ON STUDENT-TEACHER DEVIATIONS IN DISTILLATION: DOES IT PAY TO DISOBEY?

Anonymous authors
Paper under double-blind review

ABSTRACT

Knowledge distillation has been widely-used to improve the performance of a "student" network by hoping to mimic soft probabilities of a "teacher" network. Yet, for self-distillation to work, the student *must* deviate from the teacher in some manner (Stanton et al., 2021). What is the nature of these deviations, and how do they relate to the generalization gains of distillation? To investigate these questions, we first conduct a variety of experiments across image and language classification datasets. One of our key observations is that in a majority of our settings, the student underfits points that the teacher finds hard. We also find that student-teacher deviations during the *initial* phase of training are *not* crucial to see the benefits of distillation — simply switching to distillation in the middle of training can recover much of its gains. We then provide two parallel theoretical perspectives of these deviations, one casting distillation as a regularizer in eigenspace, and another as a denoiser of gradients. In both views, we argue how student-teacher deviations emerge, and how they relate to generalization in the context of our experiments. Our analysis also bridges fundamental gaps between existing theory and practice by focusing on gradient descent and avoiding label noise assumptions.

1 INTRODUCTION

Distillation (Bucilă et al., 2006; Hinton et al., 2015) has emerged as a highly effective model compression technique, wherein one trains a small "student" model to match the predicted soft label distribution of a large "teacher" model, rather than match one-hot labels. An actively developing literature has sought to explore applications of this technique to various settings (Radosavovic et al., 2018; Furlanello et al., 2018; Xie et al., 2019), design more effective variants of the above recipe (Romero et al., 2015; Anil et al., 2018; Park et al., 2019; Beyer et al., 2022), and better understand theoretically when and why distillation is effective (Lopez-Paz et al., 2016; Phuong & Lampert, 2019; Mobahi et al., 2020; Allen-Zhu & Li, 2020; Menon et al., 2021; Dao et al., 2021).

On paper, distillation intends to transfer the teacher's soft probabilities over to the student. However, Stanton et al. (2021) challenge this premise: they show there is often a mismatch of student and teacher probabilities, and in fact, that a *greater* mismatch is correlated with *better* student performance. Indeed, in the *self-distillation* setting (Furlanello et al., 2018; Zhang et al., 2019) — where the student and teacher architectures are identical — some form of deviation (in the representation, if not in the probabilities) is *necessary* for the student's generalization to supercede the teacher.

In this work, we are interested in better characterizing these deviations in probabilities, and in understanding how they play a role in the student outperforming the teacher. In the first half of the paper, we conduct experiments characterizing what kind of deviations exist between the teacher and the student, and what deviations are relevant for better generalization. In the second half, we provide two complementary theoretical perspectives on how distillation can induce such deviations, and when and how that can subsequently aid generalization. More concretely, our *key* contributions are as follows:

(i) Across various pairs of architectures (ResNet56, ResNet20, MobileNet, and RoBERTa), and image and language classification datasets (CIFAR100, TinyImageNet, CIFAR10, GLUE). we empirically demonstrate (§3.1) that in a majority of the settings, *the student tends to underfit on "hard" points for the teacher* (Fig 1a) in terms of the final probabilities learned by both models. We also find even more significant student-teacher deviations early on during training.

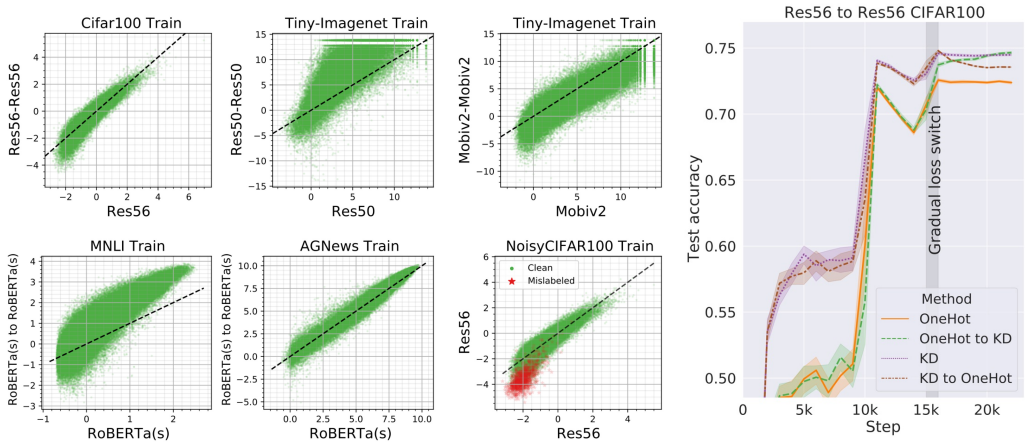

(a) Teacher-student logit plots for self-distillation.

(b) Effect of loss-switching.

Figure 1: **Left**: **Deviation in probabilities of (one-hot trained) teacher vs. (self-distilled) student:** For each training sample (x, y), we plot $\phi(p^{\mathsf{te}}_{y^{\mathsf{te}}}(x))$ versus $\phi(p^{\mathsf{st}}_{y^{\mathsf{te}}}(x))$ for logit transformation $\phi(u) = \log[u/(1-u)]$ and teacher predicted label y^{te}. We consistently find that the distilled student predictions deviate from the $X = Y$ line (dashed) with "hard" points (small X) being *underfit* ($Y \leq X$). This hints at distillation acting as a regularizer. **Right**: **Effect of late loss-switching:** In a CIFAR100 ResNet56 self-distillation setup, we switch the loss (gradually over the course of a few steps) late during training and find that switching to distillation (OneHot to KD line) recovers nearly all the gains of distillation (KD). This suggests that the initial phase of training is not critical for distillation to help.

