

Understanding Representation Learnability of Nonlinear Self-Supervised Learning

Ruofeng Yang¹, Xiangyuan Li¹, Bo Jiang¹, Shuai Li^{1*}

¹Shanghai Jiao Tong University

wanshuiyin@sjtu.edu.cn, lixiangyuan19@sjtu.edu.cn, bjiang@sjtu.edu.cn, shuaili8@sjtu.edu.cn

Abstract

Self-supervised learning (SSL) has empirically shown its data representation learnability in many downstream tasks. There are only a few theoretical works on data representation learnability, and many of those focus on final data representation, treating the nonlinear neural network as a “black box”. However, the accurate learning results of neural networks are crucial for describing the data distribution features learned by SSL models. Our paper is the first to analyze the learning results of the nonlinear SSL model accurately. We consider a toy data distribution that contains two features: the label-related feature and the hidden feature. Unlike previous linear setting work that depends on closed-form solutions, we use the gradient descent algorithm to train a 1-layer nonlinear SSL model with a certain initialization region and prove that the model converges to a local minimum. Furthermore, different from the complex iterative analysis, we propose a new analysis process which uses the **exact version of Inverse Function Theorem** to accurately describe the features learned by the local minimum. With this local minimum, we prove that the nonlinear SSL model can capture the label-related feature and hidden feature at the same time. In contrast, the nonlinear supervised learning (SL) model can only learn the label-related feature. We also present the learning processes and results of the nonlinear SSL and SL model via simulation experiments.

1 Introduction

In recent years, self-supervised learning has become an important paradigm in machine learning because it can use datasets without expensive target labels to learn useful data representations for many downstream tasks (Devlin et al. 2018; Radford et al. 2019; Wu et al. 2020).

At present, contrastive learning, a common self-supervised learning method, has shown superior performance in learning data representations and outperformed supervised learning in some downstream tasks (He et al. 2020; Chen and He 2021; Grill et al. 2020; Caron et al. 2020; Wang et al. 2022). Contrastive learning methods usually form a dual pair of siamese networks (Bromley et al. 1993) and use data augmentations for each datapoint. They treat two augmented datapoints of the same datapoint as positive pairs and maximize the similarity between positive pairs to learn data representations.

However, the siamese networks often collapse to a trivial solution during the training process, rendering the learned representation meaningless.

To avoid the above problem, earlier contrastive learning methods such as MoCo (He et al. 2020) and SimCLR (Chen et al. 2020) treat augmented datapoints from different datapoints as negative pairs and prevent model collapse by the trade-off between positive and negative pairs. However, obtaining high-quality negative pairs is difficult (Khosla et al. 2020), which in turn requires additional changes to the model. Recently, other classes of the SSL model, such as BYOL (Grill et al. 2020) and SimSiam (Chen and He 2021), which do not use negative pairs, have been studied. These models will not collapse to a trivial solution because they construct subtle asymmetry in the structure of the siamese network and create a dynamic buffer area (Tian, Chen, and Ganguli 2021). SimSiam further simplifies the structure of BYOL and only retains the core asymmetry. The simplified model makes training and analysis more convenient while obtaining competitive and meaningful data representations.

Despite the empirical success of SSL (He et al. 2020; Chen et al. 2020; Chen and He 2021; Zhong et al. 2022), there are only a few works that focus on data representation learnability (Arora et al. 2019; Tosh, Krishnamurthy, and Hsu 2021; Lee et al. 2021; HaoChen et al. 2021, 2022; Tian 2022a,b; Wen and Li 2021; Liu et al. 2021). However, studying the learnability is helpful in understanding why SSL models can obtain meaningful data representations. Many of the above works used final data representation to study the data representation learnability. Arora et al. (2019) obtained the data representation function by minimizing the empirical SSL loss in a special data representation function class. HaoChen et al. (2021) and HaoChen et al. (2022) studied final data representation by closed-form solutions. They viewed the nonlinear neural network as a “black box” and ignored the learning result of the nonlinear neural network. Thus their results do not describe the features accurately captured by SSL models and explain the encoding process of neural networks.

Wen and Li (2021) and Tian et al. (2020) tried to understand the learning results of nonlinear SSL models by analyzing a relatively overparameterized neural network. However, their results do not provide an accurate answer to whether SSL models could exactly capture the important features of data distribution or just capture a mixture of features.

*Corresponding author

Copyright © 2023, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

Liu et al. (2021) studied the learning results of SSL models, and it is the most relevant work to us. They proved that SSL models could learn the label-related features and hidden features at the same time. However, their work is a linear framework, and their results depend on the closed-form solutions of the learning results. When considering a nonlinear SSL model, we can not get closed-form solutions due to the nonconvexity. Therefore, which features can be exactly learned by nonlinear SSL models remains an important open question. We need a new analysis process to analyze the specific learning results of the nonlinear SSL model.

In this work, for the first time, we use gradient descent to train a **nonlinear SSL model** and analyze the data representation learnability by using the learning results of neural networks. We accurately describe the data distribution features captured by the SSL model. Specifically, we accomplish:

1. With a designed data distribution, we use gradient descent (GD) to train a 1-layer nonlinear SSL model and prove that the model can converge to a local minimum under a certain initialization region. Using locally strong convexity, we also obtain the convergence rate of the algorithm.
2. We describe the properties of the local minimum using **the exact version of Inverse Function Theorem**. Using these properties, we prove that the SSL model learns the label-related feature and hidden feature at the same time.
3. We prove that the nonlinear SL model can only learn the label-related feature. In other words, SSL is superior to SL in learning data representation. We verify the correctness of the above results through simulation experiments.

2 Related Work

Theoretical analyses for final data representation. For the analysis of the data representation learnability, many works focus on the final data representation (the optimal solution of the pretext task) and measure the quality of the final data representation in the downstream tasks by using a linear classifier (HaoChen et al. 2021, 2022; Arora et al. 2019; Lee et al. 2021; Tosh, Krishnamurthy, and Hsu 2021). The main difference in this line of work is how to obtain the final data representation. Arora et al. (2019) assumed that the data representation function class contains a function with low SSL loss and minimized the empirical SSL loss in this class. HaoChen et al. (2021) constructed the population positive-pair graph with augmented datapoints as vertices and the correlation of augmented datapoints as edge weights. Then they proved that the closed-form solutions of the data representations are approximately equivalent to the eigenvectors of the adjacency matrix of the above graph. Lee et al. (2021) used the nonlinear canonical correlation analysis (CCA) method to obtain the final data representation. The above works viewed the nonlinear neural network as a “black box” and ignored the learning results of the neural network. However, the learning results are crucial for analyzing which features are exactly captured by SSL methods. Hence we need to propose a new method to analyze the learning results.

