

GENERALIZATION AND ESTIMATION ERROR BOUNDS FOR MODEL-BASED NEURAL NETWORKS

Avner Shultzman¹, Eyar Azar¹, Miguel R. D. Rodrigues² & Yonina C. Eldar¹ *

Faculty of Mathematics and Computer Science, Weizmann Institute of Science, Israel¹

Faculty of Engineering Science, University College London, UK²

{avner.shultzman, eyar.azar, yonina.eldar}@weizmann.ac.il
m.rodrigues@ucl.ac.uk

ABSTRACT

Model-based neural networks provide unparalleled performance for various tasks, such as sparse coding and compressed sensing problems. Due to the strong connection with the sensing model, these networks are interpretable and inherit prior structure of the problem. In practice, model-based neural networks exhibit higher generalization capability compared to ReLU neural networks. However, this phenomenon was not addressed theoretically. Here, we leverage complexity measures including the global and local Rademacher complexities, in order to provide upper bounds on the generalization and estimation errors of model-based networks. We show that the generalization abilities of model-based networks for sparse recovery outperform those of regular ReLU networks, and derive practical design rules that allow to construct model-based networks with guaranteed high generalization. We demonstrate through a series of experiments that our theoretical insights shed light on a few behaviours experienced in practice, including the fact that ISTA and ADMM networks exhibit higher generalization abilities (especially for small number of training samples), compared to ReLU networks.

1 INTRODUCTION

Model-based neural networks provide unprecedented performance gains for solving sparse coding problems, such as the learned iterative shrinkage and thresholding algorithm (ISTA) (Gregor & LeCun, 2010) and learned alternating direction method of multipliers (ADMM) (Boyd et al., 2011). In practice, these approaches outperform feed-forward neural networks with ReLU nonlinearities.

These neural networks are usually obtained from algorithm unrolling (or unfolding) techniques, which were first proposed by Gregor and LeCun (Gregor & LeCun, 2010), to connect iterative algorithms to neural network architectures. The trained networks can potentially shed light on the problem being solved. For ISTA networks, each layer represents an iteration of a gradient-descent procedure. As a result, the output of each layer is a valid reconstruction of the target vector, and we expect the reconstructions to improve with the network’s depth. These networks capture original problem structure, which translates in practice to a lower number of required training data (Monga et al., 2021). Moreover, the generalization abilities of model-based networks tend to improve over regular feed-forward neural networks (Behboodi et al., 2020; Schnoor et al., 2021).

Understanding the generalization of deep learning algorithms has become an important open question. The generalization error of machine learning models measures the ability of a class of estimators to generalize from training to unseen samples, and avoid overfitting the training (Jakubovitz et al., 2019). Surprisingly, various deep neural networks exhibit high generalization abilities, even for increasing networks’ complexities (Neyshabur et al., 2015b; Belkin et al., 2019). Classical machine learning

*This project has received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research, the innovation programme (grant agreement No. 101000967), and the Israel Science Foundation under Grant 536/22. Y. C. Eldar and M. R. D. Rodrigues are supported by The Weizmann-UK Making Connections Programme (Ref. 129589). M. R. D. Rodrigues is also supported by the Alan Turing Institute. The authors wish to thank Dr. Gholamali Aminian from the Alan Turing Institute, UK, for his contribution to the proofs’ correctness.

measures such as the Vapnik-Chervonenkis (VC) dimension (Vapnik & Chervonenkis, 1991) and Rademacher complexity (RC) (Bartlett & Mendelson, 2002), predict an increasing generalization error (GE) with the increase of the models’ complexity, and fail to explain the improved generalization observed in experiments. More advanced measures consider the training process and result in tighter bounds on the estimation error (EE), were proposed to investigate this gap, such as the local Rademacher complexity (LRC) (Bartlett et al., 2005). To date, the EE of model based networks using these complexity measures has not been investigated to the best of our knowledge.

1.1 OUR CONTRIBUTIONS

In this work, we leverage existing complexity measures such as the RC and LRC, in order to bound the generalization and estimation errors of learned ISTA and learned ADMM networks.

- We provide new bounds on the GE of ISTA and ADMM networks, showing that the GE of model-based networks is lower than that of the common ReLU networks. The derivation of the theoretical guarantees combines existing proof techniques for computing the generalization error of multilayer networks with new methodology for bounding the RC of the soft-thresholding operator, that allows a better understanding of the generalization ability of model based networks.
- The obtained bounds translate to practical design rules for model-based networks which guarantee high generalization. In particular, we show that a nonincreasing GE as a function of the network’s depth is achievable, by limiting the weights’ norm in the network. This improves over existing bounds, which exhibit a logarithmic increase of the GE with depth (Schnoor et al., 2021). The GE bounds of the model-based networks suggest that under similar restrictions, learned ISTA networks generalize better than learned ADMM networks.
- We also exploit the LRC machinery to derive bounds on the EE of feed-forward networks, such as ReLU, ISTA, and ADMM networks. The EE bounds depend on the data distribution and training loss. We show that the model-based networks achieve lower EE bounds compared to ReLU networks.
- We focus on the differences between ISTA and ReLU networks, in term of performance and generalization. This is done through a series of experiments for sparse vector recovery problems. The experiments indicate that the generalization abilities of ISTA networks are controlled by the soft-threshold value. For a proper choice of parameters, ISTA achieves lower EE along with more accurate recovery. The dependency of the EE as a function of λ and the number of training samples can be explained by the derived EE bounds.

1.2 RELATED WORK

Understanding the GE and EE of general deep learning algorithms is an active area of research. A few approaches were proposed, which include considering networks of weights matrices with bounded norms (including spectral and $L_{2,1}$ norms) (Bartlett et al., 2017; Sokolić et al., 2017), and analyzing the effect of multiple regularizations employed in deep learning, such as weight decay, early stopping, or drop-outs, on the generalization abilities (Neyshabur et al., 2015a; Gao & Zhou, 2016; Amjad et al., 2021). Additional works consider global properties of the networks, such as a bound on the product of all Frobenius norms of the weight matrices in the network (Golowich et al., 2018). However, these available bounds do not capture the GE behaviour as a function of network depth, where an increase in depth typically results in improved generalization. This also applies to the bounds on the GE of ReLU networks, detailed in Section 2.3.

Recently, a few works focused on bounding the GE specifically for deep iterative recovery algorithms (Behboodi et al., 2020; Schnoor et al., 2021). They focus on a broad class of unfolded networks for sparse recovery, and provide bounds which scale logarithmically with the number of layers (Schnoor et al., 2021). However, these bounds still do not capture the behaviours experienced in practice.

Much work has also focused on incorporating the networks’ training process into the bounds. The LRC framework due to Bartlett, Bousquet, and Mendelson (Bartlett et al., 2005) assumes that the training process results in a smaller class of estimation functions, such that the distance between the estimator in the class and the empirical risk minimizer (ERM) is bounded. An additional related framework is the effective dimensionality due to Zhang (Zhang, 2002). These frameworks result in

different bounds, which relate the EE to the distance between the estimators. These local complexity measures were not applied to model-based neural networks.

Throughout the paper we use boldface lowercase and uppercase letters to denote vectors and matrices respectively. The L_1 and L_2 norms of a vector \mathbf{x} are written as $\|\mathbf{x}\|_1$ and $\|\mathbf{x}\|_2$ respectively, and the L_∞ (which corresponds to the maximal L_1 norm over the matrix's rows) and spectral norms of a matrix \mathbf{X} , are denoted by $\|\mathbf{X}\|_\infty$ and $\|\mathbf{X}\|_\sigma$ respectively. We denote the transpose operation by $(\cdot)^T$. For any function f and class of functions \mathcal{H} , we define $f \circ \mathcal{H} = \{x \mapsto f \circ h(x) : h \in \mathcal{H}\}$.

2 PRELIMINARIES

2.1 NETWORK ARCHITECTURE

We focus on model-based networks for sparse vector recovery, applicable to the linear inverse problem

$$\mathbf{y} = \mathbf{A}\mathbf{x} + \mathbf{e} \quad (1)$$

where $\mathbf{y} \in \mathbb{R}^{n_y}$ is the observation vector with n_y entries, $\mathbf{x} \in \mathbb{R}^{n_x}$ is the target vector with n_x entries, with $n_x > n_y$, $\mathbf{A} \in \mathbb{R}^{n_y \times n_x}$ is the linear operator, and $\mathbf{e} \in \mathbb{R}^{n_y}$ is additive noise. The target vectors are sparse with sparsity rate ρ , such that at most $\lfloor \rho n_x \rfloor$ entries are nonzero. The inverse problem consists of recovering the target vector \mathbf{x} , from the observation vector \mathbf{y} .

Given that the target vector is assumed to be sparse, recovering \mathbf{x} from \mathbf{y} in (1) can be formulated as an optimization problem, such as least absolute shrinkage and selection operator (LASSO) (Tibshirani & Ryan, 2013), that can be solved with well-known iterative methods including ISTA and ADMM. To address more complex problems, such as an unknown linear mapping \mathbf{A} , and to avoid having to fine tune parameters, these algorithms can be mapped into model-based neural networks using unfolding or unrolling techniques (Gregor & LeCun, 2010; Boyd et al., 2011; Monga et al., 2021; Yang et al., 2018). The network's architecture imitates the original iterative method's functionality and enables to learn the models' parameters with respect to a set of training examples.

We consider neural networks with L layers (referred to as the network's depth), which corresponds to the number of iterations in the original iterative algorithm. The layer outputs of an unfolded ISTA network \mathbf{h}_I^l , $l \in [1, L]$, are defined by the following recurrence relation, shown in Fig. 1:

$$\mathbf{h}_I^l = \mathcal{S}_\lambda(\mathbf{W}^l \mathbf{h}_I^{l-1} + \mathbf{b}), \quad \mathbf{h}_I^0 = \mathcal{S}_\lambda(\mathbf{b}) \quad (2)$$

where $\mathbf{W}^l \in \mathbb{R}^{n_x \times n_x}$, $l \in [1, L]$ are the weights matrices corresponding to each of the layers, with bounded norms $\|\mathbf{W}^l\|_\infty \leq B_l$. We further assume that the L_2 norm of \mathbf{W}^1 is bounded by B_1 . The vector $\mathbf{b} = \mathbf{A}^T \mathbf{y}$ is a constant bias term that depends on the observation \mathbf{y} , where we assume that the initial values are bounded, such that $\|\mathbf{h}_I^0\|_1 \leq B_0$. In addition, $\mathcal{S}_\lambda(\cdot)$ is the elementwise soft-thresholding operator

$$\mathcal{S}_\lambda(\mathbf{h}) = \text{sign}(\mathbf{h}) \max(|\mathbf{h}| - \lambda, \mathbf{0}) \quad (3)$$

where the functions $\text{sign}(\cdot)$ and $\max(\cdot)$ are applied elementwise, and $\mathbf{0}$ is a vector of zeros. As $\mathcal{S}_\lambda(\cdot)$ is an elementwise function it preserves the input's dimension, and can be applied on scalar or vector inputs. The network's prediction is given by the last layer in the network $\hat{\mathbf{x}} = \mathbf{h}_I^L(\mathbf{y})$. We note that the estimators are functions mapping \mathbf{y} to $\hat{\mathbf{x}}$, $\mathbf{h}_I^L : \mathbb{R}^{n_y} \rightarrow \mathbb{R}^{n_x}$, characterized by the weights, i.e. $\mathbf{h}_I^L = \mathbf{h}_I^L(\{\mathbf{W}^l\}_{l=1}^L)$. The class of functions representing the output at depth L in an ISTA network, is $\mathcal{H}_I^L = \{\mathbf{h}_I^L(\{\mathbf{W}^l\}_{l=1}^L) : \|\mathbf{W}^l\|_\infty \leq B_l, l \in [1, L], \|\mathbf{W}^1\|_2 \leq B_1\}$.

Similarly, the l th layer of unfolded ADMM is defined by the following recurrence relation

$$\begin{aligned} \mathbf{h}_A^l &= \mathbf{W}^l (\mathbf{z}^{l-1} + \mathbf{u}^{l-1}) + \mathbf{b} \\ \mathbf{z}^l &= \mathcal{S}_\lambda(\mathbf{h}_A^l - \mathbf{u}^{l-1}), \quad \mathbf{z}^0 = \mathbf{0} \\ \mathbf{u}^l &= \mathbf{u}^{l-1} - \gamma (\mathbf{h}_A^l - \mathbf{z}^l), \quad \mathbf{u}^0 = \mathbf{0} \end{aligned} \quad (4)$$

where $\mathbf{0}$ is a vector of zeros, $\mathbf{b} = \mathbf{A}^T \mathbf{y}$ is a constant bias term, and $\gamma > 0$ is the step size derived by the original ADMM algorithm, as shown in Fig. 1. The estimators satisfy $\mathbf{h}_A^L : \mathbb{R}^{n_y} \rightarrow \mathbb{R}^{n_x}$, and the class of functions representing the output at depth L in an ADMM network, is $\mathcal{H}_A^L = \{\mathbf{h}_A^L(\{\mathbf{W}^l\}_{l=1}^L) : \|\mathbf{W}^l\|_\infty \leq B_l, l \in [1, L], \|\mathbf{W}^1\|_2 \leq B_1\}$, where we impose the same assumptions on the weights matrices as ISTA networks.

For a depth- L ISTA or ADMM network, the learnable parameters are the weight matrices $\{\mathbf{W}^l\}_{l=1}^L$. The weights are learnt by minimizing a loss function \mathcal{L} on a set of m training examples $S = \{(\mathbf{x}_i, \mathbf{y}_i)\}_{i=1}^m$, drawn from an unknown distribution \mathcal{D} , consistent with the model in (1). We consider the case where the per-example loss function is obtained by averaging over the example per-coordinate losses:

$$\mathcal{L}(\mathbf{h}(\mathbf{y}), \mathbf{x}) = \frac{1}{n_x} \sum_{j=1}^{n_x} \ell(\mathbf{h}_j(\mathbf{y}), \mathbf{x}_j) \quad (5)$$

where $\mathbf{h}_j(\mathbf{y})$ and \mathbf{x}_j denote the j th coordinate of the estimated and true targets, and $\ell : \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}_+$ is 1-Lipschitz in its first argument. This requirement is satisfied in many practical settings, for example with the p -power of L_p norms, and is also required in a related work (Xu et al., 2016). The loss of an estimator $\mathbf{h} \in \mathcal{H}$ which measures the difference between the true value \mathbf{x} and the estimation $\mathbf{h}(\mathbf{y})$, is denoted for convenience by $\mathcal{L}(\mathbf{h}) = \mathcal{L}(\mathbf{h}(\mathbf{y}), \mathbf{x})$.

There exists additional forms of learned ISTA and ADMM networks, which include learning an additional set of weight matrices affecting the bias terms (Monga et al., 2021). Also, the optimal value of λ generally depends on the target vector sparsity level. Note however that for learned networks, the value of λ at each layer can also be learned. However, here we focus on a more basic architecture with fixed λ in order to draw theoretical conclusions.

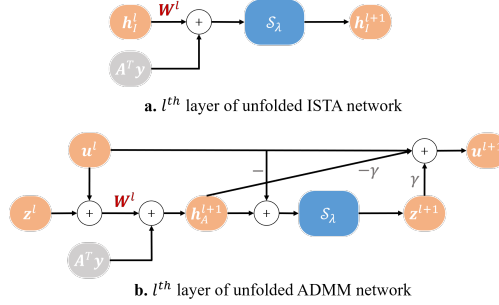


Figure 1: A single layer in the unfolded networks. **a.** Unfolded ISTA. **b.** Unfolded ADMM. The learnable parameters are the weight matrices (marked in red).

2.2 GENERALIZATION AND ESTIMATION ERRORS

In this work, we focus on upper bounding the GE and EE of the model-based neural networks of Fig. 1. The GE of a class of estimation functions $\mathbf{h} \in \mathcal{H}$, such that $\mathbf{h} : \mathbb{R}^{n_y} \rightarrow \mathbb{R}^{n_x}$, is defined as

$$\mathcal{G}(\mathcal{H}) = \mathbb{E}_S \sup_{\mathbf{h} \in \mathcal{H}} L_{\mathcal{D}}(\mathbf{h}) - L_S(\mathbf{h}) \quad (6)$$

where $L_{\mathcal{D}}(\mathbf{h}) = \mathbb{E}_{\mathcal{D}} \mathcal{L}(\mathbf{h})$ is the expected loss with respect to the data distribution \mathcal{D} (the joint probability distribution of the targets and observations), $L_S(\mathbf{h}) = \frac{1}{m} \sum_{i=1}^m \mathcal{L}(\mathbf{h}(\mathbf{y}_i), \mathbf{x}_i)$ is the average empirical loss with respect to the training set S , and \mathbb{E}_S is the expectation over the training datasets. The GE is a global property of the class of estimators, which captures how the class of estimators \mathcal{H} is suited to the learning problem. Large GE implies that there are hypotheses in \mathcal{H} for which $L_{\mathcal{D}}$ deviates much from L_S , on average over S . However, the GE in (6) does not capture how the learning algorithm chooses the estimator $\mathbf{h} \in \mathcal{H}$.