(ii) We find (§3.2) that it is possible to switch from one-hot loss in the middle of training to distillation loss and (a) still recover a considerable fraction of distillation's gains (Fig 1b), (b) and also recover the final-epoch underfitting behavior on TinyImageNet and CIFAR100. Thus, we conclude that any student-teacher deviations unique to the early phase of training are *not* by themselves adequate to explaining the success of distillation.

(iii) We extend the seminal result of Mobahi et al. (2020) — which demonstrates distillation as a regularizer in a *non-gradient-descent* setting — to the gradient descent setting. Concretely, we formalize distillation as an "eigenspace sparsifier" for early-stopped, continuous-time gradient-descent trained linear models (Theorem 4.1). We propose this view as a way to understand the empirically observed underfitting in distillation. Besides providing a much simpler proof and a more practically relevant version of Mobahi et al. (2020), our result also formalizes existing empirical intuition about the importance of early-stopping in distillation (Dong et al., 2019; Cho & Hariharan, 2019; Ji & Zhu, 2020).

(iv) As a complementary viewpoint, we also formalize distillation as a gradient denoiser in the presence of class similarities (Theorem 4.2). We propose this view as a way to understand our empirical observations on loss-switching. Importantly, unlike prior work (Menon et al. (2021)), we show how denoising can occur even when the data is perfectly classifiable and has no inherent label noise.

(v) We conclude with a discussion unifying these two perspectives, thus painting a more coherent picture of two disjoint lines of existing theories (Mobahi et al. (2020) vs. Menon et al. (2021)) these views build upon.

Overall, we hope that our discussion helps bridge the gap between existing theoretical understanding and empirics in distillation by (a) making more practical assumptions than existing theories, and (b) making connections to various empirical observations. Our findings also suggest that *not* matching the teacher probabilities exactly can be a good thing, which future empirical work on distillation may want to be mindful of.

2 BACKGROUND AND NOTATION

Our interest in this paper is *multiclass classification* problems. This involves learning a *classifier* $h\colon \mathcal{X} \to \mathcal{Y}$ which, for input $x \in \mathcal{X}$, predicts the most likely label $h(x) \in \mathcal{Y} = [K] \doteq \{1, 2, \ldots, K\}$. Such a classifier is typically implemented by computing *logits* $\mathbf{f}\colon \mathcal{X} \to \mathbb{R}^K$ that score the plausibility of each label, and then computing $h(x) = \operatorname{argmax}_{y \in \mathcal{Y}} f_y(x)$. In neural models, these logits are parameterised as $\mathbf{f}(x) = \mathbf{W}^\top \mathbf{Z}(x)$ for learned weights $\mathbf{W} \in \mathbb{R}^{D \times K}$ and embeddings $\mathbf{Z}(x) \in \mathbb{R}^D$. One may learn such logits by minimising the *empirical loss* on a training sample $S \doteq \{(x_n, y_n)\}_{n=1}^N$:

$$R_{\text{emp}}(\mathbf{f}) \doteq \frac{1}{N} \sum_{n \in [N]} \mathbf{e}(y_n)^\top \ell(\mathbf{f}(x_n)), \tag{1}$$

where $\mathbf{e}(y) \in \{0, 1\}^K$ denotes the *one-hot encoding* of y, $\ell(\cdot) \doteq [\ell(1, \cdot), \ldots, \ell(K, \cdot)] \in \mathbb{R}^K$ denotes the *loss vector* of the predicted logits, and each $\ell(y, \mathbf{f}(x))$ is the loss of predicting logits $\mathbf{f}(x) \in \mathbb{R}^K$ when the true label is $y \in [K]$. Typically, we set ℓ to be the softmax cross-entropy $\ell(y, \mathbf{f}(x)) = -\log p_y(x)$, where $\mathbf{p}(x) \propto \exp(\mathbf{f}(x))$ is the *softmax* transformation of the logits.

Equation 1 guides the learner via one-hot targets $\mathbf{e}(y_n)$ for each input. *Distillation* (Bucilă et al., 2006; Hinton et al., 2015) instead guides the learner via a target label distribution $\mathbf{p}^{\mathsf{te}}(x_n)$ provided by a *teacher*, which are the softmax probabilities from a distinct model trained on the *same* dataset. In this context, the learned model is referred to as a *student*, and the training objective is

$$R_{\text{dist}}(\mathbf{f}) \doteq \frac{1}{N} \sum_{n \in [N]} \mathbf{p}^{\mathsf{te}}(x_n)^\top \ell(\mathbf{f}(x_n)). \tag{2}$$

One may also consider a weighted combination of R_{emp} and R_{dist} for algorithmic reasons, but we focus on the above objective in this paper since we are interested in understanding each objective individually. Compared to training on the one-hot labels, distillation often results in improved performance for the student (Hinton et al., 2015). Typically, the teacher model is of higher capacity than the student model; the performance gains of the student may thus informally be attributed to the teacher transferring rich information about the problem to the student. In such settings, distillation may be seen as a form of model compression. Intriguingly, however, even when the teacher and student are of the *same* capacity (a setting known as *self-distillation*), one may see gains from distillation (Furlanello et al., 2018; Zhang et al., 2019). The questions we explore in this paper are motivated by the self-distillation setting; however, for a well-rounded analysis, we will empirically study both the self- and cross-architecture-distillation settings.