Theoretical analyses for learning results of SSL. Liu et al. (2021) analyzed the learning results of SSL methods.

With a 1-layer linear SSL model, similar to SimSiam, they demonstrated that the SSL models could learn label-related and hidden features simultaneously. Because of the linear structure and the objective function with a designed quartic regularization, they can directly obtain the closed-form solutions of the learning results by using spectral decomposition of the matrix related to the data distribution. Tian (2022a) and Tian (2022b) dealt with the learning results of the nonlinear SSL model by analyzing an objective function similar to traditional Principal Component Analysis (PCA). However, their results were extended by a hidden neuron. Hence their results can not definitively answer which data features are captured by the model and which are ignored. Wen and Li (2021) and Tian et al. (2020) tried to understand the learning results of the nonlinear SSL by using stochastic gradient descent (SGD). However, their results relied heavily on special data augmentation and relatively overparameterized neural networks. Furthermore, their results only showed that with a large number of neurons, the neural networks contain all data features. They did not accurately characterize the learning result of each neuron. In other words, these results did not show the features exactly captured by the SSL methods.

Theoretical guarantees for supervised learning. For the analysis of the supervised learning, researchers focus on (1) How to characterize the landscape of the objective function; (2) How to converge to the local minima through algorithms (such as GD and SGD); (3) How fast the algorithm converges to the local minimum (Allen-Zhu, Li, and Song 2019; Du et al. 2017; Li and Yuan 2017; Brutzkus and Globerson 2017; Du et al. 2019). Hence they focus on characterizing the relationship between the objective function and its gradient and less on the specific form or the properties of local minima. However, the specific forms of local minima are helpful to determine whether SSL methods can capture important data distribution features.

3 Problem Formulation

In this section, we introduce the data distribution and the nonlinear SSL and SL model to be studied in this paper.

3.1 Data Distribution

The classification problem is a typical downstream task in machine learning, which can be used to measure the quality of data representation. We start with a simple binary classification and want to explore the differences in the data representations learned by SSL and SL models.

To train models, we first build the data distribution. In most cases, the data distribution contains not only label-related features but also some hidden features. These hidden features may not be helpful for the current task but may be useful for other downstream tasks. We want to determine whether the nonlinear SSL models capture hidden features, resulting in a richer data representation. At the same time, we also wonder whether the SL models only learn label-related features.

For the simplicity of analysis, we consider the label-related features as a group, represented by the feature e_1 . We also use e_2 to represent the hidden features. Inspired by previous work (Liu et al. 2021), which solved the above question in the linear

setting, we construct a data distribution containing four kinds of datapoints. The number of these four kinds of datapoints are n_1, n_2, n_3, n_4 and $n = n_1 + n_2 + n_3 + n_4$. Every time we generate a datapoint, we draw among the four kinds of datapoints with a probability of $1/4$, which means $\mathbb{E}[n_l] = n/4, \forall l \in [4]$. Let $\tau > 1, \rho > 0$ are two hyperparameters of the data distribution and $\xi_1, \dots, \xi_n \in \mathbb{R}^d$ are datapoint noise terms sampled from a Gaussian distribution $\mathcal{N}(0, I)$. Define

$$\begin{aligned} \mathcal{D}_1 &= \{x_i | x_i = e_1 + \rho \xi_i\}_{i=1}^{n_1}, \\ \mathcal{D}_2 &= \{x_i | x_i = e_1 + \tau e_2 + \rho \xi_i\}_{i=n_1+1}^{n_1+n_2}, \\ \mathcal{D}_3 &= \{x_i | x_i = -e_1 + \rho \xi_i\}_{i=n_1+n_2+1}^{n_1+n_2+n_3}, \\ \mathcal{D}_4 &= \{x_i | x_i = -e_1 + \tau e_2 + \rho \xi_i\}_{i=n_1+n_2+1}^n, \end{aligned} \quad (1)$$

as the datasets of four kinds of datapoints, where $e_1, e_2 \in \mathbb{R}^d$ are two orthogonal unit-norm vectors. Then, the data distribution in this paper is $\mathcal{D} = \mathcal{D}_1 \cup \mathcal{D}_2 \cup \mathcal{D}_3 \cup \mathcal{D}_4$.

Because labels are required during the SL model training process, we modify the data distribution. Specifically, we denote the class label by $y = \{0, 1\}$. When $x_i \in \mathcal{D}_1 \cup \mathcal{D}_2, y = 0$, otherwise $y = 1$. After the above steps, we obtain the data distribution \mathcal{D}^{SL} of the nonlinear SL model.

It is clear that the binary classification task can be completed using only the representative label-related feature e_1 . However, since $\tau \geq 1, e_2$ is also an important hidden feature.

Although this data distribution is a toy setting, it is sufficient to distinguish the learnability of the SSL and SL models. This data distribution is also representative. In Sec 4.2, we explain that the proof process can be easily extended to a more general data distribution containing many label-related and hidden features.

3.2 Model

In this section, we introduce the activation function and then the nonlinear SSL and SL model.

In this paper, we analyze nonlinear models, so it is necessary to introduce activation functions. We discuss two activation functions: sigmoid function $\sigma(x) = \frac{1}{1+e^{-x}}$ and tanh function $\sigma_2(x) = \frac{e^x - e^{-x}}{e^x + e^{-x}}$.

