian Jiang, Tarun Kathuria, Yin Tat Lee, Swati Padmanabhan, and Zhao Song. A faster interior point method for semidefinite programming. In *2020 IEEE 61st annual symposium on foundations of computer science (FOCS)*, pages 910–918. IEEE, 2020.
- Pritish Kamath, Omar Montasser, and Nathan Srebro. Approximate is good enough: Probabilistic variants of dimensional and margin complexity. In *Conference on Learning Theory*, pages 2236–2262. PMLR, 2020.

- Michael Kearns and Leslie Valiant. Cryptographic limitations on learning boolean formulae and finite automata. *Journal of the ACM (JACM)*, 41(1):67–95, 1994.
- Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- Guy Kornowski and Ohad Shamir. Oracle complexity in nonsmooth nonconvex optimization. *Advances in Neural Information Processing Systems*, 34:324–334, 2021.
- Yuanzhi Li and Yingyu Liang. Learning overparameterized neural networks via stochastic gradient descent on structured data. *Advances in neural information processing systems*, 31, 2018.
- Roi Livni, Shai Shalev-Shwartz, and Ohad Shamir. On the computational efficiency of training neural networks. In *Advances in neural information processing systems*, pages 855–863, 2014.
- Wolfgang Maass. Networks of spiking neurons: the third generation of neural network models. *Neural networks*, 10(9):1659–1671, 1997.
- Eran Malach, Pritish Kamath, Emmanuel Abbe, and Nathan Srebro. Quantifying the benefit of using differentiable learning over tangent kernels. In *International Conference on Machine Learning*, pages 7379–7389. PMLR, 2021a.
- Eran Malach, Gilad Yehudai, Shai Shalev-Schwartz, and Ohad Shamir. The connection between approximation, depth separation and learnability in neural networks. In *Conference on Learning Theory*, pages 3265–3295. PMLR, 2021b.
- Arnold F Nikiforov, Vasilii B Uvarov, Sergei K Suslov, Arnold F Nikiforov, Vasilii B Uvarov, and Sergei K Suslov. *Classical orthogonal polynomials of a discrete variable*. Springer, 1991.
- Florian A Potra and Stephen J Wright. Interior-point methods. *Journal of computational and applied mathematics*, 124(1-2):281–302, 2000.
- A. Rahimi and B. Recht. Random features for large-scale kernel machines. In *NIPS*, 2007.
- Itay Safran and Ohad Shamir. Spurious local minima are common in two-layer relu neural networks. In *International conference on machine learning*, pages 4433–4441. PMLR, 2018.
- Shai Shalev-Shwartz and Shai Ben-David. *Understanding machine learning: From theory to algorithms*. Cambridge university press, 2014.
- Shai Shalev-Shwartz, Alon Gonen, and Ohad Shamir. Large-scale convex minimization with a low-rank constraint. *arXiv preprint arXiv:1106.1622*, 2011.
- Yuandong Tian. An analytical formula of population gradient for two-layered relu network and its applications in convergence and critical point analysis. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 3404–3413. JMLR. org, 2017.
- Gal Vardi, Gilad Yehudai, and Ohad Shamir. Learning a single neuron with bias using gradient descent. *Advances in Neural Information Processing Systems*, 34:28690–28700, 2021.



Heribert Vollmer. *Introduction to circuit complexity: a uniform approach*. Springer Science & Business Media, 1999.

Huiyuan Wang and Wei Lin. Harmless overparametrization in two-layer neural networks. *arXiv preprint arXiv:2106.04795*, 2021.

Blake Woodworth, Suriya Gunasekar, Jason D Lee, Edward Moroshko, Pedro Savarese, Itay Golan, Daniel Soudry, and Nathan Srebro. Kernel and rich regimes in overparametrized models. *arXiv preprint arXiv:2002.09277*, 2020.

Gilad Yehudai and Ohad Shamir. On the power and limitations of random features for understanding neural networks. *arXiv preprint arXiv:1904.00687*, 2019.

Gilad Yehudai and Ohad Shamir. Learning a single neuron with gradient methods. *arXiv preprint arXiv:2001.05205*, 2020.

Chulhee Yun, Suvrit Sra, and Ali Jadbabaie. Small nonlinearities in activation functions create bad local minima in neural networks. *arXiv preprint arXiv:1802.03487*, 2018.

## Appendix A. Proofs from Sec. 3

**Proof** [Proof of Thm. 2] We first show that we can implement AND, OR, NEG and Id using a RedEx with  $O(1)$  layers and feature dimension of  $O(d^2)$ . Recall that by definition, we added a coordinate to the data which is constant 1. To implement Id of the  $i$ -th coordinate, we can use a matrix  $V$  where the  $i$ -th row is equal to  $\mathbf{e}_i$ , and the last row (which corresponds to the constant 1) equal to  $\mathbf{e}_{d+1}$ , this way  $((V\mathbf{x})(V\mathbf{x})^\top)_{i,d+1} = x_i$

NEG can be implemented by  $x \mapsto 1 - x$ . This can be implemented for coordinate  $i$  by having the  $i$ -th row of  $V$  equal to  $-\mathbf{e}_i + \mathbf{e}_{d+1}$ . We also need the last row of  $V$  to be  $\mathbf{e}_{d+1}$ , this way  $((V\mathbf{x})(V\mathbf{x})^\top)_{i,d+1} = 1 - x_i$ .

Now,  $\text{AND}(x_i, x_j) = x_i \cdot x_j$  can be implemented by having  $V$  with  $i$ -th row equal to  $\mathbf{e}_i$ , and  $j$ -th row equal to  $\mathbf{e}_j$ , this way  $((V\mathbf{x})(V\mathbf{x})^\top)_{i,j} = x_i \cdot x_j$ . Finally, we have  $\text{OR}(x_1, x_2) = x_1 + x_2 - x_1 \cdot x_2$ . This can be implemented by applying Id and AND on  $x_1, x_2$ , If the output of the above operations are in rows  $i, j, k$  correspondingly, then we need some row of  $V$  to be equal to  $\mathbf{e}_i + \mathbf{e}_j - \mathbf{e}_k$  and the last row of  $V$  to be equal to  $\mathbf{e}_{d+1}$ .