3 A FINE-GRAINED LOOK AT TEACHER-STUDENT DEVIATIONS

Stanton et al. (2021) found that, contrary to the premise of distillation, more accurate students are poorer at matching the teacher probabilities. They quantified the student-teacher deviation by measuring the disagreement or the KL divergence between the student and teacher probabilities, in expectation over all points. Could the average deviation in probabilities be attributed merely to an *arbitrary* lack of precision in matching the probabilities during training? In the following section, through a closer study of *per-sample* relationship between the teacher and student predictions, we do *not* find this to be the case; rather, there are discernible patterns in student-teacher deviations.

3.1 A PER-SAMPLE VIEW OF TEACHER-STUDENT FIDELITY

Suppose we have teacher and distilled student models $\mathbf{f}^{\mathsf{te}}, \mathbf{f}^{\mathsf{st}}\colon \mathcal{X} \to \mathbb{R}^K$ respectively. We seek to visualize the deviations in the corresponding predicted probability vectors $\mathbf{p}^{\mathsf{te}}(x)$ and $\mathbf{p}^{\mathsf{st}}(x)$ for each (x, y) in the train and test set, rather than in the aggregated sense as in Stanton et al. (2021). To visualize the deviations, we need a scalar summary of these vectors. An obvious candidate is the probabilities $(p^{\mathsf{te}}_{y^\star}(x), p^{\mathsf{st}}_{y^\star}(x))$ assigned to the ground truth label $y^\star$. However, since the student does not have access to the ground truth label, and is only trying to mimic the teacher, we examine deviations of probabilities of the *teacher's predicted label*, i.e., $(p^{\mathsf{te}}_{y^{\mathsf{te}}}(x), p^{\mathsf{st}}_{y^{\mathsf{te}}}(x))$ where $y^{\mathsf{te}} \doteq \operatorname{argmax}_{y' \in [K]} p^{\mathsf{te}}_{y'}(x)$. To make patterns easier to detect, we further perform a monotonic logit transformation $\phi(u) = \log[u/(1-u)]$ that produces real-values in $(-\infty, +\infty)$. Thus, we compare $\phi(p^{\mathsf{te}}_{y^{\mathsf{te}}}(x))$ and $\phi(p^{\mathsf{st}}_{y^{\mathsf{te}}}(x))$ for each train and test sample (x, y).

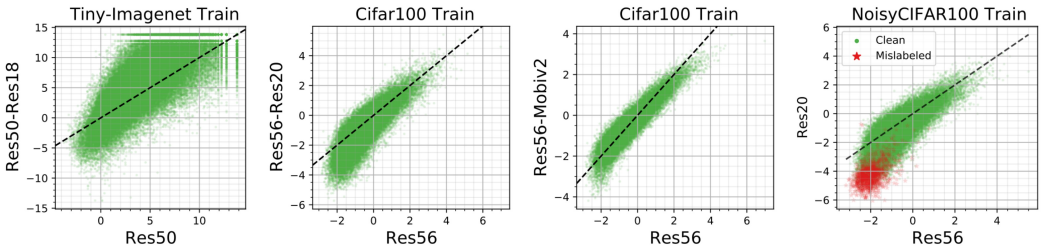

Figure 2: **Teacher-student logit-transformed probability scatter plots under cross-architecture distillation settings:** We demonstrate in the last plot that under label noise, this underfitting is clearly beneficial in deprioritizing noisy train labels, demonstrating that even a smaller student can go beyond a larger teacher.

We report a scatter plot of $\phi(p^{\mathsf{te}}_{y^{\mathsf{te}}}(x))$ (X-axis) vs. $\phi(p^{\mathsf{st}}_{y^{\mathsf{te}}}(x))$ (Y-axis) on the training set for both various self-distillation settings in Figures 1a and cross-architecture distillation settings in Fig 2. In all plots, the dashed line indicates the $X = Y$ line. All values are computed at the end of training. The tasks considered include image classification benchmarks, namely CIFAR10, CIFAR-100, Tiny-ImageNet, and text classification tasks from the GLUE benchmark (e.g., MNLI (Williams et al., 2018), AGNews (Zhang et al., 2015). See Appendix C.1 for details on the experimental hyperparameters. Many additional plots are presented in §C.2.1.

Distilled students tend to underfit hard samples. Our main finding is that in most of our cases, the student underfits a subset of points (i.e., $Y \leq X$ in the scatter plot) that are typically hard for the teacher (i.e., for small values of X). We note that this pattern is particularly strong on test data, and appears in the test data even in certain exceptions where it does not appear in the training data (see §C.2.1, e.g., CIFAR-10 plots Fig 7). We also note that in the language datasets, typically there is high deviation for harder points, involving both underfitting and overfitting, along with low deviation for easier points (i.e., large X), typically involving only overfitting (as can be seen in AGNews Fig 1 and in §C.2.1 Fig 8, Fig 9). These patterns however break down for cross-architecture distillation in language datasets. See §C.2.1 for a discussion of the exceptions.

Why is underfitting good? To better understand the significance of the aforementioned underfitting, we present plots for distillation when a portion of CIFAR100 is mislabeled in Fig 1a and Fig 2. In both cases, we observe that the noisy points are underfit by the student. This is notably the case even though the student is smaller in Fig 2 and achieves lower accuracy than the teacher (Table 2). Together, these observations hint at (a) why hard-point-underfitting may be a desirable behavior and (b) that *there is a latent sense in which smaller student models go beyond a larger teacher, one that is not apparent through mere accuracy numbers*. We provide a more formal view of the underfitting in §4.1 where we present distillation as a form of regularization. Having said that, we must note that underfitting may not always correspond to improved generalization (cf. §C.2.1, Table 2).

