A SIMPLE SPARSE DENOISING LAYER FOR ROBUST DEEP LEARNING

Anonymous authors
Paper under double-blind review

ABSTRACT

Deep models have achieved great success in many applications. However, vanilla deep models are not well-designed against the input perturbation. In this work, we take an initial step to design a simple robust layer as a lightweight plug-in for vanilla deep models. To achieve this goal, we first propose a fast sparse coding and dictionary learning algorithm for sparse coding problem with an exact k-sparse constraint or l_0 norm regularization. Our method comes with a closed-form solution for the sparse coding phase by taking advantage of a novel structured dictionary. With this handy solution, we propose a simple sparse denoising layer (SDL) as a lightweight robust plug-in. Extensive experiments on both classification and reinforcement learning tasks manifest the effectiveness of our methods.

1 INTRODUCTION

Deep neural networks have obtained a great success in many applications, including computer vision, reinforcement learning (RL) and natural language processing, etc. However, vanilla deep models are not robust to noise perturbations of the input. Even a small perturbation of input data would dramatically harm the prediction performance (Goodfellow et al., 2015).

To address this issue, there are three mainstreams of strategies: data argumentation based learning methods (Zheng et al., 2016; Ratner et al., 2017; Madry et al., 2018; Cubuk et al., 2020), loss functions/regularization techniques (Elsayed et al., 2018; Zhang et al., 2019), and importance weighting of network architecture against noisy input perturbation. Su et al. (2018) empirically investigated 18 deep classification models. Their studies found that model architecture is a more critical factor to robustness than the model size. Most recently, Guo et al. (2020) employed a neural architecture search (NAS) method to investigate the robust architectures. However, the NAS-based methods are still very computationally expensive. Furthermore, their resultant model cannot be easily adopted as a plug-in for other vanilla deep models. A handy robust plug-in for backbone models remains highly demanding.

In this work, we take an initial step to design a simple robust layer as a lightweight plug-in for the vanilla deep models. To achieve this goal, we first propose a novel fast sparse coding and dictionary learning algorithm. Our algorithm has a closed-form solution for the sparse coding phase, which is cheap to compute compared with iterative methods in the literature. The closed-form update is handy for the situation that needs fast computation, especially in the deep learning. Based on this, we design a very simple sparse denoising layer for deep models. Our SDL is very flexible, and it enables an end-to-end training. Our SDL can be used as a lightweight plug-in for many modern architecture of deep models (e.g., ResNet and DenseDet for classification and deep PPO models for RL). Our contributions are summarized as follows:

- We propose simple sparse coding and dictionary learning algorithms for both k-sparse constrained sparse coding problem and l_0-norm regularized problem. Our algorithms belong to the alternative descent optimization, for each alternative phase, our algorithms have a closed-form update, which is very simple and fast to compute.
- We introduce a simple sparse denoising layer (SDL) based on our handy update. Our SDL involves simple operations only, which is a fast plug-in layer for end-to-end training.
- Extensive experiments on both classification tasks and reinforcement learning tasks show the effeteness of our SDL.

2 RELATED WORKS

Sparse Coding and Dictionary Learning: Sparse coding and dictionary learning are widely studied in computer vision and image processing. One related popular method is K-SVD (Elad & Aharon, 2006; Rubinstein et al., 2008), it jointly learns an over-complete dictionary and the sparse representations by minimizing a l_0-norm regularized reconstruction problem. Specifically, K-SVD alternatively iterates between the sparse coding phase and dictionary updating phase. The both two steps are based on heuristic greedy methods. Despite its good performance, K-SVD is very computationally demanding. Moreover, as pointed out by Bao et al. (2013), both the sparse coding phase and dictonary updating of K-SVD use some greedy approaches that lack rigorous theoretical guarantee on its optimality and convergence. Bao et al. (2013) proposed to learn an orthogonal dictionary instead of the over-complete one. The idea is to concatenate the free parameters with predefined filters to form an orthogonal dictionary. This trick reduces the time complexity compared with K-SVD. However, their algorithm relies on the predefined filters. Furthermore, the alternative descent method heavily relies on SVD, which is not easy to extend to deep models.

In contrast, our method learns a structured over-complete dictionary, which has a simple form as a layer for deep learning. Recently, some works (Venkatakrishnan et al., 2013) employed deep neural networks to approximate alternating direction method of multipliers (ADMM) or other proximal algorithms for image denoising tasks. In (Wei et al., 2020), reinforcement learning is used to learn the hyperparameters of these deep iterative models. However, this kind of method itself needs to train a complex deep model. Thus, they are computationally expensive, which is too heavy or inflexible as a plug-in layer for backbone models in other tasks instead of image denoising tasks, e.g., reinforcement learning and multi-class classification, etc. An illustration of number of parameters of SDL, DnCNN (Zhang et al., 2017) and PnP (Wei et al., 2020) are shown in Table 1. SDL has much less parameters compared with DnCNN and PnP, and it can serve as a lightweight plug-in for other tasks, e.g., RL.

Table 1: Number of parameters of different models

	SDL	DnCNN (Zhang et al., 2017)	PnP (Wei et al., 2020)
#parameters	**630**	667,008	296,640

Robust Deep Learning: In the literature of robust deep learning, several robust losses have been studied. To achieve better generalization ability, Elsayed et al. (2018) proposed a loss function to impose a large margin of any chosen layers of a deep network. Barron (2019) proposed a general loss with a shape parameter to cover several robust losses as special cases. For the problems with noisy input perturbation, several data argumentation-based algorithms and regularization techniques are proposed (Zheng et al., 2016; Ratner et al., 2017; Cubuk et al., 2020; Elsayed et al., 2018; Zhang et al., 2019). However, the network architecture remains less explored to address the robustness of the input perturbation. Guo et al. (2020) employed NAS methods to search the robust architectures. However, the searching-based method is very computationally expensive. The resultant architectures cannot be easily used as a plug-in for other popular networks. In contrast, our SDL is based on a closed-form of sparse coding, which can be used as a handy plug-in for many backbone models.

