

Dropout Rademacher Complexity Of Deep Neural Networks

GAO Wei & ZHOU Zhi-Hua

*National Key Laboratory for Novel Software Technology, Nanjing University
Collaborative Innovation Center of Novel Software Technology and Industrialization
Nanjing 210023, China*

Received xxx, 2015; accepted xxx, 2015; published online xxx, 2015

Abstract Great successes of deep neural networks have been witnessed in various real applications. Many algorithmic and implementation techniques have been developed; however, theoretical understanding of many aspects of deep neural networks is far from clear. A particular interesting issue is the usefulness of *dropout*, which was motivated from the intuition of preventing complex co-adaptation of feature detectors. In this paper, we study the Rademacher complexity of different types of dropouts, and our theoretical results disclose that for shallow neural networks (with one or none hidden layer) dropout is able to reduce the Rademacher complexity in polynomial, whereas for deep neural networks it can amazingly lead to an exponential reduction.

Keywords artificial intelligence, machine learning, deep Learning, dropout, Rademacher complexity

Citation Gao W & Zhou Z-H. Dropout Rademacher complexity of deep neural networks. *Sci China Inf Sci*, 2015, xx: xxxxxx(12), doi: xxxxxxxxxxxxxxxx

1 Introduction

Deep neural networks [1] has become a hot wave during the past few years, and great successes have been achieved in various real applications, such as object recognition [2–4], speech recognition [5–7], video analysis [8,9], etc. Many effective algorithmic and implementation techniques [10] have been developed; however, theoretical understanding of many aspects of deep neural networks is far from clear.

It is well known that deep neural networks are complicated models with rich representations. For really deep networks, there may be millions or even billions of parameters, and thus, there are high risks of overfitting even with large-scale training data. Indeed, controlling the overfitting risk is a long-standing topic in the research of neural networks, and various techniques have been developed, such as weight elimination [11], early stopping [12], Bayesian control [13], etc.

Dropout is among the key ingredients of the success of deep neural networks. The main idea is to randomly omit some units, either hidden ones or input ones corresponding to different input features; this is executed with certain probability in the forward propagation of training phase, and the weights related to the remaining units are updated in back propagation. This technique is evidently related to overfitting control, though it was proposed with the intuition of preventing complex co-adaptations

by encouraging independent contributions from different features during training phase [14]. Extensive empirical studies [14–16] verified that dropout is able to improve the performance and reduce overfitting risk. However, theoretical understanding of dropout is far from clear.

In this paper, we study the influence on Rademacher complexity by three types of dropouts, i.e., dropout of units [14], dropout of weights [16] and dropout of both. Our theoretical results disclose that for shallow neural networks with none or one hidden layer, dropout is able to reduce the Rademacher complexity in polynomial, whereas for deep neural networks it is able to reach an exponential reduction of Rademacher complexity.

Related Work

There are several designs of dropout, such as the fast dropout [17] and adaptive dropout [18], whereas the most fundamental dropouts are the dropout of units (hidden units, or input units corresponding to input features) [14, 19] and the dropout of weights [16].

The average and regularizing properties of dropout have been studied in [20]. Wager et al. [21] showed that dropout is first-order regular equivalent to an L_2 regularizer applied after scaling the features by an estimate of the inverse diagonal Fisher information matrix. The generalization bound of dropout has been analyzed in [16, 22]. McAllester [22] presented PAC-Bayesian bounds, whereas Wan et al. [16] derived Rademacher generalization bounds. Both their results show that the reduction of complexity brought by dropout is $O(\rho)$, where ρ is the probability of keeping an element in dropout.

In contrast to previous studies [16, 22], we present better generalization bounds and disclose that dropout is able to reduce the Rademacher complexity exponentially, i.e., $O(\rho^{k+1})$ or $O(\rho^{(k+1)/2})$ for different types of dropouts, where k is the number of hidden layers within neural networks.

Extensive work [23, 24, and reference therein] studied the complexity of neural network based on VC-dimension, covering number, fat-shatter dimension, etc., and it was usually shown that these complexities are polynomial in the total number of units and weights. Note that the polynomial complexities are still very high for deep neural networks that may have million or even billions of parameters. Moreover, it is worth noting that these complexities measure the function space in the worst case, and cannot distinguish situations with/without dropouts. In contrast, we show that Rademacher complexity is proper to study the influence of dropouts, and we prove that the complexities of neural network can be bounded by the L_1 or L_2 -norm of weights, irrelevant to the number of units and weights.

This paper is organized as follows: Section 2 introduces some preliminaries. Section 3 presents general Rademacher generalization bounds for dropout. Section 4 analyzes the usefulness of different types of dropouts on shallow as well as deep neural networks. Section 5 concludes.

2 Preliminaries

Let $\mathcal{X} \subset \mathbb{R}^d$ and \mathcal{Y} be the input and output space, respectively, where $\mathcal{Y} \subset \mathbb{R}$ for regression and $\mathcal{Y} = \{+1, -1\}$ for binary classification. Throughout this paper, we restrict our attention on regression and binary classification, and it is easy to make similar analysis for multi-class tasks. Let \mathcal{D} be an unknown (underlying) distribution over $\mathcal{X} \times \mathcal{Y}$.

Let \mathcal{W} be the weight space for neural network, and denote $f(\mathbf{w}, \mathbf{x})$ the general output of a neural network with respect to input $\mathbf{x} \in \mathcal{X}$ and weight $\mathbf{w} \in \mathcal{W}$. Here f depends on the structure of neural network. During training neural network, dropout randomly omits hidden units, input units corresponding to input features, and connected weights with certain probability; therefore, it is necessary to introduce another space

$$\mathcal{R} = \{\mathbf{r} = (r_1, r_2, \dots, r_s) : r_i \in \{0, 1\}\}$$

where s depends on different neural networks and different types of dropouts, and $r_i = 0$ implies dropping out some hidden unit, input unit and weight. Here each r_i is drawn independently and identically from a Bernoulli distribution with parameter ρ , denoted by $\text{Bern}(\rho)$. Further, we denote $f(\mathbf{w}, \mathbf{x}, \mathbf{r})$ the dropout

output of a neural network, and write

$$\mathcal{F}_{\mathcal{W}} = \{f(\mathbf{w}, \mathbf{x}, \mathbf{r}) : \mathbf{w} \in \mathcal{W}\},$$

as the function space for dropout. Here we just present a general output $f(\mathbf{w}, \mathbf{x}, \mathbf{r})$ for dropout, and detailed expressions will be given for specific neural network in Section 4.

An objective function (or loss function) ℓ is introduced to measure the performance of output of neural network. For example, least square loss and cross entropy are used for regression and binary classification, respectively. We define the expected risk for dropout as

$$R(\mathbf{w}) = E_{\mathbf{r}, (\mathbf{x}, y)}[\ell(f(\mathbf{w}, \mathbf{x}, \mathbf{r}), y)].$$

The goal is to find a $\mathbf{w}^* \in \mathcal{W}$ so as to minimize the expected risk, i.e., $\mathbf{w}^* \in \arg \min_{\mathbf{w} \in \mathcal{W}} R(\mathbf{w})$. Notice that the distribution \mathcal{D} is unknown, but it is demonstrated by a training sample

$$S_n = \{(\mathbf{x}_1, y_1), (\mathbf{x}_2, y_2), \dots, (\mathbf{x}_n, y_n)\}$$

which are drawn i.i.d. from distribution \mathcal{D} . Given sample S_n and $RS_n = \{\mathbf{r}_1, \mathbf{r}_2, \dots, \mathbf{r}_n\}$, we define the empirical risk for dropout as

$$\hat{R}(\mathbf{w}, S_n, RS_n) = \frac{1}{n} \sum_{i=1}^n \ell(f(\mathbf{w}, \mathbf{x}_i, \mathbf{r}_i), y_i).$$

In this paper, we try to study on generalization bounds for dropouts, i.e., the gap between $R(\mathbf{w})$ and $\hat{R}(\mathbf{w}, S_n, RS_n)$. Rademacher complexity has always been an efficient measure for function space [25, 29]. For function space \mathcal{H} , the classical Rademacher complexity is defined by

$$\hat{\mathfrak{R}}_n(\mathcal{H}) = E \left[\sup_{h \in \mathcal{H}} \frac{1}{n} \sum_{i=1}^n \epsilon_i h(\mathbf{x}_i) \right] \tag{1}$$

where $\epsilon_1, \dots, \epsilon_n$ are independent random variables uniformly chosen from $\{+1, -1\}$, and they are referred as Rademacher variables. Rademacher complexity has been used to develop data-dependent generalization bounds in diverse learning tasks [30–32].

