On the Unreasonable Effectiveness of Knowledge Distillation: Analysis in the Kernel Regime – Long Version

Arman Rahbar
Department of Computer Science and Engineering
Chalmers University of Technology
Gothenburg, Sweden
armanr@chalmers.se

Ashkan Panahi
Department of Computer Science and Engineering
Chalmers University of Technology
Gothenburg, Sweden
ashkan.panahi@chalmers.se

Chiranjib Bhattacharyya
Department of Computer Science and Automation
Indian Institute of Science
Karnataka, India
chiru@iisc.ac.in

Devdatt Dubhashi
Department of Computer Science and Engineering
Chalmers University of Technology
Gothenburg, Sweden
dubhashi@chalmers.se

Morteza Haghir Chehreghani
Department of Computer Science and Engineering
Chalmers University of Technology
Gothenburg, Sweden
morteza.chehreghani@chalmers.se

September 28, 2020

ABSTRACT

Knowledge distillation (KD), i.e. one classifier being trained on the outputs of another classifier, is an empirically very successful technique for knowledge transfer between classifiers. It has even been observed that classifiers learn much faster and more reliably if trained with the outputs of another classifier as soft labels, instead of from ground truth data. However, there has been little or no theoretical analysis of this phenomenon. We provide the first theoretical analysis of KD in the setting of extremely wide two layer non-linear networks in model and regime in [1, 2, 3]. We prove results on what the student network learns and on the rate of convergence for the student network. Intriguingly, we also confirm the *lottery ticket hypothesis* [4]in this model. To prove our results, we extend the repertoire of techniques from linear systems dynamics. We give corresponding experimental analysis that validates the theoretical results and yields additional insights.

1 Introduction

In 2014, Hinton et al. [5] made a surprising observation: they found it easier to train classifiers using the real–valued outputs of another classifier as target values than using actual ground–truth labels. They introduced the term *knowledge distillation* (or distillation for short) for this phenomenon. Since then, distillation–based training has been confirmed robustly in several different types of neural networks [6, 7, 8]. It has been observed that optimization is generally more well–behaved than with label-based training, and it needs less if any regularization or specific optimization tricks. Consequently, in several fields, distillation has become a standard technique for transferring the information between classifiers with different architectures, such as from deep to shallow neural networks or from ensembles of classifiers to individual ones.

While the practical benefits of distillation are beyond doubt, its theoretical justification remains almost completely unclear [1]. Existing explanations rarely go beyond qualitative statements, e.g. claiming that learning from soft labels should be easier than learning from hard labels, or that in a multi-class setting the teacher's output provides information about how similar different classes are to each other. The recent paper [9] is the sole exception (as far as we are aware) but their analysis is in the setting of *linear* networks.

We carry out the first theoretical analysis of distillation in the setting of a simple two–layer *non–linear* neural network. Our analysis is carried out in the model underlying the exciting recent work that analyzes the dynamics of neural networks under the so–called *kernel regime*, see [1, 2, 3, 10] and others. In this series of papers, it was shown that the behaviour of training by gradient descent (GD) in the limit of very wide neural networks can be approximated by *linear system dynamics*. This is dubbed the *kernel regime* because it was shown in [11] that a fixed kernel – the *neural tangent kernel* – characterizes the behavior of fully-connected infinite width neural networks in this regime.

In the same model and regime, we prove the first theoretical results on knowledge distillation for non-linear networks. Our framework is general enough to encompass Vapnik's notion of *privileged information* and provides a unified analysis of *generalized distillation* in the paradigm of *machines teaching machines* as in [12]. We give results on both what is learnt by the student network (in section 2.5) and also on the speed of convergence (in section 2.6). Intriguingly, we also confirm theoretically, the *lottery ticket hypothesis* [4]in this model and regime as a special case of our analysis (in section 2.5.1. We introduce novel techniques to make possible this analysis. We also carry out systematic experimental evaluation that confirms the theoretical analysis.

2 Problem Formulation and Main Results

Unlike most previous work which study distillation in various complex architectures, we focus our attention on a simplified, analytically tractable, setting: the two layer non–linear model introduced in [1, 13, 2, 3]:

$$f(\mathbf{x}) = \sum_{k=1}^{m} \frac{a_k}{\sqrt{m}} \sigma(\mathbf{w}_k^T \mathbf{x}), \tag{1}$$

Here the weights $\{\mathbf{w}_k\}_{k=1}^m$ are variables of the network corresponding to m hidden units, $\sigma(.)$ is a real (nonlinear) activation function and the weights $\{a_k\}$ are fixed. In our theoretical analysis, the student network has this form while the teacher can be any classifier. In our experimental analysis both the teacher and student networks have this form; the student just has many fewer nodes in the hidden layer.

We introduce a general optimization framework for knowledge transfer formulating it as a least squares optimization problem with regularization provided by privileged knowledge. Given a dataset $\{(\mathbf{x}_i, y_i)\}_{i=1}^n$ comprising of n data samples $\{\mathbf{x}_i\}$ and their corresponding labels $\{y_i\}$, our framework is given by the following optimization problem

$$\min_{\{\mathbf{w}_k\}} \sum_i (y_i - f(\mathbf{x}_i))^2 + \lambda \sum_i \sum_k \left(\phi^{(k)}(\mathbf{x}_i) - f^{(k)}(\mathbf{x}_i) \right)^2, \tag{2}$$

where $f(.)$ is stated in (1) and $f^{(k)}(\mathbf{x}) = \sigma(\mathbf{w}_k^T \mathbf{x})$ is the corresponding k^{th} hidden feature of the student network. Here, we assume a given function $\phi^{(k)}$ for each hidden unit as privileged knowledge. We observe that (2) incorporates this knowledge as a regularization of the original nonlinear least squares (average ℓ_2 risk minimization) framework for fitting a function to the labeled data: $\min_{\{\mathbf{w}_k\}} \sum_i (y_i - f(\mathbf{x}_i))^2$. The coefficient $\lambda \geq 0$ in (2) is the regularization parameter.

A special case of this setup concerns knowledge distillation where $\{\phi^{(k)}\}$ are hidden features of a teacher network, i.e. $\phi^{(k)}(\mathbf{x}) = \sigma(\langle \mathbf{w}_k^{\text{teacher}}, \mathbf{x} \rangle)$ framing (2) as a standard student-teacher scheme for knowledge distillation.

2.1 Relation to Previous Work

To the best of our knowledge, [9] is the only previous attempt to give a theoretical analysis of distillation. In the setting of (2), their work corresponds to $\lambda = \infty$ (pure distillation), single hidden unit ($m = 1$), Sigmoid activation σ and cross-entropy replacing the square-error loss and hence leading to a convex objective. Their result concerns the final value and is expressed in terms of the weight values. Convexity allows [9] to avoid our assumption on the initial weights in Theorem 1 and 2, as the final solution is independent of initialization. Our result is applicable to a different regime with a large number m of units, high expression capacity and a non-convex formulation.

---

[1]This is the reason for our title with a play on Eugene Wigner's famous article entitled "The Unreasonable Effectiveness of Mathematics in the Natural Sciences"

To gain a deeper insight into the process of knowledge transfer by (2), we study the generic behavior of the gradient descent (GD) algorithm when applied to the optimization therein. In the spirit of the analysis in [13, 1], shortly explained in Section 3, we carry out an investigation on the dynamics for GD that answers two fundamental questions:

- What does the (student) network learn?
- How fast is the convergence by the gradient descent?

The answer to both these questions emerges from the analysis of the dynamics of GD.

2.2 General Framework

The existing analysis of dynamics for neural networks in a series of recent papers [1, 2, 3] is tied centrally to the premise that the behaviour of GD for the optimization can be approximated by a linear dynamics of finite order. To isolate the negligible effect of learning rate μ in GD, it is also conventional to study the case $\mu \to 0$ where GD is alternatively represented by an ordinary differential equation (ODE), known as the gradient flow, with a continuous "time" variable t replacing the iteration number r (being equivalent to the limit of $r\mu$). Let us denote by $\mathbf{f}(t)$ the vector of the output $f(\mathbf{x}_i, t)$ of the network at time t. Then, the theory of linear systems with a finite order suggests the following expression for the evolution of $\mathbf{f}(t)$:

$$\mathbf{f}(t) = \mathbf{f}_\infty + \boldsymbol{\delta}(t), \tag{3}$$

where $\mathbf{f}_\infty$ is a constant and

$$\boldsymbol{\delta}(t) = \mathbf{u}_1 e^{-p_1 t} + \mathbf{u}_2 e^{-p_2 t} + \ldots + \mathbf{u}_d e^{-p_d t}. \tag{4}$$

Here, d is the order of the linear system and complex-valued vectors $\mathbf{u}_1, \ldots, \mathbf{u}_m$ and nonzero complex values $p_1, p_2, \ldots, p_d$ are to be determined by the specifications of the dynamics. The constants $\{p_j \neq 0\}$ are called poles, that also correspond to the singular points of the Laplace transform $\mathbf{F}(s)$ of $\mathbf{f}(t)$ (except for 0, which corresponds to the constant $\mathbf{f}_\infty$ in our formulation). We observe that such a representation may only have a convergence (final) value at $t \to \infty$ if the poles have strictly positive real parts, in which case $\mathbf{f}_\infty$ is the final value. Moreover, the asymptotic rate of convergence is determined by the dominating term in (4), i.e. the smallest value $\Re(p_j)$ with a nonzero vector $\mathbf{u}_j$. We observe that identifying $\mathbf{f}_\infty$ and the dominating term responds to the aforementioned questions of interest. In this paper, we show that these values can be calculated as the number m of hidden units increases.

2.3 Definitions

we first need to introduce a number of definitions. Let us take $\mathbf{w}_k = \mathbf{w}_k(0)$ as the initial values of the weights and define $\mathbf{H}_k = (\sigma'(\mathbf{w}_k^T \mathbf{x}_i)\sigma'(\mathbf{w}_k^T \mathbf{x}_i)\mathbf{x}_i^T \mathbf{x}_j)$ as the k^{th} realization of the "associated gram matrix" where σ' denotes the derivative function of σ (that can be defined in the distribution sense). Further, denote by $\mathbf{f}^k(0)$ the vector of the initial values $f^k(\mathbf{x}_i) = \sigma(\mathbf{w}_k^T \mathbf{x}_i)$ of the k^{th} unit for different data points $\{\mathbf{x}_i\}$. Finally, we define $a = \sum_k \frac{a_k^2}{m}$.