In order to capture the effect of the learning algorithm, we consider local properties of the class of estimators, and focus on bounding the estimation error (EE)

$$\mathcal{E}(\mathcal{H}) = L_{\mathcal{D}}(\hat{\mathbf{h}}) - \inf_{\mathbf{h} \in \mathcal{H}} L_{\mathcal{D}}(\mathbf{h}) \quad (7)$$

where $\hat{\mathbf{h}}$ is the ERM satisfying $L_S(\hat{\mathbf{h}}) = \inf_{\mathbf{h} \in \mathcal{H}} L_S(\mathbf{h})$. We note that the ERM approximates the learned estimator $\mathbf{h}_{\text{learned}}$, which is obtained by training the network on a set of training examples S , using algorithms such as SGD. However, the estimator $\mathbf{h}_{\text{learned}}$ depends on the optimization algorithm, and can differ from the ERM. The difference between the empirical loss associated with $\hat{\mathbf{h}}$ and the empirical loss associated with $\mathbf{h}_{\text{learned}}$ is usually referred to as the optimization error.

Common deep neural network architectures have large GE compared to the low EE achieved in practice. This still holds, when the networks are not trained with explicit regularization, such as weight decay, early stopping, or drop-outs (Srivastava et al., 2014; Neyshabur et al., 2015b). This empirical phenomena is experienced across various architectures and hyper-parameter choices (Liang et al., 2019; Novak et al., 2018; Lee et al., 2018; Neyshabur et al., 2018).

2.3 RADEMACHER COMPLEXITY BASED BOUNDS

The RC is a standard tool which captures the ability of a class of functions to approximate noise, where a lower complexity indicates a reduced generalization error. Formally, the empirical RC of a class of scalar estimators \mathcal{H} , such that $h : \mathbb{R}^{n_y} \rightarrow \mathbb{R}$ for $h \in \mathcal{H}$, over m samples is

$$\mathfrak{R}_m(\mathcal{H}) = \mathbb{E}_{\{\epsilon_i\}_{i=1}^m} \sup_{h \in \mathcal{H}} \frac{1}{m} \sum_{i=1}^m \epsilon_i h(\mathbf{y}_i) \quad (8)$$

where $\{\epsilon_i\}_{i=1}^m$ are independent Rademacher random variables for which $Pr(\epsilon_i = 1) = Pr(\epsilon_i = -1) = 1/2$, the samples $\{\mathbf{y}_i\}_{i=1}^m$ are obtained by the model in (1) from m i.i.d. target vectors $\{\mathbf{x}_i\}_{i=1}^m$ drawn from an unknown distribution. Taking the expectation of the RC of $\mathcal{L} \circ \mathcal{H}$ with respect to the set of examples S presented in Section 2.1, leads to a bound on the GE

$$\mathcal{G}(\mathcal{H}) \leq 2\mathbb{E}_S \mathfrak{R}_m(\mathcal{L} \circ \mathcal{H}) \quad (9)$$

where \mathcal{L} is defined in Section 2.1 and $\mathcal{L} \circ \mathcal{H} = \{(x, \mathbf{y}) \mapsto \mathcal{L}(h(\mathbf{y}), x) : h \in \mathcal{H}\}$ (Shalev-Shwartz & Ben-David, 2014). We observe that the class of functions $\mathcal{L} \circ \mathcal{H}$ consists of scalar functions, such that $f : \mathbb{R}^{n_x} \times \mathbb{R}^{n_y} \rightarrow \mathbb{R}$ for $f \in \mathcal{L} \circ \mathcal{H}$. Therefore, in order to bound the GE of the class of functions defined by ISTA and ADMM networks, we can first bound their RC.

Throughout the paper, we compare the model-based architectures with a feed forward network with ReLU activations, given by $ReLU(\mathbf{h}) = \max(\mathbf{h}, \mathbf{0})$. In this section, we review existing bounds for the generalization error of these networks. The layers of a ReLU network \mathbf{h}_R^l , $\forall l \in [1, L]$, are defined by the following recurrence relation $\mathbf{h}_R^l = ReLU(\mathbf{W}^l \mathbf{h}_R^{l-1} + \mathbf{b})$, and $\mathbf{h}_R^0 = ReLU(\mathbf{b})$, where \mathbf{W}^l , $\forall l \in [1, L]$ are the weight matrices which satisfy the same conditions as the weight matrices of ISTA and ADMM networks. The class of functions representing the output at depth L in a ReLU network, is $\mathcal{H}_R^L = \{\mathbf{h}_R^L(\{\mathbf{W}^l\}_{l=1}^L) : \|\mathbf{W}^l\|_\infty \leq B_l \ l \in [1, L], \|\mathbf{W}^1\|_2 \leq B_1\}$. This architecture leads to the following bound on the GE.

Theorem 1 (Generalization error bound for ReLU networks (Gao & Zhou, 2016)). *Consider the class of feed forward networks of depth- L with ReLU activations, \mathcal{H}_R^L , as described in Section 2.3, and m i.i.d. training samples. Given a 1-Lipschitz loss function, its GE satisfies $\mathcal{G}(\mathcal{H}_R^L) \leq 2G_R^L$, where $G_R^L = \frac{B_0 \prod_{l=1}^L B_l}{\sqrt{m}}$.*

Proof. Follows from applying the bound of the RC of ReLU neural networks from (Gao & Zhou, 2016) and combining it with (9). \square

The bound in Theorem 1 is satisfied for any feed forward network with 1-Lipschitz nonlinear activations (including ReLU), and can be generalized for networks with activations with different Lipschitz constants. We show in Theorem 2, that the bound presented in Theorem 1 cannot be substantially improved for ReLU networks with the RC framework.

Theorem 2 (Lower Rademacher complexity bound for ReLU networks (Bartlett et al., 2017)). *Consider the class of feed forward networks of depth- L with ReLU activations, where the weight matrices have bounded spectral norm $\|\mathbf{W}^l\|_\sigma \leq B'_l$, $l \in [1, L]$. The dimension of the output layer is 1, and the dimension of each non-output layer is at least 2. Given m i.i.d. training samples, there exists a c such that $\mathfrak{R}_m(\mathcal{H}_R^{L,L}) \geq cB_0 \prod_{l=1}^L B'_l$, where $\mathcal{H}_R^{L,L} = \{\mathbf{h}_R^L(\{\mathbf{W}^l\}_{l=1}^L) : \|\mathbf{W}^l\|_\sigma \leq B'_l \ l \in [1, L]\}$.*

This result shows that using the RC framework the GE of ReLU networks behaves as the product of the weight matrices' norms $\prod_{l=1}^L B_l$, as captured in Theorem 1. Theorem 2, implies that the dependence on the weight matrices' norms, cannot be substantially improved with the RC framework for ReLU networks.

3 GENERALIZATION ERROR BOUNDS: GLOBAL PROPERTIES

In this section, we derive theoretical bounds on the GE of learned ISTA and ADMM networks. From these bounds we deduce design rules to construct ISTA and ADMM networks with a GE which does not increase exponentially with the number of layers. We start by presenting theoretical guarantees on the RC of any class of functions, after applying the soft-thresholding operation.

Soft-thresholding is a basic block that appears in multiple iterative algorithms, and therefore is used as the nonlinear activation in many model-based networks. It results from the proximal gradient of the L_1 norm (Palomar & Eldar, 2010). We therefore start by presenting the following lemma which expresses how the RC of a class of functions is affected by applying soft-thresholding to each function in the class. The proof is provided in the supplementary material.

Lemma 1 (Rademacher complexity of soft-thresholding). *Given any class of scalar functions \mathcal{H} where $h : \mathbb{R}^n \rightarrow \mathbb{R}$, $h \in \mathcal{H}$ for any integer $n \geq 1$, and m i.i.d. training samples,*

$$\mathfrak{R}_m(S_\lambda \circ \mathcal{H}) \leq \mathfrak{R}_m(\mathcal{H}) - \frac{\lambda T}{m}, \quad (10)$$

where $T = \sum_{i=1}^m T_i$ and $T_j = \mathbb{E}_{\{\epsilon_i\}_{i=2}^m} (\mathbb{1}_{h^*(\mathbf{y}_j) > \lambda \wedge h'^*(\mathbf{y}_j) < -\lambda})$ such that $h^*, h'^* = \arg \max_{h, h' \in \mathcal{H}} (h(\mathbf{y}_1) - h'(\mathbf{y}_1) - 2\lambda \mathbb{1}_{h(\mathbf{y}_1) > \lambda \wedge h'(\mathbf{y}_1) < -\lambda} + \sum_{i=2}^m \epsilon_i (\mathcal{S}_\lambda(h(\mathbf{y}_i)) + \mathcal{S}_\lambda(h'(\mathbf{y}_i))))$.

The quantity T is a non-negative value obtained during the proof, which depends on the networks' number of layers, underlying data distribution \mathcal{D} and soft-threshold value λ . As seen from (10), the value of T dictates the reduction in RC due to soft-thresholding, where a reduction in the RC can also be expected. The value of T increases as λ decreases. In the case that λT increases with λ , higher values of λ further reduce the RC of the class of functions \mathcal{H} , due to the soft-thresholding.

We now focus on the class of functions representing the output of a neuron at depth L in an ISTA network, \mathcal{H}_I^L . In the following theorem, we bound its GE using the RC framework and Lemma 1. The proof is provided in the supplementary material.

Theorem 3 (Generalization error bound for ISTA networks). *Consider the class of learned ISTA networks of depth L as described in (2), and m i.i.d. training samples. Then there exist $T^{(l)}$ for $l \in [1, L]$ in the range $T^{(l)} \in [0, \min \left\{ \frac{m B_l G_I^{l-1}}{\lambda}, m \right\}]$ such that $\mathcal{G}(\mathcal{H}_I^L) \leq 2\mathbb{E}_S G_I^L$, where*

$$G_I^l = \left(\frac{B_0 \prod_{l'=1}^l B_{l'}}{\sqrt{m}} - \frac{\lambda}{m} \sum_{l'=1}^{l-1} T^{(l')} \prod_{j=l'+1}^l B_j - \frac{\lambda T^{(l)}}{m} \right). \quad (11)$$

Next, we show that for a specific distribution, the expected value of $T^{(l)}$ is greater than 0. Under an additional bound on the expectation value of the estimators (specified in the supplementary material)

$$\mathbb{E}_S(T^{(l)}) \geq \max \left\{ m \left(1 - 2e^{-(cB_l b^{(l)} - \lambda)} \right), 0 \right\} \quad (12)$$

where $b^{(l+1)} = B_l b^{(l)} - \lambda$, $b^{(1)} = B_0$, and $c \in (0, 1]$. Increasing λ or decreasing B , will decrease the bound in (12), since crossing the threshold is less probable. Depending on λ and $\{B_l\}_{l=1}^L$ (specifically that $\lambda \leq cB_l b^{(l)} + \ln 2$), the bound in (12) is positive, and enforces a non-zero reduction in the GE. Along with Theorem 3, this shows the expected reduction in GE of ISTA networks compared to ReLU networks. The reduction is controlled by the value of the soft threshold.

To obtain a more compact relation, we can choose the maximal matrices' norm $B = \max_{l \in [1, L]} B_l$, and denote $T = \min_{l \in [1, L]} T^{(l)} \in [0, m]$ which leads to $\mathcal{G}(\mathcal{H}_I^L) \leq 2 \left(\frac{B_0 B^L}{\sqrt{m}} - \frac{\lambda \mathbb{E}_S(T)}{m} \frac{B^L - 1}{B - 1} \right)$.

Comparing this bound with the GE bound for ReLU networks presented in Theorem 1, shows the expected reduction due to the soft thresholding activation. This result also implies practical rules for designing low generalization error ISTA networks. We note that the network's parameters such as the soft-threshold value λ and number of samples m , are predefined by the model being solved (for example, in ISTA, the value of λ is chosen according to the singular values of \mathbf{A}).

We derive an implicit design rule from (3), for a nonincreasing GE, as detailed in Section A.2. This is done by restricting the matrices' norm to satisfy $B \leq 1 + \frac{\lambda \mathbb{E}_S(T)}{\sqrt{m} B_0}$. Moreover, these results can

be extended to convolutional neural networks. As convolution operations can be expressed via multiplication with a convolution matrix, the presented results are also satisfied in that case.

Similarly, we bound the GE of the class of functions representing the output at depth L in an ADMM network, \mathcal{H}_A^L . The proof and discussion are provided in the supplementary material.

Theorem 4 (Generalization error bound for ADMM networks). *Consider the class of learned ADMM networks of depth L as described in (4), and m i.i.d. training samples. Then there exist $T^{(l)}$ for $l \in [1, L-1]$ in the interval $T^{(l)} \in \left[0, \min \left\{ \frac{m\tilde{B}_l G_A^{l-1}}{\tilde{\lambda}}, m \right\} \right]$ where $G_A^l = \frac{B_0 \prod_{l'=1}^{l-1} \tilde{B}_{l'}}{\sqrt{m}} - \frac{\tilde{\lambda}}{m} \sum_{l'=1}^{l-2} T^{(l')} \prod_{j=l'+1}^{l-1} \tilde{B}_j - \frac{\tilde{\lambda} T^{(l-1)}}{m}$, and $\tilde{\lambda} = (1+\gamma)\lambda$, $\tilde{B}_l = (1+2\gamma)(B_l+2)$, $l \in [1, L]$, such that $\mathcal{G}(\mathcal{H}_A^L) \leq 2\tilde{B}_L \mathbb{E}_S G_A^{L-1}$.*

We compare the GE bounds for ISTA and ADMM networks, to the bound on the GE of ReLU networks presented in Theorem 1. We observe that both model-based networks achieve a reduction in the GE, which depends on the soft-threshold, the underlying data distribution, and the bound on the norm of the weight matrices. Following the bound, we observe that the soft-thresholding nonlinearity is most valuable in the case of small number of training samples. The soft-thresholding nonlinearity is the key that enables reducing the GE of the ISTA and ADMM networks compared to ReLU networks. Next, we focus on bounding the EE of feed-forward networks based on the LRC framework.

4 ESTIMATION ERROR BOUNDS: LOCAL PROPERTIES

To investigate the model-based networks' EE, we use the LRC framework (Bartlett et al., 2005). Instead of considering the entire class of functions \mathcal{H} , the LRC considers only estimators which are close to the optimal estimator $\mathcal{H}_r = \left\{ \mathbf{h} \in \mathcal{H} : \mathbb{E}_{\mathcal{D}} \|\mathbf{h}(\mathbf{y}) - \mathbf{h}^*(\mathbf{y})\|_2^2 \leq r \right\}$, where \mathbf{h}^* is such that $L_{\mathcal{D}}(\mathbf{h}^*) = \inf_{\mathbf{h} \in \mathcal{H}} L_{\mathcal{D}}(\mathbf{h})$. It is interesting to note that the class of estimators \mathcal{H}_r only restricts the distance between the estimators themselves, and not between their corresponding losses. Following the LRC framework, we consider target vectors ranging in $[-1, 1]^{n_x}$. Therefore, we adapt the networks' estimations by clipping them to lie in the interval $[-1, 1]^{n_x}$. In our case we consider the restricted classes of functions representing the output of a neuron at depth L in ISTA, ADMM, and ReLU networks. Moreover, we denote by \mathbf{W}^l and $\mathbf{W}^{l,*}$, $l \in [1, L]$ the weight matrices corresponding to $\mathbf{h} \in \mathcal{H}_r$ and \mathbf{h}^* , respectively. Based on these restricted class of functions, we present the following assumption and theorem (the proof is provided in the supplementary material).

Assumption 1. *There exists a constant $C \geq 1$ such that for every probability distribution \mathcal{D} , and estimator $\mathbf{h} \in \mathcal{H}$, $\mathbb{E}_{\mathcal{D}} \sum_{j=1}^{n_x} (\mathbf{h}_j - \mathbf{h}_j^*)^2 \leq C \mathbb{E}_{\mathcal{D}} \sum_{j=1}^{n_x} (\ell(\mathbf{h}_j) - \ell(\mathbf{h}_j^*))$, where \mathbf{h}_j and \mathbf{h}_j^* denote the j th coordinate of the estimators.*

As pointed out in (Bartlett & Mendelson, 2002), this condition usually follows from a uniform convexity condition on the loss function ℓ . For instance, if $|h(\mathbf{y}) - \mathbf{x}| \leq 1$ for any $\mathbf{h} \in \mathcal{H}$, $\mathbf{y} \in \mathbb{R}^{n_y}$ and $\mathbf{x} \in \mathbb{R}^{n_x}$, then the condition is satisfied with $C = 1$ (Yousefi et al., 2018).