The SSL model. We focus on a variant of SimSiam (Chen and He 2021). SimSiam has shown impressive performance in various downstream experiments using only positive pairs and has become a representative SSL model. Fig. 1 shows the structure of the model in this paper. The datapoint x_i is augmented by data augmentation ξ_{aug} and ξ'_{aug} to obtain two augmented datapoints x'_i and x''_i . The data representations z'_i and z''_i of x'_i and x''_i are obtained through the nonlinear encoder $\sigma(Wx'_i)$ and $\sigma(Wx''_i)$. We use inner product $\langle z'_i, z''_i \rangle$ to measure the similarity between z'_i and z''_i . Tian, Chen, and Ganguli (2021) showed that a regularizer is essential for the existence of the non-collapsed solution. Hence, $\alpha \|W\|_F^2$ is added in L . The objective function L is defined as

$$\begin{aligned} \min_W L &= \alpha \|W\|_F^2 \\ &- \frac{1}{n} \sum_{i=1}^n \mathbb{E}_{\xi_{\text{aug}}, \xi'_{\text{aug}}} \left[\langle \sigma(W(x_i + \xi_{\text{aug}})), \sigma(W(x_i + \xi'_{\text{aug}})) \rangle \right], \end{aligned} \quad (2)$$

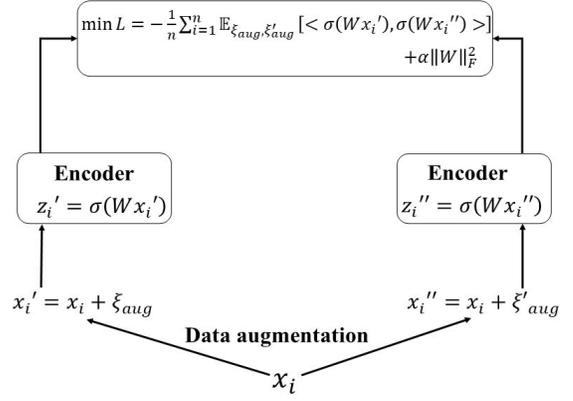


Figure 1: The structure of SSL model.

where α is the coefficient of regularizer, $W = [w_1, w_2]^T \in \mathbb{R}^{2 \times d}$ and $\xi_{\text{aug}}, \xi'_{\text{aug}} \sim \mathcal{N}(0, \rho^2 I)$. W is the weight matrix of the encoder containing two neurons, and the parameters of the encoder are the same on both sides.

Note that when Liu et al. (2021) took expectation over $\xi_{\text{aug}}, \xi'_{\text{aug}}$, they canceled the effect of $\xi_{\text{aug}}, \xi'_{\text{aug}}$ due to its linear framework. In other words, the variance of $\xi_{\text{aug}}, \xi'_{\text{aug}}$ can be arbitrarily large. However, Jing et al. (2021) showed that strong augmentation causes dimensional collapse. Hence it is necessary to consider the variance of the data augmentation. In our formulation, ξ_{aug} and ξ'_{aug} can not be canceled due to the nonlinear model. Thus our setting is more reasonable and more in line with the models in practice. To deal with data augmentation operation, we adopt $\xi_{\text{aug}}, \xi'_{\text{aug}} \sim \mathcal{N}(0, \rho^2 I)$.

The SL model. We consider a simple two-layer nonlinear SL model to deal with the above 2-classification problem.

Define $f_{F, W^{\text{SL}}}(x) \triangleq F \sigma(W^{\text{SL}}x)$ with $F \in \mathbb{R}^{1 \times 2}$ as the projection matrix and $W^{\text{SL}} \triangleq [w_1^{\text{SL}}, w_2^{\text{SL}}]^T \in \mathbb{R}^{2 \times d}$ as the weight matrix of the feature extractor. The usual process is to use sigmoid function to transform $f_{F, W^{\text{SL}}}(x_i)$ to $\hat{y}_i \in (0, 1), \forall i \in [n]$. Then, a binary cross-entropy loss function can be constructed with \hat{y}_i and label information $y_i, \forall i \in [n]$.

However, in this paper, we focus on the performance of the feature extractor $\sigma(W^{\text{SL}}x)$. Therefore an objective function that minimizes the norm of feature extractor matrix W^{SL} with margin constraint is used:

$$\begin{aligned} \min_{W^{\text{SL}}} L_{\text{SL}} &= \|w_1^{\text{SL}}\|_2^2 + \|w_2^{\text{SL}}\|_2^2, \\ \text{s.t. } &\sigma((w_{y+1}^{\text{SL}})^T x) - \sigma((w_{y'}^{\text{SL}})^T x) \\ &\geq \sigma(2) - \sigma(-2) - 5\rho d^{\frac{1}{10}}, \forall (x, y) \in \mathcal{D}^{\text{SL}}, y \neq y'. \end{aligned} \quad (3)$$

The SL objective function in this paper is similar to the linear SL model in Liu et al. (2021). We set this SL objective function mainly out of intuition: If a supervised learning model is good enough to complete the classification task, it should satisfy the above margin constraint.

Definitions and notations. To characterize the objective functions, we give the following definitions and notations.

Definition 1 (Locally strong convexity and smooth on B_0). Function $f : \mathbb{R}^d \rightarrow \mathbb{R}$ is locally μ -strongly convex and L_m -smooth if

$$\mu I \preceq \nabla^2 f(x) \preceq L_m I, \forall x \in B_0, \quad (4)$$

where $B_0 := \{x : \|x - x^*\|_2 \leq \|x^{(0)} - x^*\|_2\}$ and $x^* \in \operatorname{argmin}_{x \in \mathcal{X}} f(x)$.

Definition 2 (L_H -Lipschitz continuous Hessian). Function $f : \mathbb{R}^d \rightarrow \mathbb{R}$ is L_H -Lipschitz continuous Hessian if

$$\|\nabla^2 f(x) - \nabla^2 f(y)\|_2 \leq L_H \|x - y\|_2, \forall x, y \in \mathbb{R}^d. \quad (5)$$

Notations. For $x \in \mathbb{R}^d$, we denote by $\|x\|_2$ the vector's Euclidean norm. For $A \in \mathbb{R}^{d \times d}$, we denote by $\|A\|_F$ the standard Frobenius norm and define $\|A\|_2 = \sqrt{\lambda}$ where λ is the largest eigenvalue of $A^\top A$. For $x \in \mathbb{R}^d$ and $\nabla^3 f(x) \in \mathbb{R}^{d \times d \times d}$, we give an upper bound of $\|\nabla^3 f(x)\|_2$ by considering $\nabla^3 f(x)$ as a matrix-vector. Each element of the matrix-vector is $\frac{\partial \nabla^2 f(x)}{\partial x_i} \in \mathbb{R}^{d \times d}$. It is clear that $\|\nabla^3 f(x)\|_2^2 \leq \sum_{i=1}^d \left\| \frac{\partial \nabla^2 f(x)}{\partial x_i} \right\|_F^2$. We denote by $O(\cdot)$ standard Big-O notations, only hiding constants. We denote by $z^{(k)}$ the k -th element of $z \in \mathbb{R}^d$ and $z(t)$ the t -th iteration of the gradient descent algorithm.