2.4 Assumptions

Our analysis will also be built upon a number of assumptions:

Assumption 1. Nonzero eigenvalues of the matrices $\{\mathbf{H}_k\}$ are all distinct. Note that they are always strictly positive as $\{\mathbf{H}_k\}$ are by construction positive semi-definite (psd).

Assumption 2. -1 is the eigenvalue of the matrix

$$\mathbf{T}(s) = \sum_k \frac{a_k^2}{m}(s\mathbf{I} + \lambda\mathbf{H}_k)^{-1}\mathbf{H}_k$$

at exactly $d = n \times m$ distinct strictly negative values of $s = -p_1, -p_2, \ldots, -p_d$, where p_is are all different to the eigenvalues of $\{\mathbf{H}_k\}$, with $\mathbf{v}^1, \mathbf{v}^2, \ldots, \mathbf{v}^d$ and $\mathbf{u}^1, \mathbf{u}^2, \ldots, \mathbf{u}^d$ being the corresponding right and left eigenvectors.

Assumption 3. The function σ and its derivative σ' is Lipschitz continuous.

Assumption 4. We assume $m \to \infty$ and $\|\boldsymbol{\phi}_k - \mathbf{f}_k(0)\| = O(1/\sqrt{m})$ and $\|\mathbf{y} - \sum_k \frac{\mathbf{a}_k \boldsymbol{\phi}_k}{\sqrt{m}}\| = O(1)$. Moreover, $\|\mathbf{x}_i\|$s and a_ks are bounded.

3

2.5 What does the student learn?

This result pertains to the first question above, concerning the final value of $\mathbf{f}$. For this, we prove the following result:

Theorem 1. *Suppose that Assumption 1-4 holds true. Then,* $\lim\limits_{t \to \infty} \mathbf{f}(t) = \mathbf{f}_\infty$ *where*

$$\mathbf{f}_\infty = \frac{1}{a+\lambda}\left(a\mathbf{y} + \lambda \sum_k \frac{a_k \phi_k}{\sqrt{m}}\right) \tag{5}$$

To gain a deeper insight, we specialize this result to the case where the teacher is a well-trained network with the similar structure as in (1) with $\bar{m} \gg m$ units and coefficients $\{\bar{\mathbf{w}}_l, \frac{q\bar{a}_l}{\sqrt{m}}\}_{l=1}^{\bar{m}}$. In this setup, the constant $q = \sqrt{\frac{m}{\bar{m}}}$ is a correction term and the output of the teacher at $\mathbf{x}_i$ is exactly y_i for $i = 1, 2, \ldots, n$ (e.g. by the results in [13]). Then, the privileged knowledge is extracted by randomly selecting indices $l_1, l_2, \ldots, l_m$ and setting $\mathbf{w}_k^{\text{teacher}} = \bar{\mathbf{w}}_{l_k}$ and $a_k = \bar{a}_{l_k}$. Then, Theorem 1 leads to the the following result:

Theorem 2. *Suppose that the student is initialized by* $\mathbf{w}_k = \mathbf{w}_k^{\text{teacher}}$, *selected by the above randomized scheme and* $\sum\limits_l \frac{\bar{a}_l^2}{\bar{m}} = 1$. *Moreover,* $\sum\limits_{l,i} \frac{\bar{a}_l^2}{\bar{m}}\sigma^2(\bar{\mathbf{w}}_l^T \mathbf{x}_i)$ *is bounded with* $\bar{m}$. *Then, the output f after training satisfies:*

$$\sum_{i=1}^n (f(\mathbf{x}_i) - y_i)^2 \leq \left(1 - \frac{m}{\bar{m}}\right) \times$$

$$C\left(\frac{\lambda}{1+\lambda}\right)^2 \left[\sum_{i,l} \frac{q^2 \bar{a}_l^2}{\bar{m}}\sigma^2(\bar{\mathbf{w}}_l^T \mathbf{x}_i)\right]\log(\bar{m}) \tag{6}$$

with a probability higher than $1 - \frac{1}{m^{10}}$, *where C is a universal constant.*

In simple words, the final error in the student will be proportional to $1 - \frac{m}{\bar{m}}$, confirming the intuition that the error grows with a smaller student.

2.5.1 Relation to Lottery Ticket Hypothesis

We further observe an interesting connection between the setup in Theorem 2 and the lottery ticket hypothesis in [4]: When $\lambda = 0$ (no distillation), the training procedure on the student network can be interpreted as the retraining of the randomly selected features of the teacher. This coincides with the lottery ticket setup. As such, our analysis in Theorem 2 with $\lambda = 0$ shows that re-training in the kernel limit with a fixed fraction $m/\bar{m}$ of features leads to zero training loss, confirming the lottery ticket hypothesis.

2.6 How fast does the student learn?

Now, we turn our attention to the question of the speed of convergence, for which we have the following result:

Theorem 3. *Define*

$$\mathbf{f}_\infty^{(k)} = \frac{a_k}{\lambda\sqrt{m}}(\mathbf{y} - \mathbf{f}_\infty) + \phi_k.$$

Under Assumption 1-4, the dynamics of $\mathbf{f}$ can be written as[2]

$$\mathbf{f}(t) = \sum_{j=1}^d e^{-p_j t}\alpha_j \mathbf{u}^j + O_m\left(\frac{1}{\sqrt{m}}\right)$$

where

$$\alpha_j = \sum_k \frac{a_k^2}{m}\left\langle \mathbf{v}^j, \mathbf{H}_k(p_j\mathbf{I} - \lambda\mathbf{H}_k)^{-1}(\mathbf{f}_\infty^{(k)} - \mathbf{f}^{(k)}(0))\right\rangle$$

From the definition, α_j can be interpreted as an average overlap between a combination of the label vector $\mathbf{y}$ and the knowledge vectors ϕ_k, and the "spectral" structure of the data as reflected by the vectors $(p_j\mathbf{I} - \lambda\mathbf{H}_k)^{-1}\mathbf{H}_k\mathbf{v}^j$. This is a generalization of the geometric argument in [1] in par with the "data geometry" concept introduced in [9]. We will later use this result in our experiments to improve distillation by modifying the data geometry of α_j coefficients.

---

[2]We clarify that $O(1/\sqrt{m})$ holds in L_1 sense, i.e. $a(t) = b(t) + O(1/\sqrt{m})$ means $\int\limits_0^\infty \|a(t) - b(t)\|_2 \mathrm{d}t = O(1/\sqrt{m})$.

2.7 Further Remarks

The above two results have a number of implications on the distillation process. First, note that the case $\lambda = 0$ reproduces the results in []: The final value simply becomes $\mathbf{y}$ while the poles will become the singular values of the matrix $\mathbf{H} = \sum_k \frac{a_k^2}{m} \mathbf{H}_k$. The other extreme case of $\lambda = \infty$ corresponds pure distillation, where the effect of the first term in (2) becomes negligible and hence the optimization boils down to individually training each hidden unit by ϕ_k. One may then expect the solution of this case to be $f_k = \phi_k$. However, the conditions of the above theorems become difficult to verify, but we shortly present experiments that numerically investigate the corresponding dynamics.

We also observe that for a finite value of λ, the final value $\mathbf{f}_\infty$ is a weighted average, depending on the quality of ϕ^k. Defining $e = \|\mathbf{f}_\infty - \mathbf{y}\|$ as the final error, we simply conclude that $e = \frac{\lambda}{a+\lambda} \|\sum_k \frac{a_k \phi_k}{m} - \mathbf{y}\|$, where the term $\|\sum_k \frac{a_k \phi_k}{m} - \mathbf{y}\|$ reflects the quality of teacher in representing the labels. Also for an imperfect teacher the error e monotonically increases with λ, while larger λ generally has a positive effect on the speed of convergence as it shifts the poles to become larger. For λ, this leads to a trade-off between speed and quality. The result of Theorem 1 also stems from the probabilistic analysis of e in the lottery ticket setup.

The analysis of Theorem 3 gives us another way to assess the effect of the teacher. If the teachers have a small overlap with the "eigen-basis" $\mathbf{H}_k(p_j - \lambda \mathbf{H}_k)^{-1}\mathbf{v}_j$ corresponding to small values of p_j, then the dynamics is mainly identified by the large poles p_j, speeding up the convergence properties.

Finally, the assumption $\|\phi_k - \mathbf{f}^{(k)}(0)\| = O(1/m)$ can be simply satisfied in the distillation case with single hidden layer, where we have $\phi_k(\mathbf{x}) = \sigma(\langle \mathbf{w}_k^{\text{teacher}}, \mathbf{x} \rangle)$ and initializing the weights of the student by that of the teacher $\mathbf{w}_k(0) = \mathbf{w}_k^{\text{teacher}}$ leads to $\phi_k = \mathbf{f}^{(k)}(0)$. We further numerically investigate the consequences of violating this assumption.

We observe that the above elements can be simply investigated in the context of distillation with the lottery ticket principle in [14]. This is what we mainly study in the numerical experiments.

3 Analysis and Insights

The study in [13] on the dynamics of backpropagation serves as our main source of inspiration, which we review first. The point of departure in this work is to represent the dynamics of BP or gradient descent (GD) for the standard ℓ_2 risk minimization, as in (2) and (1) with $\lambda = 0$. In this case, the associated ODE to GD reads:

$$\frac{d\mathbf{w}_k}{dt}(t) = \frac{a_k}{\sqrt{m}} \mathbf{L}(\mathbf{w}_k(t))(\mathbf{y} - \mathbf{f}(t)), \tag{7}$$

where $\mathbf{y}, \mathbf{f}(t)$ are respectively the vectors of $\{y_k\}$ and $f(\mathbf{x}_k)$, calculated in (1) by replacing $\mathbf{w}_k = \mathbf{w}_k(t)$. Moreover, the matrix $\mathbf{L}(\mathbf{w})$ consists of $\sigma'(\mathbf{w}^T \mathbf{x}_k)\mathbf{x}_k$ as its k^{th} column. While the dynamics in (7) is generally difficult to analyze, we identify two simplifying ingredients in the study of [13]. First, it turns the attention from the dynamics of weights to the dynamics of the function, as reflected by the following relation:

$$\frac{d\mathbf{f}}{dt}(t) = \sum_k \frac{a_k}{\sqrt{m}} \mathbf{L}_k^T \frac{d\mathbf{w}_k}{dt} = \mathbf{H}(t)(\mathbf{y} - \mathbf{f}(t)) \tag{8}$$

where $\mathbf{L}_k = \mathbf{L}_k(t)$ is a short-hand notation for $\mathbf{L}(\mathbf{w}_k(t))$ and $\mathbf{H}(t) = \sum_k a_k^2 \mathbf{L}_k^T \mathbf{L}_k / m$. The second element in the proof can be formulated as follows:

Kernel Hypothesis (KH): In the asymptotic case of $m \to \infty$, the dynamics of $\mathbf{H}(t)$ has a negligible effect, such that it may be replaced by $\mathbf{H}(0)$, resulting to a linear dynamics.