Note that if in the process of the quadratic expansion of RedEx we added extra coordinates which are not needed, in the next layer we can use zero rows for the unnecessary coordinates to zero them out. This way, the application of the dimension reduction method would delete those unnecessary coordinates since their output is constant zero. Note that each operation above was implemented using at most 2-layer RedEx, hence the feature dimension is at most  $d^2$  where  $d$  is the dimension of the input. For a general intermediate layer, we can bound its input by the total size of the target binary circuit, hence we can bound the feature dimension by  $O(T^2)$ . In addition, since each operation can be implemented by a RedEx of depth  $O(1)$ , the total depth of the RedEx which implements the Boolean circuit is  $O(T)$ . ■

## Appendix B. Proofs from Sec. 4

### B.1. Proof of Thm. 3

We first need the following lemma:

**Lemma 12** *If  $-R \preceq A \preceq R$  then  $\ker(R) \subset \ker(A)$*

**Proof** Since  $A$  is symmetric and by the assumption of the lemma, it has a (non-unique) decomposition as  $A = A_+ + A_-$  where  $A_+$  is positive semi-definite with  $A_+ \preceq R$  and  $A_-$  is negative semi-definite with  $-R \preceq A_-$ . Let  $0 \neq \mathbf{x} \in \ker(R)$ , then  $0 \leq \langle \mathbf{x}, A_+\mathbf{x} \rangle \leq \langle \mathbf{x}, R\mathbf{x} \rangle = 0$ . Since  $A_+$  is PSD it has an orthogonal diagonalization with orthonormal eigenvectors  $\mathbf{v}_i$  and corresponding eigenvalues  $\lambda_i > 0$ . We can expand  $\mathbf{x}$  in this basis  $\mathbf{x} = \sum_i \alpha_i \mathbf{v}_i$ . Now we have that:

$$0 = \langle \mathbf{x}, A_+\mathbf{x} \rangle = \left\langle \sum_i \alpha_i \mathbf{v}_i, A_+ \sum_i \alpha_i \mathbf{v}_i \right\rangle = \left\langle \sum_i \alpha_i \mathbf{v}_i, \sum_i \lambda_i \alpha_i \mathbf{v}_i \right\rangle = \sum_i \alpha_i^2 \lambda_i .$$

Hence, for every  $i$  either  $\alpha_i = 0$  or  $\lambda_i = 0$ , in particular,  $\mathbf{x}$  is in the kernel of  $A_+$ . Using a similar argument we get that  $\mathbf{x}$  is in the kernel of  $A_-$ , hence it is in the kernel of  $A$ . ■

We are now ready to prove the main theorem:

**Proof** [Thm. 3] In the proof, for ease of notations we use the notion of RedEx norm, see Definition 9. We begin with the first item. Let  $h_{V,\mathbf{P}} \in \mathcal{H}_M$ , we can write:

$$\begin{aligned} h_{V,\mathbf{P}}(\mathbf{x}) &= \left( \text{Tr} \left( P_1(V\mathbf{x})(V\mathbf{x})^\top \right), \dots, \text{Tr} \left( P_k(V\mathbf{x})(V\mathbf{x})^\top \right) \right) \\ &= \left( \text{Tr} \left( (V\mathbf{x})(V\mathbf{x})^\top P_1^\top \right), \dots, \text{Tr} \left( (V\mathbf{x})(V\mathbf{x})^\top P_k^\top \right) \right) \\ &= \left( \text{Tr} \left( (V\mathbf{x})(P_1 V\mathbf{x})^\top \right), \dots, \text{Tr} \left( (V\mathbf{x})(P_k V\mathbf{x})^\top \right) \right) \\ &= (\langle V\mathbf{x}, P_1 V\mathbf{x} \rangle, \dots, \langle V\mathbf{x}, P_k V\mathbf{x} \rangle) \\ &= \left( \langle \mathbf{x}, V^\top P_1 V\mathbf{x} \rangle, \dots, \langle \mathbf{x}, V^\top P_k V\mathbf{x} \rangle \right) \end{aligned}$$

It is therefore enough to show that for  $A_i := V^\top P_i V$  we have  $\|\mathbf{A}\|_{\text{Rx}} \leq M$ . Since  $\|P_1\|_{\text{sp}} \leq 1$  we have  $-I \preceq P_i \preceq I$ . Hence, also  $-V^\top V \preceq V^\top P_i V \preceq V^\top V$ . This implies that

$$\|\mathbf{A}\|_{\text{Rx}} \leq \text{Tr}(V^\top V) = \text{Tr}(VV^\top) \leq M.$$

For the other direction, let  $h^{R,\mathbf{A}} \in \mathcal{H}^M$  with  $\|\mathbf{A}\|_{\text{Rx}} \leq M$ . Let  $R \succeq 0$  be a matrix which satisfies  $-R \preceq A_i \preceq R$  for every  $i \in [k]$ , and  $\text{Tr}(R) \leq M$ . By diagonalizing  $R$  we can write  $R = U^\top D U$ , where  $U$  is unitary and  $D$  is PSD. Define:

$$V = \sqrt{D}U, \quad P_i = (V^\dagger)^\top A_i (V^\dagger)$$

Here  $\sqrt{D}$  is the diagonal matrix equal to  $\sqrt{D_{i,i}} = \sqrt{D_{i,i}}$ , and  $V^\dagger$  is the pseudo-inverse of  $V$ . By its definition,  $R$  has orthogonal rows. We also have that:

$$\begin{aligned} \|P_i\| &= \|(V^\dagger)^\top A_i (V^\dagger)\| \\ &= \left\| \left( (\sqrt{D}U)^\dagger \right)^\top A_i (\sqrt{D}U)^\dagger \right\| \\ &\leq \|D^\dagger A_i\| \leq 1 \end{aligned}$$

where we used that  $U$  is orthogonal and  $A_i \preceq R$ . Now, we have that:

$$\|V\|_{\text{fr}}^2 = \text{Tr}(V^\top V) = \text{Tr}((\sqrt{D}U)^\top \sqrt{D}U) = \text{Tr}(DU^\top U) = \text{Tr}(D) = \text{Tr}(R) \leq M$$

where we used that  $U$  is orthogonal. We have shown that for our definitions of  $V$  and  $\mathbf{P}$  we have that  $h_{V,\mathbf{P}} \in \mathcal{H}_M$ , it is left to show that for every  $\mathbf{x} \in \mathbb{R}^d$  we have  $h_{R,\mathbf{A}}(\mathbf{x}) = h_{V,\mathbf{P}}(\mathbf{x})$ .