Evolution of teacher-student deviations. While we have looked at deviations between the *final* probabilities learned by the models, it is also instructive to look at how these deviations evolve as student training progresses. Concretely, in Fig 3a (and §C.2.3 Fig 13) we present a series of scatter plots for CIFAR100 (and Tiny-Imagenet) settings where the X axis is always fixed to be the *final* (logit-transformed) teacher probabilities while the Y axis corresponds to snapshots of the distillation-trained and the one-hot trained model at various instances. Here, we find significant differences between the scatter plots early on during training. In particular, the distilled model has converged much faster to its final values than the one-hot model. We also see from the left-most snapshot that the student has prioritized its easier points (i.e., points with large X are close to the $X = Y$ line), while for the one-hot model, the easiest points are not as prioritized (i.e., points with large X are farther from $X = Y$ compared to other points).

On the face of it, these observations appear consistent with some previously proposed hypotheses behind the success of distillation: it conditions the network into a more favorable representation/basin early on (Jha et al., 2020), or induces faster convergence (Phuong & Lampert, 2019) or perhaps, induces a favorable learning order. However, we question whether this initial phase is critical in the following section.

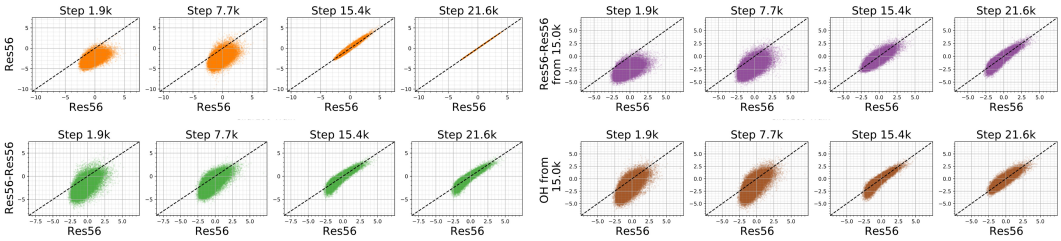

(a) One-hot and self-distillation.

(b) Loss-switching to distillation/one-hot at $15k$ steps.

Figure 3: **Evolution of teacher-student logit plots over various steps of training for CIFAR100 ResNet56 self-distillation setup:** On the **left**, we present plots for one-hot training (**top**) and distillation (**bottom**). Observe that there is a stark difference in the plots early on during training (at $1.9k$ steps): the student has converged to the $X = Y$ line much faster, fitting "easier" points more substantially. On the **right**, we present similar plots for experiments from Section 3.2, with the loss switched to distillation (**top**) and one-hot (**bottom**) at $15k$ steps. From the last two visualized plots in each, observe that switching to distillation "unlearns" the hard points, while switching to one-hot curiously "fixes" the existing underfitting.

3.2 The effect of loss switching

Next, we conduct *late distillation* experiments where we train a network with one-hot loss for a fair part of training (say about the first half of training), and then switch to distillation loss (gradually). If distillation relied solely on the "early phase" deviations, we must observe that late distillation on a one-hot-pretrained network should not result in substantial improvements.

Yet, from Figures 4 (and Sec C.2.3 Fig 14), we observe that late distillation still recovers considerable fractions of standard distillation's gains in our CIFAR100 and TinyImagenet setting. Although this does not recover all of distillation's gains, this implies that the early phase deviations are not a *necessary* requirement for distillation to help.

Conversely, we also switch from distillation to one-hot in Fig 15, and find the story to be nuanced. First, switching to one-hot typically preserves the gains of distillation — and in the rare case of ResNet50 self-distillation on TinyImagenet even outperforms distillation — which supports earlier findings (Cho & Hariharan, 2019; Zhou et al., 2020; Jafari et al., 2021) that advocate a soft final switch to one-hot training. Overall, this seems to suggest that, the early phase of distillation may be a *sufficient* requirement to gain all the benefits of distillation. *However*, we also find some cases where the switch to one-hot hurts (Fig 1b), especially if continued for an extended period of time (Fig 16).

We also find some curious connections between these experiments and the underfitting behavior from Section 3.1. In particular, even though early-phase one-hot training has already fit all examples, subsequent late-distillation "unlearns" the harder examples to force the typical underfitting behavior (See Fig 3b, top). Conversely, switching to one-hot fixes the underfitting already induced by an initial distillation phase. Thus, while our observations (a) de-emphasize the necessity of early-phase student-teacher deviations, they also (b) emphasize the relevance of the deviations in the final probabilities. Overall, these observations suggest that, to explain the workings of distillation, we also need to look for effects decoupled from the initial phase of training. In Section 4.2, we formalize one such effect.

4 Formalizing student-teacher deviations: Two perspectives

To understand the above empirical observations regarding underfitting and loss-switching, and their connections to the benefits of distillation, we provide two complementary perspectives of distillation: an *eigenspace* perspective and a *gradient-space* view. These perspectives also aim to provide key clarifications about existing intuition in literature, and paint a more coherent picture of disjoint threads of theoretical research.