3 FAST SPARSE CODING AND DICTIONARY LEARNING

In this section, we present our fast sparse coding and dictionary learning algorithm for the k-sparse problem and the l_0-norm regularized problem in Section 3.1 and Section 3.2, respectively. Both algorithms belong to the alternative descent optimization framework.

3.1 K-SPARSE CODING

We first introduce the optimization problem for sparse coding with a k-sparse constraint. Mathematically, we aim at optimizing the following objective

$$\begin{aligned} &\min_{\boldsymbol{Y},\boldsymbol{D}} \|\boldsymbol{X} - \boldsymbol{D}\boldsymbol{Y}\|_F^2 \\ &\text{subject to } \|\boldsymbol{Y}_i\|_0 \leq k, \forall i \in \{1, \cdots, N\} \\ &\qquad\quad \mu(\boldsymbol{D}) \leq \lambda \\ &\qquad\quad \|\boldsymbol{D}_j\|_2 = 1, \forall j \in \{1, \cdots, M\}, \end{aligned} \tag{1}$$

where $\boldsymbol{D} \in \mathbb{R}^{d\times M}$ is the dictionary, and $\boldsymbol{D}_i$ denotes the ith column of matrix $\boldsymbol{D}$. $\boldsymbol{Y}_i$ denotes the ith column of the matrix $\boldsymbol{Y} \in \mathbb{R}^{M\times N}$, and $\mu(\cdot)$ denotes the mutual coherence that is defined as

$$\mu(\boldsymbol{D}) = \max_{i\neq j} \frac{|\boldsymbol{D}_i^\top \boldsymbol{D}_j|}{\|\boldsymbol{D}_i\|_2 \|\boldsymbol{D}_j\|_2}. \tag{2}$$

The optimization problem (1) is non-convex and non-smooth, which is very difficult to optimize. To alleviate this problem, we employ a structured dictionary as

$$\boldsymbol{D} = \boldsymbol{R}^\top \boldsymbol{B}. \tag{3}$$

We require that $\boldsymbol{R}^\top \boldsymbol{R} = \boldsymbol{R}\boldsymbol{R}^\top = \boldsymbol{I}_d$ and $\boldsymbol{B}\boldsymbol{B}^\top = \boldsymbol{I}_d$, and each column vector of matrix $\boldsymbol{B}$ has a constant l_2-norm, i.e., $\|\boldsymbol{B}_i\|_2 = c$. The benefit of the structured dictionary is that it enables a fast update algorithm with a closed-form solution for the sparse coding phase.

3.1.1 Construction of Structured Matrix $\boldsymbol{B}$

Now, we show how to design a structured matrix $\boldsymbol{B}$ that satisfies the requirements. First, we construct $\boldsymbol{B}$ by concatenating the real and imaginary parts of rows of a discrete Fourier matrix. The proof of the following theorems regarding the properties of $\boldsymbol{B}$ can be found in Appendix.

Without loss of generality, we assume that $d = 2m, M = 2n$. Let $\boldsymbol{F} \in \mathbb{C}^{n\times n}$ be a $n \times n$ discrete Fourier matrix. $\boldsymbol{F}_{k,j} = e^{\frac{2\pi \boldsymbol{i} kj}{n}}$ is the $(k, j)^{th}$entry of $\boldsymbol{F}$, where $\boldsymbol{i} = \sqrt{-1}$. Let $\Lambda = \{k_1, k_2, ..., k_m\} \subset \{1, ..., n-1\}$ be a subset of indexes.

The structured matrix $\boldsymbol{B}$ can be constructed as Eq.(4).

$$\boldsymbol{B} = \frac{1}{\sqrt{n}} \begin{bmatrix} \mathrm{Re}\boldsymbol{F}_\Lambda & -\mathrm{Im}\boldsymbol{F}_\Lambda \\ \mathrm{Im}\boldsymbol{F}_\Lambda & \mathrm{Re}\boldsymbol{F}_\Lambda \end{bmatrix} \in \mathbb{R}^{d\times N} \tag{4}$$

where Re and Im denote the real and imaginary parts of a complex number, and $\boldsymbol{F}_\Lambda$ in Eq. (5) is the matrix constructed by m rows of $\boldsymbol{F}$

$$\boldsymbol{F}_\Lambda = \begin{bmatrix} e^{\frac{2\pi \boldsymbol{i} k_1 1}{n}} & \cdots & e^{\frac{2\pi \boldsymbol{i} k_1 n}{n}} \\ \vdots & \ddots & \vdots \\ e^{\frac{2\pi \boldsymbol{i} k_m 1}{n}} & \cdots & e^{\frac{2\pi \boldsymbol{i} k_m n}{n}} \end{bmatrix} \in \mathbb{C}^{m\times n}. \tag{5}$$

Theorem 1. *Suppose $d = 2m, M = 2n$. Construct matrix $\boldsymbol{B}$ as in Eq.(4). Then $\boldsymbol{B}\boldsymbol{B}^\top = \boldsymbol{I}_d$ and $\|\boldsymbol{B}_j\|_2 = \sqrt{\frac{m}{n}}, \forall j \in \{1, \cdots, M\}$.*

Theorem 1 shows that the structured construction $\boldsymbol{B}$ satisfies the orthogonal constraint and constant norm constraint. One thing remaining is how to construct $\boldsymbol{B}$ to achieve a small mutual coherence.