Theorem 5 (Estimation error bound of ISTA, ADMM, and ReLU networks). *Consider the class of functions represented by depth- L ISTA networks \mathcal{H}_I^L as detailed in Section 2.1, m i.i.d training samples, and a per-coordinate loss satisfying Assumption 1 with a constant C . Let $\|\mathbf{W}^l - \mathbf{W}^{l,*}\|_{\infty} \leq \alpha\sqrt{r}$ for some $\alpha > 0$. Moreover, $B \geq \max\{\alpha\sqrt{r}, 1\}$. Then there exists T in the interval $T \in \left[0, \min \left\{ \frac{\sqrt{m}B_0B^{L-1}2^L}{\lambda\eta}, m \right\} \right]$, where $\eta = \frac{LB^{L-1}(B-1)-B^{L+1}}{(B-1)^2}$, such that for any $s > 0$ with probability at least $1 - e^{-s}$,*

$$\mathcal{E}(\mathcal{H}_I^L) \leq 41r^* + \frac{17C^2 + 48C}{mn_x} s \quad (13)$$

where $r^* = C^2\alpha^2 \left(\frac{B_0B^{L-1}2^L}{\sqrt{m}} - \frac{\lambda T}{m}\eta \right)^2$. The bound is also satisfied for the class of functions represented by depth- L ADMM networks \mathcal{H}_A^L , with $r^* = C^2\alpha^2 \left(\frac{B_0\tilde{B}^{L-2}2^{L-1}}{\sqrt{m}} - \frac{\tilde{\lambda}T}{m}\tilde{\eta} \right)^2$, where $\tilde{\lambda} = (1+\gamma)\lambda$, $\tilde{B} = (1+2\gamma)(B+2)$, and $\tilde{\eta} = \frac{(L-1)\tilde{B}^{L-2}(\tilde{B}-1)-\tilde{B}^{L+1}+1}{(\tilde{B}-1)^2}$, and for the class of functions represented by depth- L ReLU networks \mathcal{H}_R^L , with $r^* = C^2\alpha^2 \left(\frac{B_0B^{L-1}2^L}{\sqrt{m}} \right)^2$.

From Theorem 5, we observe that the EE decreases by a factor of $\mathcal{O}(1/m)$, instead of a factor of $\mathcal{O}(1/\sqrt{m})$ obtained for the GE. This result complies with previous local bounds which yield faster convergence rates compared to global bounds (Blanchard et al., 2007; Bartlett et al., 2005). Also, the value of α relates the maximal distance between estimators in \mathcal{H}_r denoted by r , to the distance between their corresponding weight matrices $\|\mathbf{W}^l - \mathbf{W}^{l,*}\|_\infty \leq \alpha\sqrt{r}$, $l \in [1, L]$. Tighter bounds on the distance between the weight matrices allow us to choose a smaller value for α , resulting in smaller values r^* which improve the EE bounds. The value of α could depend on the network’s nonlinearity, underlying data distribution \mathcal{D} , and number of layers.

Note that the bounds of the model-based architectures depend on the soft-thresholding through the value of $\lambda\mathbb{E}_S(T)$. As $\lambda\mathbb{E}_S(T)$ increases, the bound on the EE decreases, which emphasizes the nonlinearity’s role in the network’s generalization abilities. Due to the soft-thresholding, ISTA and ADMM networks result in lower EE bounds compared to the bound for ReLU networks. It is interesting to note, that as the number of training samples m increases, the difference between the bounds on the model-based and ReLU networks is less significant. In the EE bounds of model-based networks, the parameter B_0 relates the bound to the sparsity level ρ , of the target vectors. Lower values of ρ result in lower EE bounds, as demonstrated in the supplementary.

5 NUMERICAL EXPERIMENTS

In this section, we present a series of experiments that concentrate on how a particular model-based network (ISTA network) compares to a ReLU network, and showcase the merits of model-based networks. We focus on networks with 10 layers (similar to previous works (Gregor & LeCun, 2010)), to represent realistic model-based network architectures. The networks are trained on a simulated dataset to solve the problem in (1), with target vectors uniformly distributed in $[-1, 1]$. The linear mapping \mathbf{A} is constructed from the real part of discrete Fourier transform (DFT) matrix rows (Ong et al., 2019), where the rows are randomly chosen. The sparsity rate is $\rho = 0.15$, and the noise’s standard deviation is 0.1. To train the networks we used the SGD optimizer with the L_1 loss over all neurons of the last layer. The target and noise vectors are generated as element wise independently from a uniform distribution ranging in $[-1, 1]$. All results are reproducible through (Authors, 2022) which provides the complete code to execute the experiments presented in this section.

We concentrate on comparing the networks’ EE, since in practice, the networks are trained with a finite number of examples. In order to empirically approximate the EE of a class of networks \mathcal{H} , we use an empirical approximation of h^* (which satisfy $\mathcal{L}_{\mathcal{D}}(h^*) = \inf_{h \in \mathcal{H}} \mathcal{L}_{\mathcal{D}}(h)$) and the ERM \hat{h} , denoted by h_{emp}^* and \hat{h}_{emp} respectively. The estimator h_{emp}^* results from the trained network with 10^4 samples, and the ERM is approximated by a network trained using SGD with m training samples (where $m \leq 10^4$). The empirical EE is given by their difference $\mathcal{L}_{\mathcal{D}}(\hat{h}_{emp}) - \mathcal{L}_{\mathcal{D}}(h_{emp}^*)$.

In Fig. 2, we compare between the ISTA and ReLU networks, in terms of EE and L_1 loss, for networks trained with different number of samples (between 10 and 10^4 samples). We observe that for small number of training samples, the ISTA network substantially reduces the EE compared to the ReLU network. This can be understood from Theorem 5, which results in lower EE bounds on the ISTA networks compared to ReLU networks, due to the term $\lambda\mathbb{E}_S(T)$. However, for large number of samples the EE of both networks decreases to zero, which is also expected from Theorem 5. This highlights that the contribution of the soft-thresholding nonlinearity to the generalization abilities of the network is more significant for small number of training samples. Throughout the paper, we considered networks with constant bias terms. In this section, we also consider learned bias terms, as detailed in the supplementary material. In Fig. 2, we present the experimental results for networks with constant and learned biases. The experiments indicate that the choice of constant or learned bias is less significant to the EE or the accuracy, compared to the choice of nonlinearity, emphasizing the relevance of the theoretical guarantees. The cases of learned and constant biases have different optimal estimators, as the networks with learned biases have more learned parameters. As a result, it is plausible a network with more learnable parameters (the learnable bias) exhibits a lower estimation error since the corresponding optimal estimator also exhibits a lower estimation error. To analyze the effect of the soft-thresholding value on the generalization abilities of the ISTA network, we show in Fig. 3, the empirical EE for multiple values of λ . The experimental results demonstrate that for small number of samples, increasing λ reduces the EE. As expected from the EE bounds, for a large number of training samples, this dependency on the nonlinearity vanishes, and the EE is similar for all values

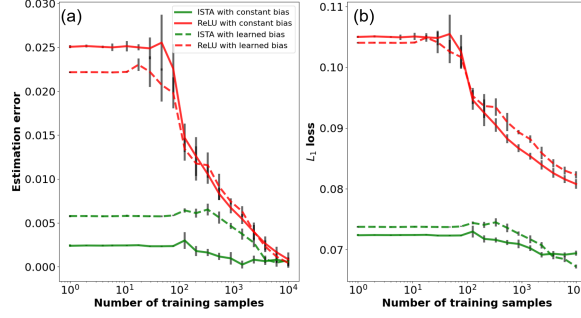


Figure 2: Comparing the EE of ISTA and ReLU networks with 10 layers. **a.** The EE of the networks as a function of the training samples. **b.** The L_1 of the networks as a function of the training samples. The ISTA network achieves lower EE compared to the ReLU network, along with lower losses.

of λ . In Fig. 3b, we show the L_1 loss of the ISTA networks for different values of λ . We observe that low estimation error does not necessarily lead to low loss value. For $m \lesssim 100$, increasing λ reduced the EE. These results suggest that given ISTA networks with different values of λ such that all networks achieve similar accuracy, the networks with higher values of λ provide lower EE. These results are also valid for additional networks’ depths. In Section B, we compare the EE for networks with different number of layers, and show that they exhibit a similar behaviour.

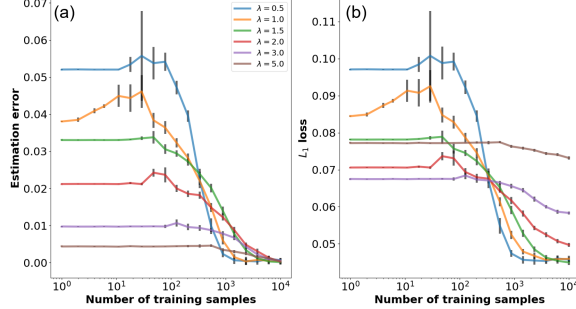


Figure 3: Estimation error and loss of ISTA networks with 10 layers, as a function of the soft-threshold’s value λ . The experiments indicate that the estimation error decreases as λ increases, which demonstrates a way to control the networks’ generalization abilities through λ .

6 CONCLUSION

We derived new GE and EE bounds for ISTA and ADMM networks, based on the RC and LRC frameworks. Under suitable conditions, the model-based networks’ GE is nonincreasing with depth, resulting in a substantial improvement compared to the GE bound of ReLU networks. The EE bounds explain EE behaviours experienced in practice, such as ISTA networks demonstrating higher estimation abilities, compared to ReLU networks, especially for small number of training samples. Through a series of experiments, we show that the generalization abilities of ISTA networks are controlled by the soft-threshold value, and achieve lower EE along with a more accurate recovery compared to ReLU networks which increase the GE and EE with the networks’ depths.

It is interesting to consider how the theoretical insights can be harnessed to enforce neural networks with high generalization abilities. One approach is to introduce an additional regularizer during the training process that is rooted in the LRC, penalizing networks with high EE (Yang et al., 2019).

REFERENCES

- J. Amjad, Z. Lyu, and M.R. Rodrigues. Deep learning model-aware regularization with applications to inverse problems. *IEEE Transactions on Signal Processing*, 69:6371–6385, 2021.
- Anonymous Authors. Model-based neural networks for sparse vector recovery. https://github.com/NeurIPS2022-Model-based-NN/Model_based_Network_Sparse_vector_recovery, 2022.
- P.L. Bartlett and S. Mendelson. Rademacher and gaussian complexities: Risk bounds and structural results. *Journal of Machine Learning Research*, 3(3):463–482, 2002.
- P.L. Bartlett, O. Bousquet, and S. Mendelson. Local rademacher complexities. *The Annals of Statistics*, 33(4):1497–1537, 2005.
- P.L. Bartlett, D.J. Foster, and M.J. Telgarsky. Spectrally-normalized margin bounds for neural networks. *Advances in neural information processing systems*, 30, 2017.
- A. Behboodi, H. Rauhut, and E. Schnoor. Compressive sensing and neural networks from a statistical learning perspective. *arXiv preprint arXiv:2010.15658*, 2020.
- M. Belkin, D. Hsu, S. Ma, and S. Mandal. Reconciling modern machine-learning practice and the classical bias–variance trade-off. *Proceedings of the National Academy of Sciences*, 116(32):15849–15854, 2019.
- G. Blanchard, O. Bousquet, and L. Zwald. Statistical properties of kernel principal component analysis. *Machine Learning*, 66(2):259–294, 2007.
- S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Found. Trends Mach. Learn.*, 3(1):1–122, 2011.
- W. Gao and Z.H. Zhou. Dropout rademacher complexity of deep neural networks. *Science China Information Sciences*, 59(7):1–12, 2016.
- N. Golowich, A. Rakhlin, and O. Shamir. Size-independent sample complexity of neural networks. *Conference on Learning Theory*, pp. 297–299, 2018.
- K. Gregor and Y. LeCun. Learning fast approximations of sparse coding. *Int. Conf. Machine Learning*, pp. 399–406, 2010.
- D. Jakubovitz, R. Giryes, and M.R. Rodrigues. Generalization error in deep learning. *Compressed sensing and its applications*, pp. 153–193, 2019.
- M. Ledoux and M. Talagrand. *Probability in Banach Spaces: isoperimetry and processes*. Springer Science Business Media, 1991.
- J. Lee, J. Sohl-dickstein, R. Pennington, J. Novak, S. Schoenholz, and Y. Bahri. Deep neural networks as gaussian processes. *International Conference on Learning Representations*, 2018.
- T. Liang, T. Poggio, A. Rakhlin, and J. Stokes. Fisher-rao metric, geometry, and complexity of neural networks. *The 22nd international conference on artificial intelligence and statistics*, pp. 888–896, 2019.
- S. Mendelson. Improving the sample complexity using global data. *IEEE transactions on Information Theory*, 48(7):1977–1991, 2002.
- V. Monga, Y. Li, and Y.C. Eldar. Algorithm unrolling: Interpretable, efficient deep learning for signal and image processing. *IEEE Signal Processing Magazine*, 38(2):18–44, 2021.
- B. Neyshabur, R. Tomioka, and N. Srebro. Norm-based capacity control in neural networks. *Conference on Learning Theory*, pp. 1376–1401, 2015a.
- B. Neyshabur, R. Tomioka, and N. Srebro. In search of the real inductive bias: On the role of implicit regularization in deep learning. *International Conference on Learning Representations workshop track*, 2015b.

- B. Neyshabur, Z. Li, S. Bhojanapalli, Y. LeCun, and N. Srebro. Towards understanding the role of over-parametrization in generalization of neural networks. *arXiv preprint arXiv:1805.12076*, 2018.
- R. Novak, Y. Bahri, D.A. Abolafia, J. Pennington, and J. Sohl-Dickstein. Sensitivity and generalization in neural networks: an empirical study. *International Conference on Learning Representations*, 2018.
- F. Ong, R. Heckel, and K. Ramchandran. A fast and robust paradigm for fourier compressed sensing based on coded sampling. *International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 5117–5121, 2019.
- D.P. Palomar and Y.C Eldar. *Convex optimization in signal processing and communications*. Cambridge university press, 2010.
- E. Schnoor, A. Behboodi, and H. Rauhut. Generalization error bounds for iterative recovery algorithms unfolded as neural networks. *arXiv preprint arXiv:2112.04364*, 2021.
- Shai Shalev-Shwartz and Shai Ben-David. *Understanding Machine Learning: From Theory to Algorithms*. Cambridge University Press, USA, 2014. ISBN 1107057132.
- J. Sokolić, R. Giryes, G. Sapiro, and M.R. Rodrigues. Robust large margin deep neural networks. *IEEE Transactions on Signal Processing*, 65(16):4265–4280, 2017.
- N. Srivastava, G.E. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov. Dropout: a simple way to prevent neural networks from overfitting. *Journal of machine learning research*, 15(1):1929–2958, 2014.
- Tibshirani and J. Ryan. The lasso problem and uniqueness. *Electronic Journal of Statistics*, 7(1):1465–1490, 2013.
- V.N. Vapnik and A. Chervonenkis. The necessary and sufficient conditions for consistency in the empirical risk minimization method. *Pattern Recognition and Image Analysis*, 1(3):260284, 1991.
- C. Xu, T. Liu, D. Tao, and C. Xu. Local rademacher complexity for multi-label learning. *IEEE Transactions on Image Processing*, 25(3):1495–1507, 2016.
- Y. Yang, J. Sun, H. Li, and Z. Xu. Admm-csnet: A deep learning approach for image compressive sensing. *IEEE transactions on pattern analysis and machine intelligence*, 42(3):521–538, 2018.
- Y. Yang, J. Yu, X. Li, J. Huan, and T.S. Huang. An empirical study on regularization of deep neural networks by local rademacher complexity. *arXiv preprint arXiv:1902.00873*, 2019.
- N. Yousefi, Y. Lei, M. Kloft, M. Mollaghasemi, and G.C. Anagnostopoulos. Local rademacher complexity-based learning guarantees for multi-task learning. *The Journal of Machine Learning Research*, 19(1):1385–1431, 2018.
- T. Zhang. Effective dimension and generalization of kernel learning. *Advances in Neural Information Processing Systems*, 15, 2002.