For notational simplicity, we denote $[n] = \{1, 2, \dots, n\}$ for integer $n > 0$. The inner product between $\mathbf{w} = (w_1, \dots, w_d)$ and $\mathbf{x} = (x_1, \dots, x_d)$ is given by $\langle \mathbf{w}, \mathbf{x} \rangle = \sum_{i=1}^d w_i x_i$, and write $\|\mathbf{w}\| = \|\mathbf{w}\|_2 = \sqrt{\langle \mathbf{w}, \mathbf{w} \rangle}$ and $\|\mathbf{w}\|_1 = \sum_{i=1}^d |w_i|$. Further, the entrywise product (also called Schur product or Hadamard product) is defined as $\mathbf{w} \odot \mathbf{x} = (w_1 x_1, \dots, w_d x_d)$.

3 General Rademacher Generalization Bounds

In conventional studies, the generalization performance is mostly affected by training sample [33], and the standard Rademacher complexity is defined on training sample only (as shown in Eq. 1). For dropout, however, the generalization performance is not only relevant to training sample, but also dropout randomization; thus, we generalize the Rademacher complexity as follows:

Definition 1. For spaces \mathcal{Z} and \mathcal{R} , let $\mathcal{H} : \mathcal{Z} \times \mathcal{R} \rightarrow \mathbb{R}$ be a real-valued function space. For $S_n = \{\mathbf{z}_1, \dots, \mathbf{z}_n\}$ and $RS_n = \{\mathbf{r}_1, \dots, \mathbf{r}_n\}$, the empirical Rademacher complexity of \mathcal{H} is defined to be

$$\hat{\mathfrak{R}}_n(\mathcal{H}, S_n, RS_n) = E_{\epsilon} \left[\sup_{h \in \mathcal{H}} \left(\frac{1}{n} \sum_{i=1}^n \epsilon_i h(\mathbf{z}_i, \mathbf{r}_i) \right) \right]$$

where $\epsilon = (\epsilon_1, \dots, \epsilon_n)$ are Rademacher variables. Further, we define the Rademacher complexity of \mathcal{H} as

$$\mathfrak{R}_n(\mathcal{H}) = E_{S_n, RS_n} [\hat{\mathfrak{R}}_n(\mathcal{H}, S_n, RS_n)].$$

Based on this definition, it is easy to get a useful lemma as follows:

Lemma 1. For function space \mathcal{H} , define $\text{absconv}(\mathcal{H}) = \{\sum \alpha_i h_i : h_i \in \mathcal{H} \text{ and } \sum |\alpha_i| = 1\}$. Then, we have

$$\hat{\mathfrak{R}}_n(\mathcal{H}, S_n, RS_n) = \hat{\mathfrak{R}}_n(\text{absconv}(\mathcal{H}), S_n, RS_n).$$

Given a set \mathcal{W} , we denote composite function space for dropout as

$$\ell \circ \mathcal{F}_{\mathcal{W}} := \{((\mathbf{x}, y), \mathbf{r}) \rightarrow \ell(f(\mathbf{w}, \mathbf{x}, \mathbf{r}), y), \mathbf{w} \in \mathcal{W}\}.$$

Based on the generalized Rademacher complexity, we present the general Rademacher generalization bounds for dropout as follows:

Theorem 1. Let $S_n = \{(\mathbf{x}_1, y_1), (\mathbf{x}_2, y_2), \dots, (\mathbf{x}_n, y_n)\}$ be a sample chosen i.i.d. according to distribution \mathcal{D} , and let $RS_n = \{\mathbf{r}_1, \mathbf{r}_2, \dots, \mathbf{r}_n\}$ be random variable sample for dropout. If the loss function ℓ is bounded by $B > 0$, for every $\delta > 0$ and $\mathbf{w} \in \mathcal{W}$, the following holds with probability at least $1 - \delta$,

$$R(\mathbf{w}) \leq \hat{R}(\mathbf{w}, S_n, RS_n) + 2\mathfrak{R}_n(\ell \circ \mathcal{F}_{\mathcal{W}}) + B\sqrt{\ln(2/\delta)/n}, \tag{2}$$

$$R(\mathbf{w}) \leq \hat{R}(\mathbf{w}, S_n, RS_n) + 2\hat{\mathfrak{R}}_n(\ell \circ \mathcal{F}_{\mathcal{W}}, S_n, RS_n) + 3B\sqrt{\ln(2/\delta)/n}. \tag{3}$$

Proof. The proof is motivated from the techniques in [25]. For every $\mathbf{w} \in \mathcal{W}$, it is easy to observe

$$R(\mathbf{w}) \leq \hat{R}(\mathbf{w}, S_n, RS_n) + \sup_{\mathbf{w}} [R(\mathbf{w}) - \hat{R}(\mathbf{w}, S_n, RS_n)],$$

and we further denote

$$\Phi(S_n, RS_n) = \sup_{\mathbf{w}} [R(\mathbf{w}) - \hat{R}(\mathbf{w}, S_n, RS_n)] = \sup_{\mathbf{w}} \left[R(\mathbf{w}) - \frac{1}{n} \sum_{i=1}^n \ell(f(\mathbf{w}, \mathbf{x}_i, \mathbf{r}_i), y_i) \right].$$

Let $S_n^{i, (\mathbf{x}'_i, y'_i)} = \{(\mathbf{x}_1, y_1), \dots, (\mathbf{x}'_i, y'_i), \dots, (\mathbf{x}_n, y_n)\}$ be the sample whose i -th example (\mathbf{x}_i, y_i) in S_n is replaced by (\mathbf{x}'_i, y'_i) , and $RS_n^{i, \mathbf{r}'_i} = \{\mathbf{r}_1, \dots, \mathbf{r}'_i, \dots, \mathbf{r}_n\}$ be the random variable vector with i -th variable \mathbf{r}_i replaced by \mathbf{r}'_i . For bounded loss $|\ell| < B$, we have

$$|\Phi(S_n, RS_n) - \Phi(S_n, RS_n^{i, \mathbf{r}'_i})| \leq B/m \text{ and } |\Phi(S_n, RS_n) - \Phi(S_n^{i, (\mathbf{x}'_i, y'_i)}, RS_n)| \leq B/m.$$

Based on McDiarmid's inequality [26], it holds that with probability at least $1 - \delta$,

$$\Phi(S_n, RS_n) \leq E_{S_n, RS_n} [\Phi(S_n, RS_n)] + B\sqrt{\ln(2/\delta)/n}.$$

Define a ghost sample $\tilde{S}_n = \{(\tilde{\mathbf{x}}_1, \tilde{y}_1), \dots, (\tilde{\mathbf{x}}_n, \tilde{y}_n)\}$ and a ghost random variable vector $\widetilde{RS}_n = \{\tilde{\mathbf{r}}_1, \dots, \tilde{\mathbf{r}}_1\}$. By using the fact

$$\Phi(S_n, RS_n) = \sup_{\mathbf{w}} [E_{\tilde{S}_n, \widetilde{RS}_n} [\hat{R}(\mathbf{w}, \tilde{S}_n, \widetilde{RS}_n) - \hat{R}(\mathbf{w}, S_n, RS_n)]],$$

we have

$$\begin{aligned} E_{S_n, RS_n} [\Phi(S_n, RS_n)] &\leq E \left[\sup_{\mathbf{w}} \left[\hat{R}(\mathbf{w}, \tilde{S}_n, \widetilde{RS}_n) - \hat{R}(\mathbf{w}, S_n, RS_n) \right] \right] \\ &= E \left[\sup_{\mathbf{w}} \left[\frac{\sum_{i=1}^n \ell(f(\mathbf{w}, \tilde{\mathbf{x}}_i, \tilde{\mathbf{r}}_i), \tilde{y}_i) - \ell(f(\mathbf{w}, \mathbf{x}_i, \mathbf{r}_i), y_i)}{n} \right] \right] \\ &\leq 2E \left[\sup_{\mathbf{w}} \frac{1}{n} \sum_{i=1}^n \epsilon_i \ell(f(\mathbf{w}, \mathbf{x}_i, \mathbf{r}_i), y_i) \right] = 2\mathfrak{R}_n(\ell \circ \mathcal{F}_{\mathcal{W}}) \end{aligned}$$

which completes the proofs of Eq. 2. Again, we apply McDiarmid's inequality to $\hat{\mathfrak{R}}_n(\mathcal{W}, S_n, RS_n)$, and we have

$$\mathfrak{R}_n(\ell \circ \mathcal{F}_{\mathcal{W}}) \leq \hat{\mathfrak{R}}_n(\ell \circ \mathcal{F}_{\mathcal{W}}, S_n, RS_n) + B\sqrt{\ln(2/\delta)/n}$$

which completes the proof of Eq. 3.