To achieve this goal, we can leverage the coordinate descent method in (Lyu, 2017) to construct the index set Λ. For a prime number n such that m divides $n-1$, i.e., $m|(n-1)$, we can employ a closed-form construction. Let g denote a primitive root modulo n. We construct the index $\Lambda = \{k_1, k_2, ..., k_m\}$ as

$$\Lambda = \{g^0, g^{\frac{n-1}{m}}, g^{\frac{2(n-1)}{m}}, \cdots, g^{\frac{(m-1)(n-1)}{m}}\} \bmod n. \tag{6}$$

The resulted structured matrix $\boldsymbol{B}$ has a bounded mutual coherence, which is shown in Theorem 2.

Theorem 2. *Suppose $d = 2m, M = 2n$, and n is a prime such that $m|(n-1)$. Construct matrix $\boldsymbol{B}$ as in Eq.(4) with index set Λ as Eq.(6). Let mutual coherence $\mu(\boldsymbol{B}) := \max_{i\neq j} \frac{|\boldsymbol{B}_i^\top \boldsymbol{B}_j|}{\|\boldsymbol{B}_i\|_2 \|\boldsymbol{B}_j\|_2}$. Then $\mu(\boldsymbol{B}) \leq \frac{\sqrt{n}}{m}$.*

Remark: The bound of mutual coherence in Theorem 2 is non-trivial when $n < m^2$. For the case $n \geq m^2$, we can use the coordinate descent method in (Lyu, 2017) to minimize the mutual coherence.

Now, we show that the structured dictionary $\boldsymbol{D} = \boldsymbol{R}^\top \boldsymbol{B}$ satisfies the constant norm constraint and has a bounded mutual coherence. The results are summarized in Theorem 3.

Theorem 3. *Let $\boldsymbol{D} = \boldsymbol{R}^\top \boldsymbol{B}$ with $\boldsymbol{R}^\top \boldsymbol{R} = \boldsymbol{R}\boldsymbol{R}^\top = \boldsymbol{I}_d$. Construct matrix $\boldsymbol{B}$ as in Eq.(4) with index set Λ as Eq.(6). Then $\mu(\boldsymbol{D}) = \mu(\boldsymbol{B}) \leq \frac{\sqrt{n}}{m}$ and $\|\boldsymbol{D}_j\|_2 = \|\boldsymbol{B}_j\|_2 = \sqrt{\frac{m}{n}}$, $\forall j \in \{1, \cdots, M\}$.*

Theorem 3 shows that, for any orthogonal matrix $\boldsymbol{R}$, each column vector of the structured dictionary $\boldsymbol{D}$ has a constant l_2-norm. Moreover, it remains a constant mutual coherence $\mu(\boldsymbol{D}) = \mu(\boldsymbol{B})$. Thus, given a fixed matrix $\boldsymbol{B}$, we only need to learn matrix $\boldsymbol{R}$ for the dictionary learning without undermining the low mutual coherence property.

3.1.2 Joint Optimization for Dictionary Learning and Sparse Coding

With the structured matrix $\boldsymbol{B}$, we can jointly optimize $\boldsymbol{R}$ and $\boldsymbol{Y}$ for the optimization problem (7).

$$\begin{aligned} &\min_{\boldsymbol{Y},\boldsymbol{R}} \|\boldsymbol{X} - \boldsymbol{R}^\top \boldsymbol{B}\boldsymbol{Y}\|_F^2 \\ &\text{subject to } \ \|\boldsymbol{Y}_i\|_0 \leq k, \forall i \in \{1, \cdots, N\} \\ &\qquad\qquad \boldsymbol{R}^\top \boldsymbol{R} = \boldsymbol{R}\boldsymbol{R}^\top = \boldsymbol{I}_d \end{aligned} \tag{7}$$

This problem can be solved by the alternative descent method. For a fixed $\boldsymbol{R}$, we show the sparse representation $\boldsymbol{Y}$ has a closed-form solution thanks to the structured dictionary. For the fixed sparse codes $\boldsymbol{Y}$, dictionary parameter $\boldsymbol{R}$ also has a closed-form solution. By alternatively optimizing $\boldsymbol{R}$ and $\boldsymbol{Y}$, we can monotonically decrease the objective.

Fix $\boldsymbol{R}$, optimize $\boldsymbol{Y}$: Since the constraints of $\boldsymbol{Y}$ is column separable, i.e., $\|\boldsymbol{Y}_i\| \leq k$, and the objective (8) is also decomposable,

$$\|\boldsymbol{X} - \boldsymbol{R}^\top \boldsymbol{B}\boldsymbol{Y}\|_F^2 = \sum_{i=1}^{N} \|\boldsymbol{X}_i - \boldsymbol{R}^\top \boldsymbol{B}\boldsymbol{Y}_i\|_F^2. \tag{8}$$

It is sufficient to optimize the sparse code $\boldsymbol{Y}_i \in \mathbb{R}^M$ for each point $\boldsymbol{X}_i \in \mathbb{R}^d$ separately.

Without loss of generality, for any input $\boldsymbol{x} \in \mathbb{R}^d$, we aim at finding the optimal sparse code $\boldsymbol{y} \in \mathbb{R}^M$ such that $\|\boldsymbol{y}\|_0 \leq k$. Since $\boldsymbol{R}^\top \boldsymbol{R} = \boldsymbol{R}\boldsymbol{R}^\top = \boldsymbol{I}_d$ and $\boldsymbol{B}\boldsymbol{B}^\top = \boldsymbol{I}_d$, we have