Define  $Q := V^\dagger V$ , this is the projection on the range of  $V^\top$ , which contains the range of  $R$ . Hence,  $I - Q$  is the projection on the orthogonal complement of the range of  $V$ , which is contained in the orthogonal complement of the range of  $R$ , which is the kernel of  $R$ . By Lemma 12 we have that  $\ker(R) \subseteq \ker(A_i)$  for every  $i \in [k]$ . Hence,  $A_i(I - Q) = (I - Q)A_i = 0$ , this implies that:

$$\begin{aligned} A_i &= (I - Q + Q)A_i(I - Q + Q) \\ &= (I - Q)A_i(I - Q) + (I - Q)A_iQ + QA_i(I - Q) + QA_iQ \\ &= QA_iQ \end{aligned}$$

Finally, using that  $P_i = (V^\dagger)^\top A_i V^\dagger$  we have for every  $\mathbf{x} \in \mathbb{R}^d$ :

$$\begin{aligned} h_{V, \mathbf{P}}(\mathbf{x}) &= \left( \langle \mathbf{x}, V^\top P_1 V \mathbf{x} \rangle, \dots, \langle \mathbf{x}, V^\top P_k V \mathbf{x} \rangle \right) \\ &= \left( \langle \mathbf{x}, V^\top (V^\dagger)^\top A_1 V^\dagger V \mathbf{x} \rangle, \dots, \langle \mathbf{x}, V^\top (V^\dagger)^\top A_k V^\dagger V \mathbf{x} \rangle \right) \\ &= \left( \langle \mathbf{x}, Q^\top A_1 Q \mathbf{x} \rangle, \dots, \langle \mathbf{x}, Q^\top A_k Q \mathbf{x} \rangle \right) \\ &= (\langle \mathbf{x}, A_1 \mathbf{x} \rangle, \dots, \langle \mathbf{x}, A_k \mathbf{x} \rangle) \\ &= h_{R, \mathbf{A}}(\mathbf{x}) \end{aligned}$$

This finishes the first part of the proof. For the second part of the theorem, we use the following algorithm to compute  $V$  and  $\mathbf{P}$  given  $\mathbf{A}$ :

1. Find  $R \succeq 0$  such that  $\|\mathbf{A}\|_{\text{Rx}} = \text{Tr}(R)$  and  $R \preceq A_i \preceq R$ .
2. Compute an orthogonal diagonalization  $R = U^\top D U$ .
3. Output  $V = \sqrt{D} U$  and  $P_i = (V^\dagger)^\top A_i V^\dagger$ .

The first step can be completed in polynomial time using SDP solvers since this is a convex problem with linearly many constraints, see e.g. [Jiang et al. \(2020\)](#). The second step can also be done in polynomial time as it only consists of diagonalizing a symmetric matrix.  $\blacksquare$

### Appendix C. Proof of Theorem 7

We first introduce some notation. For  $\mathbf{x} \in \{\pm 1\}^d$  denote  $\mathbf{x}(\mathcal{I}) = \sum_{i \in \mathcal{I}} x_i$  and  $P_i(\mathbf{x}) = p_i(\mathbf{x}(\mathcal{I}))$ . Denote  $A_i^t = (V^t)^\top P_i^t V^t$ ,  $R^t = (V^t)^\top V^t$ ,  $\Psi_t = \Psi_{V^t} \circ \dots \circ \Psi_{V^{(1)}}$  and  $h_i^t(\mathbf{x}) = \Psi_t(\mathbf{x})^\top A_i^t \Psi_t(\mathbf{x})$ . Denote also  $\Psi_0(\mathbf{x}) = (c, \mathbf{x}) \in \mathbb{R}^{d+1}$  and will refer to the first coordinate in  $\Psi_0(\mathbf{x})$  (the constant coordinate) as the 0'th coordinate (instead of 1'th). We note that  $\Psi_t$  computes a polynomial of degree  $\leq 2^t$ . Hence,  $h_i^t$  is orthogonal to  $p_{2i}(\mathbf{x}(\mathcal{I}))$  for  $2i > 2^{t+1}$ . Thus, for  $i \geq 2^t$ , the optimal solution to the  $i$ 'th coordinate of the  $t$ 'th layer is 0. This observation motivated the definition of the  $t$ -truncated loss given by

$$\text{Val}_{\Psi_{t-1}}^t(\mathbf{A}, R) = \sum_{j=0}^{2^{t-1}} \mathbb{E}_{\mathbf{x}} \left( p_{2j}(\mathbf{x}(\mathcal{I})) - \Psi_{t-1}(\mathbf{x})^\top A_j \Psi_{t-1}(\mathbf{x}) \right)^2$$

We also denote

$$\text{Val}_{\Psi_{t-1}}^t = \inf_{(\mathbf{A}, R)} \text{Val}_{\Psi_{t-1}}^t(\mathbf{A}, R)$$

We denote by  $P_{\mathcal{I}} : \mathbb{R}^{d+1} \rightarrow \mathbb{R}^{d+1}$  the projection on the coordinate in  $\mathcal{I}$  and the first (constant) coordinate. That is,  $(P_{\mathcal{I}} \mathbf{x})_j = \begin{cases} x_j & j \in \mathcal{I} \text{ or } j = 0 \\ 0 & \text{otherwise} \end{cases}$ .