A SUPPLEMENTARY MATERIAL

In this supplementary material we provide the proofs of the presented lemmas and theorems.

A.1 RADEMACHER COMPLEXITY OF THE SOFT-THRESHOLDING OPERATOR

In this section, we prove Lemma 1 which examines how the RC is affected by the soft-thresholding operator.

Lemma 1 (Rademacher complexity of soft-thresholding):

Given any class of scalar functions \mathcal{H} where $h : \mathbb{R}^n \rightarrow \mathbb{R}$, $h \in \mathcal{H}$ for any integer $n \geq 1$, there exists a value T in the interval

$$T \in \left(0, \min \left\{ \frac{m\mathfrak{R}_m(\mathcal{H})}{\lambda}, m \right\} \right] \quad (14)$$

such that

$$\mathfrak{R}_m(S_\lambda \circ \mathcal{H}) \leq \mathfrak{R}_m(\mathcal{H}) - \frac{\lambda T}{m}. \quad (15)$$

Proof. We start by denoting

$$R = m\mathfrak{R}_m(S_\lambda \circ \mathcal{H}) = \mathbb{E}_{\{\epsilon_i\}_{i=1}^m} \sup_{h \in \mathcal{H}} \sum_{i=1}^m \epsilon_i \mathcal{S}_\lambda(h(\mathbf{y}_i)) \quad (16)$$

and we explicitly compute the expectation on ϵ_1 :

$$\begin{aligned} R &= \mathbb{E}_{\{\epsilon_i\}_{i=1}^m} \sup_{h \in \mathcal{H}} \sum_{i=1}^m \epsilon_i \mathcal{S}_\lambda(h(\mathbf{y}_i)) \\ &= \mathbb{E}_{\{\epsilon_i\}_{i=1}^m} \sup_{h \in \mathcal{H}} \epsilon_1 \mathcal{S}_\lambda(h(\mathbf{y}_1)) + \sum_{i=2}^m \epsilon_i \mathcal{S}_\lambda(h(\mathbf{y}_i)) \\ &= \mathbb{E}_{\{\epsilon_i\}_{i=2}^m} \left[\frac{1}{2} \sup_{h \in \mathcal{H}} \left(\mathcal{S}_\lambda(h(\mathbf{y}_1)) + \sum_{i=2}^m \epsilon_i \mathcal{S}_\lambda(h(\mathbf{y}_i)) \right) + \frac{1}{2} \sup_{h' \in \mathcal{H}} \left(-\mathcal{S}_\lambda(h'(\mathbf{y}_1)) + \sum_{i=2}^m \epsilon_i \mathcal{S}_\lambda(h'(\mathbf{y}_i)) \right) \right] \\ &= \frac{1}{2} \mathbb{E}_{\{\epsilon_i\}_{i=2}^m} \sup_{h \in \mathcal{H}} \left(\mathcal{S}_\lambda(h(\mathbf{y}_1)) + \sum_{i=2}^m \epsilon_i \mathcal{S}_\lambda(h(\mathbf{y}_i)) \right) + \sup_{h' \in \mathcal{H}} \left(-\mathcal{S}_\lambda(h'(\mathbf{y}_1)) + \sum_{i=2}^m \epsilon_i \mathcal{S}_\lambda(h'(\mathbf{y}_i)) \right) \\ &= \frac{1}{2} \mathbb{E}_{\{\epsilon_i\}_{i=2}^m} \sup_{h, h' \in \mathcal{H}} \left(\mathcal{S}_\lambda(h(\mathbf{y}_1)) - \mathcal{S}_\lambda(h'(\mathbf{y}_1)) + \sum_{i=2}^m \epsilon_i (\mathcal{S}_\lambda(h(\mathbf{y}_i)) + \mathcal{S}_\lambda(h'(\mathbf{y}_i))) \right). \end{aligned} \quad (17)$$

To simplify the notations, we denote

$$s_\epsilon(h, h') = \sum_{i=2}^m \epsilon_i (\mathcal{S}_\lambda(h(\mathbf{y}_i)) + \mathcal{S}_\lambda(h'(\mathbf{y}_i))) \quad (18)$$

and observe that $s_\epsilon(h, h') = s_\epsilon(h', h)$.

Next, we show that the soft-thresholding on the estimations of the first sample, reduces the RC, and results in

$$R \leq \frac{1}{2} \mathbb{E}_{\{\epsilon_i\}_{i=2}^m} \sup_{h, h' \in \mathcal{H}} (h(\mathbf{y}_1) - h'(\mathbf{y}_1) - 2\lambda \mathbb{1}_{h(\mathbf{y}_1) > \lambda \wedge h'(\mathbf{y}_1) < -\lambda} + s_\epsilon(h, h')). \quad (19)$$

The soft-thresholding function is piecewise linear with three pieces, and can be written as

$$\mathcal{S}_\lambda(x) = \begin{cases} x + \lambda, & x < -\lambda \\ 0, & x \in [-\lambda, \lambda] \\ x - \lambda, & x > \lambda \end{cases}. \quad (20)$$

We focus on the terms $h(\mathbf{y}_1)$ and $h'(\mathbf{y}_1)$, in the different regions determined by (20). To capture these regions, we define the following classes of functions

$$\begin{aligned}\mathcal{H}^{-1} &= \{h \in \mathcal{H} : h(\mathbf{y}_1) < -\lambda\} \\ \mathcal{H}^0 &= \{h \in \mathcal{H} : h(\mathbf{y}_1) \in [-\lambda, \lambda]\} \\ \mathcal{H}^1 &= \{h \in \mathcal{H} : h(\mathbf{y}_1) > \lambda\}.\end{aligned}\tag{21}$$

We can then write

$$r = \max_{\alpha, \beta \in \{-1, 0, 1\}} r_{\alpha, \beta}\tag{22}$$

where

$$r_{\alpha, \beta} = \sup_{h \in \mathcal{H}^\alpha, h' \in \mathcal{H}^\beta} (\mathcal{S}_\lambda(h(\mathbf{y}_1)) - \mathcal{S}_\lambda(h'(\mathbf{y}_1)) + s_\epsilon(h, h')).\tag{23}$$

We now divide into cases for the different values of α and β :

1. For $\alpha, \beta = 1$ or $\alpha, \beta = -1$, we obtain that $\mathcal{S}_\lambda(h(\mathbf{y}_1)) - \mathcal{S}_\lambda(h'(\mathbf{y}_1)) = h(\mathbf{y}_1) - h'(\mathbf{y}_1)$. As a result,

$$r_{1,1} = \sup_{h \in \mathcal{H}^1, h' \in \mathcal{H}^1} (h(\mathbf{y}_1) - h'(\mathbf{y}_1) + s_\epsilon(h, h'))\tag{24}$$

and

$$r_{-1,-1} = \sup_{h \in \mathcal{H}^{-1}, h' \in \mathcal{H}^{-1}} (h(\mathbf{y}_1) - h'(\mathbf{y}_1) + s_\epsilon(h, h')).\tag{25}$$

2. For $\alpha = 1$ and $\beta = 0$, we obtain that $\mathcal{S}_\lambda(h(\mathbf{y}_1)) - \mathcal{S}_\lambda(h'(\mathbf{y}_1)) = h(\mathbf{y}_1) - \lambda$. In this case, $h'(\mathbf{y}_1) \leq \lambda$, which results in

$$\begin{aligned}r_{1,0} &= \sup_{h \in \mathcal{H}^1, h' \in \mathcal{H}^0} (h(\mathbf{y}_1) - \lambda + s_\epsilon(h, h')) \\ &\leq \sup_{h \in \mathcal{H}^1, h' \in \mathcal{H}^0} (h(\mathbf{y}_1) - h'(\mathbf{y}_1) + s_\epsilon(h, h')).\end{aligned}\tag{26}$$

3. For $\alpha = 0$ and $\beta = -1$, we obtain that $\mathcal{S}_\lambda(h(\mathbf{y}_1)) - \mathcal{S}_\lambda(h'(\mathbf{y}_1)) = -h'(\mathbf{y}_1) - \lambda$. In this case, $h(\mathbf{y}_1) \geq -\lambda$, which results in

$$\begin{aligned}r_{0,-1} &= \sup_{h \in \mathcal{H}^0, h' \in \mathcal{H}^{-1}} (-h'(\mathbf{y}_1) - \lambda + s_\epsilon(h, h')) \\ &\leq \sup_{h \in \mathcal{H}^0, h' \in \mathcal{H}^{-1}} (h(\mathbf{y}_1) - h'(\mathbf{y}_1) + s_\epsilon(h, h')).\end{aligned}\tag{27}$$

4. For $\alpha = 1$ and $\beta = -1$, we obtain that $\mathcal{S}_\lambda(h(\mathbf{y}_1)) - \mathcal{S}_\lambda(h'(\mathbf{y}_1)) = h(\mathbf{y}_1) - h'(\mathbf{y}_1) - 2\lambda$, which results in

$$r_{1,-1} = \sup_{h \in \mathcal{H}^1, h' \in \mathcal{H}^{-1}} (h(\mathbf{y}_1) - h'(\mathbf{y}_1) - 2\lambda + s_\epsilon(h, h')).\tag{28}$$

In this case, we obtain a reduction by 2λ in the complexity.

5. For $\alpha, \beta = 0$, we obtain that $\mathcal{S}_\lambda(h(\mathbf{y}_1)) - \mathcal{S}_\lambda(h'(\mathbf{y}_1)) = 0$, which leads to

$$r_{0,0} = \sup_{h \in \mathcal{H}^0, h' \in \mathcal{H}^0} s_\epsilon(h, h').\tag{29}$$

We bound this term to obtain a similar expression to the previous cases. Any pair of estimators $h, h' \in \mathcal{H}^0$ satisfies that $h(\mathbf{y}_1) \geq h'(\mathbf{y}_1)$ or $h(\mathbf{y}_1) \leq h'(\mathbf{y}_1)$. Using the fact that $s_\epsilon(h, h') = s_\epsilon(h', h)$, there exists a pair $h, h' \in \mathcal{H}^0$ satisfying $h(\mathbf{y}_1) - h'(\mathbf{y}_1) + s_\epsilon(h, h') \geq s_\epsilon(h, h')$ (otherwise, interchanging between the estimators satisfies this requirement since $h'(\mathbf{y}_1) - h(\mathbf{y}_1) + s_\epsilon(h, h') \geq s_\epsilon(h, h') = s_\epsilon(h', h)$, where the estimators satisfy $h', h \in \mathcal{H}^0$). This means that the term $r_{0,0}$ is bounded by

$$r_{0,0} \leq \sup_{h \in \mathcal{H}^0, h' \in \mathcal{H}^0} (h(\mathbf{y}_1) - h'(\mathbf{y}_1) + s_\epsilon(h, h')).\tag{30}$$

We will now show that the additional cases that were not presented above, are dominated by other cases (i.e. upper bounded), and therefore will not be considered in (22):

1. For $\alpha = -1$ and $\beta = 1$, we obtain that $\mathcal{S}_\lambda(h(\mathbf{y}_1)) - \mathcal{S}_\lambda(h'(\mathbf{y}_1)) = h(\mathbf{y}_1) - h'(\mathbf{y}_1) + 2\lambda$, which results in:

$$r_{-1,1} = \sup_{h \in \mathcal{H}^{-1}, h' \in \mathcal{H}^1} (h(\mathbf{y}_1) - h'(\mathbf{y}_1) + 2\lambda + s_\epsilon(h, h')). \quad (31)$$

We now show that $r_{-1,1} \leq r_{1,-1}$, and therefore the case of $\alpha = -1$ and $\beta = 1$ is upper bounded by the case of $\alpha = 1$ and $\beta = -1$. For any pair of estimators in $h \in \mathcal{H}^{-1}$ and $h' \in \mathcal{H}^1$, there exists a pair of estimators in $h \in \mathcal{H}^1$ and $h' \in \mathcal{H}^{-1}$, which achieve a higher value (showing that $r_{-1,1} \leq r_{1,-1}$). Since we consider the case where $\alpha = -1$ and $\beta = 1$, there exist $\delta(h), \delta'(h') > 0$ such that $h(\mathbf{y}_1) = -\delta(h) - \lambda$ and $h'(\mathbf{y}_1) = \delta'(h') + \lambda$, resulting in:

$$r_{-1,1} = \sup_{h \in \mathcal{H}^{-1}, h' \in \mathcal{H}^1} c_{-1,1}(h, h') \quad (32)$$

where

$$c_{-1,1}(h, h') = -\delta(h) - \delta'(h') + s_\epsilon(h, h'). \quad (33)$$

Similarly, for the case of $\alpha = 1$ and $\beta = -1$, we can write

$$r_{1,-1} = \sup_{h \in \mathcal{H}^1, h' \in \mathcal{H}^{-1}} c_{1,-1}(h, h') \quad (34)$$

such that

$$\begin{aligned} c_{1,-1}(h, h') &= \lambda + \delta(h) + \lambda + \delta'(h') - 2\lambda + s_\epsilon(h, h') \\ &= \delta(h) + \delta'(h') + s_\epsilon(h, h'). \end{aligned} \quad (35)$$

Since $s_\epsilon(h', h) = s_\epsilon(h, h')$, for any estimators $h_{-1} = -\delta_{-1}(h_{-1}) - \lambda \in \mathcal{H}^{-1}$ and $h_1 = \delta_1(h_1) + \lambda \in \mathcal{H}^1$

$$\begin{aligned} c_{1,-1}(h_1, h_{-1}) &= \delta_1(h_1) + \delta_2(h_{-1}) + s_\epsilon(h_1, h_{-1}) \\ &\geq -\delta_1(h_1) - \delta_2(h_{-1}) + s_\epsilon(h_{-1}, h_1) \\ &= c_{-1,1}(h_{-1}, h_1) \end{aligned} \quad (36)$$

implying that

$$r_{-1,1} = \sup_{h \in \mathcal{H}^{-1}, h' \in \mathcal{H}^1} c_{-1,1}(h, h') \leq \sup_{h' \in \mathcal{H}^1, h \in \mathcal{H}^{-1}} c_{1,-1}(h', h) = r_{1,-1}. \quad (37)$$

2. For $\alpha = -1$ and $\beta = 0$, we obtain that $\mathcal{S}_\lambda(h(\mathbf{y}_1)) - \mathcal{S}_\lambda(h'(\mathbf{y}_1)) = h(\mathbf{y}_1) + \lambda$, which results in:

$$r_{-1,0} = \sup_{h \in \mathcal{H}^{-1}, h' \in \mathcal{H}^1} c_{-1,0}(h, h') \quad (38)$$

where

$$c_{-1,0}(h, h') = h(\mathbf{y}_1) + \lambda + s_\epsilon(h, h'). \quad (39)$$

We will show that $r_{-1,0} \leq r_{0,-1}$, and therefore the case of $\alpha = -1$ and $\beta = 0$ is dominated (upper bounded) by the case of $\alpha = 0$ and $\beta = -1$. Since we consider the case where $\alpha = -1$ and $\beta = 0$, there exist $\delta(h) > 0$ such that $h(\mathbf{y}_1) = -\delta(h) - \lambda$, resulting in:

$$c_{-1,0}(h, h') = -\delta(h) + s_\epsilon(h, h'). \quad (40)$$

Similarly, for $\alpha = 0$ and $\beta = -1$, we can write

$$r_{0,-1} = \sup_{h \in \mathcal{H}^0, h' \in \mathcal{H}^{-1}} c_{0,-1}(h, h') \quad (41)$$

where

$$\begin{aligned} c_{0,-1}(h, h') &= \lambda + \delta(h) - \lambda + s_\epsilon(h', h) \\ &= \delta(h) + s_\epsilon(h, h'). \end{aligned} \quad (42)$$

Again, we consider any estimators $h_{-1} \in \mathcal{H}^{-1}$ and $h_0 \in \mathcal{H}^0$

$$\begin{aligned} c_{0,-1}(h_0, h_{-1}) &= \delta_{-1}(h_{-1}) + s_\epsilon(h_0, h_{-1}) \\ &\geq -\delta_{-1}(h_{-1}) + s_\epsilon(h_{-1}, h_0) \\ &= c_{-1,0}(h_{-1}, h_0). \end{aligned} \quad (43)$$

This implies that for any pair of estimators in \mathcal{H}^{-1} and \mathcal{H}^0 , there exists a pair of estimators in \mathcal{H}^0 and \mathcal{H}^{-1} , which achieve a higher value, proving that

$$r_{-1,0} = \sup_{h \in \mathcal{H}^{-1}, h' \in \mathcal{H}^0} c_{-1,0}(h, h') \leq \sup_{h' \in \mathcal{H}^0, h \in \mathcal{H}^{-1}} c_{0,-1}(h', h) = r_{0,-1}. \quad (44)$$

Following the same chain of thought, one can prove that $r_{0,1} \leq r_{1,0}$, which makes the case of $\alpha = 0$ and $\beta = 1$ dominated by that of $\alpha = 1$ and $\beta = 0$.