$$\begin{aligned} \|\boldsymbol{x} - \boldsymbol{R}^\top \boldsymbol{B}\boldsymbol{y}\|_F^2 &= \|\boldsymbol{R}\boldsymbol{x} - \boldsymbol{R}\boldsymbol{R}^\top \boldsymbol{B}\boldsymbol{y}\|_F^2 \\ &= \|\boldsymbol{R}\boldsymbol{x} - \boldsymbol{B}\boldsymbol{y}\|_F^2 \\ &= \|\boldsymbol{B}\boldsymbol{B}^\top \boldsymbol{R}\boldsymbol{x} - \boldsymbol{B}\boldsymbol{y}\|_F^2 \\ &= \|\boldsymbol{B}(\boldsymbol{B}^\top \boldsymbol{R}\boldsymbol{x} - \boldsymbol{y})\|_F^2 \\ &= \|\boldsymbol{B}(\boldsymbol{h} - \boldsymbol{y})\|_F^2. \end{aligned} \tag{9}$$

where $\boldsymbol{h} = \boldsymbol{B}^\top \boldsymbol{R}\boldsymbol{x}$ is a dense code. Rewrite Eq.(9) into a summation form, we have

$$\|\boldsymbol{B}(\boldsymbol{h} - \boldsymbol{y})\|_F^2 = \sum_{j=1}^{M} (\boldsymbol{h}_j - \boldsymbol{y}_j)^2 \|\boldsymbol{B}_j\|_2^2. \tag{10}$$

Note that $\|\boldsymbol{B}_j\|_2^2 = \frac{m}{n}$, it follows that

$$\|\boldsymbol{B}(\boldsymbol{h} - \boldsymbol{y})\|_F^2 = \frac{m}{n} \sum_{j=1}^{M} (\boldsymbol{h}_j - \boldsymbol{y}_j)^2. \tag{11}$$

Now, we consider the sparse constraint $\|\boldsymbol{y}\|_0 \leq k$. Because each term $(\boldsymbol{h}_j - \boldsymbol{y}_j)^2 \geq 0$ is minimized when $\boldsymbol{y}_j = \boldsymbol{h}_j$, we know that Eq.(11) under sparse constraints is minimized when all the non-zero term setting as $\boldsymbol{y}_j = \boldsymbol{h}_j$. Otherwise we can set a non-zero term $\boldsymbol{y}_j$ to $\boldsymbol{y}_j = \boldsymbol{h}_j$ to further reduce term $(\boldsymbol{h}_j - \boldsymbol{y}_j)^2$ to zero.

Now, the problem is to find the index set of the non-zero term to minimize Eq.(12).

$$\sum_{j=1}^{M} (\boldsymbol{h}_j - \boldsymbol{y}_j)^2 = \sum_{j=1}^{M} \boldsymbol{h}_j^2 - \sum_{i \in S, |S| \leq k} \boldsymbol{h}_i^2 \tag{12}$$

where $S := \{j | \boldsymbol{y}_j \neq 0\}$. We can see that Eq.(12) is minimized when S consists of the index of the k largest elements $\boldsymbol{y}_i^2$ of $\boldsymbol{y}^2 \in \mathbb{R}^M$.

Fix $\boldsymbol{Y}$, Optimize $\boldsymbol{R}$: For a fixed $\boldsymbol{Y}$, we know that

$$\|\boldsymbol{X} - \boldsymbol{R}^\top \boldsymbol{B}\boldsymbol{Y}\|_F^2 = \|\boldsymbol{X}\|_F^2 + \|\boldsymbol{B}\boldsymbol{Y}\|_F^2 - 2\text{tr}(\boldsymbol{R}^\top \boldsymbol{B}\boldsymbol{Y}\boldsymbol{X}^\top) \tag{13}$$

This is the nearest orthogonal matrix problem, which has a closed-form solution as shown in (Schönemann, 1966; Gong et al., 2012). Let $\boldsymbol{B}\boldsymbol{Y}\boldsymbol{X}^\top = \boldsymbol{U}\Gamma\boldsymbol{V}^\top$ obtained by singular value decomposition (SVD), where $\boldsymbol{U}, \boldsymbol{V}$ are orthgonal matrix. Then, Eq.(13) is minimized by

$$\boldsymbol{R} = \boldsymbol{U}\boldsymbol{V}^\top \tag{14}$$

3.2 l_0-NORM REGULARIZATION

We employ the structured dictionary $\boldsymbol{D} = \boldsymbol{R}^\top \boldsymbol{B}$ same as in Section 3.1. The optimization problem with l_0-norm regularization is defined as

$$\begin{aligned} &\min_{\boldsymbol{Y},\boldsymbol{R}} \|\boldsymbol{X} - \boldsymbol{R}^\top \boldsymbol{B}\boldsymbol{Y}\|_F^2 + \lambda\|\boldsymbol{Y}\|_0 \\ &\text{subject to } \ \boldsymbol{R}^\top \boldsymbol{R} = \boldsymbol{R}\boldsymbol{R}^\top = \boldsymbol{I}_d \end{aligned} \tag{15}$$

This problem can be solved by the alternative descent method. For a fixed $\boldsymbol{R}$, we show $\boldsymbol{Y}$ has a closed-form solution thanks to the structured dictionary. For fixed the sparse codes $\boldsymbol{Y}$, dictionary parameter $\boldsymbol{R}$ also has a closed-form solution. By alternatively optimizing $\boldsymbol{R}$ and $\boldsymbol{Y}$, we can monotonically decrease the objective.