4 SSL is Superior to SL in Learning Representation

In this section, we show that the nonlinear SSL model can capture the label-related feature and the hidden feature of data distribution at the same time. In contrast, the nonlinear SL model can only learn the label-related feature. For simplicity, we assume $e_1 = (1, 0, \dots, 0)^\top$, $e_2 = (0, 1, \dots, 0)^\top \in \mathbb{R}^d$.

4.1 The Learning Abilities of SSL and SL

For the convenience, we define $D_1(\tau) = \{\vec{x} \in \mathbb{R}^d | x^{(1)} \in (3.1, 3.9), \tau x^{(2)} \in (8.5, 9), x^{(k)} \in (-\frac{3}{d^{0.49}}, \frac{3}{d^{0.49}}), \forall k \in [3, d]\}$ and $D_2(\tau) = \{\vec{x} \in \mathbb{R}^d | x^{(1)} \in (-3.9, -3.1), \tau x^{(2)} \in (8.5, 9), x^{(k)} \in (-\frac{3}{d^{0.49}}, \frac{3}{d^{0.49}}), \forall k \in [3, d]\}$ as the initialization region of w_1 and w_2 in Theorem 1.

Theorem 1. For $\alpha = 1/800$, $\tau = \max\{7, d^{\frac{1}{10}}\}$, $\rho = 1/d^{1.5}$ and $n = d^2$, with probability $1 - O(e^{-d^{\frac{1}{10}}})$, the SSL objective function L exists a local minimum $W^* = (w_1^*, w_2^*)^\top$:

$$\begin{aligned} \|w_1^* - \tilde{w}_1^*\|_2 &\leq O(d^{-\frac{1}{2}}), \\ \|w_2^* - \tilde{w}_2^*\|_2 &\leq O(d^{-\frac{1}{2}}), \end{aligned}$$

where $\tilde{w}_1^{*(1)} \in [3.1, 3.9]$, $\tilde{w}_1^{*(1)} = -\tilde{w}_2^{*(1)}$, $\tau \tilde{w}_1^{*(2)} = \tau \tilde{w}_2^{*(2)} \geq 9$, $\tilde{w}_1^{*(k)} = \tilde{w}_2^{*(k)} = 0, \forall k \in [3, d]$.

Furthermore, when $(w_1(0), w_2(0)) \in D_1(\tau) \times D_2(\tau)$, using the gradient descent algorithm and choosing learning rate $\eta = \frac{2}{4\alpha + \tau^2 + 1.5}$, $\kappa = 1 + \frac{\tau^2 + 1.5 + 2d^{-0.1}}{2\alpha - d^{-0.1}}$, we have

$$\begin{aligned} \|w_1(t) - w_1^*\|_2 &\leq \left(\frac{\kappa - 1}{\kappa + 1}\right)^t \|w_1(0) - w_1^*\|_2, \\ \|w_2(t) - w_2^*\|_2 &\leq \left(\frac{\kappa - 1}{\kappa + 1}\right)^t \|w_2(0) - w_2^*\|_2. \end{aligned}$$

The projection of e_1 and e_2 on the space spanned by w_1^* and w_2^* is very close to 1, i.e.,

$$\begin{aligned} |\Pi e_1| &\geq 1 - O(\tau^3 d^{-\frac{1}{2}}), \\ |\Pi e_2| &\geq 1 - O(\tau^3 d^{-\frac{1}{2}}). \end{aligned}$$

Theorem 1 shows that using GD to train the nonlinear SSL model under a certain initialization region $D_1(\tau) \times D_2(\tau)$, the model can converge to a local minimum (w_1^*, w_2^*) . Further, the projection of e_1, e_2 on the space spanned by w_1^*, w_2^* is almost 1. In other words, the nonlinear SSL model has simultaneously learned e_1 and e_2 , which are the label-related and hidden features.

Theorem 2. Let $w_1^{SL,*}$ and $w_2^{SL,*}$ be the optimal solution of L_{SL} . Then with probability $1 - O(e^{-d^{\frac{1}{10}}})$,

$$\left(w_1^{SL,*(2)}\right)^2 + \left(w_2^{SL,*(2)}\right)^2 \leq O\left(\rho d^{\frac{1}{10}}\right).$$

When $\rho = 1/d^{1.5}$, $\left(w_1^{SL,*(2)}\right)^2 + \left(w_2^{SL,*(2)}\right)^2 \leq O(1/d^{1.4})$.

From Theorem 2, we show that $(w_1^{SL,*(2)})^2 + (w_2^{SL,*(2)})^2$ is very small, which means SL model can only learn label-related feature and some noise terms.

Note that the previous works (Tian 2017; Zhang et al. 2019; Li and Liang 2018) used the gradient-based algorithm to analyze the SL model with one hidden layer and obtained asymptotic convergence guarantees. They did not analyze the specific form of the learning results. Hence, Theorem 2 is different compared with the previous results. We obtain the bounds of each dimension of the learning results by constructing margin constraints. These bounds accurately describe the features learned by the SL model and help to characterize the representation learnability of the SL model.

Finally, Theorem 1 and Theorem 2 show that the nonlinear SSL model is superior to the nonlinear SL model in capturing important data features, which means SSL can obtain a more competitive data representation than SL.

4.2 Discussion

The extension to more general data distributions. As described in Sec 3.1, we treat label-related features as a group, represented by e_1 (e_2 represents hidden features), and obtain Theorem 1. In this part, we show that Theorem 1 can be extended to data distributions with many label-related and hidden features. Suppose there are P label-related features $E^L = \{e_1, \dots, e_P\}$ and $Q + 1$ hidden features $E^H = \{e_{P+1}, \dots, e_{P+Q}, \vec{0}\}$, where $E = \{E^L, E^H\}$ is column-orthogonal matrix. Each datapoint consists of a label-related feature and a hidden feature, $x_i = z_i e_i^L + \tau e_i^H$, where $P(z_i = 1) = P(z_i = -1) = 1/2$. e_i^L and e_i^H are features in E^L and E^H . This general distribution only considers the relationship between label-related features and hidden features. Hence, the gradient can be decoupled, and the method of this paper can be applied. Finally, we can know that if W contains $P + Q$ neurons, the learning results of W will span the space spanned by $\{e_1, \dots, e_{P+Q}\}$. This conclusion can be regarded as the general version of Theorem 1.

The challenges for nonlinear models. Since the objective function is non-convex and nonlinear, it is difficult to get a closed-form solution with a similar process of Liu et al. (2021). We need to use an optimization algorithm such as GD to converge to a local minimum (w_1^*, w_2^*) and determine features captured by (w_1^*, w_2^*) .