The reason for our terminology of the KH is that under this assumption, the dynamics of BP resembles that of a kernel regularized least squares problem. The investigation in [13] further establishes KH under mild assumptions and further notes that for random initialization of weights $\mathbf{H}(0)$ concentrates on its mean value, denoted by $\mathbf{H}^\infty$.

3.1 Dynamics of Knowledge Transfer

Following the methodology of [13], we proceed by providing the dynamics of the GD algorithm for the optimization problem in (2) with $\lambda > 0$. Direct calculation of the gradient leads us to the following associated ODE for GD:

$$\frac{d\mathbf{w}_k}{dt} = \mathbf{L}_k \left[\frac{a_k}{\sqrt{m}} (\mathbf{y} - \mathbf{f}(t)) + \lambda(\boldsymbol{\phi}^{(k)} - \mathbf{f}^{(k)}) \right], \tag{9}$$

where $\mathbf{L}_k, \mathbf{y}, \mathbf{f}(t)$ are similar to the previous case in (7). Furthermore, $\mathbf{f}^{(k)}, \boldsymbol{\phi}^{(k)}$ are respectively the vectors of $\{f^{(k)}(\mathbf{x}_i)\}_i$ and $\{\phi^{(k)}(\mathbf{x}_i)\}_i$. We may now apply the methodology of [13] to obtain the dynamics of the features. We also observe that unlike this work, the hidden features $\{\mathbf{f}^{(k)}\}$ explicitly appear in the dynamics:

$$\frac{d\mathbf{f}^{(k)}}{dt}(t) = \mathbf{L}_k^T \frac{d\mathbf{w}_k}{dt} =$$
$$\mathbf{H}_k(t) \left[\frac{a_k}{\sqrt{m}} (\mathbf{y} - \mathbf{f}(t)) + \lambda(\boldsymbol{\phi}^{(k)} - \mathbf{f}^{(k)}(t)) \right], \tag{10}$$

where $\mathbf{H}_k(t) = \mathbf{L}_k^T \mathbf{L}_k$ and

$$\mathbf{f}(t) = \sum_k \frac{a_k}{\sqrt{m}} \mathbf{f}^{(k)}(t). \tag{11}$$

This relation will be central in our analysis and we may slightly simplify it by introducing $\boldsymbol{\delta}^{(k)} = \mathbf{f}^{(k)} - \mathbf{f}_\infty^{(k)}$ and $\boldsymbol{\delta} = \mathbf{f} - \mathbf{f}_\infty$. In this case the dynamics in (10) and (11) simplifies to:

$$\frac{d\boldsymbol{\delta}^{(k)}}{dt}(t) = -\mathbf{H}_k(t) \left[\frac{a_k}{\sqrt{m}} \boldsymbol{\delta}(t) + \lambda \boldsymbol{\delta}^{(k)}(t) \right],$$
$$\boldsymbol{\delta}(t) = \sum_k \frac{a_k}{\sqrt{m}} \boldsymbol{\delta}^{(k)}(t). \tag{12}$$

Finally, we give a more abstract view on the relation in (12) by introducing the block vector $\boldsymbol{\eta}(t)$ where the k^{th} block is given by $\boldsymbol{\delta}^{(k)}(t)$. Then, we may write (12) as

$$\frac{d\boldsymbol{\eta}}{dt}(t) = -\bar{\mathbf{H}}(t)\boldsymbol{\eta}(t), \tag{13}$$

where $\bar{\mathbf{H}}(t)$ is a block matrix with $\mathbf{H}_k(t) \left(\frac{a_k a_l}{m} + \lambda \delta_{k,l} \right)$ as its k, l block ($\delta_{k,l}$ denotes the Kronecker delta function).

3.2 Dynamics Under Kernel Hypothesis

Now, we follow [13] by simplifying the relation in (10) under the kernel hypothesis, which in this case assumes the matrices $\bar{\mathbf{H}}(t)$ to be fixed to its initial value $\bar{\mathbf{H}} = \bar{\mathbf{H}}(0)$, leading again to a linear dynamics:

$$\boldsymbol{\eta}(t) = e^{-\bar{\mathbf{H}}t} \boldsymbol{\eta}(0). \tag{14}$$

Despite similarities with the case in [13, 1], the relation in (14) is not simple to analyze due to the asymmetry in $\bar{\mathbf{H}}$ and the complexity of its eigen-structure. For this reason, we proceed by taking the Laplace transform of (12) (assuming $\mathbf{H}_k = \mathbf{H}_k(t) = \mathbf{H}_k(0)$) which after straightforward manipulations gives:

$$\boldsymbol{\Delta}^{(k)}(s) = (s\mathbf{I} + \lambda \mathbf{H}_k)^{-1} \left[\boldsymbol{\delta}^{(k)}(0) - \frac{a_k}{\sqrt{m}} \mathbf{H}_k \boldsymbol{\Delta}(s) \right]$$
$$\boldsymbol{\Delta}(s) = (\mathbf{I} + \mathbf{T}(s))^{-1} \sum_k (s\mathbf{I} + \lambda \mathbf{H}_k)^{-1} \boldsymbol{\delta}^{(k)}(0), \tag{15}$$

where $\boldsymbol{\Delta}^{(k)}(s), \boldsymbol{\Delta}(s)$ are respectively the Laplace transform of $\boldsymbol{\delta}^{(k)}(y), \boldsymbol{\delta}(t)$. Hence, $\boldsymbol{\delta}(t)$ is given by taking the inverse Laplace transform of $\boldsymbol{\Delta}(s)$. Note that by construction, $\boldsymbol{\Delta}(s)$ is a rational function, which shows the finite order of the dynamics. To find the inverse Laplace transform, we only need to find the poles of $\boldsymbol{\Delta}(s)$. These poles can only be either among the eigenvalues of $\mathbf{H}_k$ or the values $-p_k$ where the matrix $\mathbf{I} + \mathbf{T}(s)$ becomes rank deficient. Under Assumption 1 and 2, we may conclude that the poles are only $-p_k$, which gives the result in Theorem 1 and 2. More details of this approach can be found in the supplement, where the kernel hypothesis for this case is also rigorously proved.

4 Experimental Results

We perform our numerical analysis on a commonly-used dataset for validating deep neural models, i.e., CIFAIR-10. This dataset is used for the experiments in [1]. As in [1], we only look at the first two classes and set the label $y_i = +1$ if image i belongs to the first class and $y_i = -1$ if it belongs to the second class. The images $\{x_i\}_{i=1}^n$ are normalized such that $||\boldsymbol{x}_i||_2 = 1$ for all $i = 1, \ldots, n$.

The weights in our model are initialized as follows:

$$\boldsymbol{w}_i \sim \mathcal{N}(0, k^2 \boldsymbol{\mathcal{I}}),\ a_r \sim \text{Unif}(\{-1, 1\}), \forall r \in [m]. \tag{16}$$

For optimization, we use (full batch) gradient descent with the learning rate η. In our experiments we set $k = 10^{-2}, \eta = 2 \times 10^{-4}$ similar to [1].

In all of our experiments we use 100 hidden neurons for the teacher network and 20 hidden neurons for the student network.

4.1 Dynamics of knowledge transfer

In this section we study knowledge transfer and distillation in different settings. We first consider a finite regularization in Eq. 2 by setting $\lambda = 0.01$. Figure 1 shows the dynamics of the results in different settings, i) no teacher, i.e., the student is independently trained without access to a teacher, ii) student training, where the student is trained by both the teacher and the true labels according to Eq. 2, and iii) the teacher, trained by only the true labels. For each setting, we illustrate the training loss (Figure 1(a)) and the test loss (Figure 1(b)). Note that true labels are the same for the teacher and the students. Teacher shows the best performance because of its considerably larger capacity. On the other hand, we observe, i) for the student with access to the teacher its performance is better than the student without access to the teacher. This corresponds to the result in Theorem 1 and the discussion in section 2.7, where the final performance of the student is shown to be improved by the teacher. ii) the convergence rate of the optimization is significantly faster for the student with teacher compared to the other alternatives. This confirms the prediction of Theorem 3. This experiment implies the importance of a proper knowledge transfer to the student network via the information from the teacher.

In the following, we study two special cases of the generic formulation in Eq. 2 where $\lambda \to 0$ and $\lambda \to \infty$. As discussed in section 2.5.1, the case $\lambda = 0$ corresponds to the lottery ticket setup. Figure 2 compares these two extreme cases with the student with $\lambda = 0.01$ and the teacher w.r.t. training loss (Figure 2(a)) and test loss (Figure 2(b)). We observe that the student with a finite regularization ($\lambda = 0.01$) outperforms the two other students in terms of both convergence rate (optimization speed) and the quality of the results. In particular, when the student is trained with $\lambda \to \infty$ and it is initialized with the weights of the teacher, then the generic loss in Eq. 2 equals 0. This renders the student network to keep its weights unchanged for $\lambda \to \infty$ and the performance remains equal to that of the privileged knowledge $\sum_k \frac{a_k}{\sqrt{m}} \phi_k$ without the label input.

4.2 Dynamics of knowledge transfer with imperfect teacher

In this section, we study the impact of the quality of the teacher on the student network. We consider the student-teacher scenario in three different settings, i) perfect teacher where the student is initialized with the final weights of the teacher and uses the final teacher outputs in Eq. 2, ii) imperfect teacher where the student is initialized with the intermediate (early) weights of the teacher network and uses the respective intermediate teacher outputs in Eq. 2, and iii) no student initialization where the student is initialized randomly but uses the final teacher outputs. In all the settings, we assume $\lambda = 0.01$.

Figure 3 shows the results for these three settings, respectively w.r.t. training loss (Figure 3(a)) and test loss (Figure 3(b)). We observe that initializing and training the student with the perfect (fully trained) teacher yields the best results in terms of both quality (training and test loss) and convergence rate (optimization speed). This observation verifies our theoretical analysis on the importance of initialization of the student with fully trained teacher, as the student should be very close to the teacher.