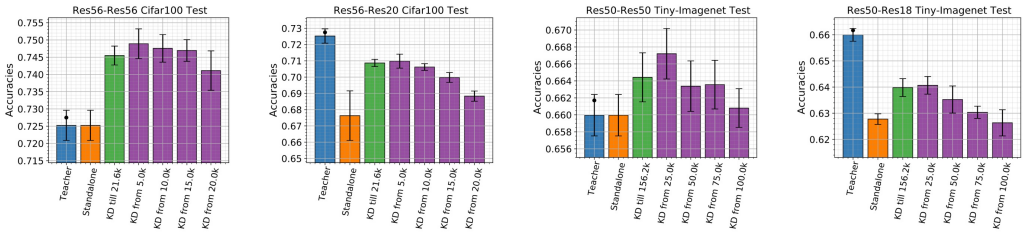

Figure 4: **Results of *loss-switching***: For each panel, we report the performance of the teacher, one-hot student (Standalone; which is the same as teacher in self-distillation), and distilled student (denoted by KD upto T steps where T is the total number of steps) and late-distilled student (denoted by KD from T_{switch} where T_{switch} is the step where the loss begins to switch gradually). We find that such switching generally yields significant gains over one-hot training of the student, and can largely recover the gains of regular distillation.

4.1 THE EIGENSPACE VIEW: DISTILLATION AS A REGULARIZER

The underfitting of hard points possibly suggests that distillation induces a form of regularization that deprioritizes certain complex features. Indeed, Mobahi et al. (2020) have proven that self-distillation acts as a form of sparsifier that focuses on simpler basis functions while ignoring more complex ones. However, their result is agnostic to the biases of gradient descent.

We bridge the gap between this intuition and practice, by showing that a similar sparsification effect emerges under gradient descent training of linear models, thanks to the inherent bias of gradient descent in converging at varied rates along various eigendirections. Concretely, we analyze a continuous-flow gradient descent model on a linear regression setting with early-stopping, a typical design choice in distillation practice (Dong et al., 2019; Cho & Hariharan, 2019; Ji & Zhu, 2020). Consider an $n \times p$ dataset $\mathbf{X}$ (where n is the number of samples, p the number of parameters) with target labels $\mathbf{y}$. Assume that the Gram matrix $\mathbf{X}\mathbf{X}^\top$ is invertible. Note that this setting includes overparameterized scenarios ($p > n$) such as when $\mathbf{X}$ corresponds to the linearized (NTK) features of neural networks (Jacot et al., 2018; Lee et al., 2019). Then, a standard calculation reveals that the weights learned at time t under gradient descent can be written as:

$$\boldsymbol{\beta}(t) = \mathbf{X}^\top(\mathbf{X}\mathbf{X}^\top)^{-1}\mathbf{A}(t)\mathbf{y} \qquad \text{where } \mathbf{A}(t) := \mathbf{I} - e^{-t\mathbf{X}\mathbf{X}^\top}. \tag{3}$$

Intuitively, $\mathbf{A}$ is a "sparsifying" matrix that skews down the weight assigned to an eigendirection of eigenvalue λ by the value $1 - e^{-\lambda t}$. As $t \to \infty$ this factor goes to 1 for all directions, thus becoming irrelevant; but for any finite $t > 0$, the topmost direction would have a larger factor than the rest, thus producing some sort of sparsification. Our argument is that, *under gradient-descent-based self-distillation, this sparsification is amplified*. Consider a setting where the teacher is trained to time T^{te}. Through a simple calculation, the weights of the student model can be written as:

$$\tilde{\boldsymbol{\beta}}(\tilde{t}) = \mathbf{X}^\top(\mathbf{X}\mathbf{X}^\top)^{-1}\tilde{\mathbf{A}}(\tilde{t})\mathbf{y} \qquad \text{where } \tilde{\mathbf{A}}(\tilde{t}) := \mathbf{A}(t)\mathbf{A}(T^{\mathsf{te}}). \tag{4}$$

Akin to Mobahi et al. (2020), one can argue that the sparsifier $\tilde{\mathbf{A}}$ corresponding to the student is more skewed towards the top eigenvectors than the teacher. Notably, a more aggressive early-stopping of the teacher *cannot* induce an equivalent effect to distillation. More formally:

Theorem 4.1. *Let $\alpha_k(t), \tilde{\alpha}_k(t)$ be the eigenvalues of the k'th top eigendirection in $\mathbf{A}(t)$ and $\tilde{\mathbf{A}}(t)$ respectively. Let λ_k be the corresponding eigenvalue of the Gram matrix $\mathbf{X}\mathbf{X}^\top$. For any two indices $k_1 < k_2$ such that $\lambda_{k_1} > \lambda_{k_2}$ (assuming such a pair exists), and for any $\alpha^\star_{k_1} \in (0, 1)$ (a required "convergence lower bound" on the higher eigenvector), there exists a choice of distillation hyperparameters $T^{\mathsf{te}}, \tilde{t}$ such that (a) $\tilde{\alpha}_{k_1}(\tilde{t}) \geq \alpha^\star_{k_1}$, and (b) for **any** choice of time t for the standalone early-stopped model satisfying $\alpha_{k_1}(t) \geq \alpha^\star_{k_1}$, the components along the lower eigenvectors satisfy:*

$$\underbrace{\tilde{\alpha}_{k_2}(\tilde{t})}_{\substack{\textit{Student's component} \\ \textit{along lower eigenvector}}} < \underbrace{\alpha_{k_2}(t).}_{\substack{\textit{Teacher's component} \\ \textit{along lower eigenvector}}} \tag{5}$$

The result says that the student relies less on the bottom eigendirections than the teacher under the condition that both models have converged well along the top eigendirections (i.e., they have sufficiently fit the "nice" directions). A more aggressive early-stopping of the teacher may have less reliance on the bottom eigenvectors but at the cost of violating the convergence criterion along the nice direction.