The main benefit of dropout lies in the sharp reduction on Rademacher complexities of $\mathfrak{R}_n(\mathcal{F}_W)$ as will be shown in Section 4; on the other hand, extensive experiments show that dropout decreases the empirical risk $\hat{R}(\mathbf{w}, S_n, RS_n)$ [14–16] because dropout intuitively prevents complex co-adaptations by encouraging independent contributions from different features during training phase [14]. This paper tries to present theoretical analysis on the the former, and leave the latter to future work.

To efficiently estimate $\mathfrak{R}_n(\ell \circ \mathcal{F}_W)$, we introduce a concentration as follows:

Lemma 2. [27] Let \mathcal{H} be a bounded real-valued function space from some space \mathcal{Z} and $\mathbf{z}_1, \dots, \mathbf{z}_n \in \mathcal{Z}$. Let $\phi: \mathbb{R} \rightarrow \mathbb{R}$ be Lipschitz with constant L and $\phi(0) = 0$. Then, we have

$$E_\epsilon \sup_{h \in \mathcal{H}} \frac{1}{n} \sum_{i \in [n]} \epsilon_i \phi(h(\mathbf{z}_i)) \leq LE_\epsilon \sup_{h \in \mathcal{H}} \frac{1}{n} \sum_{i \in [n]} \epsilon_i h(\mathbf{z}_i).$$

Based on this lemma, we have

Lemma 3. If $\ell(\cdot, \cdot)$ is Lipschitz with the first argument and constant L , we have

$$\mathfrak{R}_n(\ell \circ \mathcal{F}_W) \leq L\mathfrak{R}_n(\mathcal{F}_W).$$

Proof. We first write $\ell'(\cdot, \cdot) = \ell(\cdot, \cdot) - \ell(0, \cdot)$, and it is easy to get $\mathfrak{R}_n(\ell \circ \mathcal{F}_W) = \mathfrak{R}_n(\ell' \circ \mathcal{F}_W)$. This lemma holds by applying Lemma 2 to ℓ' .

For classification, we always use the entropy loss as the loss function in neural network as follows:

$$\ell(f(\mathbf{w}, \mathbf{x}, \mathbf{r}), y) = y \ln(y/\phi(f(\mathbf{w}, \mathbf{x}, \mathbf{r}))) + (1 - y) \ln((1 - y)/(1 - \phi(f(\mathbf{w}, \mathbf{x}, \mathbf{r}))),$$

where $\phi(t) = 1/(1 + e^{-t})$. It is easy to find that $\partial \ell(f(\mathbf{w}, \mathbf{x}, \mathbf{r}), y)/\partial f(\mathbf{w}, \mathbf{x}, \mathbf{r}) \in [-1, 1]$, and thus $\ell(\cdot, \cdot)$ is a Lipschitz function with the first argument.

For regression, we always use the square loss as the loss function in neural network as follows:

$$\ell(f(\mathbf{w}, \mathbf{x}, \mathbf{r}), y) = (y - f(\mathbf{w}, \mathbf{x}, \mathbf{r}))^2.$$

For bounded $f(\mathbf{w}, \mathbf{x}, \mathbf{r})$ and y , it is easy to find that $\ell(\cdot, \cdot)$ is a Lipschitz function with the first argument.

Based on Lemma 3, it is easy to estimate $\mathfrak{R}_n(\ell \circ \mathcal{F}_W)$ from $\mathfrak{R}_n(\mathcal{F}_W)$; therefore, we will focus on how to estimate $\mathfrak{R}_n(\mathcal{F}_W)$ for different types of dropouts and different neural networks in the subsequent section.

Finally, we introduce a useful lemma as follows:

Lemma 4. Let $\mathbf{r}_1 = (r_{11}, \dots, r_{1d})$ and $\mathbf{r}_2 = (r_{21}, \dots, r_{2d})$ be two random variable vectors, and each element in \mathbf{r}_1 and \mathbf{r}_2 is drawn i.i.d. from distribution $\text{Bern}(\rho)$. For $\mathbf{x} \in \mathcal{X}$, we have

$$E_{\mathbf{r}_1}[\langle \mathbf{x} \odot \mathbf{r}_1, \mathbf{x} \odot \mathbf{r}_1 \rangle] = \rho \langle \mathbf{x}, \mathbf{x} \rangle, \tag{4}$$

$$E_{\mathbf{r}_1, \mathbf{r}_2}[\langle \mathbf{x} \odot \mathbf{r}_1 \odot \mathbf{r}_2, \mathbf{x} \odot \mathbf{r}_1 \odot \mathbf{r}_2 \rangle] = \rho^2 \langle \mathbf{x}, \mathbf{x} \rangle. \tag{5}$$

Further, let $\mathbf{r} = (r_1, \dots, r_k)$ be k random variables drawn i.i.d. from distribution $\text{Bern}(\rho)$. We have

$$E_{\mathbf{r}_1, \mathbf{r}} \left\langle \mathbf{x} \odot \mathbf{r}_1 \prod_{i=1}^k r_i, \mathbf{x} \odot \mathbf{r}_1 \prod_{i=1}^k r_i \right\rangle = \rho^{k+1} \langle \mathbf{x}, \mathbf{x} \rangle, \tag{6}$$

$$E_{\mathbf{r}, \mathbf{r}_1, \mathbf{r}_2} \left\langle \mathbf{x} \odot \mathbf{r}_1 \odot \mathbf{r}_2 \prod_{i=1}^k r_i, \mathbf{x} \odot \mathbf{r}_1 \odot \mathbf{r}_2 \prod_{i=1}^k r_i \right\rangle = \rho^{k+2} \langle \mathbf{x}, \mathbf{x} \rangle. \tag{7}$$

Proof. Let $\mathbf{x} = (x_1, \dots, x_d)$. From the definitions of inner product and entrywise product, Eq. 4 holds from

$$E_{\mathbf{r}_1}[\langle \mathbf{x} \odot \mathbf{r}_1, \mathbf{x} \odot \mathbf{r}_1 \rangle] = E_{\mathbf{r}_1} \left[\sum_{j=1}^d x_j x_j r_{1j}^2 \right] = \rho \sum_{j=1}^d x_j x_j$$

where we use the fact $E_{r_{1j}}[r_{1j}^2] = \rho$ since r_{1j} is drawn i.i.d. from distribution $\text{Bern}(\rho)$. In a similar manner, Eqs. 5-7 hold from $E_{r_{1j}, r_{2j}}[r_{1j}^2 r_{2j}^2] = \rho^2$, $E_{\mathbf{r}, r_{1j}}[r_{1j} \prod_{i=1}^k (r_i)^2] = \rho^{1+k}$ and $E_{\mathbf{r}, r_{1j}, r_{2j}}[r_{1j}^2 r_{2j}^2 \prod_{i=1}^k (r_i)^2] = \rho^{2+k}$, respectively. This lemma follows as desired.

4 Dropouts on Different Types of Network

We study the two types of most fundamental dropouts: dropout units [14] and dropout weights [16]. In addition, we also study dropout both units and weights. For $\rho \in [0, 1]$, these types of dropouts are defined as:

- Type I (Drp^(I)): randomly drop out each unit including input unit (corresponding to input feature) with probability $1 - \rho$.
- Type II (Drp^(II)): randomly drop out each weight with probability $1 - \rho$.
- Type III (Drp^(III)): randomly drop out each weight and unit including input unit (corresponding to input feature) with probability $1 - \rho$.

We assume that a full-connected neural network has k hidden layers, and the i th hidden layer has m_i hidden units. The general output for this neural network is given by

$$f(\mathbf{w}, \mathbf{x}) = \langle \mathbf{w}_1^{[k]}, \Psi_k \rangle \text{ with } \Psi_i = (\sigma(\langle \mathbf{w}_1^{[i-1]}, \Psi_{i-1} \rangle), \dots, \sigma(\langle \mathbf{w}_{m_i}^{[i-1]}, \Psi_{i-1} \rangle)) \text{ for } i \in [k]$$

and $\Psi_0 = \mathbf{x}$, where $\mathbf{w} = (\mathbf{w}_1^{[k]}, \mathbf{w}_1^{[k-1]}, \dots, \mathbf{w}_{m_k}^{[k-1]}, \dots, \mathbf{w}_1^{[0]}, \dots, \mathbf{w}_{m_1}^{[0]})$ in which each $\mathbf{w}_i^{[j]}$ has the same size as Ψ_j and σ is an activation function.

Throughout this work, we assume that activation function σ is Lipschitz with constant L and $\sigma(0) = 0$, and many commonly used activation functions satisfy such assumptions, e.g., *tanh*, *center sigmoid*, *relu* [34], etc.