Before proceeding to the main body of the proof, we specialize theorem 6 for the square loss  $\ell_y(\hat{y}) = \|\hat{y} - y\|^2$ . While the square loss is not globally Lipschitz, it is Lipschitz on any bounded domain. Specifically, we have  $\nabla_{\hat{y}} \ell_y(\hat{y}) = 2(\hat{y} - y)$ . We also have that  $\|\mathbf{x}^\top \mathbf{A} \mathbf{x}\| \leq \sqrt{k} M \|\mathbf{x}\|^2$ . Hence, if  $\Pr_{(x,y) \sim \mathcal{D}}(\|y\| \leq M_2) = 1$  then we have the the square loss is  $(2M_2 + 2\sqrt{k} M_1^2 M)$ -Lipchitz in the relevant domain. Hence,

**Corollary 13** Assume that the dataset is an i.i.d. sample from a distribution  $\mathcal{D}$  on  $B_{M_1}^d \times B_{M_2}^k$  and that the loss is the square loss. Let  $(\mathbf{A}, R)$  be the output of algorithm 2. Then

$$\mathbb{E}_S \text{Val}_{\mathcal{D}}(\mathbf{A}, R) \leq \text{Val}_{\mathcal{D}, M} + \lambda(k+1)M^2 + \frac{M_1^4(2M_2 + 2\sqrt{k}M_1^2M)^2}{\lambda m}$$

In our case, Eq. (6) implies that  $|p_i(x)| \leq k^i$ . Hence, the output of the learned function is in  $B_{k^k}^{1+k/2}$ . Since  $\lambda = \frac{1}{\sqrt{m}}$  we get

$$\mathbb{E}_S \text{Val}_{\mathcal{D}}(\mathbf{A}, R) \leq \text{Val}_{\mathcal{D}, M} + \frac{(k+1)M^2 + M_1^4(2k^k + 2\sqrt{k}M_1^2M)^2}{\sqrt{m}}$$

### C.1. First layer

Let  $\mathbf{1}_{\mathcal{I}}, \mathbf{1}_0 \in \mathbb{R}^{d+1}$  be the indicator vectors of  $\mathcal{I}$  and  $\{0\}$ . Denote by  $J, I_{\mathcal{I}}$  and  $I_0$  the  $(d+1) \times (d+1)$  matrices given by

$$J = \frac{1}{k} \mathbf{1}_{\mathcal{I}} \mathbf{1}_{\mathcal{I}}^\top, \quad I_0 = \mathbf{1}_0 \mathbf{1}_0^\top, \quad I_{\mathcal{I}} = \frac{1}{k} (P_{\mathcal{I}} - I_0)$$

We note that

$$A_0 = \frac{1}{2} I_0, \quad A_1 = \frac{1}{\sqrt{2-2/k}} J - \frac{1}{2\sqrt{2-2/k}} I_0, \quad R = \frac{1}{2\sqrt{2-2/k}} J + \frac{1}{2} I_0$$

Is a solution to the first layer with zero 1-truncated loss. Thus, by corollary 13 we will have  $\text{Val}_{\Psi_0}^1(\mathbf{A}^1, R^1) \leq \frac{2(k+1)+(d+1)^2(2k^k+4\sqrt{k}(d+1))^2}{\delta\sqrt{m}} =: \epsilon_1$  w.p.  $1 - \delta$ . The following lemma shows that in this case it holds that  $\|V^1(I - P_{\mathcal{I}})\|_{\text{sp}} \leq (8d\epsilon_1)^{\frac{1}{4}} =: \epsilon$ .

**Lemma 14** If  $\text{Val}_{\Psi_0}^1(\mathbf{A}^1, R^1) \leq \epsilon^2$  then  $\|V^1(I - P_{\mathcal{I}})\|_{\text{sp}} \leq (8d)^{1/4} \sqrt{\epsilon}$

**Proof** We have that  $\|h_0^1 - P_0\|_2^2 + \|h_1^1 - P_2\|_2^2 \leq \epsilon^2$ . Hence, there are  $(d+1) \times (d+1)$  matrices  $A'_0, A'_1$  with  $\|A_0^1 - A'_0\|_F^2 + \|A_1^1 - A'_1\|_F^2 \leq \epsilon^2$  such that  $P_0(\mathbf{x}) = \Psi_0(\mathbf{x})^\top A'_0 \Psi_0(\mathbf{x})$  and  $P_2(\mathbf{x}) = \Psi_0(\mathbf{x})^\top A'_1 \Psi_0(\mathbf{x})$ . Since  $\|A'_i - A_i^1\|_{\text{Tr}} \leq \sqrt{d} \|A'_i - A_i^1\|_F \leq \sqrt{d}\epsilon$ , there is a PSD matrix  $R'$  such that  $-R' \preceq A'_i \preceq R'$  and  $\text{Tr}(R') \leq \text{Tr}(R^1) + \sqrt{2d}\epsilon \leq \frac{1}{2} + \frac{1}{\sqrt{2-2/k}} + \sqrt{2d}\epsilon$ .

Now, consider the matrices  $A''_i, R''$  obtained by zeroing  $(A'_i)_{jj}, R'_{jj}, R'_{jl}$  and  $R'_{lj}$  for any  $j \in [d] \setminus \mathcal{I}$  and  $0 \leq l \leq d$ , and adding  $\frac{1}{2} \sum_{j \in [d] \setminus \mathcal{I}} (A'_i)_{jj}$  to  $(A'_i)_{00}$  as well as  $\frac{1}{2} \sum_{j \in [d] \setminus \mathcal{I}} R'_{jj}$  to  $R'_{00}$ . We have  $\Psi_0(\mathbf{x})^\top A''_i \Psi_0(\mathbf{x}) = \Psi_0(\mathbf{x})^\top A'_i \Psi_0(\mathbf{x})$ ,  $-R'' \preceq A''_i \preceq R''$  and  $\text{Tr}(R'') = \text{Tr}(R') - \frac{1}{2} \sum_{j \in [d] \setminus \mathcal{I}} R'_{jj}$ . lemma 15 now implies that

$$\frac{1}{2} + \frac{1}{\sqrt{2-2/k}} \leq \text{Tr}(R'') = \text{Tr}(R') - \frac{1}{2} \sum_{j \in [d] \setminus \mathcal{I}} R'_{jj} \leq \frac{1}{2} + \frac{1}{\sqrt{2-2/k}} + \sqrt{2d}\epsilon - \frac{1}{2} \sum_{j \in [d] \setminus \mathcal{I}} R'_{jj}$$

Hence,

$$\sum_{j \in [d] \setminus \mathcal{I}} R'_{jj} \leq \sum_{j \in [d] \setminus \mathcal{I}} R'_{jj} \leq \sqrt{8d}\epsilon$$