4.3 Kernel embedding

In order to provide the teacher and the student with more relevant information and to study the role of the data geometry, as stated by Theorem 3, we can use properly designed kernel embeddings. Specifically, instead of using the original features for the networks, we could first learn an optimal kernel which is highly aligned with the labels in training

7
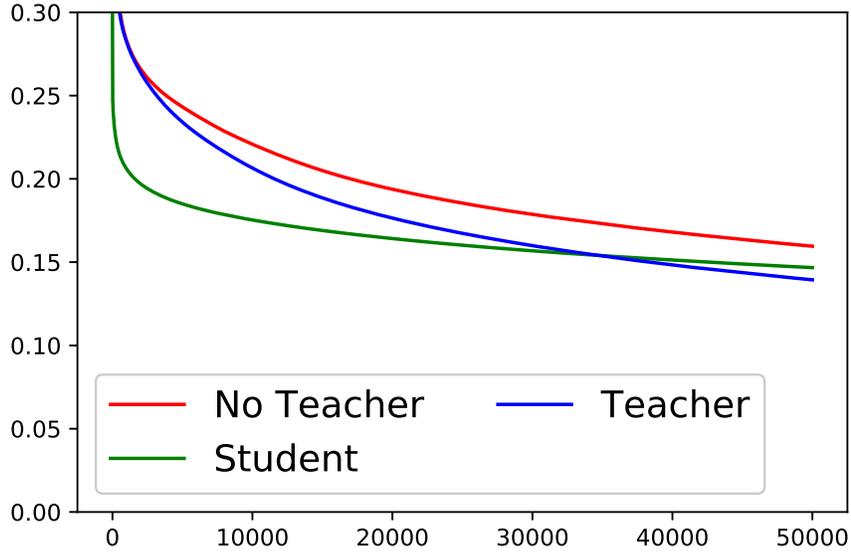
(a) Training loss
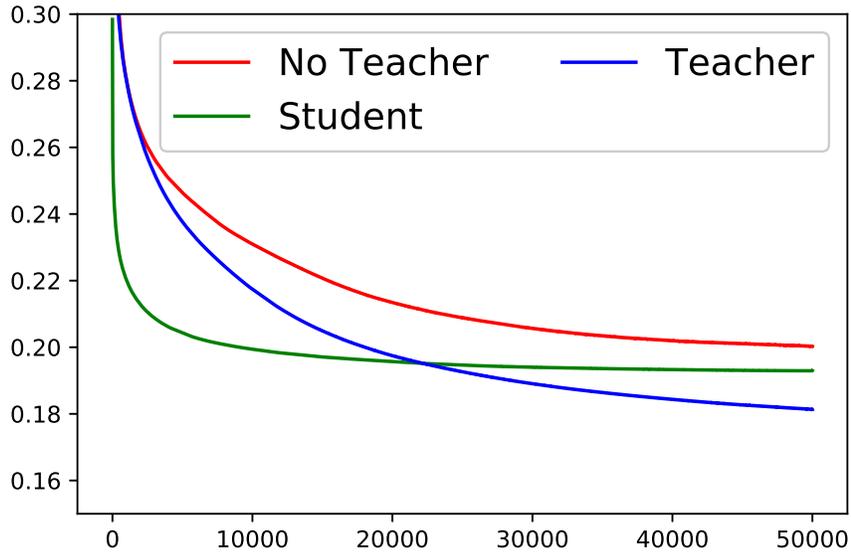
(b) Test loss

Figure 1: Dynamics of knowledge transfer from a teacher to the student, where the regularization parameter is set by $\lambda = 0.01$. The student supported by the knowledge from the teacher performs significantly better than the student with no teacher in terms of both the quality and the convergence rate of the optimization.

data, implicitly improving the combination of $\{\alpha_j\}$ in Theorem 3 and then we feed the features induced by that kernel embedding into the networks (both student and teacher).

For this purpose, we employ the method proposed in [15] that develops an algorithm to learn a new kernel from a group of kernels according to a similarity measure between the kernels, namely centered alignment. Then, the
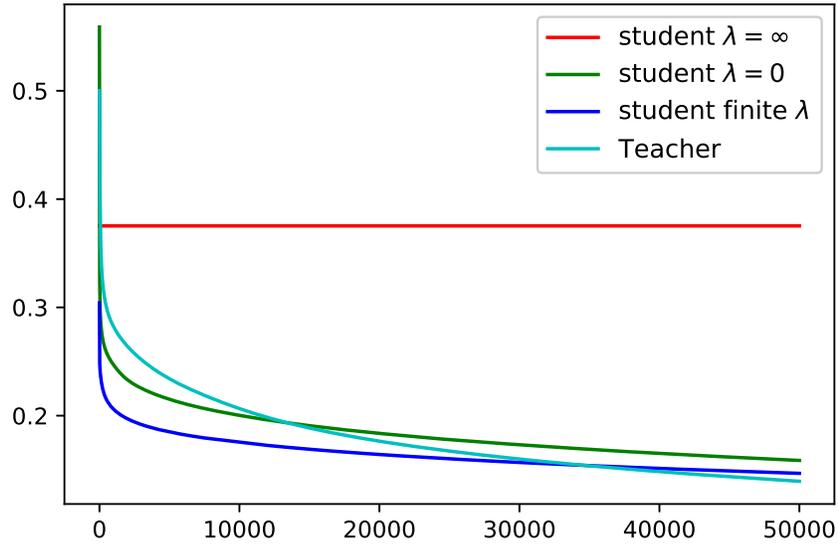
(a) Train loss
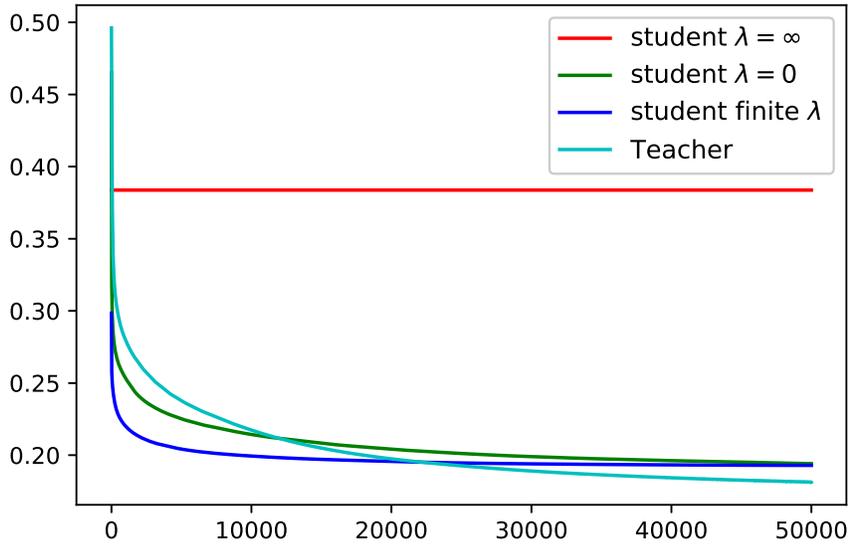
(b) Test loss

Figure 2: Performance of the student networks with different regularisation λ. The finite regularization (i.e., $\lambda = 0.01$) yields the best results for both training and test in terms of quality and convergence rate.

problem of learning a kernel with a maximum alignment between the input data and the labels is formulated as a quadratic programming (QP) problem. The respective algorithm is known as *alignf* [15].

Let us denote by K^c the centered variant of a kernel matrix K. To obtain the optimal combination of the kernels (i.e., a weighted combination of some base kernels), [15] suggests the objective function to be centered alignment between
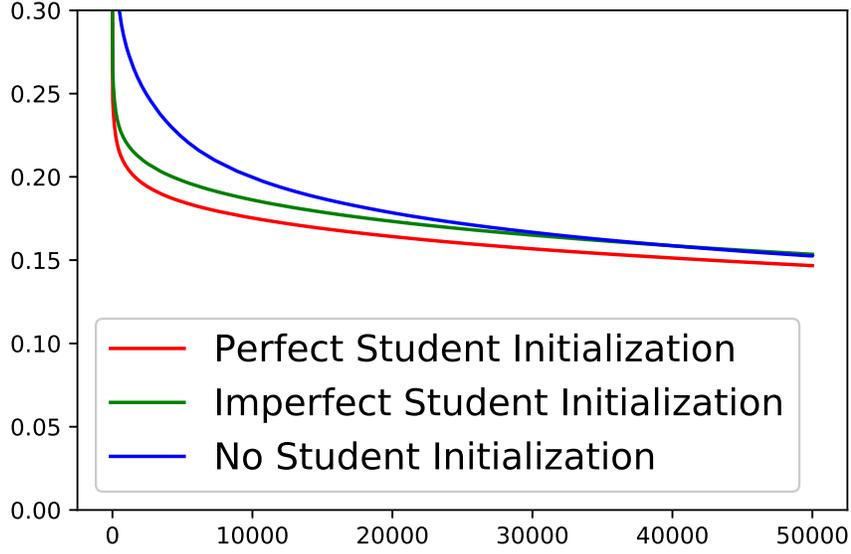
(a) Train loss
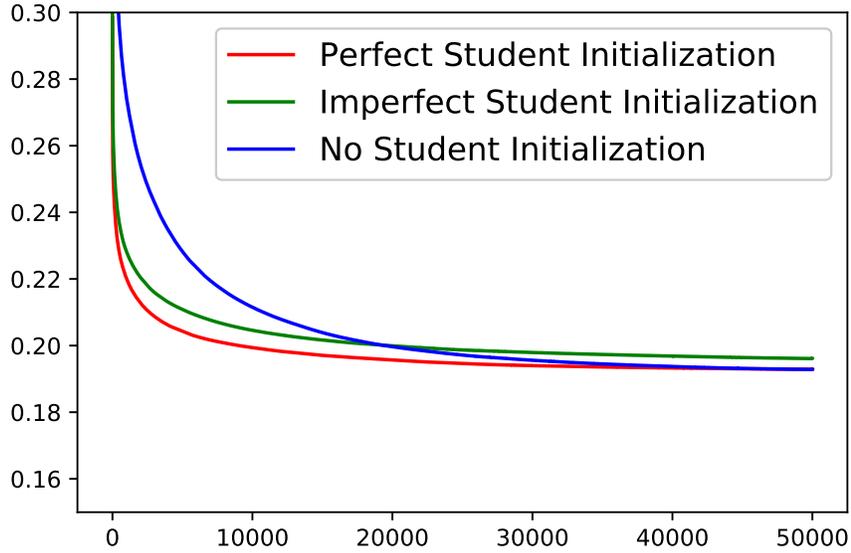
(b) Test loss

Figure 3: Dynamics of knowledge transfer with perfect and imperfect teacher w.r.t. the quality of the teacher.

the combination of the kernels and yy^T, where y is the true labels vector. By restricting the weights to be non-negative, a QP can be formulated as

$$minimize \; v^T M v - 2v^T a \; w.r.t. \; v \in R_+^P \tag{17}$$

P is the number of the base kernels and $M_{kl} = \langle K_k^c, K_l^c \rangle_F$ for $k, l \in [1, P]$, and finally a is a vector wherein $a_i = \langle K_i^c, yy^T \rangle_F$ for $i \in [1, P]$. If v^* is the solution of the QP, then the vector of kernel weights is given by [15, 16]

$$\mu^* = v^*/\|v^*\| \tag{18}$$

Using this algorithm we learn an optimal kernel based on seven different Gaussian kernels. Then, we need to approximate the kernel embeddings. To do so, we use the Nyström method [17]. Then we feed the approximated embeddings to the neural networks. The results in Figure 4 clearly show that using the kernel embeddings as inputs to the neural networks, helps both teacher and student networks in terms of training loss (Figure 4(a)) and test loss (Figure 4(b)).