Connection to underfitting. The connection between how distillation amplifies regularization and the underfitting in §3.1 is best verified through our experiments on noisy CIFAR100 experiments from Fig 1a and Fig 2. First, noisy labels in the training data typically correspond to lower eigendirections. Indeed, early-stopped gradient-descent is said to have the implicit bias to underfit noisy data (Li et al., 2020; Dong et al., 2019; Arpit et al., 2017). This is also evidenced by the fact that the mislabeled points come lowest in terms of teacher probabilities. But more interestingly, we see that the student essentially *exaggerates* this ordering between points, by assigning further lower probabilities to the mislabeled data. This exaggeration must be the manifestation of how distillation amplifies the sparsifying effect of early-stopping.

4.2 The gradient-space view: distillation as a gradient denoiser

While the previous perspective gives us an abstract theory of how distillation can induce regularization, it may not be obvious *when* this regularization can be helpful besides in settings with label noise. In fact, if ignoring the bottom eigenvectors led to ignoring non-noisy "tail" datapoints, this can even be hurtful (Feldman, 2019). Below, we formalize how distillation can help by extending a line of work analyzing distillation on noisy datasets with class similarities (Menon et al., 2021; Dao et al., 2021; Ren et al., 2022). But crucially, we show how distillation can help by *denoising* even in a perfectly classifiable dataset with class similarities, and *no explicit label noise*.

A concrete example. To make our discussion easier, consider a concrete K-class classification dataset, where the ith datapoint's features can be written as a K-"channel" input $\mathbf{x}_i = (\mathbf{x}_1^{(i)}, \ldots, \mathbf{x}_K^{(i)})$ where each channel $\mathbf{x}_k^{(i)} \in \mathbb{R}^D$ is D-dimensional. Assuming a uniform distribution over K classes, given a label y, we generate $\mathbf{x}_y$ from a Gaussian $\mathcal{N}(0, \mathbf{I}/D)$ truncated to the support $\mathbf{x}_y \cdot \boldsymbol{\mu}_{y_i}^\star = \tau$ for a pre-defined $\tau > 0$ and a "ground truth" class vector $\boldsymbol{\mu}_k^\star$. We also assume that for point $\mathbf{x}_i$, we pick a non-target "similar" class z_i for which we pick $\mathbf{x}_{z_i}$ from a Gaussian $\mathcal{N}(0, \mathbf{I}/D)$ truncated to the support $\mathbf{x}_{z_i} \cdot \boldsymbol{\mu}_{z_i}^\star = \tau/2$. All other co-ordinates are zero for $\mathbf{x}_i$. We consider a simplified linear architecture whose K-dimensional output is of the form $\mathbf{f}(\mathbf{x}) = (\mathbf{w}_1 \cdot \mathbf{x}_1, \ldots, \mathbf{w}_K \cdot \mathbf{x}_K)$ where the k'th node acts only on the k'th channel. Observe that this is a perfectly classifiable dataset with deterministic labels since if we set the model weights to be $\mathbf{w}_k = \boldsymbol{\mu}_k^\star$, we obtain $\arg\max_k f_k(\mathbf{x}) = y^\star$. In other words, the ground truth probabilities are one-hot, which can be recovered by computing $\text{softmax}(\{\alpha \cdot \mathbf{f}_k(\mathbf{x})\}_{k=1}^K)$ as $\alpha \to \infty$. Thus, it may seem that training on the one-hot loss is the wisest choice in this setting.

However, we show that distillation's gradients (a) are more optimal than one-hot gradients and also (b) more optimal than the teacher's weights themselves. To formalize this, let cos-sim$(\cdot, \cdot)$ denote the cosine-similarity between two vectors. Then:

Theorem 4.2. ***(informal; see Thm B.3)*** *Consider a student whose weights satisfy $\mathbf{w}_k = \alpha \boldsymbol{\mu}_k^\star$. Consider an imperfect teacher with weights such that $\|\mathbf{w}_k^{te}\| = \alpha_k^{te}$ and cos-sim$(\mathbf{w}_k^{te}, \boldsymbol{\mu}_k^\star) \cdot \geq 1 - \epsilon$. Let $\mathbf{v}_k$ be the one-hot update on node k of the student, and $\tilde{\mathbf{v}}_k$ the distillation update, both under cross-entropy loss. Then, for $\alpha \ll \alpha^{te}$ and $\alpha^{te} \leq O(\frac{1}{\tau} \log K)$, and sufficiently large τ, D:*

$$\underbrace{\textit{cos-sim}(\tilde{\mathbf{v}}_k, \tilde{\boldsymbol{\mu}}_k)}_{\textit{Quality of distillation update}} > \max \left(\underbrace{\textit{cos-sim}(\mathbf{v}_k, \boldsymbol{\mu}_k^\star)}_{\textit{Quality of one-hot update}}, \underbrace{\textit{cos-sim}(\mathbf{w}_k^{te}, \boldsymbol{\mu}_k^\star)}_{\textit{Quality of teacher weights}} \right) \tag{6}$$

There are three salient aspects to this result. First, this demonstrates that one-hot gradients can be inherently sub-optimal at any timestep in training *even if the model is initialized at an optimal solution*. Conversely, distillation can help improve these gradients *even if the teacher is imperfect*. Finally, distillation gradients can themselves be "more perfect" than the imperfect teacher, thus formalizing how the student can go beyond the teacher.

To understand the intuition behind this result, imagine that at some point of training, we have recovered a good $\mathbf{f}$ such that $\mathbf{f} \approx \alpha \cdot \mathbf{f}^\star$ (for some finite scaling factor α). Now as we train longer,

we ideally hope that (a) α continues increasing towards ∞ to recover the one-hot ground truth probabilities, and (b) the approximation of the ground truth logits continues to hold throughout.