Finally, we have

$$\begin{aligned}
 \|V^1(I - P_{\mathcal{I}})\|_{\text{sp}}^2 &= \|(I - P_{\mathcal{I}})(V^1)^\top V^1(I - P_{\mathcal{I}})\|_{\text{sp}} \\
 &= \|(I - P_{\mathcal{I}})R^1(I - P_{\mathcal{I}})\|_{\text{sp}} \\
 &\leq \|(I - P_{\mathcal{I}})R^1(I - P_{\mathcal{I}})\|_{\text{Tr}} \\
 &= \sum_{j \in [d] \setminus \mathcal{I}} R_{jj}^1 \\
 &\leq \sqrt{8d}\epsilon
 \end{aligned}$$

■

**Lemma 15** *Let  $A_0, A_1, R$  be  $(d+1) \times (d+1)$  matrices such that  $-R \preceq A_i \preceq R$  and  $P_{2i}(\mathbf{x}) = (\Psi_0(\mathbf{x}))^\top A_i \Psi_0(\mathbf{x})$ . Then,  $\text{Tr}(R) \geq \frac{1}{2} + \frac{1}{\sqrt{2+2/k}}$*

**Proof** We assume w.l.o.g. that  $A_0, A_1, R$  minimizes  $\text{Tr}(R)$  under the above constraints. It is not hard to verify that  $A_0$  and  $A_1$  are linear combination of the PSD matrices  $J, I_0, \tilde{I}_{\mathcal{I}} := \frac{k}{k-1}(I_{\mathcal{I}} - \frac{1}{k}J)$  and  $I_{\mathcal{I}^c} := \frac{1}{d-k}(I - P_{\mathcal{I}})$ . Write

$$A_i = a_i^1 J + a_i^2 I_0 + a_i^3 \tilde{I}_{\mathcal{I}} + a_i^4 I_{\mathcal{I}^c}$$

Since the the matrices  $J, I_0, \tilde{I}_{\mathcal{I}}$  and  $I_{\mathcal{I}^c}$  are supported on orthogonal spaces, the minimal trace of a PSD matrix  $R$  with  $-R \preceq A_i \preceq R$  is  $\sum_{j=1}^4 \max(|a_0^j|, |a_1^j|)$ . Now, zeroing  $a_0^4$  and  $a_1^4$  while adding  $\frac{1}{2}a_0^4$  and  $\frac{1}{2}a_1^4$  to  $a_0^1$  and  $a_1^1$  will not alter the functions computed by  $A_0$  and  $A_1$  and will not increase  $\sum_{j=1}^4 \max(|a_0^j|, |a_1^j|)$ . Thus, we can assume that  $a_0^4 = a_1^4 = 0$ .

Likewise, zeroing  $a_0^3$  and  $a_1^3$  while adding  $-\frac{a_i^3}{k-1}$  to  $a_i^1$ , and  $\frac{k}{k-1} \frac{a_i^3}{2}$  to  $a_i^2$  will not alter the functions computed by  $A_0$  and  $A_1$  and will not increase  $\sum_{j=1}^4 \max(|a_0^j|, |a_1^j|)$ . Thus, we can assume that  $a_0^3 = a_1^3 = 0$ . This implies that  $A_0 = \frac{1}{2}I_0$  and  $A_1 = \frac{1}{\sqrt{2-2/k}}J - \frac{1}{2\sqrt{2-2/k}}I_0$  ■

## C.2. Remaining Layers and conclusion of the proof

Suppose that the  $t$ 'th layer has  $t$ -truncated error at most  $\epsilon_t$  and that  $\|V^1(I - P_{\mathcal{I}})\|_{\text{sp}} \leq \epsilon$ . Lemma 16 below implies that there is a solution for the  $t$ 'th layer with Frobenius norm at most  $2^{k/2}$  and  $t$ -truncated error at most  $\epsilon_t + dk^{O(k^2)}\epsilon^2$ . Lemma 18 below now implies that there is a solution for the  $(t+1)$ 'th layer with  $(t+1)$ -truncated error at most  $k^{O(k^2)}\epsilon_t + dk^{O(k^2)}\epsilon^2$ . By lemma 16 and corollary 13 we have that w.p.  $1 - \delta$ , the  $(t+1)$ 'th layer has  $(t+1)$ -truncated error at most  $k^{O(k^2)}\epsilon_t + dk^{O(k^2)}(\epsilon^2 + 1/\delta\sqrt{m})$ . By induction, we conclude that w.p.  $1 - t\delta$ , the truncated error of the  $t$ 'th layer is  $k^{tO(k^2)}\epsilon_1 + dk^{tO(k^2)}(\epsilon^2 + 1/\delta\sqrt{m})$ . Hence, w.p.  $1 - \log(k)\delta$ , the error of the final layer is

$$k^{O(k^2 \log(k))}\epsilon_1 + dk^{O(k^2 \log(k))}(\epsilon^2 + 1/\delta\sqrt{m}) = \frac{d^4 2^{O(k^2 \log^2(k))}}{\delta m^{1/4}}$$

which concludes the proof.

**Lemma 16** *If  $\|V^1(I - P_{\mathcal{I}})\|_{\text{sp}} \leq \epsilon$  then, for any  $\mathbf{x} \in [-1, 1]^d$  we have,*

1.  $\|\Psi_t(\mathbf{x})\|_F \leq M^{2^t-2}(M|\mathcal{I}| + \epsilon d)^{2^t} \leq M^{3 \cdot 2^t}$
2.  $\|\Psi_t(\mathbf{x}) - \Psi_t(P_{\mathcal{I}}\mathbf{x})\|_F \leq d2^t M^{2^t+3 \cdot (2^t-1)} \epsilon$
3. *For the projection  $P$  on  $\text{span}(V^t \Psi_{t-1}(P_{\mathcal{I}}\{\pm 1\}^d))$  and any symmetric  $A$  with  $\|A\|_{\text{sp}} \leq 1$  and  $\mathbf{x} \in \{\pm 1\}^d$  we have*

$$|\langle A, \Psi_t(\mathbf{x}) \rangle - \langle PAP, \Psi_t(\mathbf{x}) \rangle| \leq d2^t M^{2^t+3 \cdot (2^t-1)} \epsilon \text{ and } \|PAP\|_F \leq 2^{|\mathcal{I}|/2}$$

**Proof** We first prove item 1. by induction on  $t$ . For  $t = 1$  we have

$$\begin{aligned} \|\Psi_1(\mathbf{x})\|_F &= \|V^1 \mathbf{x}\|^2 \\ &\leq (\|V^1(I - P_{\mathcal{I}})\|_{\text{sp}} \|\mathbf{x}\| + \|V^1\|_{\text{sp}} \|P_{\mathcal{I}}\mathbf{x}\|)^2 \\ &\leq (\epsilon \sqrt{d} + M \sqrt{|\mathcal{I}|})^2 \end{aligned}$$