Formally, three types of dropouts for the full-connected network are defined as:

- The output for Drp^(I) (first type) is given by

$$\begin{aligned} f^{(I)}(\mathbf{w}, \mathbf{x}, \mathbf{r}) &= \langle \mathbf{w}_1^{[k]}, \Psi_k \odot \mathbf{r}^{[k]} \rangle \text{ with} \\ \Psi_i &= (\sigma(\langle \mathbf{w}_1^{[i-1]}, \Psi_{i-1} \odot \mathbf{r}^{[i-1]} \rangle), \dots, \sigma(\langle \mathbf{w}_{m_i}^{[i-1]}, \Psi_{i-1} \odot \mathbf{r}^{[i-1]} \rangle)) \end{aligned} \quad (8)$$

for $i \in [k]$ and $\Psi_0 = \mathbf{x}$. Here $\mathbf{r} = (\mathbf{r}^{[k]}, \dots, \mathbf{r}^{[1]}, \mathbf{r}^{[0]})$, each $\mathbf{r}^{[i]}$ has the same size with Ψ_i and each element in $\mathbf{r}^{[i]}$ is drawn i.i.d. from $\text{Bern}(\rho)$.

- The output for Drp^(II) (second type) is given by

$$\begin{aligned} f^{(II)}(\mathbf{w}, \mathbf{x}, \mathbf{r}) &= \langle \mathbf{w}_1^{[k]} \odot \mathbf{r}_1^{[k]}, \Psi_k \rangle \text{ with} \\ \Psi_i &= (\sigma(\langle \mathbf{w}_1^{[i-1]} \odot \mathbf{r}_1^{[i-1]}, \Psi_{i-1} \rangle), \dots, \sigma(\langle \mathbf{w}_{m_i}^{[i-1]} \odot \mathbf{r}_{m_i}^{[i-1]}, \Psi_{i-1} \rangle)) \end{aligned} \quad (9)$$

for $i \in [k]$ and $\Psi_0 = \mathbf{x}$. Here $\mathbf{r} = \{\mathbf{r}_1^{[k]}, \mathbf{r}_1^{[k-1]}, \dots, \mathbf{r}_{m_k}^{[k-1]}, \dots, \mathbf{r}_1^{[0]}, \dots, \mathbf{r}_{m_1}^{[0]}\}$, and for $0 \leq j \leq k$, $\mathbf{r}_i^{[j]}$ has the same size with Ψ_j , and each element in $\mathbf{r}_i^{[j]}$ is drawn i.i.d. from $\text{Bern}(\rho)$.

- The output for Drp^(III) (third type) is given by

$$\begin{aligned} f^{(III)}(\mathbf{w}, \mathbf{x}, \mathbf{r}) &= \langle \mathbf{w}_1^{[k]} \odot \mathbf{r}_1^{[k]}, \Psi_k \odot \mathbf{r}_2^{[k]} \rangle \text{ with} \\ \Psi_i &= (\sigma(\langle \mathbf{w}_1^{[i-1]} \odot \mathbf{r}_1^{[i-1]}, \Psi_{i-1} \odot \mathbf{r}_{m_i+1}^{[i-1]} \rangle), \dots, \sigma(\langle \mathbf{w}_{m_i}^{[i-1]} \odot \mathbf{r}_{m_i}^{[i-1]}, \Psi_{i-1} \odot \mathbf{r}_{m_i+1}^{[i-1]} \rangle)) \end{aligned} \quad (10)$$

for $i \in [k]$ and $\Psi_0 = \mathbf{x}$. Here $\mathbf{r} = (\mathbf{r}_1^{[k]}, \mathbf{r}_2^{[k]}, \dots, \mathbf{r}_1^{[0]}, \dots, \mathbf{r}_{m_1+1}^{[0]})$, and for $0 \leq j \leq k$, $\mathbf{r}_j^{[i]}$ has the same size with Ψ_j , and each element in $\mathbf{r}_j^{[i]}$ is drawn i.i.d. from $\text{Bern}(\rho)$.

Given a set \mathcal{W} , we denote

$$\mathcal{F}_{\mathcal{W}}^{(I)} = \{f^{(I)}(\mathbf{w}, \mathbf{x}, \mathbf{r}) : \mathbf{w} \in \mathcal{W}\} \quad (11)$$

$$\mathcal{F}_{\mathcal{W}}^{(II)} = \{f^{(II)}(\mathbf{w}, \mathbf{x}, \mathbf{r}) : \mathbf{w} \in \mathcal{W}\} \quad (12)$$

$$\mathcal{F}_{\mathcal{W}}^{(III)} = \{f^{(III)}(\mathbf{w}, \mathbf{x}, \mathbf{r}) : \mathbf{w} \in \mathcal{W}\} \quad (13)$$

where $f^{(I)}(\mathbf{w}, \mathbf{x}, \mathbf{r})$, $f^{(II)}(\mathbf{w}, \mathbf{x}, \mathbf{r})$ and $f^{(III)}(\mathbf{w}, \mathbf{x}, \mathbf{r})$ are defined in Eqs. 8-10.

We will focus on full-connected neural networks, both shallow ones (with none or one hidden layer) and deep ones (with more hidden layers).

4.1 Shallow Network without Hidden Layer

We first consider the shallow network without hidden layer, and therefore, the output is a linear function, i.e., $f(\mathbf{w}, \mathbf{x}) = \langle \mathbf{w}, \mathbf{x} \rangle$. Further, the outputs for $\text{Drp}^{(I)}$, $\text{Drp}^{(II)}$ and $\text{Drp}^{(III)}$ are given, respectively, by

$$\begin{aligned} f^{(I)}(\mathbf{w}, \mathbf{x}, \mathbf{r}) &= \langle \mathbf{w}, \mathbf{x} \odot \mathbf{r} \rangle \\ f^{(II)}(\mathbf{w}, \mathbf{x}, \mathbf{r}) &= \langle \mathbf{w} \odot \mathbf{r}, \mathbf{x} \rangle \\ f^{(III)}(\mathbf{w}, \mathbf{x}, (\mathbf{r}_1, \mathbf{r}_2)) &= \langle \mathbf{w} \odot \mathbf{r}_1, \mathbf{x} \odot \mathbf{r}_2 \rangle \end{aligned}$$

where \mathbf{r} , \mathbf{r}_1 and \mathbf{r}_2 are of size d , and each element in \mathbf{r} , \mathbf{r}_1 and \mathbf{r}_2 is drawn i.i.d. from $\text{Bern}(\rho)$. The following theorem shows the Rademacher complexity for three types of dropouts:

Theorem 2. Let $\mathcal{W} = \{\mathbf{w}: \|\mathbf{w}\| < B_1\}$, $\mathcal{X} = \{\mathbf{x}: \|\mathbf{x}\| \leq B_2\}$, and $\mathcal{F}_{\mathcal{W}}^{(I)}$, $\mathcal{F}_{\mathcal{W}}^{(II)}$ and $\mathcal{F}_{\mathcal{W}}^{(III)}$ are defined in Eqs. 11-13. Then, we have

$$\mathfrak{R}_n(\mathcal{F}_{\mathcal{W}}^{(1)}) = \mathfrak{R}_n(\mathcal{F}_{\mathcal{W}}^{(2)}) \leq B_1 B_2 \sqrt{\rho/n} \quad \text{and} \quad \mathfrak{R}_n(\mathcal{F}_{\mathcal{W}}^{(3)}) \leq B_1 B_2 \rho / \sqrt{n}.$$

If we do not drop out any weights and input units (corresponding to input features), i.e., $\rho = 1$, the above theorem gives a similar estimation for the Rademacher complexity of linear function space as stated in [28, Theorem 3]. Also, these complexities are independent to feature dimension, and thus can be applied to high-dimensional data. In addition, such result has independent interests in missing feature problems.