4.4 Spectral analysis

Here, we investigate the overlap parameter of different networks, where we compute a simplified but conceptually consistent variant of the overlap parameter α_j in theorem 3. For a specific network, we consider the normalized columns of matrix ϕ (as defined in Eq. 2) corresponding to the nonlinear outputs of the hidden neurons, and compute the dot product of each column with the top eigenvectors of $\mathbf{H}^\infty$, and take the average. We repeat this for all the columns and depict the histogram. For a small value of λ, the resulting values are approximately equal to $\{\alpha_j\}$ in Theorem 3.

Figure 5 shows such histograms for two settings. In Figure 5(a) we compare the overlap parameter for two teachers, one trained partially (imperfect teacher) and the other trained fully (perfect teacher). We observe that the overlap parameter is larger for the teacher trained perfectly, i.e., there is more consistency between its outputs ϕ and the matrix $\mathbf{H}^\infty$. This analysis is consistent with the results in Figure 3 which demonstrates the important of fully trained (perfect) teacher. In Figure 5(b), we show that this improvement is transferred to the student via distillation.

We confirm this point further: if the teacher learns better representations via kernels, this is also transferred by distillation to the student.

4.5 Kernel spectral analysis

Finally, we perform a similar analysis to section 4.4, where we compute the overlap parameter for a teacher (Figure 6(a)) and a student (Figure 6(b)) trained with the representations from an optimal kernel described in section 4.3. The histograms show that the overlap parameter for a teacher trained with kernel embeddings is larger than the teacher trained with original features. We also observe that this better overlap parameter, i.e., higher consistency between ϕ and H^∞, is transferred to the student trained with the same kernel representations.

5 Related Work

In its current and most widely known form, distillation was introduced by Hinton et al. [5] with the aim of compressing neural networks. Since then, distillation has quickly gained popularity among practitioners and established its place in deep learning folklore. It has been found to work well across a wide range of applications, including

In contrast to its empirical success, a rigorous theoretical understanding of the principles underlying the effectiveness of distillation have largely remained a mystery. Until recently, Lopez-Paz et al. [12] was the only work that tried to examine distillation from a theoretical perspective. It casts distillation as a form of Vapnik's notion of learning using privileged information, a learning setting in which additional per-instance information is available at training time but not at test time. Their paper is more a heuristic argument for the effectiveness of distillation with repect to generalization error rather than a rigorous analysis. Very recently, Phuong and Lampert [9] made the first attempt to analyze distillation in a simple model. In their setting, both the teacher and the student are *linear* classifiers (although the student's weight vector is allowed a over-parametrized representation as a product of matrices). They give conditions under which the student's weight vector converges (approximately) to that of the teacher and derive consequences for generalization error. Crucially, their analysis is limited to linear networks while we analyze distillation in the context of non-linear networks.

A series of recent papers [1, 2, 3] achieved breakthroughs in understanding how (infinitely) wide neural network training behaves in the so-called *kernel regime*. In this regime, the dynamics of training by gradient descent can be approximated by the dynamics of a linear system. In this paper, we extend the repertoire of the methods that can be applied in such settings.
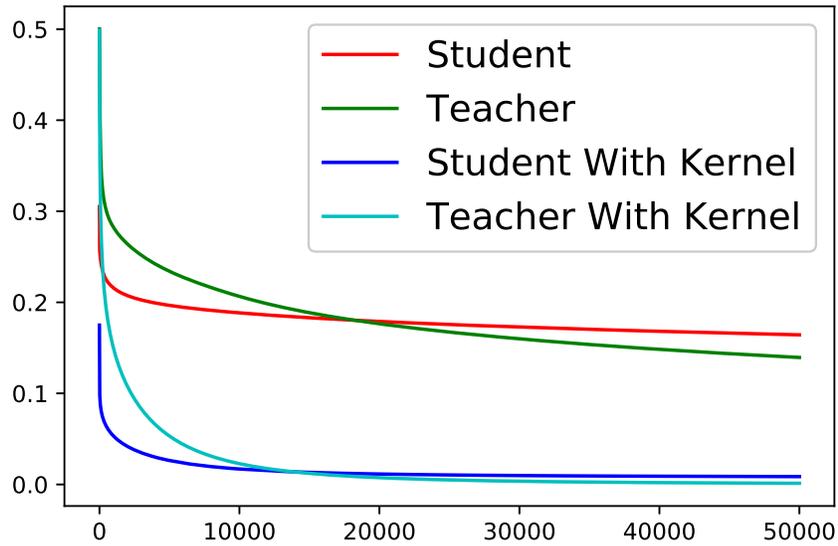
(a) Train loss
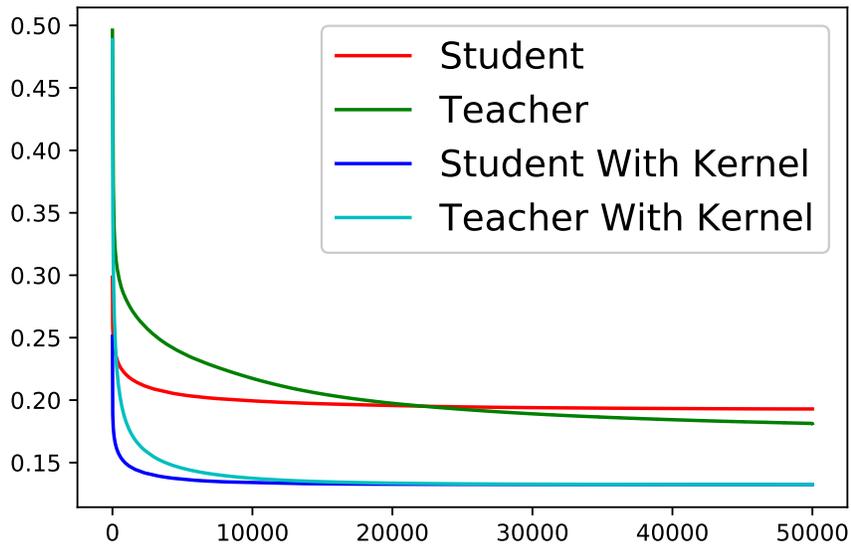
(b) Test loss

Figure 4: Dynamics of knowledge transfer with teacher trained using kernel embedding. The kernel embeddings help both teacher and student networks in terms of training loss (Figure 4(a)) and test loss.

6   Conclusions

We give the first theoretical analysis of knowledge distillation for non–linear neural networks in the model and regime of [1, 2, 3]. We provide results for both what is learnt by the student and on the speed of convergence. As an intriguing
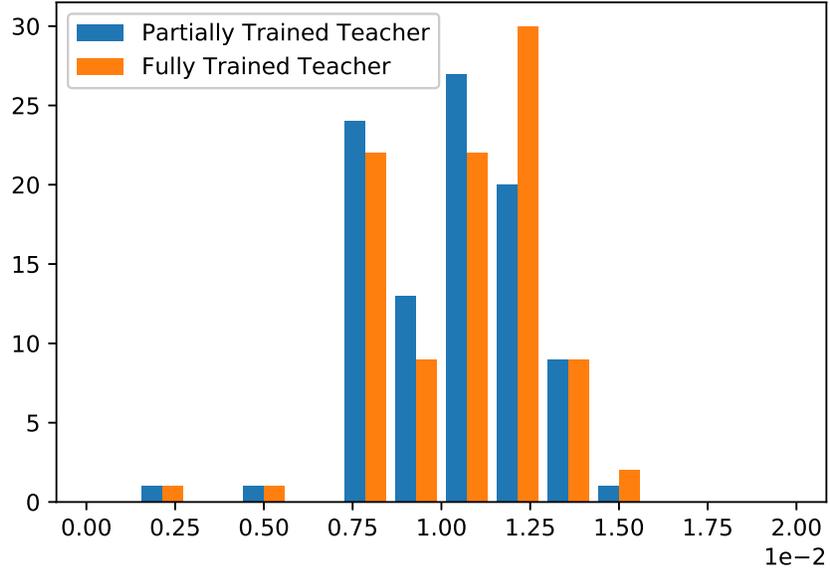
(a) Overlap parameter, partial vs fully trained teacher
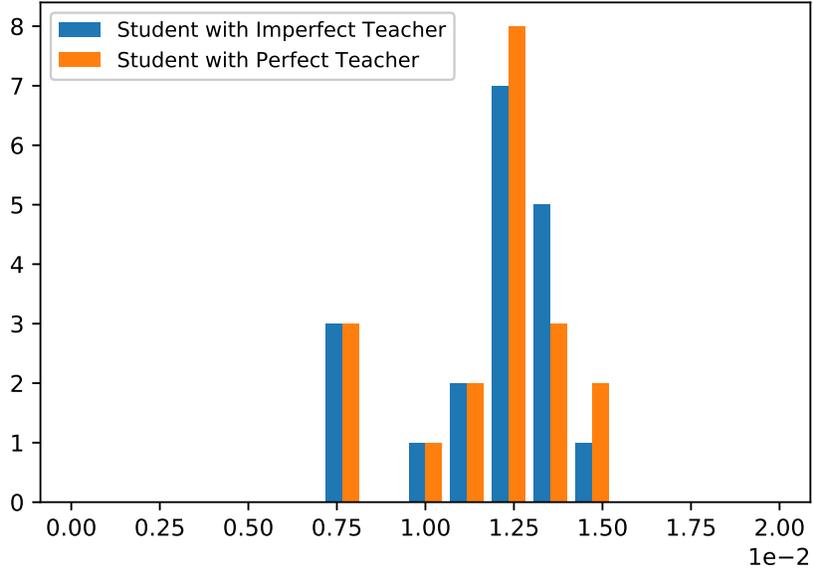
(b) Overlap parameter, student with imperfect and perfect teacher