Taking into account the dominant cases, we can re-write (22) as

$$r = \max_{\alpha, \beta \in \{-1, 0, 1\}} r_{\alpha, \beta} = \max_{(\alpha, \beta) \in \mathcal{I}} r_{\alpha, \beta} \quad (45)$$

where

$$\mathcal{I} = \{(1, 1), (1, 0), (1, -1), (0, 0), (0, -1)\} \quad (46)$$

is the set of dominant cases. We observe, that all the dominant cases, $(\alpha, \beta) \in \mathcal{I}$, are bounded by

$$r_{\alpha, \beta} \leq \sup_{h \in \mathcal{H}^\alpha, h' \in \mathcal{H}^\beta} (h(\mathbf{y}_1) - h'(\mathbf{y}_1) - 2\lambda \mathbb{1}_{h(\mathbf{y}_1) > \lambda \wedge h'(\mathbf{y}_1) < -\lambda} + s_\epsilon(h, h')) \quad (47)$$

where the indicator takes into account the reduction in complexity obtained for the case of $\alpha = 1$ and $\beta = -1$. Since $\mathcal{H}^\alpha \subseteq \mathcal{H}$, $\alpha \in \{-1, 0, 1\}$, the above is bounded by

$$r_{\alpha, \beta} \leq \sup_{h \in \mathcal{H}, h' \in \mathcal{H}} (h(\mathbf{y}_1) - h'(\mathbf{y}_1) - 2\lambda \mathbb{1}_{h(\mathbf{y}_1) > \lambda \wedge h'(\mathbf{y}_1) < -\lambda} + s_\epsilon(h, h')) \quad (48)$$

which leads to

$$r = \max_{(\alpha, \beta) \in \mathcal{I}} r_{\alpha, \beta} \leq \sup_{h \in \mathcal{H}, h' \in \mathcal{H}} (h(\mathbf{y}_1) - h'(\mathbf{y}_1) - 2\lambda \mathbb{1}_{h(\mathbf{y}_1) > \lambda \wedge h'(\mathbf{y}_1) < -\lambda} + s_\epsilon(h, h')) \quad (49)$$

proving the bound in (19).

Next, we aim to separate the indicator in (19) from the rest of the expression. Let us denote by h^* and h'^* the estimators that achieve the supremum in (19), which depend on the value of the Rademacher random variables ϵ_i , $i \in [2, m]$. Then,

$$\begin{aligned} R &\leq \frac{1}{2} \mathbb{E}_{\{\epsilon_i\}_{i=2}^m} (h^*(\mathbf{y}_1) - h'^*(\mathbf{y}_1) - 2\lambda \mathbb{1}_{h^*(\mathbf{y}_1) > \lambda \wedge h'^*(\mathbf{y}_1) < -\lambda} + s_\epsilon(h^*, h'^*)) \\ &= \frac{1}{2} \mathbb{E}_{\{\epsilon_i\}_{i=2}^m} (h^*(\mathbf{y}_1) - h'^*(\mathbf{y}_1) + s_\epsilon(h^*, h'^*)) - \lambda \mathbb{E}_{\{\epsilon_i\}_{i=2}^m} \mathbb{1}_{h^*(\mathbf{y}_1) > \lambda \wedge h'^*(\mathbf{y}_1) < -\lambda} \\ &= \frac{1}{2} \mathbb{E}_{\{\epsilon_i\}_{i=2}^m} (h^*(\mathbf{y}_1) - h'^*(\mathbf{y}_1) + s_\epsilon(h^*, h'^*)) - \lambda T_1 \end{aligned} \quad (50)$$

where we denoted $T_1 = \mathbb{E}_{\{\epsilon_i\}_{i=2}^m} \mathbb{1}_{h^*(\mathbf{y}_1) > \lambda \wedge h'^*(\mathbf{y}_1) < -\lambda} \in [0, 1]$. The RC is a nonnegative quantity, which translates to restrictions on the value of T_1 . Note, that the supremum over the first term in (50), increases its value, leading to

$$R \leq \frac{1}{2} \mathbb{E}_{\{\epsilon_i\}_{i=2}^m} \sup_{h, h' \in \mathcal{H}} (h(\mathbf{y}_1) - h'(\mathbf{y}_1) + s_\epsilon(h, h')) - \lambda T_1. \quad (51)$$

Following the same chain of reasoning, we introduce a Rademacher random variable into (51) and express the above as

$$R \leq \mathbb{E}_{\{\epsilon_i\}_{i=1}^m} \sup_{h \in \mathcal{H}} \left(\epsilon_1 h(\mathbf{y}_1) + \sum_{i=2}^m \epsilon_i \mathcal{S}_\lambda(h(\mathbf{y}_i)) \right) - \lambda T_1. \quad (52)$$

Applying the same chain of thoughts on each sample, results in similar reduction in the RC by quantities $T_i \in [0, 1]$, $i \in [2, m]$, which are obtained by repeating the process in (52) for all m samples, similarly to (50). Repeating the process for all $i \in [2, m]$, leads to

$$R \leq \mathbb{E}_{\{\epsilon_i\}_{i=1}^m} \sup_{h \in \mathcal{H}} \sum_{i=1}^m \epsilon_i h(\mathbf{y}_i) - \lambda T \quad (53)$$

where $T = \sum_{i=1}^m T_i \in [0, m]$. Substituting the definition of the RC in (8), results in

$$\mathfrak{R}_m(S_\lambda \circ \mathcal{H}) = \frac{R}{m} \leq \mathfrak{R}_m(\mathcal{H}) - \frac{\lambda T}{m}. \quad (54)$$

The RC is nonnegative, leading to

$$T \leq \min \left\{ \frac{m\mathfrak{R}_m(\mathcal{H})}{\lambda}, m \right\}, \quad (55)$$

which completes the proof. \square

A.2 GENERALIZATION ERROR BOUNDS

In this section we bound the GE of ISTA and ADMM networks. We start by presenting a few useful lemmas, starting by Talagrand's contraction lemma.

Lemma 2 (Talagrand's contraction lemma (Ledoux & Talagrand, 1991)). *Let ϕ be a G -Lipschitz function. Then*

$$\mathbb{E}_{\{\epsilon_i\}_{i=1}^m} \sup_{h \in \mathcal{H}} \frac{1}{m} \sum_{i=1}^m \epsilon_i \phi(h(\mathbf{y}_i)) \leq G \mathbb{E}_{\{\epsilon_i\}_{i=1}^m} \sup_{h \in \mathcal{H}} \frac{1}{m} \sum_{i=1}^m \epsilon_i h(\mathbf{y}_i). \quad (56)$$

Lemma 3. *For any class of functions \mathcal{H} , define $\text{absconv}(\mathcal{H}) = \{\sum \alpha_i h_i : h_i \in \mathcal{H}, \sum |\alpha_i| = 1\}$. Then (Bartlett & Mendelson, 2002),*

$$\mathfrak{R}_m(\mathcal{H}) = \mathfrak{R}_m(\text{absconv}(\mathcal{H})). \quad (57)$$

Lemma 4. *Consider a class of functions \mathcal{H} such that $h : \mathbb{R}^{n_x} \rightarrow \mathbb{R}$, $h \in \mathcal{H}$. Denote by $\mathbf{h}(\mathbf{y})$ the vector of estimators such that any entry results from an estimator in the same class of functions $\mathbf{h}_j \in \mathcal{H}$, $j \in [1, n_x]$. If $\mathbf{w} \in \mathbb{R}^{n_x}$ such that $\|\mathbf{w}\|_1 < B$, then*

$$\mathbb{E}_{\{\epsilon_i\}_{i=1}^m} \sup_{\mathbf{w}, \mathbf{h}} \frac{1}{m} \sum_{i=1}^m \epsilon_i \mathbf{w}^T \mathbf{h}(\mathbf{y}_i) \leq B \mathbb{E}_{\{\epsilon_i\}_{i=1}^m} \sup_{h \in \mathcal{H}} \frac{1}{m} \sum_{i=1}^m \epsilon_i h(\mathbf{y}_i). \quad (58)$$

Proof. Computing the LHS explicitly

$$\begin{aligned} & \mathbb{E}_{\{\epsilon_i\}_{i=1}^m} \sup_{\|\mathbf{w}\|_1 \leq B, \mathbf{h}} \frac{1}{m} \sum_{i=1}^m \epsilon_i \mathbf{w}^T \mathbf{h}(\mathbf{y}_i) \\ &= \mathbb{E}_{\{\epsilon_i\}_{i=1}^m} \sup_{\|\mathbf{w}\|_1 \leq B, \mathbf{h}} \|\mathbf{w}\|_1 \frac{1}{m} \sum_{i=1}^m \epsilon_i \frac{\mathbf{w}^T}{\|\mathbf{w}\|_1} \mathbf{h}(\mathbf{y}_i) \\ &\leq B \mathbb{E}_{\{\epsilon_i\}_{i=1}^m} \sup_{\|\mathbf{w}\|_1 \leq B, \mathbf{h}} \frac{1}{m} \sum_{i=1}^m \epsilon_i \frac{\mathbf{w}^T}{\|\mathbf{w}\|_1} \mathbf{h}(\mathbf{y}_i) \\ &= B \mathbb{E}_{\{\epsilon_i\}_{i=1}^m} \sup_{\|\tilde{\mathbf{w}}\|_1 = 1, \mathbf{h}} \frac{1}{m} \sum_{i=1}^m \epsilon_i \tilde{\mathbf{w}}^T \mathbf{h}(\mathbf{y}_i) \end{aligned} \quad (59)$$

where $\tilde{\mathbf{w}} = \mathbf{w} / \|\mathbf{w}\|_1$. Since $\|\tilde{\mathbf{w}}\|_1 = 1$, applying Lemma 3 results in

$$\begin{aligned} \mathbb{E}_{\{\epsilon_i\}_{i=1}^m} \sup_{\|\tilde{\mathbf{w}}\|_1 = 1, \mathbf{h}} \frac{1}{m} \sum_{i=1}^m \epsilon_i \tilde{\mathbf{w}}^T \mathbf{h}(\mathbf{y}_i) &= \mathbb{E}_{\{\epsilon_i\}_{i=1}^m} \sup_{\|\tilde{\mathbf{w}}\|_1 = 1, \mathbf{h}} \frac{1}{m} \sum_{i=1}^m \epsilon_i \sum_{j=1}^{n_x} \tilde{w}_j \mathbf{h}_j(\mathbf{y}_i) \\ &= \mathbb{E}_{\{\epsilon_i\}_{i=1}^m} \sup_{h \in \text{absconv}(\mathcal{H})} \frac{1}{m} \sum_{i=1}^m \epsilon_i h(\mathbf{y}_i) \\ &= \mathbb{E}_{\{\epsilon_i\}_{i=1}^m} \sup_{h \in \mathcal{H}} \frac{1}{m} \sum_{i=1}^m \epsilon_i h(\mathbf{y}_i), \end{aligned} \quad (60)$$

where the last equality results from Lemma 3. Combining (59) and (60), completes the proof. \square

Theorem 3 (Generalization error bound for ISTA networks):

Consider the class of learned ISTA networks of depth L as described in (2), and m i.i.d. samples. Then there exist $T^{(l)}$ for $l \in [1, L]$ in the range

$$T^{(l)} \in \left[0, \min \left\{ \frac{m B_l G_I^{l-1}}{\lambda}, m \right\} \right] \quad (61)$$

where

$$G_I^l = \frac{B_0 \prod_{l'=1}^l B_{l'}}{\sqrt{m}} - \frac{\lambda}{m} \sum_{l'=1}^{l-1} T^{(l')} \prod_{j=l'+1}^l B_j - \frac{\lambda T^{(l)}}{m}, \quad (62)$$

satisfying

$$\mathcal{G}(\mathcal{H}_I^L) \leq 2\mathbb{E}_S G_I^L. \quad (63)$$

Proof. To bound the GE of \mathcal{H}_I^L , we rely on the relation with the RC of \mathcal{H}_I^L , described in (9). Using the assumption that the loss function \mathcal{L} can be written as an average of the per-coordinate losses, leads to

$$\begin{aligned} \mathfrak{R}_m(\mathcal{L} \circ \mathcal{H}_I^L) &= \mathbb{E}_{\{\epsilon_i\}_{i=1}^m} \sup_{\mathbf{h}^L \in \mathcal{H}_I^L} \frac{1}{m} \sum_{i=1}^m \epsilon_i \mathcal{L}(\mathbf{h}^L(\mathbf{y}_i), \mathbf{x}_i) \\ &= \mathbb{E}_{\{\epsilon_i\}_{i=1}^m} \sup_{\mathbf{h}^L \in \mathcal{H}_I^L} \frac{1}{m} \sum_{i=1}^m \epsilon_i \frac{1}{n_x} \sum_{j=1}^{n_x} \ell(\mathbf{h}_j^L(\mathbf{y}_i), \mathbf{x}_{i,j}) \end{aligned} \quad (64)$$

where $\mathbf{x}_{i,j}$ denotes the j th coordinate of the i th sample. We relax the above supremum, by taking the supremum on each coordinate separately

$$\mathfrak{R}_m(\mathcal{L} \circ \mathcal{H}_I^L) \leq \mathbb{E}_{\{\epsilon_i\}_{i=1}^m} \sup_{\mathbf{h}_j^L \in \mathcal{H}_I^L, j \in [1, n_x]} \frac{1}{m} \sum_{i=1}^m \epsilon_i \frac{1}{n_x} \sum_{j=1}^{n_x} \ell(\mathbf{h}_j^L(\mathbf{y}_i), \mathbf{x}_{i,j}) \quad (65)$$

where \mathcal{H}_I^L is the class of scalar function that represents a single neuron at the output layer. We observe that each neuron in the output layer is given by the same class of functions \mathcal{H}_I^L . Since the supremum applies on a separable function, (65) reads

$$\mathfrak{R}_m(\mathcal{L} \circ \mathcal{H}_I^L) \leq \frac{1}{n_x} \sum_{j=1}^{n_x} \mathbb{E}_{\{\epsilon_i\}_{i=1}^m} \sup_{\mathbf{h}_j^L \in \mathcal{H}_I^L} \frac{1}{m} \sum_{i=1}^m \epsilon_i \ell(\mathbf{h}_j^L(\mathbf{y}_i), \mathbf{x}_{i,j}). \quad (66)$$

Applying a 1-Lipschitz loss (with respect to its first entry) on the network's prediction, satisfies the same bound, as detailed in (Gao & Zhou, 2016) (the proof relies on Lemma 2). As a result

$$\mathfrak{R}_m(\mathcal{L} \circ \mathcal{H}_I^L) \leq \frac{1}{n_x} \sum_{j=1}^{n_x} \mathbb{E}_{\{\epsilon_i\}_{i=1}^m} \sup_{\mathbf{h}_j^L \in \mathcal{H}_I^L} \frac{1}{m} \sum_{i=1}^m \epsilon_i \mathbf{h}_j^L(\mathbf{y}_i). \quad (67)$$

We observe that the supremum in (67) is repeated for each coordinate, although the function is taken over the same class of scalar functions \mathcal{H}_I^L . The bound in (67) can be replaced with

$$\mathfrak{R}_m(\mathcal{L} \circ \mathcal{H}_I^L) \leq \mathbb{E}_{\{\epsilon_i\}_{i=1}^m} \sup_{\mathbf{h}^L \in \mathcal{H}_I^L} \frac{1}{m} \sum_{i=1}^m \epsilon_i \mathbf{h}^L(\mathbf{y}_i). \quad (68)$$

For the rest of the proof, we focus on bounding the RC of a single neuron in the ISTA network, denoted by

$$R_I^L := \mathbb{E}_{\{\epsilon_i\}_{i=1}^m} \sup_{\mathbf{h}^L \in \mathcal{H}_I^L} \frac{1}{m} \sum_{i=1}^m \epsilon_i \mathbf{h}^L(\mathbf{y}_i). \quad (69)$$

Consider the change of the RC between classes of functions representing networks with a consecutive number of layers. The RC of \mathcal{H}_I^{l+1} is

$$R_I^{l+1} = \mathfrak{R}_m(\mathcal{H}_I^{l+1}) = \mathbb{E}_{\{\epsilon_i\}_{i=1}^m} \sup_{\mathbf{w}^{l+1}, \mathbf{h}^l} \frac{1}{m} \sum_{i=1}^m \epsilon_i \mathcal{S}_\lambda(\mathbf{w}^{l+1, T} \mathbf{h}^l(\mathbf{y}_i) + b) \quad (70)$$

where \mathbf{w}^{l+1} is a row in the weights matrix of the corresponding layer (such that $\|\mathbf{w}^{l+1}\|_1 \leq B_{l+1}$), and \mathbf{h}^l is a vector of estimators such that any entry results from an estimator in the class of functions \mathcal{H}_I^l , which results from previous layer in the network. In addition, b corresponds to a single entry of the bias term \mathbf{b} .