Proof. From $\langle \mathbf{w} \odot \mathbf{r}, \mathbf{x} \rangle = \langle \mathbf{w}, \mathbf{x} \odot \mathbf{r} \rangle$, it is easy to prove $\mathfrak{R}_n(\mathcal{F}_{\mathcal{W}}^{(I)}) = \mathfrak{R}_n(\mathcal{F}_{\mathcal{W}}^{(II)})$. For $S_n = \{\mathbf{x}_1, \dots, \mathbf{x}_n\}$ and $RS_n = \{\mathbf{r}_1, \dots, \mathbf{r}_n\}$, we have

$$\hat{\mathfrak{R}}_n(\mathcal{F}_{\mathcal{W}}^{(I)}, S_n, RS_n) = \frac{1}{n} E_{\epsilon} \sup_{\mathbf{w} \in \mathcal{W}} \sum_{i=1}^n \epsilon_i \langle \mathbf{w}, \mathbf{x}_i \odot \mathbf{r}_i \rangle$$

where $\epsilon = (\epsilon_1, \dots, \epsilon_n)$ are rademacher variables, and this yields that

$$\hat{\mathfrak{R}}_n(\mathcal{F}_{\mathcal{W}}^{(I)}, S_n, RS_n) = \frac{1}{n} E_{\epsilon} \sup_{\mathbf{w} \in \mathcal{W}} \left\langle \mathbf{w}, \sum_{i=1}^n \epsilon_i \mathbf{x}_i \odot \mathbf{r}_i \right\rangle.$$

By using the Cauchy-Schwartz inequality $\langle a, b \rangle \leq \|a\| \|b\|$ and $\|\mathbf{w}\| \leq B_1$, we have

$$\begin{aligned} \hat{\mathfrak{R}}_n(\mathcal{F}_{\mathcal{W}}^{(I)}, S_n, RS_n) &\leq \frac{B_1}{n} E_{\epsilon} \left\| \sum_{i=1}^n \epsilon_i \mathbf{x}_i \odot \mathbf{r}_i \right\| \\ &= \frac{B_1}{n} E_{\epsilon} \left(\sum_{i=1}^n \sum_{j=1}^n \epsilon_i \epsilon_j \langle \mathbf{x}_i \odot \mathbf{r}_i, \mathbf{x}_j \odot \mathbf{r}_j \rangle \right)^{1/2} \\ &\leq \frac{B_1}{n} \left(\sum_{i=1}^n \sum_{j=1}^n E_{\epsilon_i, \epsilon_j} \epsilon_i \epsilon_j \langle \mathbf{x}_i \odot \mathbf{r}_i, \mathbf{x}_j \odot \mathbf{r}_j \rangle \right)^{1/2} \end{aligned}$$

where the last inequality holds from Jensen's inequality. Since $E_{\epsilon_i, \epsilon_j} \epsilon_i \epsilon_j = 0$ for $i \neq j$ and $E_{\epsilon_i} \epsilon_i \epsilon_i = 1$ for rademacher variables, we have

$$\hat{\mathfrak{R}}_n(\mathcal{F}_{\mathcal{W}}^{(I)}, S_n, RS_n) \leq \frac{B_1}{n} \left(\sum_{i=1}^n \langle \mathbf{x}_i \odot \mathbf{r}_i, \mathbf{x}_i \odot \mathbf{r}_i \rangle \right)^{1/2}. \tag{14}$$

Based on the above inequality, it holds that

$$\begin{aligned} \mathfrak{R}_n(\mathcal{F}_{\mathcal{W}}^{(I)}) &= E_{S_n, RS_n} [\hat{\mathfrak{R}}_n(\mathcal{F}_{\mathcal{W}}^{(I)}, S_n, RS_n)] \\ &\leq \frac{B_1}{n} E_{S_n, RS_n} \left(\sum_{i=1}^n \langle \mathbf{x}_i \odot \mathbf{r}_i, \mathbf{x}_i \odot \mathbf{r}_i \rangle \right)^{1/2} \end{aligned}$$

$$\leq \frac{B_1}{n} E_{S_n} \left(\sum_{i=1}^n E_{\mathbf{r}_i} \langle \mathbf{x}_i \odot \mathbf{r}_i, \mathbf{x}_i \odot \mathbf{r}_i \rangle \right)^{1/2}$$

where the second inequality holds from Jensen's inequality. Finally, we have $\mathfrak{R}_n(\mathcal{F}_{\mathcal{W}}^{(I)}) \leq B_1 B_2 \sqrt{\rho/n}$ from Eq. 4 and $\|\mathbf{x}_i\| \leq B_2$.

In a similar manner, we have $\mathfrak{R}_n(\mathcal{F}_{\mathcal{W}}^{(III)}) \leq B_1 B_2 \rho / \sqrt{n}$ from $\langle \mathbf{w} \odot \mathbf{r}_1, \mathbf{x} \odot \mathbf{r}_2 \rangle = \langle \mathbf{w}, \mathbf{x} \odot \mathbf{r}_1 \odot \mathbf{r}_2 \rangle$ and Eq. 5. This theorem follows as desired.

4.2 Shallow Network with One Hidden Layer

We consider the shallow network with only one hidden layer, and assume that the hidden layer has m hidden units. The output for such network is given by

$$f(\mathbf{w}, \mathbf{x}) = \langle \mathbf{w}^{[1]}, \Psi(\mathbf{w}_1^{[0]}, \dots, \mathbf{w}_m^{[0]}, \mathbf{x}) \rangle$$

with

$$\Psi(\mathbf{w}_1^{[0]}, \dots, \mathbf{w}_m^{[0]}, \mathbf{x}) = (\sigma(\langle \mathbf{w}_1^{[0]}, \mathbf{x} \rangle), \dots, \sigma(\langle \mathbf{w}_m^{[0]}, \mathbf{x} \rangle)) \tag{15}$$

where $\mathbf{w} = (\mathbf{w}^{[1]}, \mathbf{w}_1^{[0]}, \dots, \mathbf{w}_m^{[0]})$, and $\mathbf{w}^{[1]}$ and $\mathbf{w}_i^{[0]}$ ($i \in [m]$) are of size m and d , respectively.

From Eqs. 8-10, the outputs for $\text{Drp}^{(I)}$, $\text{Drp}^{(II)}$ and $\text{Drp}^{(III)}$ are given, respectively, by

$$f^{(I)}(\mathbf{w}, \mathbf{x}, \mathbf{r}) = \langle \mathbf{w}^{[1]}, \mathbf{r}_1^{[1]} \odot \Psi(\mathbf{w}_1^{[0]}, \dots, \mathbf{w}_m^{[0]}, \mathbf{x} \odot \mathbf{r}_1^{[0]}) \rangle$$

$$f^{(II)}(\mathbf{w}, \mathbf{x}, \mathbf{r}) = \langle \mathbf{w}^{[1]} \odot \mathbf{r}_1^{[1]}, \Psi(\mathbf{w}_1^{[0]} \odot \mathbf{r}_1^{[0]}, \dots, \mathbf{w}_m^{[0]} \odot \mathbf{r}_m^{[0]}, \mathbf{x}) \rangle$$

and

$$f^{(III)}(\mathbf{w}, \mathbf{x}, \mathbf{r}) = \langle \mathbf{w}^{[1]} \odot \mathbf{r}_1^{[1]}, \mathbf{r}_2^{[1]} \odot \Psi(\mathbf{w}_1^{[0]} \odot \mathbf{r}_1^{[0]}, \dots, \mathbf{w}_m^{[0]} \odot \mathbf{r}_m^{[0]}, \mathbf{x} \odot \mathbf{r}_{m+1}^{[0]}) \rangle$$

where Ψ is defined in Eq. 15. Here $\mathbf{r}_i^{[1]}$ and $\mathbf{r}_j^{[0]}$ are of size m and d , respectively, and each element in $\mathbf{r}_i^{[1]}$ and $\mathbf{r}_j^{[0]}$ is drawn i.i.d. from $\text{Bern}(\rho)$. The following theorem shows the Rademacher complexity for three types of dropouts.

Theorem 3. Let $\mathcal{W} = \{(\mathbf{w}^{[1]}, \mathbf{w}_1^{[0]}, \dots, \mathbf{w}_m^{[0]}) : \|\mathbf{w}^{[1]}\|_1 \leq B_1, \|\mathbf{w}_i^{[0]}\| \leq B_0\}$, $\mathcal{X} = \{\mathbf{x} \in \mathbb{R}^d : \|\mathbf{x}\| \leq \hat{B}\}$ and $\mathcal{F}_{\mathcal{W}}^{(I)}$, $\mathcal{F}_{\mathcal{W}}^{(II)}$, $\mathcal{F}_{\mathcal{W}}^{(III)}$ are defined in Eqs. 11-13. Suppose that the activation σ is Lipschitz with constant L and $\sigma(0) = 0$. Then, we have

$$\mathfrak{R}_n(\mathcal{F}_{\mathcal{W}}^{(I)}) \leq \mathfrak{R}_n(\mathcal{F}_{\mathcal{W}}^{(II)}) \leq L B_1 B_0 \hat{B} \rho / \sqrt{n} \quad \text{and} \quad \mathfrak{R}_n(\mathcal{F}_{\mathcal{W}}^{(III)}) \leq L B_1 B_0 \hat{B} \rho^2 / \sqrt{n}.$$

Proof. We first have

$$\mathfrak{R}_n(\mathcal{F}_{\mathcal{W}}^{(I)}) \leq \mathfrak{R}_n(\mathcal{F}_{\mathcal{W}}^{(II)})$$

from $f^{(II)}(\mathbf{w}, \mathbf{x}, \mathbf{r}) = f^{(I)}(\mathbf{w}, \mathbf{x}, \mathbf{r}')$ by selecting $\mathbf{r}_1^{[0]} = \dots = \mathbf{r}_m^{[0]} = \mathbf{r}'^{[0]}$ and $\mathbf{r}^{[1]} = \mathbf{r}'^{[1]}$. In the following, we will estimate $\mathfrak{R}_n(\mathcal{F}_{\mathcal{W}}^{(II)})$.