Figure 5: Spectral analysis of teachers (5(a)) and students (5(b)) in different settings w.r.t. the overlap parameters. Perfect (full) training improves the overlap compared with the imperfect (partial) training.

side result, we also confirm the "lottery ticket hypothesis" [4] in this model and regime. Our numerical studies further confirm our theoretical findings on the role of data geometry and distillation in the final performance of student.
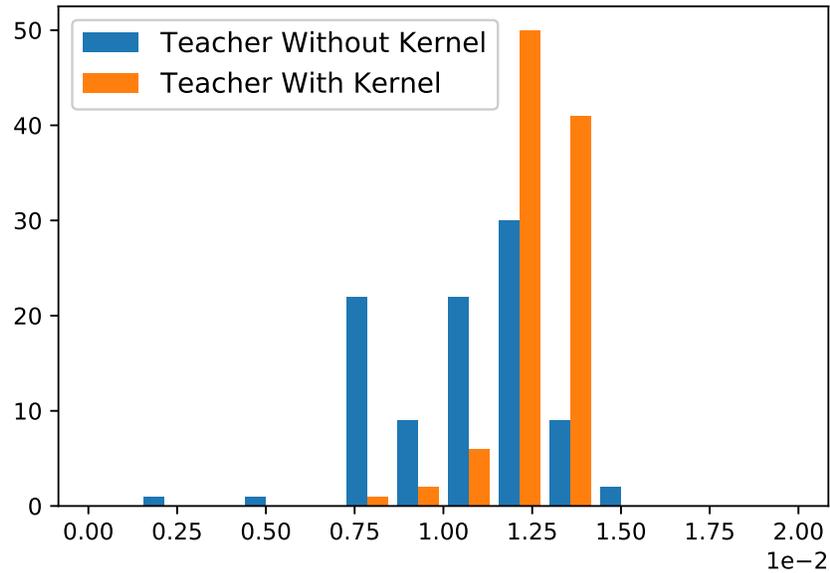
(a) Overlap parameter, teacher with and without kernel
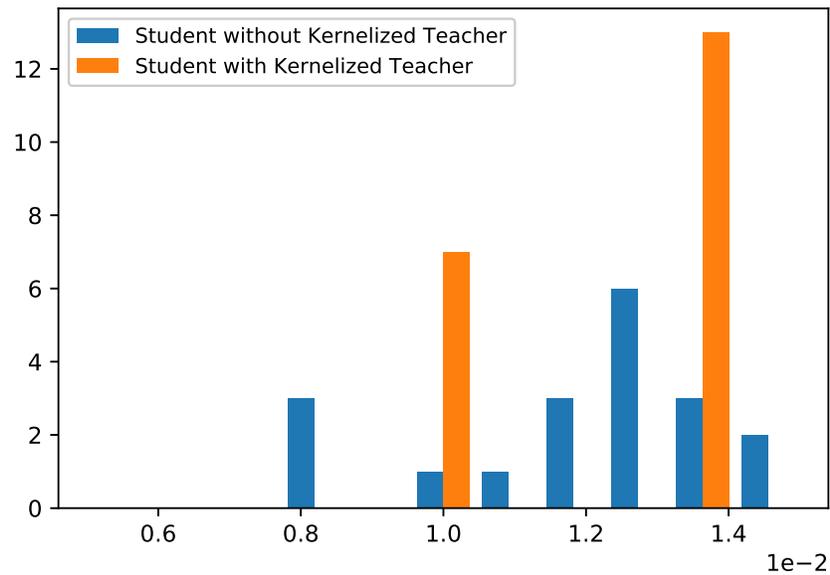
(b) Overlap parameter, student with and without kernelized teacher

Figure 6: Spectral analysis of teachers (6(a)) and students (6(b)) in different settings w.r.t. the overlap parameters. Using kernel embeddings improves the overlap compared with normal training.

Acknowledgement

References

[1] Sanjeev Arora, Simon S. Du, Wei Hu, Zhiyuan Li, and Ruosong Wang. Fine-grained analysis of optimization and generalization for overparameterized two-layer neural networks. In *Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA*, pages 322–332, 2019.

[2] Simon S. Du and Wei Hu. Width provably matters in optimization for deep linear neural networks. In *Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA*, pages 1655–1664, 2019.

[3] Yuan Cao and Quanquan Gu. Generalization bounds of stochastic gradient descent for wide and deep neural networks. In *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, 8-14 December 2019, Vancouver, BC, Canada*, pages 10835–10845, 2019.

[4] Jonathan Frankle and Michael Carbin. The lottery ticket hypothesis: Finding sparse, trainable neural networks. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*, 2019.

[5] Geoffrey E. Hinton, Oriol Vinyals, and Jeffrey Dean. Distilling the knowledge in a neural network. *CoRR*, abs/1503.02531, 2015.

[6] Guobin Chen, Wongun Choi, Xiang Yu, Tony Han, and Manmohan Chandraker. Learning efficient object detection models with knowledge distillation. In *Advances in Neural Information Processing Systems*, pages 742–751, 2017.

[7] Junho Yim, Donggyu Joo, Jihoon Bae, and Junmo Kim. A gift from knowledge distillation: Fast optimization, network minimization and transfer learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4133–4141, 2017.

[8] Ruichi Yu, Ang Li, Vlad I Morariu, and Larry S Davis. Visual relationship detection with internal and external linguistic knowledge distillation. In *Proceedings of the IEEE international conference on computer vision*, pages 1974–1982, 2017.

[9] Mary Phuong and Christoph Lampert. Towards understanding knowledge distillation. In Kamalika Chaudhuri and Ruslan Salakhutdinov, editors, *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pages 5142–5151, Long Beach, California, USA, 09–15 Jun 2019. PMLR.

[10] Song Mei, Theodor Misiakiewicz, and Andrea Montanari. Mean-field theory of two-layers neural networks: dimension-free bounds and kernel limit. *arXiv preprint arXiv:1902.06015*, 2019.

[11] Arthur Jacot, Franck Gabriel, and Clement Hongler. Neural tangent kernel: Convergence and generalization in neural networks. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing Systems 31*, pages 8571–8580. Curran Associates, Inc., 2018.

[12] David Lopez-Paz, Léon Bottou, Bernhard Schölkopf, and Vladimir Vapnik. Unifying distillation and privileged information. In *4th International Conference on Learning Representations, ICLR 2016, San Juan, Puerto Rico, May 2-4, 2016, Conference Track Proceedings*, 2016.

[13] Simon S Du, Xiyu Zhai, Barnabas Poczos, and Aarti Singh. Gradient descent provably optimizes overparameterized neural networks. *arXiv preprint arXiv:1810.02054*, 2018.

[14] Jonathan Frankle and Michael Carbin. The lottery ticket hypothesis: Finding sparse, trainable neural networks. *arXiv preprint arXiv:1803.03635*, 2018.

[15] Corinna Cortes, Mehryar Mohri, and Afshin Rostamizadeh. Algorithms for learning kernels based on centered alignment. *J. Mach. Learn. Res.*, 13:795–828, March 2012.

[16] Mehmet Gönen and Ethem Alpayd. Multiple kernel learning algorithms. *J. Mach. Learn. Res.*, 12:2211–2268, July 2011.

[17] Christopher K. I. Williams and Matthias Seeger. Using the nyström method to speed up kernel machines. In T. K. Leen, T. G. Dietterich, and V. Tresp, editors, *Advances in Neural Information Processing Systems 13*, pages 682–688. MIT Press, 2001.

7 Supplementary Material

7.1 Errata

In this long version we have made few clarifications compared to the submitted version according to the following corrections.

After (6 in paper), $\frac{1}{n^{10}}$ must be replaced by $\frac{1}{m^{10}}$.

In Assumption 2, the number n of values p_i must be changed to $d = n \times m$.

In Assumption 3, we must also have the derivative function σ' to be also Lipschitz.

In assumption 4, $O(1/m)$ can be changed to $O(1/\sqrt{m})$. We also assume that $\|\mathbf{y} - \sum_k \frac{\mathbf{a}_k \phi_k}{\sqrt{m}}\|$ remains bounded.

Two more assumptions $\|\mathbf{x}_i\|$s and a_ks are bounded.

We clarify that in theorem 2, $+O(1/\sqrt{m})$ holds in L_1 sense, i.e. $a(t) = b(t) + O(1/\sqrt{m})$ means $\int_0^\infty \|a(t) - b(t)\|_2 dt = O(1/\sqrt{m})$.

In Assumption 2, we found out that $\mathbf{T}(s)$ is symmetric. Hence $\mathbf{u}^j = \mathbf{v}^j$, but we keep the expression unchanged.

Theorem 2 also needs the assumption that $\sum_k \frac{\bar{a}_k^2}{\bar{m}} \|\bar{\phi}_k\|^2$ is bounded (it can be incorporated in C, but the theorem does not reflect the dependency on $\bar{m}$).

7.2 Notation

$\sigma_{\max}$ is the largest singular value and $\|\|$ means 2-norm of vectors.