For nonlinear SSL models, previous work (Wen and Li 2021) used SGD to update the model step-by-step and observed the learning result during the iteration. However, the step-by-step process is complex, and it is easy to ignore the change process. Therefore, this procedure makes it difficult to analyze the learning results of local minima accurately.

Different from the complex iterative analysis of the previous work, we propose a new analysis process. We first obtain the approximate region and properties of the local minimum from the simplified objective function \tilde{L} and then extend it to the original complex objective function L . For the transformation from \tilde{L} to L , we use the **exact version of Inverse Function Theorem** as a bridge, avoiding the direct analysis of the local minimum of L . In the remainder of this part, we demonstrate the intuitions and techniques for each part.

Non-convex and nonlinear objective function. At this step, we consider the structure of the objective function, ignore noise terms, and take expectation over data distribution:

$$\min \tilde{L} = -\mathbb{E}_{\tilde{x}}[\langle \sigma(W\tilde{x}), \sigma(W\tilde{x}) \rangle] + \alpha \|W\|_F^2,$$

where \tilde{x}_i is the datapoint without noise term $\rho\xi_i$. We use the intermediate value principle, locally strong convexity of \tilde{L} , and the properties of activation function carefully to prove the existence of the local minimum $(\tilde{w}_1^*, \tilde{w}_2^*)$ of \tilde{L} .

The exact version of Inverse Function Theorem. There are many noise terms in L , such as $\rho\xi_i, \forall i \in [n]$ (datapoint noise), ξ_{aug} (data augmentation noise), and the error terms due to the expectation operation over the data distribution. After obtaining the upper bound of these noise terms (Sec. 4.3), we need a bridge to deal with the transformation from \tilde{L} to L . Since $(\tilde{w}_1^*, \tilde{w}_2^*)$ is local minimum of \tilde{L} and noise terms are bounded, L should be μ -strongly convex and L_m -smooth in the neighborhood of $(\tilde{w}_1^*, \tilde{w}_2^*)$. With these properties, it is clear that $\frac{\partial L}{\partial w_1}$ is one-to-one in a small neighborhood of \tilde{w}_1^* using the origin Inverse Function Theorem (Rudin et al. 1976). However, we need a exact neighborhood to guarantee that the solution w_1^* of $\frac{\partial L}{\partial w_1} = 0$ is in the one-to-one region. Hence we introduce Lipschitz continuous Hessian constant L_H to build an open ball centered at w_1^* with radius $r = \frac{1}{2\mu L_H}$ as the exact neighborhood and modify the Inverse Function Theorem to complete our proof.

4.3 Proof Sketch of Main Theorem

Proof sketch of SSL. For the sake of discussion, we respectively define $\tilde{D}_1(\tau) = \{\tilde{x} \in \mathbb{R}^d | x^{(1)} \in [3.1, 3.9], \tau x^{(2)} \in [9, +\infty), x^{(k)} = 0, \forall k \in [3, d]\}$ and $\tilde{D}_2(\tau) = \{\tilde{x} \in \mathbb{R}^d | x^{(1)} \in [-3.9, -3.1], \tau x^{(2)} \in [9, +\infty), x^{(k)} = 0, \forall k \in [3, d]\}$ as the region of \tilde{w}_1^* and \tilde{w}_2^* .

As a beginning, we focus on \tilde{L} . To obtain the solution of $\frac{\partial \tilde{L}}{\partial w_1} = 0$, we first solve $\frac{\partial \tilde{L}}{\partial w_1^{(k)}} = 0, \forall k \in [2]$ separately in

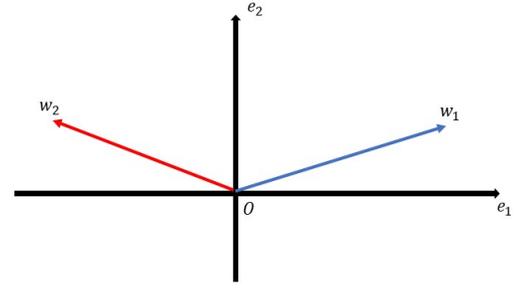


Figure 2: Theoretical Results of Theorem 1

$\tilde{D}_1(\tau)$. Subsequently, we use the intermediate value principle twice to prove the existence of \tilde{w}_1^* . Finally, we use the Hessian matrix to prove that \tilde{w}_1^* is a local minimum. We demonstrate that \tilde{L} is $\tilde{\mu}$ -strongly convexity and \tilde{L}_m -smooth in the region around \tilde{w}_1^* .

To prove Theorem 1, we need to deal with the noise terms in L . Due to the activation function, we cannot use the noise matrix to treat the noise terms as in Liu et al. (2021). Hence, we use the Lagrange's Mean Value Theorem to separate $\xi_i, \xi'_{\text{aug}}, \xi''_{\text{aug}}$ from the activation function and bound these noise terms using the tail bound of Gaussian variable. There are also some error terms due to the expectation operation over data distribution. With the intuition that $n_i, \forall i \in [4]$ can not be far away from $n/4$, we bound these error terms.

After obtaining the upper bound of the above noise terms, we characterize the landscape of L by using the Matrix Eigenvalue Perturbation Theory (Kahan 1975). We sum up the properties of L when w_1 around \tilde{w}_1^* as follows.

1. $\frac{\partial L}{\partial w_1} |_{w_1=\tilde{w}_1^*}$ is very close to 0.
2. L is μ -strongly convex and L_m -smooth. Specifically, we show that $\tilde{\mu} - \epsilon_1 \leq \mu \leq \tilde{\mu}$ and $\tilde{L}_m \leq L_m \leq \tilde{L}_m + \epsilon_1$ where ϵ_1 is a small term related to ρ and d .
3. L is L_H -Lipschitz continuous Hessian.

Combined with these properties, we use the exact version of Inverse Function Theorem to prove the existence of the local minimum (w_1^*, w_2^*) of L . Finally, we show that with good initialization, specifically initialization around the local minimum, $w_1(0)$ converges to w_1^* using the gradient descent algorithm (Bubeck et al. 2015). We remark that the above process only analyzes w_1 , we can get w_2 through symmetry.

Proof sketch of SL. The proof sketch of SL is similar to the proof of the linear SL model in Liu et al. (2021). However, because of the nonlinear SL model in this paper, we need to perform finer scaling to get a high probability guarantee.