Given $S_n = \{\mathbf{x}_1, \dots, \mathbf{x}_n\}$ and $RS_n = \{\mathbf{r}_1, \dots, \mathbf{r}_n\}$, it holds that

$$\begin{aligned} \hat{\mathfrak{R}}_n(\mathcal{F}_{\mathcal{W}}^{(II)}, S_n, RS_n) &= \frac{1}{n} E_\epsilon \left[\sup_{\mathbf{w}} \langle \mathbf{w}^{[1]}, \sum_{i=1}^n \epsilon_i \mathbf{r}_i^{[1]} \odot \Delta_i \rangle \right] \\ &\leq B_1 E_\epsilon \left[\sup_{\mathbf{w}} \left\langle \frac{\mathbf{w}^{[1]}}{\|\mathbf{w}^{[1]}\|_1}, \frac{1}{n} \sum_{i=1}^n \epsilon_i \mathbf{r}_i^{[1]} \odot \Delta_i \right\rangle \right] \end{aligned}$$

where $\Delta_i = \Psi(\mathbf{w}_1^{[0]} \odot \mathbf{r}_{i1}^{[0]}, \dots, \mathbf{w}_m^{[0]} \odot \mathbf{r}_{im}^{[0]}, \mathbf{x}_i)$ and Ψ is defined by Eq. 15 and the inequality holds from $\|\mathbf{w}^{[1]}\|_1 \leq B$. From Lemma 1, we have

$$\hat{\mathfrak{R}}_n(\mathcal{F}_{\mathcal{W}}^{(II)}, S_n, RS_n) \leq \frac{B_1}{n} E_\epsilon \sup_{\mathbf{w}_1^{[0]}} \sum_{i=1}^n \epsilon_i r_{i1}^{[1]} \sigma(\langle \mathbf{w}_1^{[0]} \odot \mathbf{r}_{i1}^{[0]}, \mathbf{x}_i \rangle). \tag{16}$$

From $\sigma(0) = 0$ and $r_{i1}^{[1]} \in \{0, 1\}$, we have $r_{i1}^{[1]}\sigma(t) = \sigma(r_{i1}^{[1]}t)$. Since $\sigma(\cdot)$ is Lipschitz with constant L , Lemma 2 gives

$$E_\epsilon \sup_{\mathbf{w}_1^{[0]}} \sum_{i=1}^n \epsilon_i r_{i1}^{[1]} \sigma(\langle \mathbf{w}_1^{[0]} \odot \mathbf{r}_{i1}^{[0]}, \mathbf{x}_i \rangle) \leq LE_\epsilon \sup_{\mathbf{w}_1^{[0]}} \sum_{i=1}^n \epsilon_i \langle \mathbf{w}_1^{[0]} \odot \mathbf{r}_{ij}^{[0]}, r_{i1}^{[1]} \mathbf{x}_i \rangle.$$

Similarly to the proof of Eq. 14, we have

$$E_\epsilon \sup_{\mathbf{w}_1^{[0]}} \sum_{i=1}^n \epsilon_i \langle \mathbf{w}_1^{[0]}, r_{i1}^{[1]} \mathbf{x}_i \odot \mathbf{r}_{i1}^{[0]} \rangle = B_0 \left(\sum_{i=1}^n \langle r_{i1}^{[1]} \mathbf{x}_i \odot \mathbf{r}_{i1}^{[0]}, r_{i1}^{[1]} \mathbf{x}_i \odot \mathbf{r}_{i1}^{[0]} \rangle \right)^{\frac{1}{2}}.$$

Combining with the previous analysis, we have

$$\begin{aligned} \mathfrak{R}_n(\mathcal{F}_{\mathcal{W}}^{(\text{II})}) &= E[\hat{\mathfrak{R}}_n(\mathcal{F}_{\mathcal{W}}^{(\text{II})}, S_n, RS_n)] \\ &\leq \frac{LB_1B_0}{n} E_{S_n, RS_n} \left(\sum_{i=1}^n \langle r_{i1}^{[1]} \mathbf{x}_i \odot \mathbf{r}_{i1}^{[0]}, r_{i1}^{[1]} \mathbf{x}_i \odot \mathbf{r}_{i1}^{[0]} \rangle \right)^{\frac{1}{2}} \\ &\leq \frac{LB_1B_0}{n} E_{S_n} \sum_{i=1}^n E_{RS_n} \langle r_{i1}^{[1]} \mathbf{x}_i \odot \mathbf{r}_{i1}^{[0]}, r_{i1}^{[1]} \mathbf{x}_i \odot \mathbf{r}_{i1}^{[0]} \rangle \leq LB_1B_0\hat{B}\rho/\sqrt{n}, \end{aligned}$$

where the last inequality holds from Eq. 6 and $\|\mathbf{x}_i\| \leq \hat{B}$.

In a similar way, we can prove $\mathfrak{R}_n(\mathcal{F}_{\mathcal{W}}^{(\text{III})}) \leq B_1B_0\hat{B}\rho^2/\sqrt{n}$ by combining with Eqs. 6-7, and $\langle \mathbf{w} \odot \mathbf{r}_1, \mathbf{x} \odot \mathbf{r}_2 \rangle = \langle \mathbf{w}, \mathbf{x} \odot \mathbf{r}_1 \odot \mathbf{r}_2 \rangle$. This completes the proof.

4.3 Deep Network with k Hidden Layers

Now we consider the neural network with k ($k \geq 1$) hidden layers, and the i th layer has m_i hidden units ($i \in [k]$). The output for this neural network is given by

$$\begin{aligned} f(\mathbf{w}, \mathbf{x}) &= \langle \mathbf{w}_1^{[k]}, \Psi_k \rangle \text{ with } \Psi_0 = \mathbf{x}, \text{ and for } i \in [k] \\ \Psi_i &= (\sigma(\langle \mathbf{w}_1^{[i-1]}, \Psi_{i-1} \rangle), \dots, \sigma(\langle \mathbf{w}_{m_i}^{[i-1]}, \Psi_{i-1} \rangle)), \end{aligned}$$

and three types of dropouts $\text{Drp}^{(\text{I})}$, $\text{Drp}^{(\text{II})}$ and $\text{Drp}^{(\text{III})}$ are defined by Eqs. 8-10. The following theorem shows the Rademacher complexity for three types of dropouts:

Theorem 4. Let $\mathcal{W} = \{(\mathbf{w}_1^{[k]}, \mathbf{w}_1^{[k-1]}, \dots, \mathbf{w}_{m_k}^{[k-1]}, \dots, \mathbf{w}_1^{[0]}, \dots, \mathbf{w}_{m_2}^{[0]}) : \|\mathbf{w}_i^{[0]}\| \leq B_0, \|\mathbf{w}_i^{[j]}\|_1 \leq B_j \text{ for } j \geq 1\}$, $\mathcal{X} = \{\mathbf{x} \in \mathbb{R}^d : \|\mathbf{x}\| \leq \hat{B}\}$, and $\mathcal{F}_{\mathcal{W}}^{(\text{I})}$, $\mathcal{F}_{\mathcal{W}}^{(\text{II})}$, and $\mathcal{F}_{\mathcal{W}}^{(\text{III})}$ are defined in Eqs. 11-13. Suppose that the activation σ is Lipschitz with constant L and $\sigma(0) = 0$. Then, we have

$$\mathfrak{R}_n(\mathcal{F}_{\mathcal{W}}^{(\text{I})}) \leq \mathfrak{R}_n(\mathcal{F}_{\mathcal{W}}^{(\text{II})}) \leq \frac{\rho^{(k+1)/2}}{\sqrt{n}} L^k \hat{B} \prod_{j=0}^k B_j \quad \text{and} \quad \mathfrak{R}_n(\mathcal{F}_{\mathcal{W}}^{(\text{III})}) \leq \frac{\rho^{(k+1)}}{\sqrt{n}} L^k \hat{B} \prod_{j=0}^k B_j.$$

Here the Lipschitz constant L is dependent on the activation function, e.g., $L \leq 1$ if we choose the center sigmoid and relu [34] as activation function, and this follows $L\rho < 1$ for $\rho < 1$, which could improve the generalization bounds.