Applying the bound on the RC due to the soft-thresholding from Lemma 1 leads to

$$\begin{aligned} R_I^{l+1} &\leq \mathbb{E}_{\{\epsilon_i\}_{i=1}^m} \sup_{\mathbf{w}^{l+1}, \mathbf{h}^l} \frac{1}{m} \sum_{i=1}^m \epsilon_i (\mathbf{w}^{l+1} \mathbf{h}^l(\mathbf{y}_i) + b) - \frac{\lambda T^{(l+1)}}{m} \\ &= \mathbb{E}_{\{\epsilon_i\}_{i=1}^m} \sup_{\mathbf{w}^{l+1}, \mathbf{h}^l} \frac{1}{m} \sum_{i=1}^m \epsilon_i \mathbf{w}^{l+1} \mathbf{h}^l(\mathbf{y}_i) - \frac{\lambda T^{(l+1)}}{m} \end{aligned} \quad (71)$$

where $T^{(l+1)}$ results from Lemma 1 and corresponds to the l th layer in the network. The last equality follows from the symmetry of the Rademacher random variables which cancels the dependence on b , since it is a constant scalar. We observe that the assumption of a constant bias is used in (71). Since it is the only place where we used this assumption, obtaining the result in (71) without the restriction of a constant bias, will allow to relax this assumption.

Applying Lemma 4 implies

$$R_I^{l+1} \leq B_{l+1} \mathbb{E}_{\{\epsilon_i\}_{i=1}^m} \sup_{\mathbf{h}^l} \frac{1}{m} \sum_{i=1}^m \epsilon_i h^l(\mathbf{y}_i) - \frac{\lambda T^{(l+1)}}{m}. \quad (72)$$

Substituting into the RC definition from (8), leads to a recurrence relation on the RC of consecutive layers in the network

$$R_I^{l+1} \leq B_{l+1} R_I^l - \frac{\lambda T^{(l+1)}}{m}. \quad (73)$$

Moreover, the first layer of the network is obtained by applying a linear mapping, followed by the soft-thresholding operator. The bound on the RC of linear predictors, $\mathcal{H}_L = \{\mathbf{x} \mapsto \mathbf{w}\mathbf{x} : \|\mathbf{w}\|_2 \leq B_1\}$, is given by

$$R_{\mathcal{H}_L} \leq \frac{B_1 B_0}{\sqrt{m}} \quad (74)$$

where B_0 bounds the network's initialization as detailed in Section 2.1 (Shalev-Shwartz & Ben-David, 2014). Therefore, the RC of the first layer is obtained by applying the bound in (74), followed by Lemma 1

$$R_I^1 \leq \frac{B_0 B_1}{\sqrt{m}} - \frac{\lambda T^{(1)}}{m}. \quad (75)$$

Applying the recurrence relation L times results in

$$R_I^L \leq G_I^L = \frac{B_0 \prod_{l=1}^L B_l}{\sqrt{m}} - \frac{\lambda}{m} \sum_{l=1}^{L-1} T^{(l)} \prod_{j=l+1}^L B_j - \frac{\lambda T^{(L)}}{m} \quad (76)$$

where G_I^L is obtained by the same recurrence relation as in (73). To ensure that the bound on the RC is nonnegative, we restrict the value of $T^{(l+1)}$ to

$$T^{(l+1)} \in \left[0, \min \left\{ \frac{m B_{l+1} G_I^l}{\lambda}, m \right\} \right]. \quad (77)$$

Following (9), the GE is bounded by

$$\mathcal{G}(\mathcal{H}_I^L) \leq 2\mathbb{E}_S \mathfrak{R}_m(\mathcal{L} \circ \mathcal{H}_I^L) \leq 2\mathbb{E}_S R_I^L \leq 2\mathbb{E}_S G_I^L, \quad (78)$$

which completes the proof. \square

Next, we show that for a specific distribution, the expected value of T is greater than 0, demonstrating that for specific network parameters, ISTA and ADMM networks achieve a significant reduction in the GE. In the following proposition, the quantities $T^{(l)}$, B_l , and \mathbf{w}^l , for $l \in [1, L]$, follow from Theorem 3.

Proposition 1. Consider the distribution of $T^{(l)}$ where \mathbf{b}_i is a vector of random variables, ranging in $[-B_0, B_0]$. We assume there exists a constant $c \in [0, 1]$ such that for the underlying data distribution \mathcal{D} , $\mathbb{E}_{\mathcal{D}}(e^{-\mathbf{w}^{l,*T} \mathbf{b}_1}) \leq e^{-cB_l b^{(l)}}$ and $\mathbb{E}_{\mathcal{D}}(e^{\mathbf{w}^{l,*T} \mathbf{b}_1}) \leq e^{-cB_l b^{(l)}}$, where $b^{(l+1)} = B_l b^{(l)} - \lambda$, $b^{(1)} = B_0$. In addition, the weights $\mathbf{w}^{l,*}$ and $\mathbf{w}^{l,*}$ are obtained by the optimal estimators from Lemma 1. Then the expected value of $T^{(l)}$ is lower bounded by

$$\mathbb{E}_S \left(T^{(l)} \right) \geq \max \left\{ m \left(1 - 2e^{-(cB_l b^{(l)} - \lambda)} \right), 0 \right\}. \quad (79)$$

Proof. From the definition in (50), the expected value of $T_1^{(1)}$ is:

$$\begin{aligned} \mathbb{E}_{\mathcal{D}} \left(T_1^{(1)} \right) &= \mathbb{E}_{\mathcal{D}} \left(\mathbb{E}_{\{\epsilon_i\}_{i=2}^m} \mathbb{1}_{h^{1,*}(\mathbf{y}_1) > \lambda \wedge h'^{1,*}(\mathbf{y}_1) < -\lambda} \right) \\ &= \mathbb{E}_{\{\epsilon_i\}_{i=2}^m} \left(\mathbb{E}_{\mathcal{D}} \mathbb{1}_{h^{1,*}(\mathbf{y}_1) > \lambda \wedge h'^{1,*}(\mathbf{y}_1) < -\lambda} \right) \\ &= \mathbb{E}_{\{\epsilon_i\}_{i=2}^m} \left(\mathbb{P}_{\mathcal{D}} \left(h^{1,*}(\mathbf{y}_1) > \lambda \wedge h'^{1,*}(\mathbf{y}_1) < -\lambda \right) \right) \\ &= \mathbb{E}_{\{\epsilon_i\}_{i=2}^m} \left(1 - \mathbb{P}_{\mathcal{D}} \left(h^{1,*}(\mathbf{y}_1) \leq \lambda \vee h'^{1,*}(\mathbf{y}_1) \geq -\lambda \right) \right) \\ &\geq \mathbb{E}_{\{\epsilon_i\}_{i=2}^m} \left(1 - \mathbb{P}_{\mathcal{D}} \left(h^{1,*}(\mathbf{y}_1) \leq \lambda \right) - \mathbb{P}_{\mathcal{D}} \left(h'^{1,*}(\mathbf{y}_1) \geq -\lambda \right) \right) \end{aligned} \quad (80)$$

where we replaced the expected value of the indicator function, by the probability of the event, and applied the union bound. We substitute the outcome of the first layer by $h^{1,*}(\mathbf{y}_1) = \mathbf{w}^{1,*T} \mathbf{b}_1$ and $h'^{1,*}(\mathbf{y}_1) = \mathbf{w}'^{1,*T} \mathbf{b}_1$, where \mathbf{b}_1 is the bias term defined in Section 2.1.

Following the assumption, there exists a constant $c \in [0, 1]$ such that $\mathbb{E}_{\mathcal{D}}(e^{-\mathbf{w}^{1,*T} \mathbf{b}_1}) \leq e^{-cBB_0}$. This assumption captures the relation between the optimal weights $\mathbf{w}^{1,*}$ and the inputs \mathbf{b}_1 . Since the entries of $\mathbf{w}^{1,*}$ and \mathbf{b}_1 are smaller than B_0 and B in absolute value, the constant c captures how the product of both vectors is close to their maximal value.

Applying Chernoff bound leads to

$$\mathbb{P}_{\mathcal{D}} \left(h^{1,*}(\mathbf{y}_1) \leq \lambda \right) \leq \frac{\mathbb{E}_{\mathcal{D}} \left(e^{-h^{1,*}(\mathbf{y}_1)} \right)}{e^{-\lambda}} = \frac{\mathbb{E}_{\mathcal{D}} \left(e^{-\mathbf{w}^{1,*T} \mathbf{b}_1} \right)}{e^{-\lambda}} \leq \frac{\mathbb{E}_{\mathcal{D}} \left(e^{-cBB_0} \right)}{e^{-\lambda}} = e^{\lambda - cBB_0}. \quad (81)$$

Similarly, we assume that $\mathbb{E}_{\mathcal{D}}(e^{\mathbf{w}'^{1,*T} \mathbf{b}_1}) \leq e^{-cBB_0}$, and obtain

$$\mathbb{P}_{\mathcal{D}} \left(h'^{1,*}(\mathbf{y}_1) \geq -\lambda \right) \leq \frac{\mathbb{E}_{\mathcal{D}} \left(e^{h'^{1,*}(\mathbf{y}_1)} \right)}{e^{-\lambda}} = \frac{\mathbb{E}_{\mathcal{D}} \left(e^{\mathbf{w}'^{1,*T} \mathbf{b}_1} \right)}{e^{-\lambda}} \leq \frac{\mathbb{E}_{\mathcal{D}} \left(e^{-cBB_0} \right)}{e^{-\lambda}} = e^{\lambda - cBB_0}. \quad (82)$$

As a result,

$$\mathbb{E}_{\mathcal{D}} \left(T_1^{(1)} \right) \geq \mathbb{E}_{\{\epsilon_i\}_{i=2}^m} \left(1 - 2e^{-(cBB_0 - \lambda)} \right) = 1 - 2e^{-(cBB_0 - \lambda)}. \quad (83)$$

Repeating this process for all samples, leads to the overall reduction obtained by the first layer

$$\mathbb{E}_S \left(T^{(1)} \right) \geq m \left(1 - 2e^{-(\lambda - cBB_0)} \right). \quad (84)$$

We do not consider the case of $\lambda > cBB_0$, since many of the entries will be zeroed by the soft thresholding operation. For example, when $c = 1$, only the zero estimator can be achieved, which is not relevant for learning. To obtain a meaningful bound which is greater than 0, we assume that $\lambda < cBB_0 + \ln(2)$.

Next, we consider the distribution of $T^{(l)}$ for a general layer $l \in [1, L]$. Now, the entries of the estimator $\mathbf{h}^{l,*}$ range in $[-b^{(l)}, b^{(l)}]$, where

$$b^{(l+1)} = Bb^{(l)} - \lambda \quad (85)$$

and $b^{(1)} = B_0$. Again, we assume that there exists a constant $c \geq 0$ such that $\mathbb{E}_{\mathcal{D}}(e^{-\mathbf{w}^{l,*T} \mathbf{b}_1}) \leq e^{-cBb^{(l)}}$. Following the above derivation and noticing that the bound is non-negative, the expected value of $T^{(l)}$ is bounded by

$$\mathbb{E}_S \left(T^{(l)} \right) \geq \max \left\{ m \left(1 - 2e^{-(cBb^{(l)} - \lambda)} \right), 0 \right\}, \quad (86)$$

completing the proof. \square

To obtain a meaningful bound which is greater than 0, we assume that $\lambda < cBb^{(l)} + \ln(2)$, which translates to $\lambda \leq \frac{cB^{l+1}B_0 + \ln(2)}{1 + lcB}$. The result in (86), behaves as expected with respect to the network's parameters n_x , λ , B , and B_0 . Increasing the value of the soft threshold or decreasing B , will decrease the bound in (86), since crossing the threshold is less probable.

Taking the limit of (76) for $B_l \rightarrow 1, l \in [1, L]$, the bound reduces to

$$\mathcal{G}(\mathcal{H}_I^L) \leq 2 \left(\frac{B_0}{\sqrt{m}} - \frac{\lambda}{m} \sum_{l=1}^L \mathbb{E}_S(T^{(l)}) \right) \quad (87)$$

showing that the GE decreases with the number of layers. This is in contrast to ReLU networks, where the GE increases exponentially with the network's depth, as shown in Section 2.3. This bound also improves over previously available bounds for ISTA networks, where the increase is logarithmic with the number of layers (Behboodi et al., 2020; Schnoor et al., 2021).

To obtain a more compact relation, we can choose the maximal matrices' norm $B = \max_{l \in [1, L]} B_l$, and denote $T = \min_{l \in [1, L]} T^{(l)} \in [0, m]$. The recurrence relation in (73) then reads

$$G_I^{l+1} = BG_I^l - \frac{\lambda \mathbb{E}_S(T)}{m}. \quad (88)$$

This means that

$$G_I^L = \frac{B_0 B^L}{\sqrt{m}} - \frac{\lambda \mathbb{E}_S(T)}{m} \frac{B^L - 1}{B - 1}, \quad (89)$$

which results in a simpler bound

$$\mathcal{G}(\mathcal{H}_I^L) \leq 2G_I^L = 2 \left(\frac{B_0 B^L}{\sqrt{m}} - \frac{\lambda \mathbb{E}_S(T)}{m} \frac{B^L - 1}{B - 1} \right). \quad (90)$$

To verify under what condition a nonincreasing GE is obtained, we subtract between the bounds of the GE of networks with consecutive number of layers

$$\begin{aligned} G_I^l - G_I^{l-1} &= \frac{B_0 B^{l-1}}{\sqrt{m}} (B - 1) - \frac{\lambda \mathbb{E}_S(T)}{m} \frac{B^{l-1}}{B - 1} (B - 1) \\ &= \frac{B^{l-1}}{\sqrt{m}} \left(B_0 (B - 1) - \frac{\lambda \mathbb{E}_S(T)}{m} \right). \end{aligned} \quad (91)$$

As a result, a nonincreasing GE is achievable by restricting the matrices' norm to satisfy

$$B \leq 1 + \frac{\lambda \mathbb{E}_S(T)}{\sqrt{m} B_0}. \quad (92)$$

We observe that the value of T is also dependent on B (as is seen from (50)), and therefore the design rule results in an implicit function. This result indicates how large an intermediate result of the network can be increased, as a function of $\lambda \mathbb{E}_S(T)$, without increasing the GE. Moreover, for all combinations of λ , G and m , there exists a value of B such that $B > 1$.