Fix $\boldsymbol{R}$, optimize $\boldsymbol{Y}$: Since the objective can be rewritten as Eq.(16)

$$\|\boldsymbol{X} - \boldsymbol{R}^\top \boldsymbol{B}\boldsymbol{Y}\|_F^2 + \lambda\|\boldsymbol{Y}\|_0 = \sum_{i=1}^{N} \|\boldsymbol{X}_i - \boldsymbol{R}^\top \boldsymbol{B}\boldsymbol{Y}_i\|_F^2 + \lambda\|\boldsymbol{Y}_i\|_0. \tag{16}$$

It is sufficient to optimize $\boldsymbol{Y}_i$ for each point $\boldsymbol{X}_i$ separately. Without loss of generality, for any input $\boldsymbol{x} \in \mathbb{R}^d$, we aim at finding an optimal sparse code $\boldsymbol{y} \in \mathbb{R}^M$. Since $\boldsymbol{R}^\top \boldsymbol{R} = \boldsymbol{R}\boldsymbol{R}^\top = \boldsymbol{I}_d$ and $\boldsymbol{B}\boldsymbol{B}^\top = \boldsymbol{I}_d$, following the derivation in Section 3.1.2, we have

$$\|\boldsymbol{x} - \boldsymbol{R}^\top \boldsymbol{B}\boldsymbol{y}\|_F^2 + \lambda\|\boldsymbol{y}\|_0 = \|\boldsymbol{B}(\boldsymbol{h} - \boldsymbol{y})\|_F^2 + \lambda\|\boldsymbol{y}\|_0, \tag{17}$$

where $\boldsymbol{h} = \boldsymbol{B}^\top \boldsymbol{R}\boldsymbol{x}$ is a dense code. Note that $\|\boldsymbol{B}_j\|_2^2 = \frac{m}{n}$, together with Eq.(17), it follows that

$$\|\boldsymbol{B}(\boldsymbol{h} - \boldsymbol{y})\|_F^2 + \lambda\|\boldsymbol{y}\|_0 = \frac{m}{n}\left(\sum_{j=1}^{M} (\boldsymbol{h}_j - \boldsymbol{y}_j)^2 + \frac{n\lambda}{m}\|\boldsymbol{y}_j\|_0\right). \tag{18}$$

This problem is separable for each variable $\boldsymbol{y}_j$, and each term is minimized by setting

$$\boldsymbol{y}_j = \begin{cases} \boldsymbol{h}_j & \text{if } \ \boldsymbol{h}_j^2 \geq \frac{n\lambda}{m} \\ 0 & \text{otherwise} \end{cases}. \tag{19}$$

Fix $\boldsymbol{Y}$, update $\boldsymbol{R}$: For a fixed $\boldsymbol{Y}$, minimizing the objective leads to the same nearest orthogonal matrix problem as shown in Section 3.1.2. Let $\boldsymbol{B}\boldsymbol{Y}\boldsymbol{X}^\top = \boldsymbol{U}\Gamma\boldsymbol{V}^\top$ obtained by SVD, where $\boldsymbol{U}, \boldsymbol{V}$ are orthogonal matrix. Then, the reconstruction problem is minimized by $\boldsymbol{R} = \boldsymbol{U}\boldsymbol{V}^\top$.

Remark: Problems with other separable regularization terms can be solved in a similar way. The key difference is how to achieve sparse codes $\boldsymbol{y}$. For example, for l_1-norm regularized problems, $\boldsymbol{y}$ can be obtained by a soft thresholding function, i.e., $\boldsymbol{y} = \text{sign}(\boldsymbol{y}) \odot \max\big(0, |\boldsymbol{y}| - n\lambda/(2m)\big)$.

4 SPARSE DENOISING LAYER

One benefit of our fast sparse coding algorithm is that it enables a simple closed-form reconstruction, which can be used as a plug-in layer for deep neural networks. Specifically, given an orthogonal matrix $\boldsymbol{R}$ and input vector $\boldsymbol{x}$, the optimal reconstruction of our method can be expressed as

$$\boxed{\widetilde{\boldsymbol{x}} = \boldsymbol{R}^\top \boldsymbol{B} f(\boldsymbol{B}^\top \boldsymbol{R}\boldsymbol{x})}, \tag{20}$$

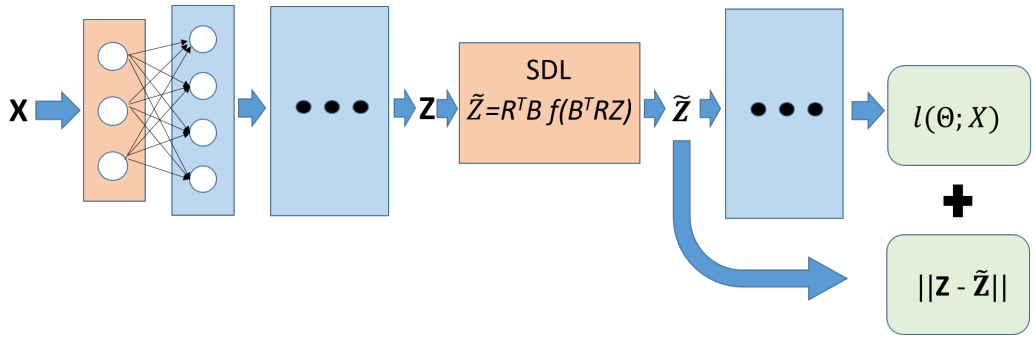

Figure 1: Illustration of the SDL Plug-in

where $f(\cdot)$ is a non-linear mapping function. For the k-sparse constrained problem, $f(\cdot)$ is a k-max pooling function (w.r.t the absolute value) as Eq.(21)

$$f(\boldsymbol{h}_j) = \begin{cases} \boldsymbol{h}_j & \text{if } \ |\boldsymbol{h}_j| \text{ is one of the k-highest values of } \ |\boldsymbol{h}| \in \mathbb{R}^M \\ 0 & \text{otherwise} \end{cases} . \tag{21}$$

For the l_0-norm regularization problem, $f(\cdot)$ is a hard thresholding function as Eq.(22)

$$f(\boldsymbol{h}_j) = \begin{cases} \boldsymbol{h}_j & \text{if } \ |\boldsymbol{h}_j| \geq \sqrt{\frac{n\lambda}{m}} \\ 0 & \text{otherwise} \end{cases} . \tag{22}$$

For the l_1-norm regularization problem, $f(\cdot)$ is a soft thresholding function as Eq.(23)

$$f(\boldsymbol{h}_j) = \text{sign}(\boldsymbol{h}_j) \times \max\left(0, |\boldsymbol{h}_j| - \frac{n\lambda}{2m}\right), \tag{23}$$

where $\text{sign}(\cdot)$ denotes the Sign function.