This theorem shows that dropout can lead to an exponential reduction of the Rademacher complexity with respect to the number of hidden layers within neural network. If we do not drop out any weights and units (including hidden units, or input units corresponding to input features), i.e., $\rho = 1$, the above theorem improves the results in [35, Lemma 26]. The Rademacher complexities are dependent on the norms of weights, but irrelevant to the number of units and weights in the network, as well as the dimension of input datasets.

Previous empirical studies showed that dropout tends to be resistant to overfitting in many real applications [7, 15, 16, 19]. This theoretical result discloses the reason that dropout makes an exponential reduction of the Rademacher complexity, i.e., dropout can reduce the model complexity so as to avoid overfitting.

Finally, this theorem is also consistent with the empirical experience of training deep network¹⁾, e.g., multi-layer network with multi-layers of dropout performs not worse than that of single layer of dropout. An interesting open problem is to study the convergence rate of learning algorithms such as SGD for multi-layer dropout network.

Proof. We first have

$$\mathfrak{R}_n(\mathcal{F}_{\mathcal{W}}^{(I)}) \leq \mathfrak{R}_n(\mathcal{F}_{\mathcal{W}}^{(II)})$$

from $f^{(II)}(\mathbf{w}, \mathbf{x}, \mathbf{r}) = f^{(I)}(\mathbf{w}, \mathbf{x}, \mathbf{r}')$ by selecting $\mathbf{r}_1^{[i]} = \dots = \mathbf{r}_{m_{i+1}}^{[i]} = \mathbf{r}'^{[i]}$ for $0 \leq i \leq k-1$ and $\mathbf{r}_1^{[k]} = \mathbf{r}'^{[k]}$.

For any $S_n = \{\mathbf{x}_1, \dots, \mathbf{x}_n\}$ and $RS_n = \{\mathbf{r}_1, \dots, \mathbf{r}_n\}$, we will prove that

$$\hat{\mathfrak{R}}_n(\mathcal{F}_{\mathcal{W}}^{(II)}, S_n, RS_n) \leq \frac{L^k}{n} E_\epsilon \left[\sup_{\mathbf{w}_1^{[0]}} \sum_{i=1}^n \epsilon_i \langle \mathbf{w}_1^{[0]} \odot \mathbf{r}_{i,1}^{[0]}, \mathbf{x}_i \prod_{s=1}^k r_{i,j_{s+1},j_s}^{[s]} \rangle \right] \prod_{j=1}^k B_j \quad (17)$$

by induction on k , i.e., the number of layers in neural network, where $j_{k+1} = 1$. It is easy to find the above holds for $k = 1$ from Eq. 16. Assume that Eq. 17 holds for neural network of $k-1$ layers ($k \geq 2$), and in the following we will prove for neural network of k layers.

For $\|\mathbf{w}_1^{[k]}\|_1 \leq B_k$, we have

$$\begin{aligned} \hat{\mathfrak{R}}_n(\mathcal{F}_{\mathcal{W}}^{(II)}, S_n, RS_n) &= \frac{1}{n} E_\epsilon \sup_{\mathbf{w}} \sum_{i=1}^n \epsilon_i \langle \mathbf{w}_1^{[k]} \odot \mathbf{r}_{i1}^{[k]}, \Psi_{ik} \rangle \\ &= \frac{1}{n} E_\epsilon \sup_{\mathbf{w}} \left\langle \mathbf{w}_1^{[k]}, \sum_{i=1}^n \epsilon_i \Psi_{ik} \odot \mathbf{r}_{i1}^{[k]} \right\rangle \leq \frac{B_k}{n} E_\epsilon \sup_{\mathbf{w}} \left\langle \frac{\mathbf{w}_1^{[k]}}{\|\mathbf{w}_1^{[k]}\|_1}, \sum_{i=1}^n \epsilon_i \Psi_{ik} \odot \mathbf{r}_{i1}^{[k]} \right\rangle \end{aligned}$$

where $\Psi_{ij} = (\sigma(\langle \mathbf{w}_1^{[j-1]} \odot \mathbf{r}_{i1}^{[j-1]}, \Psi_{i,j-1} \rangle), \dots, \sigma(\langle \mathbf{w}_{m_j}^{[j-1]} \odot \mathbf{r}_{m_j}^{[j-1]}, \Psi_{i,j-1} \rangle))$ for $j \in [k]$ and $\Psi_{i0} = \mathbf{x}_i$. From Lemma 1, we have

$$E_\epsilon \sup_{\mathbf{w}} \left\langle \frac{\mathbf{w}_1^{[k]}}{\|\mathbf{w}_1^{[k]}\|_1}, \sum_{i=1}^n \epsilon_i \Psi_{ik} \odot \mathbf{r}_{i1}^{[k]} \right\rangle \leq E_\epsilon \sup_{\mathbf{w}} \sum_{i=1}^n \epsilon_i r_{i,1,1}^{[k]} \times \sigma(\langle \mathbf{w}_1^{[k-1]} \odot \mathbf{r}_{i,1}^{[k-1]}, \Psi_{i,k-1} \rangle),$$

which yields that

$$\hat{\mathfrak{R}}_n(\mathcal{F}_{\mathcal{W}}^{(II)}, S_n, RS_n) \leq \frac{B_k}{n} E_\epsilon \sup_{\mathbf{w}} \sum_{i=1}^n \epsilon_i r_{i,1,1}^{[k]} \times \sigma(\langle \mathbf{w}_1^{[k-1]} \odot \mathbf{r}_{i,1}^{[k-1]}, \Psi_{i,k-1} \rangle)$$

Since $\sigma(0) = 0$ and σ is Lipschitz with constant L , Lemma 2 gives

$$\hat{\mathfrak{R}}_n(\mathcal{F}_{\mathcal{W}}^{(II)}, S_n, RS_n) \leq \frac{LB_k}{n} E_\epsilon \sup_{\mathbf{w}} \sum_{i=1}^n \epsilon_i r_{i,1,1}^{[k]} \langle \mathbf{w}_1^{[k-1]} \odot \mathbf{r}_{i,1}^{[k-1]}, \Psi_{i,k-1} \rangle. \quad (18)$$

By using $r_{i,1,1}^{[k]} \sigma(t) = \sigma(r_{i,1,1}^{[k]} t)$, the term

$$\frac{1}{n} E_\epsilon \sup_{\mathbf{w}} \sum_{i=1}^n \epsilon_i r_{i,1,1}^{[k]} \langle \mathbf{w}_1^{[k-1]} \odot \mathbf{r}_{i,1}^{[k-1]}, \Psi_{i,k-1} \rangle$$

can be viewed as the empirical Rademacher for another $k-1$ layers neural network with respect to sample $S'_n = \{\mathbf{x}_1 r_{i,1,1}^{[k]}, \dots, \mathbf{x}_n r_{i,1,1}^{[k]}\}$ and $RS'_n = (\mathbf{r}_1^{[k-1]}, \mathbf{r}_1^{[k-2]}, \dots, \mathbf{r}_{m_{k-2}}^{[k-2]}, \mathbf{r}_1^{[0]}, \dots, \mathbf{r}_{m_1}^{[0]})$. Therefore, by our assumption that Eq. 17 holds for any $k-1$ layers neural network, we have

$$\begin{aligned} &\frac{1}{n} E_\epsilon \sup_{\mathbf{w}} \sum_{i=1}^n \epsilon_i r_{i,1,1}^{[k]} \langle \mathbf{w}_1^{[k-1]} \odot \mathbf{r}_{i,1}^{[k-1]}, \Psi_{i,k-1} \rangle \\ &\leq \frac{L^{k-1} \prod_{i=1}^{k-1} B_i}{n} E_\epsilon \sum_{i=1}^n \epsilon_i \langle \mathbf{w}_1^{[0]} \odot \mathbf{r}_{i,1}^{[0]}, \mathbf{x}_i \prod_{s=1}^{k-1} r_{i,j_{s+1},j_s}^{[s]} \rangle \end{aligned}$$

1) This is devoted to one reviewer.

which proves that Eq. 17 holds for k layers neural network by combining with Eq.18.