Theorem 4 (Generalization error bound of ADMM networks):

Consider the class of learned ADMM networks of depth L as described in (4), and m i.i.d. samples. Then there exist $T^{(l)}$ for $l \in [1, L - 1]$ in the interval

$$T^{(l)} \in \left[0, \min \left\{ \frac{m \tilde{B}_l G_A^{l-1}}{\tilde{\lambda}}, m \right\} \right]. \quad (93)$$

where

$$G_A^l = \frac{B_0 \prod_{l'=1}^{l-1} \tilde{B}_{l'}}{\sqrt{m}} - \frac{\tilde{\lambda}}{m} \sum_{l'=1}^{l-2} T^{(l')} \prod_{j=l'+1}^{l-1} \tilde{B}_j - \frac{\tilde{\lambda} T^{(l-1)}}{m}. \quad (94)$$

where $\tilde{\lambda} = (1 + \gamma)\lambda$ and $\tilde{B}_l = (1 + 2\gamma)(B_l + 2)$, $l \in [1, L]$, satisfying

$$\mathcal{G}(\mathcal{H}_A^L) \leq 2\tilde{B}_L \mathbb{E}_S G_A^{L-1}. \quad (95)$$

Proof. The ADMM recurrence relation in (4) can be re-written as

$$\begin{aligned} \mathbf{z}^l &= \mathcal{S}_\lambda (\mathbf{b} + \mathbf{W}^L \mathbf{z}^{l-1} + (\mathbf{W}^L - \mathbf{I}) \mathbf{u}^{l-1}) \\ \mathbf{u}^l &= (\mathbf{I} - \gamma \mathbf{W}^L) \mathbf{u}^{l-1} - \gamma \mathbf{W}^L \mathbf{z}^{l-1} + \gamma \mathbf{z}^l. \end{aligned} \quad (96)$$

Following the same logic from the proof of Theorem 3, we focus on a single neuron in the layers. We denote $R_z^l = \mathbb{E}_{\{\epsilon_i\}_{i=1}^m} \sup_{\mathbf{z}^l} \frac{1}{m} \sum_{i=1}^m \epsilon_i z^{l-1}$ and $R_u^l = \mathbb{E}_{\{\epsilon_i\}_{i=1}^m} \sup_{\mathbf{u}^l} \frac{1}{m} \sum_{i=1}^m \epsilon_i u^l$, the RC of a single entry in the \mathbf{z}^l and \mathbf{u}^l , respectively. Then

$$R_z^l = \mathbb{E}_{\{\epsilon_i\}_{i=1}^m} \sup_{\mathbf{w}^l, \mathbf{z}^{l-1}, \mathbf{u}^{l-1}} \frac{1}{m} \sum_{i=1}^m \epsilon_i \mathcal{S}_\lambda (\mathbf{w}^l \mathbf{z}^{l-1} + (\mathbf{w}^l - \mathbf{I}) \mathbf{u}^{l-1} + \mathbf{b}). \quad (97)$$

From Lemma 1, the soft-thresholding leads to

$$R_z^l \leq \mathbb{E}_{\{\epsilon_i\}_{i=1}^m} \sup_{\mathbf{w}^l, \mathbf{z}^{l-1}, \mathbf{u}^{l-1}} \frac{1}{m} \sum_{i=1}^m \epsilon_i (\mathbf{w}^l \mathbf{z}^{l-1} + (\mathbf{w}^l - \mathbf{I}) \mathbf{u}^{l-1}) - \frac{\lambda T^{(l)}}{m}. \quad (98)$$

By splitting the supremum and applying Lemma 4, the above reads

$$\begin{aligned} R_z^l &\leq B_l \mathbb{E}_{\{\epsilon_i\}_{i=1}^m} \sup_{\mathbf{z}^{l-1}} \frac{1}{m} \sum_{i=1}^m \epsilon_i z^{l-1} + (1 + B_l) \mathbb{E}_{\{\epsilon_i\}_{i=1}^m} \sup_{\mathbf{u}^{l-1}} \frac{1}{m} \sum_{i=1}^m \epsilon_i u^{l-1} - \frac{\lambda T^{(l)}}{m} \\ &\leq B_l R_z^{l-1} + (1 + B_l) R_u^{l-1} - \frac{\lambda T^{(l)}}{m}. \end{aligned} \quad (99)$$

Similarly for R_u^l we get,

$$\begin{aligned} R_u^l &\leq (1 + \gamma B_l) R_u^{l-1} + \gamma B_l R_z^{l-1} + \gamma R_z^l \\ &\leq (1 + \gamma(2B_l + 1)) R_u^{l-1} + 2\gamma B_l R_z^{l-1} - \frac{\gamma \lambda T^{(l)}}{m}. \end{aligned} \quad (100)$$

Applying Lemma 1, and using the fact that the RC of a sum can be bounded by the sum of the individual RCs (Bartlett & Mendelson, 2002), results in

$$\begin{aligned} R_z^l &\leq B_l R_z^{l-1} + (1 + B_l) R_u^{l-1} - \frac{\lambda T^{(l)}}{m} \\ R_u^l &\leq (1 + \gamma B_l) R_u^{l-1} + \gamma B_l R_z^{l-1} + \gamma R_z^l \\ &\leq (1 + \gamma(2B_l + 1)) R_u^{l-1} + 2\gamma B_l R_z^{l-1} - \frac{\gamma \lambda T^{(l)}}{m}. \end{aligned} \quad (101)$$

As a result,

$$\begin{aligned} R_z^l + R_u^l &\leq B_l R_z^{l-1} + (1 + B_l) R_u^{l-1} - \frac{\lambda T^{(l)}}{m} + (1 + \gamma(2B_l + 1)) R_u^{l-1} + 2\gamma B_l R_z^{l-1} - \frac{\gamma \lambda T^{(l)}}{m} \\ &= (1 + 2\gamma) B_l R_z^{l-1} + ((1 + 2\gamma) B_l + 2 + \gamma) R_u^{l-1} - \frac{(1 + \gamma) \lambda T^{(l)}}{m} \\ &\leq (1 + 2\gamma) B_l R_z^{l-1} + (1 + 2\gamma) (B_l + 2) R_u^{l-1} - \frac{(1 + \gamma) \lambda T^{(l)}}{m} \\ &\leq (1 + 2\gamma) (B_l + 2) (R_z^{l-1} + R_u^{l-1}) - \frac{(1 + \gamma) \lambda T^{(l)}}{m}. \end{aligned} \quad (102)$$

We obtain a recurrence relation on the sum of RCs $R_z^l + R_u^l$. Repeating the proof of Theorem 3 for ISTA networks, and replacing λ by $\tilde{\lambda} = (1 + \gamma)\lambda$ and $\tilde{B}_l = (1 + 2\gamma)(B_l + 2)$, $\forall l \in [1, L]$, respectively, leads to

$$R_z^{L-1} + R_u^{L-1} \leq G_A^{L-1} = \frac{B_0 \prod_{l=1}^{L-1} \tilde{B}_l}{\sqrt{m}} - \frac{\tilde{\lambda}}{m} \sum_{l=1}^{L-2} T^{(l)} \prod_{j=l+1}^{L-1} \tilde{B}_j - \frac{\tilde{\lambda} T^{(L-1)}}{m}. \quad (103)$$

The value of $T^{(l)}$ are in the interval

$$T^{(l)} \in \left[0, \min \left\{ \frac{m \tilde{B}_l G_A^{l-1}}{\tilde{\lambda}}, m \right\} \right]. \quad (104)$$

Moreover, from the ADMM recurrence relation in (4), the following holds

$$R_A^l \leq \tilde{B}_l (R_z^{l-1} + R_u^{l-1}) \quad (105)$$

where $R_A^l := \mathfrak{R}_m(\mathcal{H}_A^l)$ is the RC of \mathcal{H}_A^l , leading to

$$R_A^L \leq \tilde{B}_L \left(\frac{B_0 \prod_{l=1}^{L-1} \tilde{B}_l}{\sqrt{m}} - \frac{\tilde{\lambda}}{m} \sum_{l=1}^{L-2} T^{(l)} \prod_{j=l+1}^{L-1} \tilde{B}_j - \frac{\tilde{\lambda} T^{(L-1)}}{m} \right). \quad (106)$$

Similarly to the proof of Theorem 3, the bound on the GE is obtained by applying a 1-Lipschitz loss on the network's prediction, which concludes the proof. \square

Similarly to Theorem 3, Theorem 4 results in design rules for ADMM networks with low GE. We observe that the RC bound of the ADMM network is obtained from the bound of the ISTA network, by replacing λ and B_l with $\tilde{\lambda}$ and \tilde{B}_l , $l \in [1, L]$. The relation between B_l and \tilde{B}_l sheds light on the relation between the GE of learned ISTA and ADMM. Since $\tilde{B}_l > B_l$, the GE bound on ISTA networks is potentially lower compared to ADMM networks with the same weight's norm B_l , indicating that ISTA networks have better generalization abilities compared to ADMM networks. Depending on the behaviour of T , as the number of training samples m increases, the difference between the bounds on the model-based and ReLU networks (presented in Theorem 1) might be less significant. In this case, the soft-thresholding nonlinearity is most valuable in the case of small number of training samples.

Similarly to the GE bound of ISTA networks, we define $\tilde{B} = \max_{l \in [1, L]} \tilde{B}_l$ and extract a more compact bound on the GE of ADMM networks

$$\mathcal{G}(\mathcal{H}_A^L) \leq 2\tilde{B} \left(\frac{B_0 \tilde{B}^{L-1}}{\sqrt{m}} - \frac{\tilde{\lambda} \mathbb{E}_S(T) \tilde{B}^{L-1} - 1}{m \tilde{B} - 1} \right). \quad (107)$$

A.3 ESTIMATION ERROR BOUNDS

Here we prove the EE bounds for ISTA, ADMM, and ReLU networks provided in Theorem 5, derived with the LRC machinery, presented in the following theorem.

Definition 1. A function $\psi : [0, \infty) \rightarrow [0, \infty)$ is sub-root if it is nonnegative, nondecreasing, and if $r \mapsto \psi(r)/\sqrt{r}$ is nonincreasing for $r > 0$.

Theorem 6. (Local Rademacher complexity bound on the estimation error for vector estimators (Yousefi et al., 2018)) Let \mathcal{H} be a class of functions where each coordinate ranges in $[-1, 1]$ and let \mathcal{L} be a loss function satisfying:

- There exists an estimator $\mathbf{h}^* \in \mathcal{H}$ satisfying $\mathbb{E}_{\mathcal{D}} \mathcal{L}(\mathbf{h}^*) = \inf_{\mathbf{h} \in \mathcal{H}} \mathbb{E}_{\mathcal{D}} \mathcal{L}(\mathbf{h})$.
- The loss \mathcal{L} is an averaged of 1-Lipschitz per-coordinate losses ℓ

$$\mathcal{L}(\mathbf{h}) = \frac{1}{n_x} \sum_{j=1}^{n_x} \ell(\mathbf{h}_j). \quad (108)$$

- There exists a constant $C \geq 1$ such that for every probability distribution \mathcal{D} , and estimator $\mathbf{h} \in \mathcal{H}$, such that

$$\mathbb{E}_{\mathcal{D}} \sum_{j=1}^{n_x} (\mathbf{h}_j - \mathbf{h}_j^*)^2 \leq C \mathbb{E}_{\mathcal{D}} \sum_{j=1}^{n_x} (\ell(\mathbf{h}_j) - \ell(\mathbf{h}_j^*)). \quad (109)$$

Let ψ be a sub-root function with fixed point r^* such that $\psi(r) \geq \mathfrak{R}_m^{\text{vec}}(\mathcal{H}_r), \forall r \geq r^*$ where

$$\mathfrak{R}_m^{\text{vec}}(\mathcal{H}_r) := \mathbb{E}_{\{\epsilon_{i,j}\}_{i=1,j=1}^{m,n_x}} \sup_{\mathbf{h} \in \mathcal{H}_r} \frac{1}{m} \sum_{i=1}^m \frac{1}{n_x} \sum_{j=1}^{n_x} \epsilon_{i,j} \mathbf{h}_j(\mathbf{y}_i) \quad (110)$$

where \mathcal{H}_r is defined in (4).

Then for any $s > 0$, any $K > 1$, and any $r \geq \psi(r)$ with probability at least $1 - e^{-s}$

$$\mathcal{E}(\mathcal{H}) \leq 40Kr^* + \frac{16C^2K + 48C}{mn_x} s. \quad (111)$$

In the case of proper learning (such that the target vector belongs to the class of estimators \mathcal{H} , implying that $\inf_{h \in \mathcal{H}} L_{\mathcal{D}}(h) = 0$), the requirements on the loss are satisfied for losses given by the L_p norm ($p \geq 1$), as discussed in (Mendelson, 2002). The value of C in Theorem 6, depends on the L_p norm being used as a loss.

We derive the following theorem on the EE of ISTA, ADMM, and ReLU networks.

Theorem 5 (Estimation error bound of ISTA, ADMM, and ReLU networks): Consider the class of functions represented by depth- L ISTA networks \mathcal{H}_I^L as detailed in Section 2.1, m training samples, and a loss satisfying Assumption 1 with a constant C . If $\|\mathbf{W}^l - \mathbf{W}^{l,*}\|_{\infty} \leq \alpha\sqrt{r}$ for some $\alpha > 0$. Moreover, $B \geq \max\{\alpha\sqrt{r}, 1\}$. Then there exists T in the interval

$$T \in \left[0, \min \left\{ \frac{\sqrt{m}B_0B^{L-1}2^L}{\lambda\eta}, m \right\} \right] \quad (112)$$

where $\eta = \frac{LB^{L-1}(B-1)-B^L+1}{(B-1)^2}$, such that for any $s > 0$ with probability at least $1 - e^{-s}$

$$\mathcal{E}(\mathcal{H}_I^L) \leq 41r^* + \frac{17C^2 + 48C}{mn_x} s \quad (113)$$

where

$$r^* = C^2\alpha^2 \left(\frac{B_0B^{L-1}2^L}{\sqrt{m}} - \frac{\lambda T}{m}\eta \right)^2. \quad (114)$$

For the class of functions represented by depth- L ADMM networks \mathcal{H}_A^L , the bound is satisfied with

$$r^* = C^2\alpha^2 \left(\frac{B_0\tilde{B}^{L-2}2^{L-1}}{\sqrt{m}} - \frac{\tilde{\lambda}T}{m}\tilde{\eta} \right)^2 \quad (115)$$

where $\tilde{\lambda} = (1 + \gamma)\lambda$, $\tilde{B} = (1 + 2\gamma)(B + 2)$, and $\tilde{\eta} = \frac{(L-1)\tilde{B}^{L-2}(\tilde{B}-1)-\tilde{B}^{L-1}+1}{(\tilde{B}-1)^2}$. The bound is also satisfied for the class of functions represented by depth- L ReLU networks \mathcal{H}_R^L , with

$$r^* = C^2\alpha^2 \left(\frac{B_0B^{L-1}2^L}{\sqrt{m}} \right)^2. \quad (116)$$

Proof. To derive the upper bound on the EE, we will use the LRC framework developed in (Bartlett et al., 2005), and rely on Theorem 6 derived in (Yousefi et al., 2018), assuming that the loss function satisfies Assumption 1.

We define $\mathcal{H}_{I,r}^{l,j}$ to be the class of scalar functions that represent the j th neuron (coordinate) at the l th output layer

$$\mathcal{H}_{I,r}^{l,j} = \left\{ \mathbf{h}_j : \mathbf{h} \in \mathcal{H}_{I,r}^l \right\}. \quad (117)$$

In order to apply the above theorem, one needs to bound the RC of the class of functions $\mathcal{H}_{I,r}^{l,j}$, defined in (117). The weight matrices of \mathbf{h}^* and $\mathbf{h} \in \mathcal{H}_{I,r}^l$ are denoted by $\{\mathbf{W}^{l,*}\}_{l=1}^L$ and $\{\mathbf{W}^l\}_{l=1}^L$,

respectively. The weights differences is defined as $\Delta \mathbf{W}^l = \mathbf{W}^l - \mathbf{W}^{l,*}$, satisfying $\|\Delta \mathbf{W}^l\|_\infty \leq \alpha \sqrt{r} := \Delta B$. We can now bound the RC of $\mathcal{H}_{I,r}^{l,j}$:

$$\begin{aligned} \mathfrak{R}_m \left(\mathcal{H}_{I,r}^{l,j} \right) &= \mathbb{E}_{\{\epsilon_i\}_{i=1}^m} \sup_{h^l} \frac{1}{m} \sum_{i=1}^m \epsilon_i h^l(\mathbf{y}_i) \\ &= \mathbb{E}_{\{\epsilon_i\}_{i=1}^m} \sup_{\mathbf{W}^l, h^{l-1}} \frac{1}{m} \sum_{i=1}^m \epsilon_i \mathcal{S}_\lambda \left(\mathbf{W}^l h^{l-1}(\mathbf{y}_i) + \mathbf{b} \right). \end{aligned} \quad (118)$$

Applying Lemma 1 results in

$$\begin{aligned} \mathfrak{R}_m \left(\mathcal{H}_{I,r}^{l,j} \right) &\leq \mathbb{E}_{\{\epsilon_i\}_{i=1}^m} \sup_{\mathbf{W}^l, h^{l-1}} \frac{1}{m} \sum_{i=1}^m \epsilon_i \mathbf{W}^l h^{l-1}(\mathbf{y}_i) - \frac{\lambda T^{(l)}}{m} \\ &= \mathbb{E}_{\{\epsilon_i\}_{i=1}^m} \sup_{\Delta \mathbf{W}^l, h^{l-1}} \frac{1}{m} \sum_{i=1}^m \epsilon_i \left(\mathbf{W}^{l,*} + \Delta \mathbf{W}^l \right) h^{l-1}(\mathbf{y}_i) - \frac{\lambda T^{(l)}}{m} \\ &\leq (B + \Delta B) \mathfrak{R}_m \left(\mathcal{H}_r^{l-1} \right) - \frac{\lambda T^{(l)}}{m} \end{aligned} \quad (119)$$

where the last inequality holds, by applying Lemma 4 and taking into account that $\|\mathbf{W}^{l,*} + \Delta \mathbf{W}^l\|_\infty \leq \|\mathbf{W}^{l,*}\|_\infty + \|\Delta \mathbf{W}^l\|_\infty \leq B + \Delta B$. Accumulating the contributions from all L layers, we obtain

$$\mathfrak{R}_m \left(\mathcal{H}_{I,r}^{L,j} \right) \leq \frac{B_0 \left((B + \Delta B)^L - B^L \right)}{\sqrt{m}} - \frac{\lambda T}{m} \frac{(B + \Delta B)^L - 1}{(B + \Delta B) - 1}, \quad (120)$$

where $T = \min_{l \in [1, L]} T^{(l)}$. In addition, the subtraction of the factor B^L results from the following term $\mathbf{W}^{L,*} \mathbf{W}^{L-1,*} \mathbf{W}^{L-2,*} \dots \mathbf{W}^{1,*} \mathbf{x}_i$, which does not contribute to the complexity, since these are fixed matrices.