The reconstruction in Eq.(20) can be used as a simple plug-in layer for deep networks, we named it as sparse denoising layer (SDL). It is worth noting that only the orthogonal matrix $\boldsymbol{R}$ is needed to learn. The structured matrix $\boldsymbol{B}$ is constructed as in Section 3.1.1 and fixed.

The orthogonal matrix $\boldsymbol{R}$ can be parameterized by exponential mapping or Cayley mapping (Helfrich et al., 2018) of a skew-symmetric matrix. In this work, we employ the Cayley mapping to enable gradient update using deep learning tools. Specifically, the orthogonal matrix $\boldsymbol{R}$ can be obtained by the Cayley mapping of a skew-symmetric matrix as

$$\boldsymbol{R} = (\boldsymbol{I} + \boldsymbol{W})(\boldsymbol{I} - \boldsymbol{W})^{-1}, \tag{24}$$

where $\boldsymbol{W}$ is a skew-symmetric matrix, i.e., $\boldsymbol{W} = -\boldsymbol{W}^\top \in \mathbb{R}^{d\times d}$. For a skew-symmetric matrix $\boldsymbol{W}$, only the upper triangular matrix (without main diagonal) are free parameters. Thus, the number of free parameters of SDL is $d(d-1)/2$, which is much smaller compared with the number of parameters of backbone deep networks.

For training a network with a SDL, we add a reconstruction loss term as a regularization. The optimization problem is defined as

$$\min_{\boldsymbol{W}\in\mathbb{R}^{d\times d},\boldsymbol{\Theta}} \ell(\boldsymbol{X};\boldsymbol{W},\boldsymbol{\Theta}) + \beta\|\widetilde{\boldsymbol{Z}} - \boldsymbol{Z}\|_F^2, \tag{25}$$

where $\boldsymbol{W}$ is a skew-symmetric matrix parameter of SDL. $\boldsymbol{\Theta}$ is the parameter of the backbone network. $\widetilde{\boldsymbol{Z}}$ is the reconstruction of the latent representation $\boldsymbol{Z}$ via SDL (Eq.(20)). An illustration of the SDL plug-in is shown in Figure 1. When SDL is used as the first layer, then $\boldsymbol{Z} = \boldsymbol{X}$, and $\widetilde{\boldsymbol{Z}} = \widetilde{\boldsymbol{X}}$. In this case, $\widetilde{\boldsymbol{Z}}$ is the reconstruction of the input data $\boldsymbol{X}$.

It is worth noting that the shape of the input and output of SDL are same. Thus, SDL can be used as plug-in for any backbone models without changing the input/output shape of different layers in the backbone network. With the simple SDL plug-in, backbone models can be trained from scratches.

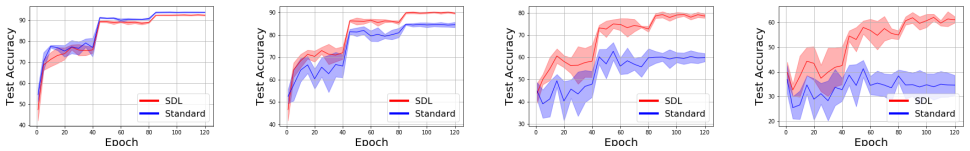

(a) Densenet-Clean (b) Densenet-Gaussian-0.1 (c) Densenet-Gaussian-0.2 (d) Densenet-Gaussian-0.3

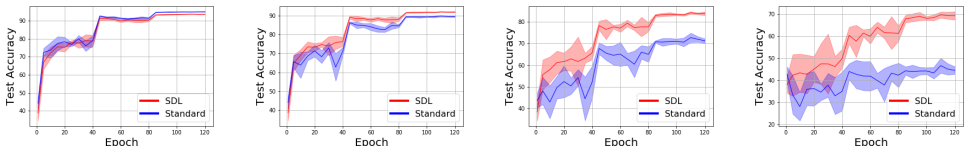

(e) Resnet-Clean (f) Resnet-Gaussian-0.1 (g) Resnet-Gaussian-0.2 (h) Resnet-Gaussian-0.3

Figure 2: Mean test accuracy $\pm$ std over 5 independent runs on CIFAR10 dataset with Gaussian noise for Densenet and Resnet with or without SDL

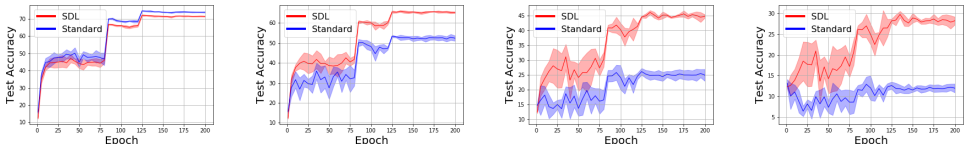

(a) Densenet-Clean (b) Densenet-Gaussian-0.1 (c) Densenet-Gaussian-0.2 (d) Densenet-Gaussian-0.3

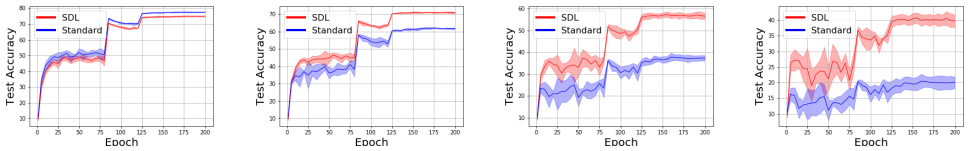

(e) Resnet-Clean (f) Resnet-Gaussian-0.1 (g) Resnet-Gaussian-0.2 (h) Resnet-Gaussian-0.3