7.3 Proof of Theorem 1 and 3 in Paper

We continue the discussion in (15 of paper). Note that the values $s = -p_i$ correspond to the points where $\det(\mathbf{I} + \mathbf{T}(s)) = 0$. We also observe that these values correspond to the negative of eigenvalues of the matrix $\bar{\mathbf{H}}(t = 0)$. We conclude that under assumption 2, the eigenvalues of $\bar{\mathbf{H}} = \bar{\mathbf{H}}(t = 0)$ are distinct and strictly positive, hence this matrix is diagonalizable. Now, we write $\bar{\mathbf{H}}(t) = \bar{\mathbf{H}} + \Delta\bar{\mathbf{H}}(t)$ and state the following lemma:

Lemma 1. *Suppose that $\bar{\mathbf{H}}$ is a diagonalizable matrix with strictly positive eigenvalues and denote its smallest eigenvalue by p. Take $\Delta\bar{\mathbf{H}}(t)$ ar a matrix valued function of the continuous valiable t such that for a given fixed value of t*

$$q = q(t) = \frac{\sup_{\tau \in [0 \; t]} \sigma_{\max}(\Delta\bar{\mathbf{H}}(\tau))}{p} < 1$$

Let $\boldsymbol{\eta}(t)$ denote the solution to (13 in paper) with $\bar{\mathbf{H}}(t) = \bar{\mathbf{H}} + \Delta\bar{\mathbf{H}}(t)$. Then,

$$\int_0^t \left\| \boldsymbol{\eta}(\tau) - e^{-\bar{\mathbf{H}}\tau} \boldsymbol{\eta}(0) \right\| d\tau \leq \frac{q\|\boldsymbol{\eta}(0)\|}{p(1-q)}$$

Proof. Consider the iteration $\boldsymbol{\eta}^{r+1} = \mathcal{T}\boldsymbol{\eta}^r$ that generates a sequence of function functions $\boldsymbol{\eta}^r(t)$ for $r = 0, 1, \ldots$ where $\boldsymbol{\eta}^r(0) = e^{-\bar{\mathbf{H}}t}\boldsymbol{\eta}(0)$ and $\boldsymbol{\eta}' = \mathcal{T}\boldsymbol{\eta}$ is the solution to

$$\frac{\mathrm{d}}{\mathrm{d}t}\boldsymbol{\eta}'(t) = -\bar{\mathbf{H}}\boldsymbol{\eta}'(t) - \Delta\bar{\mathbf{H}}(t)\boldsymbol{\eta}(t) \tag{19}$$

with $\boldsymbol{\eta}'(0) = \boldsymbol{\eta}(0)$, which can also be written as

$$\boldsymbol{\eta}'(t) = e^{-\bar{\mathbf{H}}t}\boldsymbol{\eta}(0) - \int_0^t e^{-\bar{\mathbf{H}}(t-\tau)}\Delta\bar{\mathbf{H}}(\tau)\boldsymbol{\eta}(\tau)\mathrm{d}\tau \tag{20}$$

We observe that $\mathcal{T}$ on the interval $[0\ t]$ is a contraction map under L_1 norm as we have

$$\Delta\boldsymbol{\eta}'(t) = -\int_0^t e^{-\bar{\mathbf{H}}(t-\tau)}\Delta\bar{\mathbf{H}}(\tau)\Delta\boldsymbol{\eta}(\tau)\mathrm{d}\tau \tag{21}$$

and hence by the triangle inequality on the Hilbert space of functions, we get

$$\|\Delta\boldsymbol{\eta}'(t)\|_2 \leq \sup_{\tau\in[0\ t]}\sigma_{\max}(\Delta\bar{\mathbf{H}}(\tau)) \times$$
$$\int_0^t e^{-p(t-\tau)}\|\Delta\boldsymbol{\eta}(\tau)\|_2\mathrm{d}\tau \tag{22}$$

we conclude that

$$\int_0^t \|\Delta\boldsymbol{\eta}'(\tau)\|_2\mathrm{d}t \leq$$

$$\sup_{\tau\in[0\ t]}\sigma_{\max}(\Delta\bar{\mathbf{H}}(\tau)) \times \int_0^t \frac{1-e^{-p(t-\tau)}}{p}\|\Delta\boldsymbol{\eta}(\tau)\|_2$$

$$\leq q\int_0^t \|\Delta\boldsymbol{\eta}(\tau)\|_2\mathrm{d}t$$

which shows that $\mathcal{T}$ is a contraction. Then, from Banach fixed-point theorem we conclude that $\boldsymbol{\eta}^r$ converges uniformly on the interval $[0\ t]$ to the fixed-point $\boldsymbol{\eta}$ of $\mathcal{T}$, which coincides with the solution of (13 in paper). Moreover,

$$\int_0^t \|\boldsymbol{\eta}-\boldsymbol{\eta}^0(\tau)\|_2\mathrm{d}\tau \leq \frac{\int_0^t \|\boldsymbol{\eta}^1(\tau)-\boldsymbol{\eta}^0(\tau)\|_2\mathrm{d}\tau}{1-q}$$

Now, we observe that

$$\boldsymbol{\eta}^1(t)-\boldsymbol{\eta}^0(t) = -\int_0^t e^{-\bar{\mathbf{H}}(t-\tau)}\Delta\bar{\mathbf{H}}(\tau)e^{-\bar{\mathbf{H}}(\tau)}\boldsymbol{\eta}(0)\mathrm{d}\tau$$

Hence,

$$\|\boldsymbol{\eta}^1(t)-\boldsymbol{\eta}^0(t)\| \leq \int_0^t e^{-p(t-\tau)}\sigma_{\max}(\Delta\bar{\mathbf{H}}(\tau))e^{-p\tau}\|\boldsymbol{\eta}(0)\|\mathrm{d}\tau$$
$$\leq te^{-pt}\sup_{\tau\in[0\ t]}\sigma_{\max}(\Delta\bar{\mathbf{H}}(\tau))\|\boldsymbol{\eta}(0)\|$$

and

$$\int_0^t \|\boldsymbol{\eta}^1(\tau)-\boldsymbol{\eta}^0(\tau)\|\mathrm{d}\tau \leq$$

$$\sup_{\tau\in[0\ t]}\sigma_{\max}(\Delta\bar{\mathbf{H}}(\tau))\|\boldsymbol{\eta}(0)\|\int_0^t \tau e^{-p\tau}\mathrm{d}\tau \leq \frac{q}{p}\|\boldsymbol{\eta}(0)\|$$

which completes the proof. $\qquad\square$

Now, we state two results that connect $\sigma_{\max}(\Delta\mathbf{H})$ to the change of $\mathbf{w}_k(t)$:

Lemma 2. *Under Assumption 3, the following relation holds:*

$$\sigma_{\max}(\bar{\mathbf{H}}(t)) \leq \sqrt{2}\sqrt{\lambda^2 + \left(\sum_k \frac{a_k^2}{m}\right)^2}\max_k \sigma_{\max}(\Delta\mathbf{H}_k(t)) \tag{23}$$

where $\Delta\mathbf{H}_k(t) = \mathbf{H}_k(t) - \mathbf{H}_k(0)$.

Proof. Take an arbitrary block vector $\boldsymbol{\eta} = [\boldsymbol{\delta}_k]_k$ with $\|\boldsymbol{\eta}\| = 1$ and note that

$$\|\Delta\bar{\mathbf{H}}(t)\boldsymbol{\eta}\|^2 = \sum_k \|\Delta\mathbf{H}_k(t)(\frac{a_k}{\sqrt{m}}\boldsymbol{\delta} + \lambda\boldsymbol{\delta}_k)\|^2$$

$$\leq 2 \max_k \sigma_{\max}^2(\Delta\mathbf{H}_k(t)) \sum_k \left(\frac{a_k^2}{m}\|\boldsymbol{\delta}\|^2 + \lambda\|\boldsymbol{\delta}_k\|^2\right)$$

$$= 2 \max_k \sigma_{\max}^2(\Delta\mathbf{H}_k(t)) \left(\|\boldsymbol{\delta}\|^2 \sum_k \frac{a_k^2}{m} + \lambda^2\right),$$

where $\boldsymbol{\delta} = \sum_k \frac{a_k}{\sqrt{m}}\boldsymbol{\delta}_k$. We obtain the desired result by observing that

$$\|\boldsymbol{\delta}\|^2 = \left\|\sum_k \frac{a_k}{\sqrt{m}}\boldsymbol{\delta}_k\right\|^2 \leq \left(\sum_k \frac{a_k^2}{m}\right)\sum_k \|\boldsymbol{\delta}_k\|^2$$

$$= \sum_k \frac{a_k^2}{m}.$$

$\square$

Next, we show

Lemma 3. *We have*

$$\sigma_{\max}(\Delta\mathbf{H}_k(t)) \leq L^2\sigma_x^2 \max_i \|\mathbf{x}_i\|^2 \times$$

$$\|\mathbf{w}_k(t) - \mathbf{w}_k(0)\|(\|\mathbf{w}_k(t) - \mathbf{w}_k(0)\| + 2\|\mathbf{w}_k(0)\|)$$

where σ_x is the maximal eigenvalue of the the data matrix $\mathbf{X} = [\mathbf{x}_1\ \mathbf{x}_2\ldots\mathbf{x}_n]$ and L is the largest of the Lipschitz constants of σ, σ'.

Proof. Note that since $\Delta\mathbf{H}_k$ is symmetric, we have (e.g. by eigen-decomposition)

$$\sigma_{\max}(\Delta\mathbf{H}_k(t)) = \max_{\boldsymbol{\delta}|\|\boldsymbol{\delta}\|=1} |\boldsymbol{\delta}^T \Delta\mathbf{H}_k(t)\boldsymbol{\delta}|$$

Taking an arbitrary normalized $\boldsymbol{\delta}$, we observe that

$$\boldsymbol{\delta}^T \Delta\mathbf{H}_k(t)\boldsymbol{\delta} = \left|\|\mathbf{L}(\mathbf{w}_k(t))\boldsymbol{\delta}\|^2 - \|\mathbf{L}(\mathbf{w}_k(0))\boldsymbol{\delta}\|^2\right|$$

On other hand,

$$\mathbf{L}(\mathbf{w}_k(t))\boldsymbol{\delta} = \mathbf{L}(\mathbf{w}_k(0))\boldsymbol{\delta} + \sum_i \mathbf{x}_i\sigma_i\delta_i,$$

where $\sigma_i = \sigma'(\mathbf{w}_k(t)^T\mathbf{x}_i) - \sigma'(\mathbf{w}_k(0)^T\mathbf{x}_i)$. Hence,

$$\boldsymbol{\delta}^T \Delta\mathbf{H}_k(t)\boldsymbol{\delta} \leq \left\|\sum_i \mathbf{x}_i\sigma_i\delta_i\right\|^2$$

$$+2\left\|\sum_i \mathbf{x}_i\delta_i\sigma_i\right\|\|\mathbf{L}(\mathbf{w}_k(0))\boldsymbol{\delta}\|$$

We also observe that

$$\left\|\sum_i \mathbf{x}_i\sigma_i\delta_i\right\| \leq \sigma_x\sqrt{\sum_i \delta_i^2\sigma_i^2}$$

and from Lipschitz continuity,

$$\sigma_i^2 \leq L^2\langle\mathbf{x}_i, \mathbf{w}_k(t) - \mathbf{w}_k(0)\rangle^2 \leq L^2\|\mathbf{x}_i\|^2\|\mathbf{w}_k(t) - \mathbf{w}_k(0)\|^2$$