This is however not the case with the one-hot updates. Consider any point $(\mathbf{x}, y)$ where $\mathbf{f}_k^\star(\mathbf{x})$ is high for some $k \neq y$ i.e., the point has similarities to a non-target class. On such a point, the ideal dynamics of increasing α requires that we *increase the logit on the non-target node* $f_k(\mathbf{x})$. But the one-hot update would attempt to suppress this logit, since the target probability for this node is 0. We argue that this is a *destructive gradient* since it directs the model in the opposite direction of the ideal dynamics. Notably, this manifests despite the lack of inherent label noise in the training set.

Distillation on the other hand, can flip the sign of these gradients (assuming the teacher is early-stopped), thus denoising the gradients — even if there is no explicit noise in the data. Notably, this sign-flipping provides a "fresh example" for the non-target node k, which is a privilege unavailable to the teacher. This is crucial to prove that the student is more denoised than the the teacher.

Connection to loss-switching. The gradient-space view gives us a way to make sense of the findings in §3.2. In particular, this tells us why each gradient in itself may be constructive (in the case of distillation) or destructive (in the case of one-hot). Thus, switching to distillation (or one-hot) even in the middle of training can still be helpful (or hurtful), given sufficient time after the switch.

4.3 Unifying the two perspectives

We now discuss how various themes in existing distillation theory can be understood within both of the above perspectives. Through this, we hope to piece together a more coherent and unified picture of the two lines of work we build upon: Mobahi et al. (2020) and Menon et al. (2021).

Dark knowledge and denoising. Distillation is said to benefit by implicitly introducing "dark knowledge" in the form of new information absent in the observed labels of the original datasets (Hinton et al., 2015). In both views, we can view dark knowledge as a form of recovering denoised features:

(i) In the gradient-space view, the dark knowledge comes from the non-target logits, which allows us to denoise the gradients during distillation.

(ii) In the eigenspace view, we posit that the new information corresponds to knowledge of the top eigenvector(s) of the data. Crucially, *these eigenvectors are independent of labels* and thereby independent of any label noise; hence these eigenvectors correspond to "denoised" features in the data, which distillation more readily relies on.

On the gap between early-stopping and distillation. In both our views, we see a subtle distinction between early-stopping and distillation:

(i) In the eigenspace perspective, although both induce a similar sparsification effect, distillation *amplifies* this effect.

(ii) In the gradient-space view, early-stopping would reduce the overall magnitude of destructive gradients accumulated by the teacher. But distillation amplifies this effect further by altogether flipping the sign of these gradients.

Lack of fidelity in probabilities. We can understand the underfitting phenomenon from both views:

(i) Recall that in the eigenspace view, the student's underfitting of teacher's probabilities can be seen as a consequence of ignoring the lower eigenvectors.

(ii) In the gradient-space view, the harder points of the teacher would have high similarity to other classes. Subsequently, under the one-hot loss, these points would experience heavy destructive gradients on the non-target logits. Since distillation would no longer suppress the non-target probabilities, those probabilities would increase, while the target probability would be underfit.

5 Relation to Existing Work

Distillation as a probability matching process. Distillation has been touted to be a process that benefits from matching the teacher's probabilities (Hinton et al., 2015). Indeed, many distillation algorithms have been designed in a way to more aggressively match the student and teacher functions (Czarnecki et al., 2017; Beyer et al., 2022). Theoretical analyses too rely on explaining the

benefits of distillation based on a student that obediently matches the teacher's probabilities Menon et al. (2021). However, as Stanton et al. (2021) demonstrate, an exact matching of probabilities can be *anti-correlated* with performance improvements. Our work builds on this nuance as to why we may *desire* that the student deviate from the teacher.

Understanding distillation via class similarities. A long-standing intuitive explanation for why distillation helps is that the teacher's probabilities contain "dark knowledge" about class similarities (Hinton et al., 2015; Müller et al., 2019), which is the idea we formalize in §4.2. Several works (Menon et al., 2021; Dao et al., 2021; Ren et al., 2022; Zhou et al., 2021) have formalized this by assuming that the ground truth class memberships are inherently soft, and so the observed one-hot labels are "noisy". They argue that distillation reduces the student model's variance (in an abstract non-gradient-descent setting), assuming the student exactly matches the teacher. We go beyond these works by demonstrating a form of denoising even in the absence of noisy labels, and by showing how the student can outperform the teacher under gradient descent.

Alternative theories of distillation. Some works (Furlanello et al., 2018; Yuan et al., 2020; Tang et al., 2020) have argued that the "class similarity" hypothesis cannot be the sole explanation, because distillation can help even if the student is only taught information about the target probabilities (e.g., by smoothing out all non-target probabilities). This has resulted in various alternative hypothesis such as *faster convergence* (Phuong & Lampert, 2019; Rahbar et al., 2020; Ji & Zhu, 2020) (which we question in §3.2), *feature learning* (Allen-Zhu & Li, 2020) (which would not apply in our linear settings), and regularization, either in the sense of Mobahi et al. (2020) or in the sense of *instance-specific label smoothing* (Zhang & Sabuncu, 2020; Yuan et al., 2020; Tang et al., 2020). Our eigenspace perspective in §4.1 falls under this umbrella. Finally, we also refer the reader to Lopez-Paz et al. (2016); Kaplun et al. (2022) who theoretically study distillation in orthogonal settings.

Early-stopping and knowledge distillation. Early-stopping has received much attention in the context of distillation (Liu et al., 2020; Ren et al., 2022; Dong et al., 2019; Cho & Hariharan, 2019; Ji & Zhu, 2020). We closely build on Dong et al. (2019), who argue how early-stopping a gradient-descent-trained teacher can automatically denoise the labels due to regularization in the eigenspace. However, neither Dong et al. (2019) nor any of the other works provide a formal argument for why distillation can outperform the teacher. Note that Mobahi et al. (2020) claim early-stopping has a *densification* rather than a sparsification effect; but this holds only in their non-gradient-descent setting where the function is picked from a Hilbert space with ℓ_2 regularization.