Figure 3: Mean test accuracy $\pm$ std over 5 independent runs on CIFAR100 dataset with Gaussian noise for Densenet and Resnet with or without SDL

5 EXPERIMENTS

We evaluate the performance of our SDL on both classification tasks and RL tasks. For classification, we employ both DenseNet-100 (Huang et al., 2017) and ResNet-34 (He et al., 2016) as backbone. For RL tasks, we employ deep PPO models (Schulman et al., 2017) as backbone. For all the tasks, we test the performance of backbone models with and without our SDL when adding Gaussian noise or Laplace noise. In all the experiments, we plug SDL as the first layer of deep models. We set the standard deviation of input noise as $\{0, 0.1, 0.2, 0.3\}$, respectively. (The input noise is added after input normalization) . We keep all the hyperparameters of the backbone models same, the only difference is whether plugging SDL. The parameter β for the reconstruction loss is fixed as $\beta = 100$ in all the experiments.

Classification Tasks: We test SDL on CIFAR10 and CIFAR100 dataset. We construct the structured matrix $\boldsymbol{B} \in \mathbb{R}^{432\times 866}$ by Eq.(4). In this setting, the orthogonal matrix $\boldsymbol{R}$ corresponds to the convolution parameter of $Conv2d(\cdot)$ with *kernelsize*=12×12. We set the sparse parameter of our k-sparse SDL as $k = 75$ in all the classification experiments. The average test accuracy over five independent runs on CIFAR10 and CIFAR100 dataset with Gaussian noise are shown in Fig. 2 and Fig. 3, respectively. We can observe that models with SDL obtain a similar performance compared with vanilla model on the clean input case. With an increasing variance of input noise, models with SDL outperform the vanilla models more and more significantly. The experimental results with Laplace noise are presented in Fig. 8 and Fig. 9 in the supplementary. The results on Laplace noise cases show the similar trends with Gaussian noise cases.

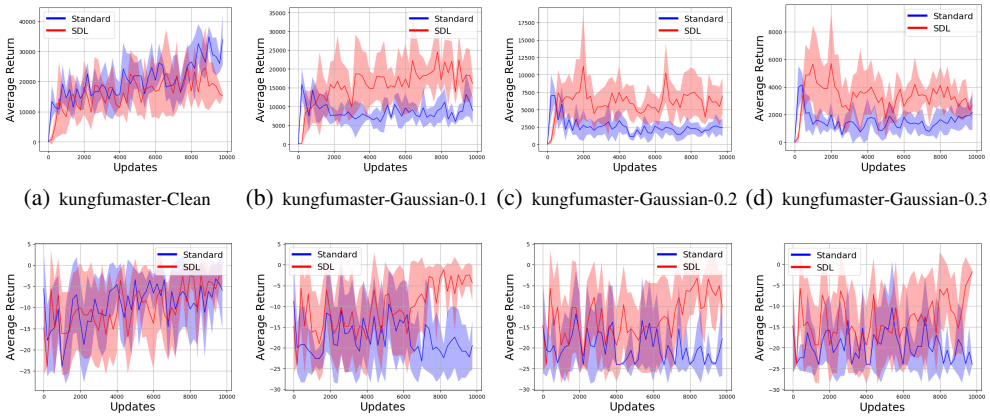

(a) kungfumaster-Clean (b) kungfumaster-Gaussian-0.1 (c) kungfumaster-Gaussian-0.2 (d) kungfumaster-Gaussian-0.3

(e) Tennis-Clean (f) Tennis-Gaussian-0.1 (g) Tennis-Gaussian-0.2 (h) Tennis-Gaussian-0.3

Figure 4: Average Return $\pm$ std over 5 independent runs on *KungfuMaster* and *Tennis* game with Gaussian noise with or without SDL

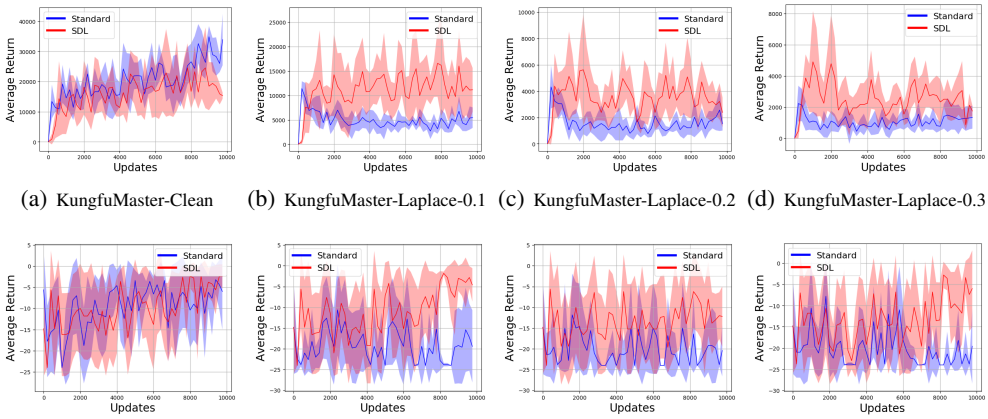

(a) KungfuMaster-Clean (b) KungfuMaster-Laplace-0.1 (c) KungfuMaster-Laplace-0.2 (d) KungfuMaster-Laplace-0.3

(e) Tennis-Clean (f) Tennis-Laplace-0.1 (g) Tennis-Laplace-0.2 (h) Tennis-Laplace-0.3

Figure 5: Average Return $\pm$ std over 5 independent runs on *KungfuMaster* and *Tennis* game with Laplace noise with or without SDL

RL Tasks: We test deep PPO model with SDL on Atari games: *KungFuMaster*, *Tennis* and *Seaquest*. The deep PPO model concatenates four frames as the input state. The size of input state is $84 \times 84 \times 4$. We construct the structured matrix $\boldsymbol{B} \in \mathbb{R}^{36\times 38}$ by Eq.(4). In this setting, the orthogonal matrix $\boldsymbol{R}$ corresponds to the convolution parameter of $Conv2d(\cdot)$ with *kernelsize*=3×3. The number of free parameters is 630. We set the sparse parameter of our k-sparse SDL as $k = 4$ in all the RL experiments.