We conclude that

$$\left\|\sum_i \mathbf{x}_i\sigma_i\delta_i\right\| \leq L\sigma_x\|\mathbf{w}_k(t) - \mathbf{w}_k(0)\| \times \max_i \|\mathbf{x}_i\|$$

Similarly, we obtain

$$\|\mathbf{L}(\mathbf{w}_k(0))\boldsymbol{\delta}\| \leq L\sigma_x \times \max_i \|\mathbf{x}_i\|$$

which completes the proof. $\square$

We finally connect the magnitude of the change $\mathbf{w}_k(t) - \mathbf{w}_k(0)$ to $\boldsymbol{\delta}_k$:

Lemma 4. *With the same definitions as in Lemma 3, we have*

$$\|\mathbf{w}_k(t) - \mathbf{w}_k(0)\| \leq \frac{|a_k|}{\sqrt{m}} L\sigma_x \max_i \|\mathbf{x}_i\| \int_0^t \|\boldsymbol{\delta}_k(\tau)\| d\tau \tag{24}$$

Proof. Note that

$$\mathbf{w}_k(t) - \mathbf{w}_k(0) = -\frac{a_k}{\sqrt{m}} \int_0^t \mathbf{L}(\mathbf{w}_k(\tau))\boldsymbol{\delta}_k(\tau)d\tau$$

and hence

$$\|\mathbf{w}_k(t) - \mathbf{w}_k(0)\| \leq \frac{|a_k|}{\sqrt{m}} \int_0^t \|\mathbf{L}(\mathbf{w}_k(\tau))\boldsymbol{\delta}_k(\tau)\| d\tau$$

From a similar argument as in Lemma 3, we have

$$\|\mathbf{L}(\mathbf{w}_k(\tau))\boldsymbol{\delta}_k(\tau)\| \leq L\sigma_x \max_i \|\mathbf{x}_i\| \|\boldsymbol{\delta}_k(\tau)\|$$

which completes the proof. $\qquad\square$

We may now proceed to the proof of Theorem 1 and 3. Define

$$T = \{t \mid \forall \tau \in [0\ t]; \ q(\tau) < \frac{1}{2}\}$$

Note that T is nonempty as $0 \in T$ and open since q is continuous. We show that for sufficiently large m, $T = [0, \infty)$. Otherwise T is an open interval $[0\ t_0)$ where $q(t_0) = \frac{1}{2}$. For any $t \in T$, we have from Lemma 1

$$A = \int_0^t \|\boldsymbol{\eta}(\tau)\| d\tau \leq \left(\frac{1 - e^{-pt}}{p} + \frac{q}{p(1-q)}\right) \|\boldsymbol{\eta}(0)\|$$

$$\leq \frac{\|\boldsymbol{\eta}(0)\|}{p(1-q)}$$

Denote $b = \max_k |a_k|$ and $B = \max_i \|\mathbf{x}_i\|$ and define

$$C = \frac{\sqrt{2}}{p} L^3 \sigma_B^3 x^3 b \sqrt{\lambda^2 + a^2}$$

Then Lemma 2,3 and 4 give us:

$$q \leq \frac{CA}{\sqrt{m}}$$

We conclude that

$$q(t) \leq \frac{C\|\boldsymbol{\eta}(0)\|}{p(1 - q(t))\sqrt{m}}$$

Note that by Assumption 4, we have

$$\boldsymbol{\delta}_k(0) = \frac{a_k}{(\lambda + a)\sqrt{m}} (\mathbf{y} - \sum_k \frac{a_k \boldsymbol{\phi}_k}{\sqrt{m}})$$

$$+ \boldsymbol{\phi}_k - \mathbf{f}_k(0) = O(1/\sqrt{m})$$

and hence $\boldsymbol{\eta}(0) = O(1)$. This shows that there exists a constant C_0 such that

$$q(t)(1 - q(t)) \leq \frac{C_0}{\sqrt{m}}$$

19

for $t \in T$. But for large values of m this is in contradiction to $q(t_0) = 1/2$. Hence, for such values t_0 does not exist and $q(t) < 1/2$ for all t. We conclude that for sufficiently large values of m we have $q(t) < 3C_0/\sqrt{m}$ for all t. Then, according to Lemma 1 and the monotone convergence theorem, we have

$$\int_0^\infty \left\| \boldsymbol{\eta}(\tau) - e^{-\bar{\mathbf{H}}\tau} \boldsymbol{\eta}(0) \right\| \mathrm{d}\tau = O(\frac{1}{\sqrt{m}})$$

This shows that

$$\lim_{t \to \infty} \| \boldsymbol{\eta}(t) - e^{-\bar{\mathbf{H}}t} \boldsymbol{\eta}(0) \| = 0$$

Note that as $\bar{\mathbf{H}}$ is diagonalizable and has strictly positive eigenvalues, we get that

$$\lim_{t \to \infty} e^{-\bar{\mathbf{H}}t} \boldsymbol{\eta}(0) = 0,$$

which further leads to

$$\lim_{t \to \infty} \boldsymbol{\eta}(t) = 0,$$

This proves Theorem 1. For Theorem 3, we see that

$$\boldsymbol{\eta}(t) = e^{-\bar{\mathbf{H}}t} \boldsymbol{\eta}(0) + O(1/\sqrt{m})$$

It suffices to show that the expression in Theorem 3 coincides with $e^{-\bar{\mathbf{H}}t} \boldsymbol{\eta}(0)$. This is simple to see through the following lemma:

Lemma 5. *According to Assumption 2, the right and left eigenvectors of $\bar{\mathbf{H}}$ corresponding to p_j are respectively given by vectors $\boldsymbol{\eta}_j^{\mathrm{r}} = [\mathbf{v}_k^j]_k$ and $\boldsymbol{\eta}_j^{\mathrm{l}} = [\mathbf{u}_k^j]_k$, where*

$$\mathbf{v}_k^j = \frac{a_k}{\sqrt{m}}(p_j \mathbf{I} - \lambda \mathbf{H}_k)^{-1} \mathbf{H}_k \mathbf{v}^j$$

and

$$\mathbf{u}_k^j = \frac{a_k}{\sqrt{m}}(p_j \mathbf{I} - \lambda \mathbf{H}_k)^{-1} \mathbf{u}^j$$

Moreover, $\mathbf{v}^j = \sum_k \frac{a_k}{\sqrt{m}} \mathbf{v}_k^j$.

Proof. According to the definition of $\bar{\mathbf{H}}$, we have that

$$\mathbf{H}_k(\frac{a_k}{\sqrt{m}}\mathbf{v} + \lambda \mathbf{v}_k^j) = p\mathbf{v}_k^j$$

where $\mathbf{v} = \sum_k \frac{a_k}{\sqrt{m}} \mathbf{v}_k^j$ which gives

$$\mathbf{v}_k^j = \frac{a_k}{\sqrt{m}}(p_j \mathbf{I} - \lambda \mathbf{H}_k)^{-1} \mathbf{H}_k \mathbf{v}$$

Replacing this expression in the definition of $\mathbf{v}$ shows that $\mathbf{v} = \mathbf{v}^j$. The case for $\mathbf{u}_k^j$ is similarly proved. $\square$

Theorem 3 simply follows by replacing the result of Lemma 5 to the eigen-decomposition of $e^{-\bar{\mathbf{H}}t}$:

$$e^{-\bar{\mathbf{H}}t} = \sum_j e^{-p_j t} \mid \boldsymbol{\eta}_j^{\mathrm{r}} \rangle \langle \boldsymbol{\eta}_j^{\mathrm{l}} \mid$$

7.4 Proof of Theorem 2 in Paper

The left hand side of (6) is given by

$$\|\mathbf{f}_\infty - \mathbf{y}\| = \left(\frac{\lambda}{\lambda + a}\right)^2 \left\| \sum_k \frac{a_k \boldsymbol{\phi}_k}{\sqrt{m}} - \mathbf{y} \right\|$$

Now note that

$$\sum_k \frac{a_k \boldsymbol{\phi}_k}{\sqrt{m}} = \sum_k \frac{\bar{a}_k \bar{\boldsymbol{\phi}}_k \alpha_k}{\sqrt{m}}$$

where $\{\alpha_k \in \{0, 1\}\}$ is a sequence of Bernoulli random variables with $\mathbb{E}[\alpha_k] = \frac{m}{\bar{m}}$. We note that

$$\mathbb{E}\left[\sum_k \frac{a_k \phi_k}{\sqrt{m}}\right] = \sum_k \frac{\bar{a}_k \bar{\phi}_k \frac{m}{\bar{m}}}{\sqrt{m}}$$

$$= \sum_k \frac{q \bar{a}_k \bar{\phi}_k}{\sqrt{\bar{m}}} = \mathbf{y}$$

Furthermore,

$$\mathbb{E}\left\|\sum_k \frac{a_k \phi_k}{\sqrt{m}} - \mathbf{y}\right\|^2 = \sum_k \frac{\bar{a}_k^2}{m} \|\bar{\phi}_k\|^2 \mathrm{Var}(\alpha_k)$$

$$= \sum_k \frac{\bar{a}_k^2}{\bar{m}} \|\bar{\phi}_k\|^2 \left(1 - \frac{m}{\bar{m}}\right)$$

Then, the Chernoff bound gives us the desired result.

7.4.1 Remark on Theorem 2 in Paper

Theorem 2 shows that the quality of knowledge distillation with the randomized scheme is proportional to $1 - \frac{m}{\bar{m}}$. This can be used to support recent results that a two-stage scheme, with knowledge being distilled to an intermediate network with size $m < m_1 < \bar{m}$ and then used as a new teacher for training the student, has a better performance than direct distillation. To support this claim, we denote $\alpha = \frac{m}{m_1}$ and $\beta = \frac{m_1}{\bar{m}}$ and note that according to Theorem 2, the performance of the two-stage scheme is proportional to $S_1 = (1 - \alpha)(1 - \beta)$ while the performance of the standard approach is proportional to $S_2 = 1 - \alpha\beta$. We show that the former is smaller. For this, we note that

$$\alpha\beta \le \sqrt{\alpha\beta} \le \frac{\alpha + \beta}{2}$$

which shows that

$$S_1 = 1 - \alpha - \beta + \alpha\beta \le 1 - \alpha\beta = S_2.$$