For  $t > 1$  we have by the induction hypothesis

$$\begin{aligned} \|\Psi_t(\mathbf{x})\|_F &= \|V^t \Psi_{t-1}(\mathbf{x})\|^2 \\ &\leq M^2 \left( M^{2^{t-1}-2} (M|\mathcal{I}| + \epsilon d)^{2^{t-1}} \right)^2 \\ &= M^{2^t-2} (M|\mathcal{I}| + \epsilon d)^{2^t} \end{aligned}$$

We next prove item 2. by induction on  $t$ . For  $t = 1$  we have

$$\begin{aligned} \|\Psi_1(\mathbf{x}) - \Psi_1(P_{\mathcal{I}}\mathbf{x})\|_F &= \|(V^1 \mathbf{x}) \otimes (V^1 \mathbf{x}) - (V^1 P_{\mathcal{I}}\mathbf{x}) \otimes (V^1 P_{\mathcal{I}}\mathbf{x})\| \\ &\leq \|(V^1 \mathbf{x}) \otimes (V^1(I - P_{\mathcal{I}})\mathbf{x})\| + \|(V^1(I - P_{\mathcal{I}})\mathbf{x}) \otimes (V^1 P_{\mathcal{I}}\mathbf{x})\| \\ &\leq 2dM\epsilon \end{aligned}$$

For  $t > 1$  we have by the induction hypothesis and item 1.

$$\begin{aligned} \|\Psi_t(\mathbf{x}) - \Psi_t(P_{\mathcal{I}}\mathbf{x})\|_F &= \|(V^t \Psi_{t-1}(\mathbf{x})) \otimes (V^t \Psi_{t-1}(\mathbf{x})) - (V^t \Psi_{t-1}(P_{\mathcal{I}}\mathbf{x})) \otimes (V^t \Psi_{t-1}(P_{\mathcal{I}}\mathbf{x}))\| \\ &\leq \|(V^t \Psi_{t-1}(\mathbf{x})) \otimes (V^t \Psi_{t-1}(\mathbf{x}) - V^t \Psi_{t-1}(P_{\mathcal{I}}\mathbf{x}))\| \\ &\quad + \|(V^t \Psi_{t-1}(\mathbf{x}) - V^t \Psi_{t-1}(P_{\mathcal{I}}\mathbf{x})) \otimes (V^t \Psi_{t-1}(P_{\mathcal{I}}\mathbf{x}))\| \\ &\leq M^2 \cdot \|\Psi_{t-1}(\mathbf{x})\| \cdot \|\Psi_{t-1}(\mathbf{x}) - \Psi_{t-1}(P_{\mathcal{I}}\mathbf{x})\| \\ &\quad + M^2 \cdot \|\Psi_{t-1}(P_{\mathcal{I}}\mathbf{x})\| \cdot \|\Psi_{t-1}(\mathbf{x}) - \Psi_{t-1}(P_{\mathcal{I}}\mathbf{x})\| \\ &\leq 2M^2 M^{3 \cdot 2^{t-1}} 2^{t-1} d M^{2(t-1)+3 \cdot (2^{t-1}-1)} \epsilon \\ &= 2^t d M^{2^t+3 \cdot (2^t-1)} \epsilon \end{aligned}$$

We now prove item 3.

$$\begin{aligned} \langle A, \Psi_t(\mathbf{x}) \rangle - \langle PAP, \Psi_t(\mathbf{x}) \rangle &= (V^t \Psi_{t-1}(\mathbf{x}))^\top A (V^t \Psi_{t-1}(\mathbf{x})) \\ &\quad - (PV^t \Psi_{t-1}(\mathbf{x}))^\top A (PV^t \Psi_{t-1}(\mathbf{x})) \\ &= (V^t \Psi_{t-1}(\mathbf{x}))^\top A (V^t \Psi_{t-1}(\mathbf{x}) - PV^t \Psi_{t-1}(\mathbf{x})) \\ &\quad + (V^t \Psi_{t-1}(\mathbf{x}) - PV^t \Psi_{t-1}(\mathbf{x}))^\top A (PV^t \Psi_{t-1}(\mathbf{x})) \end{aligned}$$



Hence, by the previous items,

$$\begin{aligned}
 |\langle A, \Psi_t(\mathbf{x}) \rangle - \langle PAP, \Psi_t(\mathbf{x}) \rangle| &\leq \|A\|_{\text{sp}} \cdot \|V^t \Psi_{t-1}(\mathbf{x}) - PV^t \Psi_{t-1}(\mathbf{x})\| \cdot (\|V^t \Psi_{t-1}(\mathbf{x})\| + \|PV^t \Psi_{t-1}(\mathbf{x})\|) \\
 &\leq 2M^{1+3 \cdot 2^{t-1}} \|V^t \Psi_{t-1}(\mathbf{x}) - PV^t \Psi_{t-1}(\mathbf{x})\| \\
 &\leq 2M^{1+3 \cdot 2^{t-1}} \|V^t \Psi_{t-1}(\mathbf{x}) - V^t \Psi_{t-1}(P_{\mathcal{I}} \mathbf{x})\| \\
 &\leq 2M^{2+3 \cdot 2^{t-1}} d 2^{t-1} M^{2(t-1)+3 \cdot (2^{t-1}-1)} \epsilon \\
 &\leq d 2^t M^{2t+3 \cdot (2^t-1)} \epsilon
 \end{aligned}$$

Finally,

$$\|PAP\|_F \leq \|PAP\|_{\text{sp}} \sqrt{\text{rank}(PAP)} \leq 1 \cdot \sqrt{\text{rank}(P)} \leq 2^{\lfloor t \rfloor / 2}$$

■

**Lemma 17** For even  $i \geq 0$  denote

$$T_i = \left\{ (j, l) : 0 \leq j \leq l \leq \max(2, 2^{\lceil \log_2(i/2) \rceil}) \text{ and } j+l \leq i \text{ and } j, l \text{ are even} \right\}$$

for odd  $i \geq 0$  denote

$$T_i = \left\{ (j, l) : 0 \leq j, l \leq \max(2, 2^{\lceil \log_2(i/2) \rceil}) \text{ and } j+l \leq i \text{ and } j \text{ is odd and } l \text{ is even} \right\}$$