The return of one episode is the sum of rewards over all steps during the whole episode. We present the average return over five independent runs on *KungFuMaster* and *Tennis* game with Gaussian noise and Laplace noise in Fig. 4 and Fig. 5, respectively. Results on *Seaquest* game are shown in Fig. 10 in the supplement due to the space limitation. We can see that models with SDL achieve a competitive average return on the clean cases. Moreover, models with SDL obtain higher average return than vanilla models when the input state is perturbed with noise.

6 Conclusion

We proposed fast sparse coding algorithms for both k-sparse problem and l_0-norm regularization problems. Our algorithms have a simple closed-form update. We proposed a sparse denoising layer as a lightweight plug-in for backbone models against noisy input perturbation based on this handy closed-form. Experiments on both ResNet/DenseNet classification model and deep PPO RL model showed the effeteness of our SDL against noisy input perturbation.

REFERENCES

Chenglong Bao, Jian-Feng Cai, and Hui Ji. Fast sparsity-based orthogonal dictionary learning for image restoration. In *Proceedings of the IEEE International Conference on Computer Vision*, pp. 3384–3391, 2013.

Jonathan T Barron. A general and adaptive robust loss function. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 4331–4339, 2019.

Jean Bourgain, Alexey A Glibichuk, and SERGEI VLADIMIROVICH KONYAGIN. Estimates for the number of sums and products and for exponential sums in fields of prime order. *Journal of the London Mathematical Society*, 73(2):380–398, 2006.

Ekin D Cubuk, Barret Zoph, Jonathon Shlens, and Quoc V Le. Randaugment: Practical automated data augmentation with a reduced search space. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pp. 702–703, 2020.

Michael Elad and Michal Aharon. Image denoising via sparse and redundant representations over learned dictionaries. *IEEE Transactions on Image processing*, 15(12):3736–3745, 2006.

Gamaleldin Elsayed, Dilip Krishnan, Hossein Mobahi, Kevin Regan, and Samy Bengio. Large margin deep networks for classification. In *Advances in neural information processing systems*, pp. 842–852, 2018.

Yunchao Gong, Svetlana Lazebnik, Albert Gordo, and Florent Perronnin. Iterative quantization: A procrustean approach to learning binary codes for large-scale image retrieval. *IEEE transactions on pattern analysis and machine intelligence*, 35(12):2916–2929, 2012.

Ian J Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples. *ICLR*, 2015.

Minghao Guo, Yuzhe Yang, Rui Xu, Ziwei Liu, and Dahua Lin. When nas meets robustness: In search of robust architectures against adversarial attacks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 631–640, 2020.

Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778, 2016.

Kyle Helfrich, Devin Willmott, and Qiang Ye. Orthogonal recurrent neural networks with scaled cayley transform. In *International Conference on Machine Learning*, pp. 1969–1978. PMLR, 2018.

Gao Huang, Zhuang Liu, Laurens Van Der Maaten, and Kilian Q Weinberger. Densely connected convolutional networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 4700–4708, 2017.

Yueming Lyu. Spherical structured feature maps for kernel approximation. In *International Conference on Machine Learning*, pp. 2256–2264, 2017.

Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. Towards deep learning models resistant to adversarial attacks. *ICLR*, 2018.

Alexander J Ratner, Henry Ehrenberg, Zeshan Hussain, Jared Dunnmon, and Christopher Ré. Learning to compose domain-specific transformations for data augmentation. In *Advances in neural information processing systems*, pp. 3236–3246, 2017.

Ron Rubinstein, Michael Zibulevsky, and Michael Elad. Efficient implementation of the k-svd algorithm using batch orthogonal matching pursuit. Technical report, Computer Science Department, Technion, 2008.

Peter H Schönemann. A generalized solution of the orthogonal procrustes problem. *Psychometrika*, 31(1):1–10, 1966.

John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.

Dong Su, Huan Zhang, Hongge Chen, Jinfeng Yi, Pin-Yu Chen, and Yupeng Gao. Is robustness the cost of accuracy?–a comprehensive study on the robustness of 18 deep image classification models. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pp. 631–648, 2018.

Singanallur V Venkatakrishnan, Charles A Bouman, and Brendt Wohlberg. Plug-and-play priors for model based reconstruction. In *2013 IEEE Global Conference on Signal and Information Processing*, pp. 945–948. IEEE, 2013.

Kaixuan Wei, Angelica Aviles-Rivero, Jingwei Liang, Ying Fu, Carola-Bibiane Schnlieb, and Hua Huang. Tuning-free plug-and-play proximal algorithm for inverse imaging problems. *ICML*, 2020.

Hongyang Zhang, Yaodong Yu, Jiantao Jiao, Eric P Xing, Laurent El Ghaoui, and Michael I Jordan. Theoretically principled trade-off between robustness and accuracy. *ICML*, 2019.

Kai Zhang, Wangmeng Zuo, Yunjin Chen, Deyu Meng, and Lei Zhang. Beyond a gaussian denoiser: Residual learning of deep cnn for image denoising. *IEEE Transactions on Image Processing*, 26 (7):3142–3155, 2017.

Stephan Zheng, Yang Song, Thomas Leung, and Ian Goodfellow. Improving the robustness of deep neural networks via stability training. In *Proceedings of the ieee conference on computer vision and pattern recognition*, pp. 4480–4488, 2016.