Different activation function. We can easily extend the results to the case where the activation function is tanh because sigmoid can be viewed as a compressed version of tanh.

To get similar results with Theorem 1, we just need to modify the region of the local minimum and the initialization region. For $\tilde{D}_1(\tau)$, we change the range of $x^{(1)}$ from $[3.1, 3.9]$ to $[2.7, 3.1]$ and the range of $x^{(2)}$ from $[9, +\infty)$ to $[6.1, +\infty)$ to obtain $\tilde{D}_1^{\sigma^2}(\tau)$. For $D_1(\tau)$, we change the

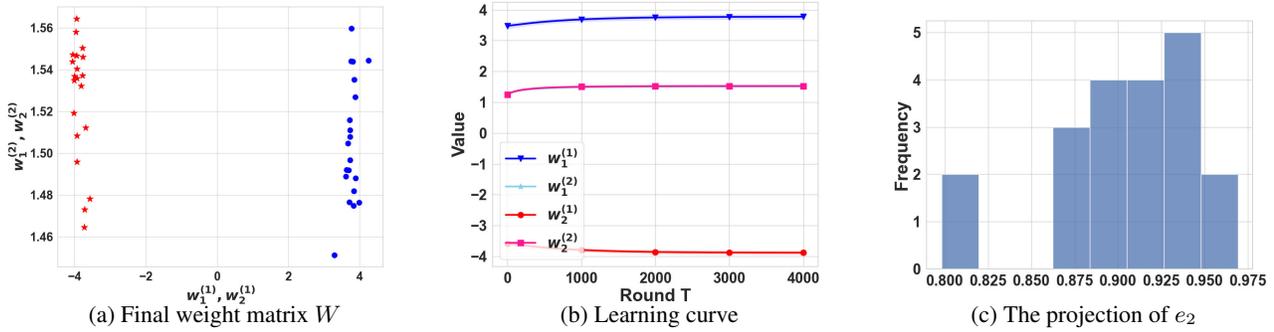


Figure 3: Experiment results of SSL model with $d = 10, \tau = 7$

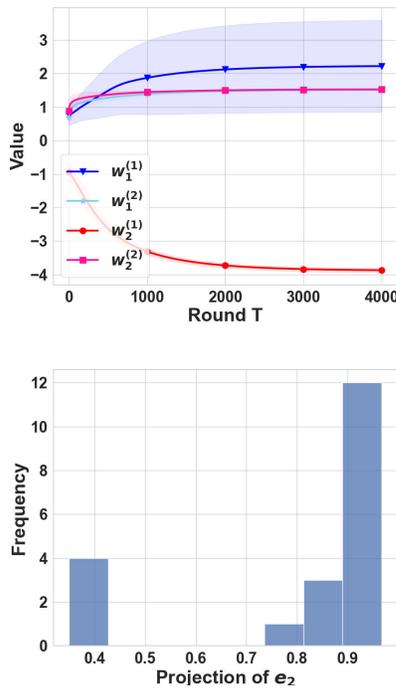


Figure 4: The experiments results with the correct sign

range of $x^{(1)}$ from (3.1, 3.9) to (2.7, 3.1) and the range of $x^{(2)}$ from (8.5, 9) to (5.75, 6.1) to obtain $D_1^{\sigma^2}(\tau)$. With similar process, we can get $\tilde{D}_2^{\sigma^2}(\tau)$ and $D_2^{\sigma^2}(\tau)$.

5 Simulation Experiments

In this section, we illustrate the correctness of Theorem 1 and Theorem 2 experimentally. We conduct experiments for the nonlinear SSL model in Sec. 5.1 and Sec. 5.2. Furthermore, we show the training process of the nonlinear SL model with projection matrix F in Sec. 5.3. In this section, we choose $\tau = 7, d = 10, \rho = 1/d^{1.5}, \alpha = \frac{1}{800}, n = d^2$ and learning rate $\eta = 0.001$ if we do not specify otherwise. Experiments are averaged over 20 random seeds, and we show the average

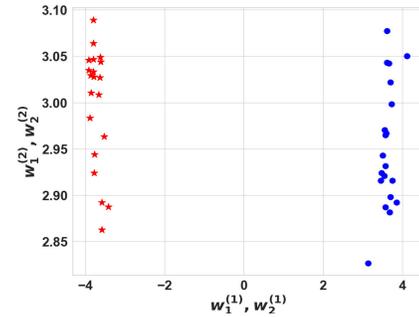


Figure 5: SSL final weight matrix W with $d = 10, \tau = 3$

results with 95% confidence interval for learning curves.

5.1 SSL Model: the Correctness of Theorem 1

In this part, we validate the correctness of Theorem 1 by strictly following the settings of the theorem. Define $T_1 = 4000$ as the number of iterations of the SSL model. Fig. 3a shows the learning result of weight matrix $W = [w_1, w_2]^T$. The blue points are learning results of $(w_1^{(1)}(T_1), w_1^{(2)}(T_1))$ and the red stars are learning results of w_2 . It is clear that $w_1(T_1)$ and $w_2(T_1)$ are almost symmetrical about the e_2 -axis, which is consistent with the theoretical result (Fig. 2). Fig. 3b shows the learning process of w_1 and w_2 . Because we initialize $w_1(0)$ and $w_2(0)$ around the local minimum, w_1 and w_2 can easily converge to (w_1^*, w_2^*) . Fig. 3c shows the projection of e_2 on the space spanned by $w_1(T_1)$ and $w_2(T_1)$. We can find the projection is almost 1. These experimental results show that W learns e_1, e_2 at the same time. The results of larger τ are similar to the results of $\tau = 7$.

5.2 SSL Model: Results Beyond Analysis

In this part, we relax requirements in Theorem 1, such as (w_1, w_2) must be initialized near (w_1^*, w_2^*) , τ must be large. We show that the SSL model still learns the label-related and hidden feature even if the requirements are relaxed.