There are coefficients  $\left\{ \alpha_{j,l}^i \right\}_{(j,l) \in T_i}$  such that

$$p_i = \sum_{(j,l) \in T_i} \alpha_{j,l}^i p_j p_l$$

furthermore,  $|\alpha_{j,l}^i| \leq (3k)^i$

**Proof** By induction on  $i$ . For  $i = 0$  we have  $p_0 = p_0 p_0$  and of  $i = 1$  we have  $p_1 = p_0 p_1$ . For any  $i + 1 \geq 1$  we have by Eq. (6)

$$p_{i+1}(x) = \frac{1}{\sqrt{(i+1)(k-i)}} x p_i(x) - \frac{\sqrt{i(k-i+1)}}{\sqrt{(i+1)(k-i)}} p_{i-1}(x)$$

By the induction hypothesis and Eq. (6)

$$\begin{aligned}
 p_{i+1}(x) &= \frac{1}{\sqrt{(i+1)(k-i)}} \sum_{(j,l) \in T_i} \alpha_{j,l}^i x p_j(x) p_l(x) + \frac{\sqrt{i(k-i+1)}}{\sqrt{(i+1)(k-i)}} \sum_{(j,l) \in T_{i-1}} \alpha_{j,l}^{i-1} p_j(x) p_l(x) \\
 &= \sum_{(j,l) \in T_i} \alpha_{j,l}^i \left( \frac{\sqrt{(j+1)(k-j)}}{\sqrt{(i+1)(k-i)}} p_{j+1}(x) + \frac{\sqrt{j(k-j+1)}}{\sqrt{(i+1)(k-i)}} p_{j-1}(x) \right) p_l(x) \\
 &\quad + \frac{\sqrt{i(k-i+1)}}{\sqrt{(i+1)(k-i)}} \sum_{(j,l) \in T_{i-1}} \alpha_{j,l}^{i-1} p_j(x) p_l(x)
 \end{aligned}$$

The lemma follows from the fact that

$$\frac{\sqrt{(j+1)(k-j)}}{\sqrt{(i+1)(k-i)}} + \frac{\sqrt{j(k-j+1)}}{\sqrt{(i+1)(k-i)}} + \frac{\sqrt{i(k-i+1)}}{\sqrt{(i+1)(k-i)}} \leq 3\sqrt{k} \leq 3k$$

■

**Lemma 18** *Let  $\Psi : \{\pm 1\}^d \rightarrow B_{M_1}^n$ . Assume that for any  $0 \leq i \leq 2^{t-1}$  there is a vector  $\mathbf{v}_i \in B_{M_2}^n$  such that for  $\tilde{P}_i(\mathbf{x}) := \langle \mathbf{v}_i, \Psi(\mathbf{x}) \rangle$  we have  $\|P_i - \tilde{P}_i\|_2 \leq \epsilon$ . Then, there are matrices  $A_i$  for  $0 \leq i \leq 2^t$  such that for  $\hat{P}_i(\mathbf{x}) := \Psi(\mathbf{x})^\top A_i \Psi(\mathbf{x})$  we have  $\|P_i - \hat{P}_i\|_2 \leq 2M_1M_2(3|\mathcal{I}|)^{|\mathcal{I}|}|\mathcal{I}|^2\epsilon$  and  $\|A_i\|_{\text{Tr}} \leq 2^t(3k)^{2^t}M_2^2$*

**Proof** Fix  $0 \leq i \leq 2^t$  and consider the matrix

$$A_i = \sum_{j,l \in \mathcal{I}_i} \alpha_{j,l}^i \mathbf{v}_j \mathbf{v}_l^\top$$

We have  $\|A_i\|_{\text{Tr}} \leq 2^t(3k)^{2^t}M_2^2$  and

$$\hat{P}_i(\mathbf{x}) = \Psi(\mathbf{x})^\top A_i \Psi(\mathbf{x}) = \sum_{j,l \in \mathcal{I}_i} \alpha_{j,l}^i \tilde{P}_j(\mathbf{x}) \tilde{P}_l(\mathbf{x})$$

Hence,

$$\begin{aligned} \|P_i - \hat{P}_i\|_2 &= \left\| \sum_{j,l \in \mathcal{I}_i} \alpha_{j,l}^i (P_j P_l - \tilde{P}_j \tilde{P}_l) \right\|_2 \\ &\leq (3|\mathcal{I}|)^{|\mathcal{I}|} \sum_{j,l \in \mathcal{I}_i} \|P_j P_l - \tilde{P}_j \tilde{P}_l\|_2 \\ &\leq (3|\mathcal{I}|)^{|\mathcal{I}|} \sum_{j,l \in \mathcal{I}_i} \|P_j P_l - P_j \tilde{P}_l\|_2 + \|P_j \tilde{P}_l - \tilde{P}_j \tilde{P}_l\|_2 \\ &\leq (3|\mathcal{I}|)^{|\mathcal{I}|} \sum_{j,l \in \mathcal{I}_i} \|P_j\|_\infty \|P_l - \tilde{P}_l\|_2 + \|\tilde{P}_l\|_\infty \|P_j - \tilde{P}_j\|_2 \\ &\leq 2M_1M_2(3|\mathcal{I}|)^{|\mathcal{I}|}|\mathcal{I}|^2\epsilon \end{aligned}$$

■

## Appendix D. Proofs and Additional Algorithms from Sec. 6

### D.1. Proof of Lemma 10

It is clear that  $\|\cdot\|_{\text{Rx}}$  is homogeneous and non-negative. It remains to show that the triangle inequality is satisfied and that  $\|\mathbf{A}\|_{\text{Rx}} > 0$  for  $\mathbf{A} \neq 0$ . For the triangle inequality we have:

$$\begin{aligned} \|\mathbf{A} + \mathbf{B}\|_{\text{Rx}} &= \min\{\text{Tr}(R) : R \succeq 0 \text{ and } \forall i, -R \preceq B_i + A_i \preceq R\} \\ &\leq \min\{\text{Tr}(R_1) + \text{Tr}(R_2) : R_1, R_2 \succeq 0 \text{ and } \forall i, -R_2 \preceq A_i \preceq R_1 \text{ and } -R_1 \preceq B_i \preceq R_2\} \\ &= \|\mathbf{A}\|_{\text{Rx}} + \|\mathbf{B}\|_{\text{Rx}} \end{aligned}$$