Similarly to the proof of Theorem 2, the term in Eq. 17 can be further bounded by

$$\begin{aligned} & E_{\epsilon} \sup_{\mathbf{w}_1^{[0]}} \sum_{i=1}^n \epsilon_i \langle \mathbf{w}_1^{[0]} \odot \mathbf{r}_{i,1}^{[0]}, \mathbf{x}_i \prod_{s=1}^k r_{i,j_{s+1},j_s}^{[s]} \rangle \\ & \leq B_0 \left(\sum_i \langle \mathbf{r}_{i,1}^{[0]} \odot \mathbf{x}_i \prod_{s=1}^k r_{i,j_{s+1},j_s}^{[s]}, \mathbf{r}_{i,1}^{[0]} \odot \mathbf{x}_i \prod_{s=1}^k r_{i,j_{s+1},j_s}^{[s]} \rangle \right)^{\frac{1}{2}} \end{aligned}$$

which yields that, from Eq. 6,

$$\mathfrak{R}_n(\mathcal{F}_{\mathcal{W}}^{(II)}) \leq \frac{1}{\sqrt{n}} L^k \rho^{(k+1)/2} \hat{B} \prod_{i=0}^k B_i.$$

In a similar manner, we can prove that

$$\mathfrak{R}_n(\mathcal{F}_{\mathcal{W}}^{(III)}) \leq \frac{1}{\sqrt{n}} L^k \rho^{k+1} \hat{B} \prod_{i=0}^k B_i.$$

by using Eq. 7, and this completes the proof.

5 Conclusion

Deep neural networks have witnessed great successes in various real applications. Many implementation techniques have been developed, however, theoretical understanding of many aspects of deep neural networks is far from clear. Dropout is an effective strategy to improve the performance as well as reduce the influence of overfitting during training of deep neural network, and it is motivated from the intuition of preventing complex co-adaptation of feature detectors. In this work, we study the Rademacher complexity of different types of dropouts, and our theoretical results disclose that for shallow neural networks (with one or none hidden layer) dropout is able to reduce the Rademacher complexity in polynomial, whereas for deep neural networks it can amazingly lead to an exponential reduction of the Rademacher complexity. An interesting future work is to present tighter generalization bounds for dropouts. In this work, we focused on very fundamental types of dropouts. Analyzing other types of dropouts is another interesting issue for future work, and we believe that the current work sheds a light on the way for the analysis.

Acknowledgements

This work was supported by the National Key Basic Research Program of China (2014CB340501), the National Science Foundation of China (61333014) and the Jiangsu Science Foundation (BK20150586).

References

- 1 Hinton G E, Salakhutdinov R R. Reducing the dimensionality of data with neural networks. *Science*, 2006, 313(5786): 504–507
- 2 Cireřan D C, Meier U, Gambardella L M, et al. Deep, big, simple neural nets for handwritten digit recognition. *Neural Computation*, 2010, 22(12): 3207–3220
- 3 Cireřan D C, Meier U, Schmidhuber J. Multi-column deep neural networks for image classification. In: *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*. Providence, RI, 2012. 3642–3649
- 4 Coates A, Huval B, Wang T, et al. Deep learning with COTS HPC systems. In: *Proceedings of the 30th International Conference on Machine Learning*. Atlanta, GA, 2013. 1337–1345
- 5 Dahl G E, Sainath T N, Hinton G E. Improving deep neural networks for lvcsr using rectified linear units and dropout. In: *Proceeding of IEEE International Conference on Acoustics, Speech and Signal Processing*. Vancouver, Canada, 2013. 8609–8613
- 6 Dahl G E, Yu D, Deng L, et al. Context-dependent pre-trained deep neural networks for large-vocabulary speech recognition. *IEEE Transactions on Audio, Speech, and Language Processing*, 2012, 20(1): 30–42

- 7 Hinton G E, Deng L, Yu D. Deep neural networks for acoustic modeling in speech recognition: The shared views of four research groups. *IEEE Signal Processing Magazine*, 2012, 29(6): 82–97
- 8 Bo L, Lai K, Ren X, et al. Object recognition with hierarchical kernel descriptors. In: *Proceeding of IEEE Conference on Computer Vision and Pattern Recognition*. Colorado Springs, CO, 2011. 1729–1736
- 9 Mobahi H, Collobert R, Weston J. Deep learning from temporal coherence in video. In: *Proceedings of the 26th International Conference on Machine Learning*. Montreal, Canada, 2009. 737–744
- 10 Liu Y J, Liu L, Tong S C. Adaptive neural network tracking design for a class of uncertain nonlinear discrete-time systems with dead-zone. *Science China Information Sciences*, 2014, 57(3): 1–12
- 11 Andreas S W, David E R, Bernardo A H. Generalization by weight-elimination with application to forecasting. *Advances in Neural Information Processing Systems 3*. Cambridge: MIT Press, 1991. 875–882
- 12 Amari S, Murata N, Muller K, et al. Asymptotic statistical theory of overtraining and cross-validation. *IEEE Transactions on Neural Networks*, 1997, 8(5): 985–996
- 13 Neal R M. *Bayesian Learning for Neural Networks*. Lecture Notes in Statistics. Springer, 1996
- 14 Hinton G E, Srivastava N, Krizhevsky A, et al. Improving neural networks by preventing co-adaptation of feature detectors. 2012, CoRR/abstract: 1207.0580
- 15 Krizhevsky A, Sutskever I, Hinton G E. Imagenet classification with deep convolutional neural networks. In: Bartlett P L, Pereira F C N, Burges C J C, et al., eds. *Advances in Neural Information Processing Systems 25*. Cambridge, MA: MIT Press, 2012. 1106–1114
- 16 Wan L, Zeiler M, Zhang S, et al. Regularization of neural networks using dropconnect. In: *Proceedings of the 30th International Conference on Machine Learning*. Atlanta, GA, 2013. 1058–1066
- 17 Wang S I, Manning C D. Fast dropout training. In: *Proceedings of the 30th International Conference on Machine Learning*. Atlanta, GA, 2013. 118–126
- 18 Ba J, Frey, B. Adaptive dropout for training deep neural networks. In: Burges C J C, Bottou L, Ghahramani Z, et al., eds. *Advances in Neural Information Processing Systems 26*. Cambridge, MA: MIT Press, 2013. 3084–3092
- 19 Srivastava N, Hinton G, Krizhevsky A. Dropout: A simple way to prevent neural networks from overfitting. *Journal of Machine Learning Research*, 2014, 15(1): 1929–1958
- 20 Baldi P, Sadowski P J. The dropout learning algorithm. *Artificial Intelligence*, 2014, 210:78–122
- 21 Wager S, Wang S, Liang P. Dropout training as adaptive regularization. In: Burges C J C, Bottou L, Ghahramani Z, et al., eds. *Advances in Neural Information Processing Systems 26*. Cambridge, MA: MIT Press, 2013. 351–359
- 22 McAllester D. A pac-bayesian tutorial with a dropout bound. 2013, CoRR/abstract 1307.2118
- 23 Karpinski M, Macintyre A. Polynomial bounds for VC dimension of sigmoidal and general pfaffian neural networks. *Journal of Computer and System Sciences*, 1997, 54(1): 169–176
- 24 Anthony M, Bartlett P L. *Neural network learning: Theoretical foundations*. Cambridge: Cambridge University Press, 2009
- 25 Bartlett P L, Mendelson S. Rademacher and gaussian complexities: Risk bounds and structural results. *Journal of Machine Learning Research*, 2002, 3:463–482
- 26 McDiarmid C. On the method of bounded differences. *Surveys in Combinatorics*. Cambridge: Cambridge University Press, 1989. 148–188
- 27 Ledoux M, Talagrand, M. *Probability in Banach spaces: Isoperimetry and processes*. Springer 2002
- 28 Kakade S M, Sridharan K, Tewari A. On the complexity of linear prediction: Risk bounds, margin bounds, and regularization. In: Koller D, Schuurmans D, Bengio Y, et al., eds. *Advances in Neural Information Processing Systems 24*. Cambridge, MA: MIT Press, 2008. 351–359
- 29 Koltchinskii V, Panchenko D. Empirical margin distributions and bounding the generalization error of combined classifiers. *Annals of Statistics*, 2002, 30(1): 1–50
- 30 Cortes C, Mohri M, Rostamizadeh A. Generalization bounds for learning kernels In: *Proceedings of the 27th International Conference on Machine Learning*. Haifa, Israel, 2010. 247–254
- 31 Maurer A. Bounds for linear multi-task learning. *Journal of Machine Learning Research*, 2006 7: 117–139
- 32 Meir R, Zhang T. Generalization error bounds for bayesian mixture algorithms. *Journal of Machine Learning Research*, 2003, 4: 839–860
- 33 Zou B, Peng Z. M, Xu Z. B. The learning performance of support vector machine classification based on Markov sampling. *Science China Information Sciences*, 2013, 56(3): 1–16
- 34 Nair V, Hinton G E. Rectified linear units improve restricted boltzmann machines. In: *Proceedings of the 27th International Conference on Machine Learning*. Haifa, Israel, 2010. 807–814
- 35 Bartlett P L. The sample complexity of pattern classification with neural networks: the size of the weights is more important than the size of the network. *IEEE Transactions on Information Theory*, 1998, 44(2): 525–536