Good enough initialization. In Theorem 1, we initialize w_1 and w_2 around (w_1^*, w_2^*) . We experimentally show

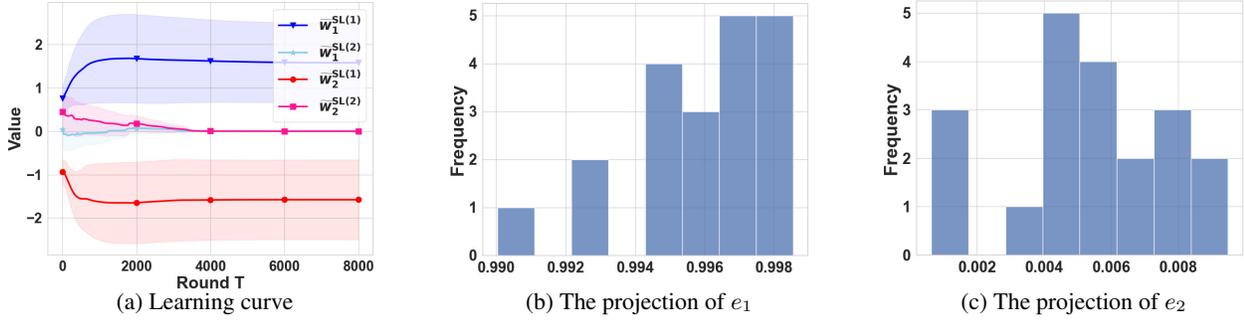


Figure 6: Experiment results of SL model with $d = 10, \tau = 7$

that initialization only need the correct sign ($w_1^{(1)}(0) > 0, w_2^{(1)}(0) < 0, w_1^{(2)}(0) > 0, w_2^{(2)}(0) > 0$) is required. Fig. 4 shows that if the initialization sign is correct, the SSL model can converge to (w_1^*, w_2^*) with high probability. ‘‘With high probability’’ means there are still a few cases where the SSL model cannot converge to (w_1^*, w_2^*) . However, compared with the learning results (Fig. 6c) of the SL model, the SSL model with the correct sign still shows the ability to learn e_2 .

Large enough τ . In the proof process of Theorem 1, we need $\tau = \max\{7, d^{\frac{1}{10}}\}$ to use the monotonicity of the solution of $\frac{\partial \tilde{L}}{\partial w_1^{(2)}}$. We experimentally show that the SSL model can get a good result even if τ does not meet this requirement. Fig. 5 shows even if $\tau = 3$, the space spanned by w_1 and w_2 is still very close to the space spanned by e_1 and e_2 .

5.3 SL Experiment Results

In Theorem 2, we mainly focus on the performance of the feature extractor W^{SL} and ignore the projection matrix F . In this section, we experimentally show that even if F is considered, W^{SL} still only learns the label-related feature. Specifically, we consider the binary cross-entropy loss function:

$$\min \tilde{L}_{\text{SL}} = -\frac{1}{n} \sum_{i=1}^n y_i \ln(\hat{y}_i) + (1 - y_i) \ln(1 - \hat{y}_i) + \beta \|W^{\text{SL}}\|_F^2 + \gamma \|F\|_2^2,$$

where $\hat{y}_i = \sigma(F\sigma(W^{\text{SL}}x_i))$, $\forall i \in [n]$, β is the coefficient of W^{SL} regularizer, and γ is the coefficient of F regularizer. In this section, we choose $\beta = \gamma = 1/800$.

Define $T_2 = 8000$ as the number of iterations of the nonlinear SL model. Fig. 6a shows the learning curve of $(\tilde{w}_1^{\text{SL}}, \tilde{w}_2^{\text{SL}})$. It is clear that $\tilde{w}_1^{\text{SL}(1)}(T_2)$ and $\tilde{w}_2^{\text{SL}(1)}(T_2)$ are the main terms, and the other terms $\tilde{w}_j^{\text{SL}(k)}, \forall j \in [2], k \in [2, d]$ will converge to 0. Fig. 6b and Fig. 6c show the projection of e_1 and e_2 on the space spanned by $\tilde{w}_1^{\text{SL}}(T_2)$ and $\tilde{w}_2^{\text{SL}}(T_2)$. It is clear that the projection of e_1 is almost 1, and the projection of e_2 is almost 0. The above experiment results mean that the nonlinear SL model can only learn label-related feature e_1 , which is consistent with the results of Theorem 2.

All experiments are conduct on a desktop with AMD Ryzen 7 5800H with Radeon Graphics 3.20 GHz

and 16 GB memory. The codes of this section are available at <https://github.com/wanshuiyin/AAAI-2023-The-Learnability-of-Nonlinear-SSL>.

6 Conclusion

Summary. Our paper is the first to analyze the data representation learnability of the nonlinear SSL model by analyzing the learning results of the neural network. We start with a 1-layer nonlinear SSL model and use GD to train this model. We prove that the model converges to a local minimum. Further, we accurately describe the properties of this local minimum and prove that the nonlinear SSL model can capture label-related features and hidden features at the same time. In contrast, the nonlinear SL model only learns label-related features. This conclusion shows that even though the nonlinear network significantly improves the learnability of the SL model, the SSL model still has a superior ability to capture important features compared with the SL model. We verify the correctness of the results through simulation experiments.

Due to the nonconvexity of the objective function and noise terms, we propose a new analysis process to describe the properties of the local minimum. This analysis process is divided into two steps. In the first step, we focus on the structure of L by ignoring all noise terms. Then we obtain the approximate region of the local minimum. In the second step, we use the exact version of Inverse Function Theorem as a bridge to connect the simplified objective function \tilde{L} and L . Finally, we prove the existence of the local minimum (w_1^*, w_2^*) and describe the properties of this local minimum.

Compared with linear SSL models, nonlinear alternatives are closer to the state-of-the-art SSL methods. The conclusions in this paper can guide us further in understanding the learning results of SSL methods and provide a theoretical basis for subsequent improvements.

Future work. This paper analyzes a 1-layer nonlinear SSL model. After that, we plan to expand the scope of the analysis to a multi-layer nonlinear network. The multi-layer network analysis requires a more refined exploration of local minima. The weight matrix of each layer needs to be uniformly processed to analyze the landscape of the objective function, which we will do in the follow-up work.

Acknowledgments

The corresponding author Shuai Li is supported by National Key Research and Development Program of China No. 2020AAA0107600, National Natural Science Foundation of China (62006151, 62076161) and Shanghai Sailing Program. The author Bo Jiang is supported by National Natural Science Foundation of China (62072302).