Let  $A \neq 0$ , then there is  $j \in [k]$  with  $A_j \neq 0$  which also means that  $\|A_j\|_{\text{Tr}} > 0$ . Note that

$$\|A\|_{\text{Rx}} \geq \max_{i \in [k]} \|A_i\|_{\text{Tr}} \geq \|A_j\|_{\text{Tr}} > 0$$

For the second part, we need to find  $R$  which minimizes:

$$\min \text{Tr}(R) \quad \text{s.t.} \quad -R \preceq A \preceq R, \quad 0 \preceq R.$$

First, assume that  $A$  is a diagonal matrix. For every unit vector  $\mathbf{e}_i$  we have that  $-\mathbf{e}_i^\top R \mathbf{e}_i \leq \mathbf{e}_i^\top A \mathbf{e}_i \leq \mathbf{e}_i^\top R \mathbf{e}_i$ , which means that  $-r_{i,i} \leq a_{i,i} \leq r_{i,i}$ . In other words, we get that  $r_{i,i} \geq |a_{i,i}|$ , and the minimum on the trace of  $R$  is achieved when  $r_{i,i} = |a_{i,i}|$  for every  $i$ . Consider some  $R$  that achieves the minimum, and assume it is not a diagonal matrix. Then, there are indices  $i \neq j$  with  $r_{i,j} \neq 0$  (and also  $r_{j,i} \neq 0$  since  $R$  is symmetric). Assume that  $r_{i,j} > 0$  and let  $\mathbf{v}$  be the vector with 1 in the  $i$ -th and  $j$ -th coordinates and 0 in every other coordinate. By the condition of  $R$  we get that  $\mathbf{v}^\top (R - A) \mathbf{v} \geq 0$ . But we have that:

$$\mathbf{v}^\top (R - A) \mathbf{v} = -2r_{i,j} < 0$$

which is a contradiction. In case  $r_{i,j} < 0$  we can take  $\mathbf{v}$  to be equal 1 in the  $i$ -th coordinate,  $-1$  in the  $j$ -th coordinate and 0 in every other coordinate. This shows that if  $A$  is diagonal, then there is a single solution for the minimization problem with a diagonal  $R$  such that  $r_{i,i} = |a_{i,i}|$  for every  $i$ .

Assume now that  $A$  is some symmetric matrix, and write  $A = U^\top D U$  where  $U$  is orthogonal and  $D$  is diagonal. Let  $R$  be some solution to the minimization problem, then  $U^\top R U$  is a solution to the same minimization problem where we replace  $A$  with  $D$ . But since  $D$  is diagonal, then there is a single solution where  $U^\top R U$  is diagonal with  $(U^\top R U) = |D|$ . Hence  $R = U^\top |D| U$  is the single solution to the minimization problem for  $A$ . Finally, we have:

$$\|A\|_{\text{Tr}} = \sum_i |D_{i,i}| = \text{Tr}(R) = \|A\|_{\text{Rx}}.$$

## D.2. Additional Algorithms

In Sec. 6 we provided several additional algorithms for training RedEx using norm constraints instead of semi-definite constraints. Here we provide the full algorithms. In Algorithm 4 we provide the 1-layer RedEx algorithm where we use the RedEx norm instead of the semi-definite constraints on  $R$ . Note that the optimization algorithm does not include  $R$ , although it does require finding  $R$  to output  $V$ . In Algorithm 5 we show how to train a 1-layer RedEx where  $k = 1$ . In this case, the RedEx norm is equivalent to the trace norm, hence this algorithm requires solving an unconstrained optimization problem. This can be solved using standard gradient methods such as GD or SGD. Note that it is also not needed to explicitly find  $R$ , since by Lemma 10 it can be calculated directly from  $A$ .

---

**Algorithm 4** Training 1 -layer RedEx with the RedEx norm

---

- 1: **Parameters:** A loss  $\ell : \mathbb{R} \times \mathcal{Y} \rightarrow [0, \infty)$ , width parameter  $\lambda_1$  and a regularization parameter  $\lambda_2$
- 2: **Input:** A dataset  $(\mathbf{x}_1, y_1), \dots, (\mathbf{x}_m, y_m) \in \mathbb{R}^d \times \mathcal{Y}$
- 3: Find symmetric  $d \times d$  matrices  $A_1, \dots, A_k$  by solving the program:

$$\min \quad \frac{1}{m} \sum_{i=1}^m \ell_{y_i} (\mathbf{x}_i^\top \mathbf{A} \mathbf{x}_i) + \lambda_1 \|\mathbf{A}\|_{\text{Rx}} + \lambda_2 \|\mathbf{A}\|_{\text{fr}}^2 \quad (9)$$

- 4: Find  $R$  that minimizes Eq. (7)
  - 5: Compute an orthogonal diagonalization  $R = U^\top D U$
  - 6: Output  $V = \sqrt{D} U$  and  $P_i = (V^\dagger)^\top A_i V^\dagger$ .
- 

---

**Algorithm 5** Training 1 -layer RedEx with output dimension 1

---

- 1: **Parameters:** A loss  $\ell : \mathbb{R} \times \mathcal{Y} \rightarrow [0, \infty)$ , width parameter  $\lambda_1$  and a regularization parameter  $\lambda_2$
- 2: **Input:** A dataset  $(\mathbf{x}_1, y_1), \dots, (\mathbf{x}_m, y_m) \in \mathbb{R}^d \times \mathcal{Y}$
- 3: Find symmetric  $d \times d$  matrix  $A$  by solving the program:

$$\min \quad \frac{1}{m} \sum_{i=1}^m \ell_{y_i} (\mathbf{x}_i^\top A \mathbf{x}_i) + \lambda_1 \|A\|_{\text{Tr}} + 2\lambda_2 \|A\|_{\text{fr}}^2 \quad (10)$$

- 4: Compute an orthogonal diagonalization  $A = U^\top D U$
  - 5: Output  $V = \sqrt{|D|} U$  and  $P = (V^\dagger)^\top A V^\dagger$ .
-