Next, we bound the expression in (120) with a term linearly dependent on ΔB . Applying Newton's binomial formula, reads

$$\begin{aligned} \mathfrak{R}_m \left(\mathcal{H}_{I,r}^{L,j} \right) &\leq \frac{B_0 \left(\sum_{k=0}^L \binom{L}{k} B^k \Delta B^{L-k} - B^L \right)}{\sqrt{m}} - \frac{\lambda T}{m} \frac{(B + \Delta B)^L - 1}{(B + \Delta B) - 1} \\ &= \frac{B_0 \Delta B \left(\sum_{k=0}^{L-1} \binom{L}{k} B^k \Delta B^{L-k-1} \right)}{\sqrt{m}} - \frac{\lambda T}{m} \frac{(B + \Delta B)^L - 1}{(B + \Delta B) - 1}. \end{aligned} \quad (121)$$

Assuming that $\Delta B < B$ the above is bounded by

$$\mathfrak{R}_m \left(\mathcal{H}_{I,r}^{L,j} \right) \leq \frac{B_0 \Delta B \left(\sum_{k=0}^{L-1} \binom{L}{k} B^{L-1-k} \right)}{\sqrt{m}} - \frac{\lambda T}{m} \frac{(B + \Delta B)^L - 1}{(B + \Delta B) - 1}. \quad (122)$$

Using the known relation $\sum_{k=0}^L \binom{L}{k} = 2^L$, leads to

$$\begin{aligned} \mathfrak{R}_m \left(\mathcal{H}_{I,r}^{L,j} \right) &\leq \frac{B_0 \Delta B B^{L-1} 2^L}{\sqrt{m}} - \frac{\lambda T}{m} \sum_{l=0}^{L-1} (B + \Delta B)^l \\ &\leq \frac{B_0 \Delta B B^{L-1} 2^L}{\sqrt{m}} - \frac{\lambda T}{m} \sum_{l=1}^{L-1} (B + \Delta B)^l \\ &\leq \frac{B_0 \Delta B B^{L-1} 2^L}{\sqrt{m}} - \frac{\lambda T}{m} \sum_{l=1}^{L-1} \sum_{k=0}^l \binom{l}{k} \Delta B^k B^{l-k} \end{aligned} \quad (123)$$

where we used Newton's binomial formula. Neglecting all terms with $k \neq 1$, reduces the above to

$$\begin{aligned} \mathfrak{R}_m \left(\mathcal{H}_{I,r}^{L,j} \right) &\leq \Delta B \left(\frac{B_0 B^{L-1} 2^L}{\sqrt{m}} - \frac{\lambda T}{m} \sum_{l=1}^{L-1} \binom{l}{1} B^{l-1} \right) \\ &= \Delta B \left(\frac{B_0 B^{L-1} 2^L}{\sqrt{m}} - \frac{\lambda T L B^{L-1} (B - 1) - B^L + 1}{m (B - 1)^2} \right). \end{aligned} \quad (124)$$

Substituting $\Delta B = \alpha\sqrt{r}$, we have

$$\mathfrak{R}_m(\mathcal{H}_{I,r}^{L,j}) \leq \alpha\sqrt{r} \left(\frac{B_0 B^{L-1} 2^L}{\sqrt{m}} - \frac{\lambda T}{m} \frac{LB^{L-1}(B-1) - B^L + 1}{(B-1)^2} \right). \quad (125)$$

Since the RC is nonnegative, the value of T is restricted to the interval

$$T \in \left[0, \min \left\{ \frac{\sqrt{m} B_0 B^{L-1} 2^L}{\lambda \eta}, m \right\} \right] \quad (126)$$

where $\eta = \frac{LB^{L-1}(B-1) - B^L + 1}{(B-1)^2}$.

To simplify our notation, let us denote

$$\beta := \left(\frac{B_0 B^{L-1} 2^L}{\sqrt{m}} - \frac{\lambda T}{m} \eta \right). \quad (127)$$

Then, for $j \in [1, n_x]$

$$\mathfrak{R}_m(\mathcal{H}_{I,r}^{L,j}) \leq \alpha\sqrt{r}\beta. \quad (128)$$

To apply Theorem 6, we observe that from Jensen's inequality $\mathfrak{R}_m^{\text{vec}}(\mathcal{H}_r)$ satisfies

$$\begin{aligned} \mathfrak{R}_m^{\text{vec}}(\mathcal{H}_r) &= \mathbb{E}_{\{\epsilon_{i,j}\}_{i=1,j=1}^{m,n_x}} \sup_{\mathbf{h} \in \mathcal{H}_{I,r}^L} \frac{1}{m} \sum_{i=1}^m \frac{1}{n_x} \sum_{j=1}^{n_x} \epsilon_{i,j} \mathbf{h}_j(\mathbf{y}_i) \\ &\leq \frac{1}{n_x} \sum_{j=1}^{n_x} \mathbb{E}_{\{\epsilon_{i,j}\}_{i=1,j=1}^{m,n_x}} \sup_{\mathbf{h} \in \mathcal{H}_{I,r}^L} \frac{1}{m} \sum_{i=1}^m \epsilon_{i,j} \mathbf{h}_j(\mathbf{y}_i) \\ &= \frac{1}{n_x} \sum_{j=1}^{n_x} \mathfrak{R}_m(\mathcal{H}_{I,r}^{L,j}) \\ &\leq \frac{1}{n_x} \sum_{j=1}^{n_x} \alpha\sqrt{r}\beta \\ &= \alpha\sqrt{r}\beta. \end{aligned} \quad (129)$$

This leads to the following sub-root function

$$\psi(r) := C\alpha\sqrt{r}\beta \geq C\mathfrak{R}_m^{\text{vec}}(\mathcal{H}_r). \quad (130)$$

Since $C, \beta > 0$ we get that ψ is a nonnegative and nondecreasing function over $r \geq 0$. In addition $\psi(r)/\sqrt{r} = C\alpha\beta$ which is a constant function and in particular is nonincreasing, implying that the conditions for Theorem 6 are satisfied.

The fixed-point of the function, such that $r^* = \psi(r^*)$, can be computed explicitly from

$$r^* = C\alpha\sqrt{r^*} \left(\frac{B_0 B^{L-1} 2^L}{\sqrt{m}} - \frac{\lambda T}{m} \eta \right), \quad (131)$$

which reads

$$r^* = C^2 \alpha^2 \left(\frac{B_0 B^{L-1} 2^L}{\sqrt{m}} - \frac{\lambda T}{m} \eta \right)^2. \quad (132)$$

By applying Theorem 6, we conclude that for any $r \geq r^*$, $K > 1$, and $s > 0$, with probability at least $1 - e^{-s}$

$$\mathcal{E}(\mathcal{H}_I^L) \leq 40Kr^* + \frac{16KC^2 + 48C}{mn_x} s. \quad (133)$$

Moreover, we choose $K = 41/40 > 1$, leading to

$$\begin{aligned} \mathcal{E}(\mathcal{H}_I^L) &\leq 41r^* + \frac{\frac{16 \cdot 41}{40} C^2 + 48C}{mn_x} s \\ &\leq 41r^* + \frac{17C^2 + 48C}{mn_x} s. \end{aligned} \quad (134)$$

Following the observation on the RC of ADMM networks, the RC bounds on ISTA networks are satisfied for ADMM networks by replacing λ, B by $\tilde{\lambda}, \tilde{B}$, respectively, as defined in Theorem 4.

Applying the same proof methodology as Theorem 5, results in the following sub-root function

$$\psi_A(r) := C\alpha\sqrt{r} \left(\frac{B_0\tilde{B}^{L-2}2^{L-1}}{\sqrt{m}} - \frac{\tilde{\lambda}T(L-1)\tilde{B}^{L-2}(\tilde{B}-1) - \tilde{B}^{L-1} + 1}{m(\tilde{B}-1)^2} \right), \quad (135)$$

leading to the fixed-point

$$r^* = C^2\alpha^2 \left(\frac{B_0\tilde{B}^{L-2}2^{L-1}}{\sqrt{m}} - \frac{\tilde{\lambda}T(L-1)\tilde{B}^{L-2}(\tilde{B}-1) - \tilde{B}^{L-1} + 1}{m(\tilde{B}-1)^2} \right)^2. \quad (136)$$

Applying Theorem 6 proves the bound.

Finally, the same proof methodology applies for ReLU networks, with the sub-root function given by

$$\psi_R(r) = C\alpha\sqrt{r} \frac{B_0B^{L-1}2^L}{\sqrt{m}} \quad (137)$$

leading to the fixed-point

$$r^* = C^2\alpha^2 \left(\frac{B_0B^{L-1}2^L}{\sqrt{m}} \right)^2. \quad (138)$$

Applying Theorem 6 concludes the proof. \square

In contrast to the GE bounds, the role of the weight's norm B is less dominant in the EE bounds presented in Theorem 5. Therefore, it is not clear from the EE bounds which model-based networks exhibit better generalization.

The parameter B_0 relates the bound to the sparsity level of the target vectors

$$B_0 = \|\mathbf{b}\|_1 = \|\mathbf{A}^T \mathbf{y}\|_1 = \|\mathbf{A}^T \mathbf{A} \mathbf{x} + \mathbf{A}^T \mathbf{e}\|_1 \leq \|\mathbf{A}^T \mathbf{A} \mathbf{x}\|_1 + \|\mathbf{A}^T \mathbf{e}\|_1. \quad (139)$$

The target vector \mathbf{x} is sparse with sparsity rate of ρ , meaning that $\|\mathbf{x}\|_1 \leq \rho n_x$, so that

$$B_0 \leq \|\mathbf{A}^T \mathbf{A}\|_1 \rho n_x + \|\mathbf{A}^T \mathbf{e}\|_1. \quad (140)$$

Therefore, as the target vector is more sparse, ρ decreases, which decreases r^* , and implies a lower bound on the EE.

B ADDITIONAL SIMULATION RESULTS

In this section, we provide additional results for the EE of ISTA and ReLU networks with varying number of layers. We show that the EE and performance of the networks behave similarly to the depth-10 networks presented in Section 5.

In our work, we considered only one set of learned weight matrices. However, a second set of weight matrices can also be learned, as understood from the following relation between consecutive layers

$$\mathbf{h}_I^l = \mathcal{S}_\lambda(\mathbf{W}_1^l \mathbf{h}_I^{l-1} + \mathbf{W}_2^l \mathbf{y}), \quad \mathbf{h}_I^0 = \mathcal{S}_\lambda(\mathbf{y}).$$

The provided comparisons show that the soft thresholding nonlinearity affects the EE of ISTA networks more significantly, compared to learned bias terms, obtained with the set of additional learned matrices. This observation emphasizes that the performed analysis on the ISTA network with constant biases, could be applicable to additional variations of the ISTA networks.




Figure 4: Comparing the EE of ISTA and ReLU networks with 2 layers.

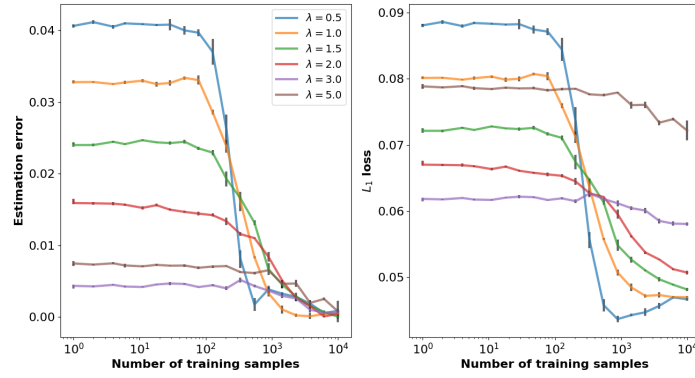
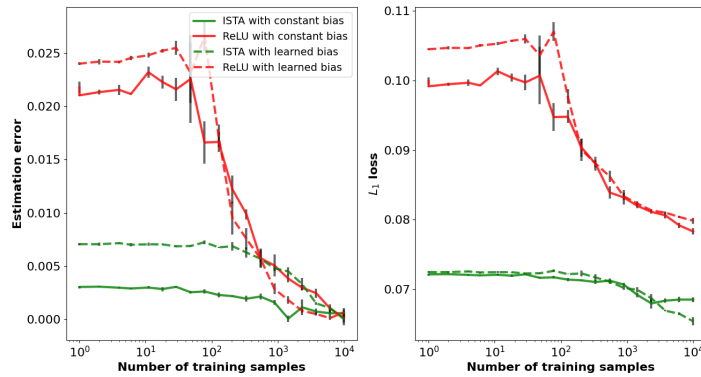
Figure 5: Estimation error and loss of ISTA networks with 2 layers, as a function of the soft-threshold's value λ .

Figure 6: Comparing the EE of ISTA and ReLU networks with 4 layers.




Figure 7: Estimation error and loss of ISTA networks with 4 layers, as a function of the soft-threshold’s value λ .




Figure 8: Comparing the EE of ISTA and ReLU networks with 6 layers.

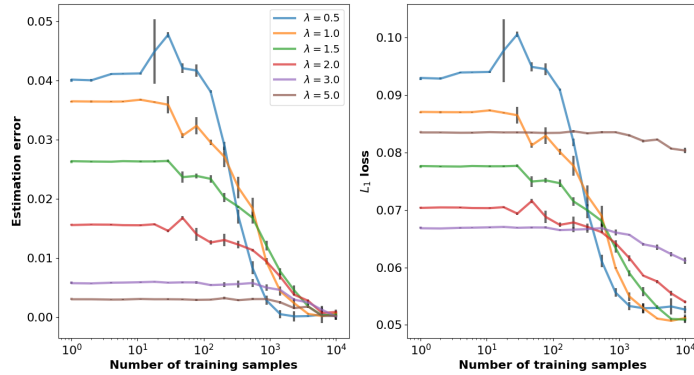


Figure 9: Estimation error and loss of ISTA networks with 6 layers, as a function of the soft-threshold’s value λ .

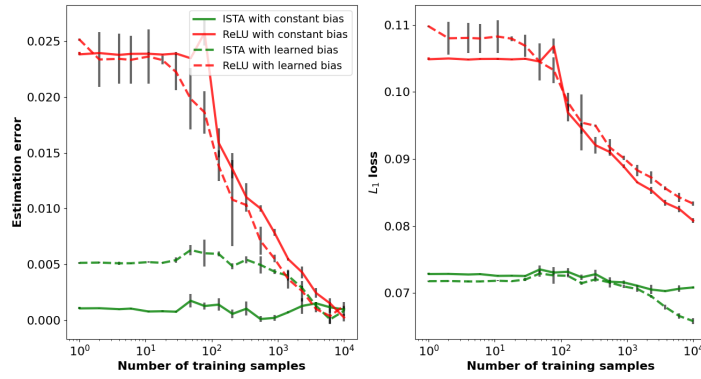
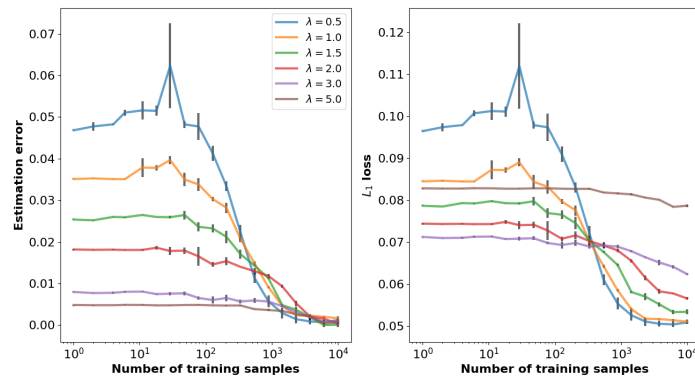


Figure 10: Comparing the EE of ISTA and ReLU networks with 8 layers.

Figure 11: Estimation error and loss of ISTA networks with 8 layers, as a function of the soft-threshold's value λ .