References

- Allen-Zhu, Z.; Li, Y.; and Song, Z. 2019. A convergence theory for deep learning via over-parameterization. In *International Conference on Machine Learning*, 242–252. PMLR.
- Arora, S.; Khandeparkar, H.; Khodak, M.; Plevrakis, O.; and Saunshi, N. 2019. A theoretical analysis of contrastive unsupervised representation learning. *arXiv preprint arXiv:1902.09229*.
- Bromley, J.; Guyon, I.; LeCun, Y.; Säckinger, E.; and Shah, R. 1993. Signature verification using a” siamese” time delay neural network. *Advances in neural information processing systems*, 6.
- Brutzkus, A.; and Globerson, A. 2017. Globally optimal gradient descent for a convnet with gaussian inputs. In *International conference on machine learning*, 605–614. PMLR.
- Bubeck, S.; et al. 2015. Convex optimization: Algorithms and complexity. *Foundations and Trends® in Machine Learning*, 8(3-4): 231–357.
- Caron, M.; Misra, I.; Mairal, J.; Goyal, P.; Bojanowski, P.; and Joulin, A. 2020. Unsupervised learning of visual features by contrasting cluster assignments. *Advances in Neural Information Processing Systems*, 33: 9912–9924.
- Chen, T.; Kornblith, S.; Norouzi, M.; and Hinton, G. 2020. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, 1597–1607. PMLR.
- Chen, X.; and He, K. 2021. Exploring Simple Siamese Representation Learning. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2021, virtual, June 19-25, 2021*, 15750–15758. Computer Vision Foundation / IEEE.
- Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Du, S.; Lee, J.; Li, H.; Wang, L.; and Zhai, X. 2019. Gradient descent finds global minima of deep neural networks. In *International conference on machine learning*, 1675–1685. PMLR.
- Du, S. S.; Jin, C.; Lee, J. D.; Jordan, M. I.; Singh, A.; and Póczos, B. 2017. Gradient descent can take exponential time to escape saddle points. *Advances in neural information processing systems*, 30.
- Grill, J.-B.; Strub, F.; Altché, F.; Tallec, C.; Richemond, P.; Buchatskaya, E.; Doersch, C.; Avila Pires, B.; Guo, Z.; Gheshlaghi Azar, M.; et al. 2020. Bootstrap your own latent—a new approach to self-supervised learning. *Advances in neural information processing systems*, 33: 21271–21284.
- HaoChen, J. Z.; Wei, C.; Gaidon, A.; and Ma, T. 2021. Provable guarantees for self-supervised deep learning with spectral contrastive loss. *Advances in Neural Information Processing Systems*, 34: 5000–5011.
- HaoChen, J. Z.; Wei, C.; Kumar, A.; and Ma, T. 2022. Beyond separability: Analyzing the linear transferability of contrastive representations to related subpopulations. *arXiv preprint arXiv:2204.02683*.
- He, K.; Fan, H.; Wu, Y.; Xie, S.; and Girshick, R. 2020. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 9729–9738.
- Jing, L.; Vincent, P.; LeCun, Y.; and Tian, Y. 2021. Understanding dimensional collapse in contrastive self-supervised learning. *arXiv preprint arXiv:2110.09348*.
- Kahan, W. 1975. Spectra of nearly Hermitian matrices. *Proceedings of the American Mathematical Society*, 48(1): 11–17.
- Khosla, P.; Teterwak, P.; Wang, C.; Sarna, A.; Tian, Y.; Isola, P.; Maschinot, A.; Liu, C.; and Krishnan, D. 2020. Supervised contrastive learning. *Advances in Neural Information Processing Systems*, 33: 18661–18673.
- Lee, J. D.; Lei, Q.; Saunshi, N.; and Zhuo, J. 2021. Predicting what you already know helps: Provable self-supervised learning. *Advances in Neural Information Processing Systems*, 34: 309–323.
- Li, Y.; and Liang, Y. 2018. Learning overparameterized neural networks via stochastic gradient descent on structured data. *Advances in neural information processing systems*, 31.
- Li, Y.; and Yuan, Y. 2017. Convergence analysis of two-layer neural networks with relu activation. *Advances in neural information processing systems*, 30.
- Liu, H.; HaoChen, J. Z.; Gaidon, A.; and Ma, T. 2021. Self-supervised learning is more robust to dataset imbalance. *arXiv preprint arXiv:2110.05025*.
- Radford, A.; Wu, J.; Child, R.; Luan, D.; Amodei, D.; Sutskever, I.; et al. 2019. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8): 9.
- Rudin, W.; et al. 1976. *Principles of mathematical analysis*, volume 3. McGraw-hill New York.
- Tian, Y. 2017. An analytical formula of population gradient for two-layered relu network and its applications in convergence and critical point analysis. In *International conference on machine learning*, 3404–3413. PMLR.
- Tian, Y. 2022a. Deep contrastive learning is provably (almost) principal component analysis. *arXiv preprint arXiv:2201.12680*.
- Tian, Y. 2022b. Understanding the Role of Nonlinearity in Training Dynamics of Contrastive Learning. *arXiv preprint arXiv:2206.01342*.
- Tian, Y.; Chen, X.; and Ganguli, S. 2021. Understanding self-supervised learning dynamics without contrastive pairs. In *International Conference on Machine Learning*, 10268–10278. PMLR.

Tian, Y.; Yu, L.; Chen, X.; and Ganguli, S. 2020. Understanding self-supervised learning with dual deep networks. *arXiv preprint arXiv:2010.00578*.

Tosh, C.; Krishnamurthy, A.; and Hsu, D. 2021. Contrastive estimation reveals topic posterior information to linear models. *J. Mach. Learn. Res.*, 22: 281–1.

Wang, Y.; Wang, H.; Shen, Y.; Fei, J.; Li, W.; Jin, G.; Wu, L.; Zhao, R.; and Le, X. 2022. Semi-Supervised Semantic Segmentation Using Unreliable Pseudo-Labels. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2022, New Orleans, LA, USA, June 18-24, 2022*, 4238–4247. IEEE.

Wen, Z.; and Li, Y. 2021. Toward understanding the feature learning process of self-supervised contrastive learning. In *International Conference on Machine Learning*, 11112–11122. PMLR.

Wu, A.; Wang, C.; Pino, J.; and Gu, J. 2020. Self-supervised representations improve end-to-end speech translation. *arXiv preprint arXiv:2006.12124*.

Zhang, X.; Yu, Y.; Wang, L.; and Gu, Q. 2019. Learning one-hidden-layer relu networks via gradient descent. In *The 22nd international conference on artificial intelligence and statistics*, 1524–1534. PMLR.

Zhong, Y.; Tang, H.; Chen, J.; Peng, J.; and Wang, Y.-X. 2022. Is Self-Supervised Learning More Robust Than Supervised Learning? *arXiv preprint arXiv:2206.05259*.