Empirical studies of distillation. Our study crucially builds on observations from (Stanton et al., 2021; Lukasik et al., 2021) demonstrating student-teacher deviations in an aggregated sense than a sample-wise sense sense. Other studies (Abnar et al., 2020; Ojha et al., 2022) investigate in what ways the student is *similar* to the teacher, beyond in terms of class probabilities e.g., out-of-distribution behavior, calibration, and so on. Deng & Zhang (2021) show how a smaller student can outperform the teacher if it was allowed better match the teacher on more data, which is orthogonal to our setting.

6 Conclusion

Via experiments across image and language classification data, we saw that the student tends to underfit points that the teacher finds hard, thus concretely characterizing the student-teacher deviations identified in Stanton et al. (2021). We also found that deviations in the initial phase of distillation are *sufficient but not necessary* to help the student. We then formalized these observations via two parallel perspectives of student-teacher deviations: one casting distillation as a regularizer in eigenspace, and another as a gradient denoiser. Our analysis bridges key gaps between theory and practice, while also unifying disjoint lines of work, namely Mobahi et al. (2020) and Menon et al. (2021).

There are several directions for future work. One question is whether we can study more practical distillation setups, such as in *semi-supervised* setups (Cotter et al., 2021). It would also be insightful to extend our eigenspace view in to multi-layered models where the eigenspace regularization effect may "compound" across layers. Another interesting direction is whether we can build on the gradient space view to design teacher-free training techniques that suppress "destructive" non-target gradients in one-hot training. Arguably, the most important implication is on current attempts to bridge the gap between student and teacher performance. Perhaps, these attempts would benefit from encouraging *some carefully-designed* gaps in the teacher and student probabilities.

REFERENCES

Samira Abnar, Mostafa Dehghani, and Willem H. Zuidema. Transferring inductive biases through knowledge distillation. abs/2006.00555, 2020. URL https://arxiv.org/abs/2006.00555.

Zeyuan Allen-Zhu and Yuanzhi Li. Towards understanding ensemble, knowledge distillation and self-distillation in deep learning. *CoRR*, abs/2012.09816, 2020. URL https://arxiv.org/abs/2012.09816.

Rohan Anil, Gabriel Pereyra, Alexandre Passos, Robert Ormandi, George E. Dahl, and Geoffrey E. Hinton. Large scale distributed neural network training through online distillation. In *International Conference on Learning Representations*, 2018.

Devansh Arpit, Stanislaw Jastrzebski, Nicolas Ballas, David Krueger, Emmanuel Bengio, Maxinder S. Kanwal, Tegan Maharaj, Asja Fischer, Aaron C. Courville, Yoshua Bengio, and Simon Lacoste-Julien. A closer look at memorization in deep networks. In *Proceedings of the 34th International Conference on Machine Learning, ICML 2017*, Proceedings of Machine Learning Research. PMLR, 2017.

Lucas Beyer, Xiaohua Zhai, Amélie Royer, Larisa Markeeva, Rohan Anil, and Alexander Kolesnikov. Knowledge distillation: A good teacher is patient and consistent. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 10925–10934, June 2022.

Cristian Bucilă, Rich Caruana, and Alexandru Niculescu-Mizil. Model compression. In *Proceedings of the 12th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '06, pp. 535–541, New York, NY, USA, 2006. ACM.

J. H. Cho and B. Hariharan. On the efficacy of knowledge distillation. In *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 4793–4801, 2019.

Andrew Cotter, Aditya Krishna Menon, Harikrishna Narasimhan, Ankit Singh Rawat, Sashank J. Reddi, and Yichen Zhou. Distilling double descent. *CoRR*, abs/2102.06849, 2021. URL https://arxiv.org/abs/2102.06849.

Wojciech M. Czarnecki, Simon Osindero, Max Jaderberg, Grzegorz Swirszcz, and Razvan Pascanu. Sobolev training for neural networks. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (eds.), *Advances in Neural Information Processing Systems 30*, pp. 4278–4287. Curran Associates, Inc., 2017.

Tri Dao, Govinda M Kamath, Vasilis Syrgkanis, and Lester Mackey. Knowledge distillation as semiparametric inference. In *International Conference on Learning Representations*, 2021.

Xiang Deng and Zhongfei Zhang. Can students outperform teachers in knowledge distillation based model compression?, 2021. URL https://openreview.net/forum?id=XZDeL25T12l.

Bin Dong, Jikai Hou, Yiping Lu, and Zhihua Zhang. Distillation $\approx$ early stopping? harvesting dark knowledge utilizing anisotropic information retrieval for overparameterized neural network, 2019.

Vitaly Feldman. Does learning require memorization? A short tale about a long tail. *CoRR*, abs/1906.05271, 2019. URL http://arxiv.org/abs/1906.05271.

Tommaso Furlanello, Zachary Chase Lipton, Michael Tschannen, Laurent Itti, and Anima Anandkumar. Born-again neural networks. In *Proceedings of the 35th International Conference on Machine Learning, ICML 2018*, pp. 1602–1611, 2018.

Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Identity mappings in deep residual networks. In Bastian Leibe, Jiri Matas, Nicu Sebe, and Max Welling (eds.), *Computer Vision - ECCV 2016 - 14th European Conference, Amsterdam, The Netherlands, October 11-14, 2016, Proceedings, Part IV*, volume 9908 of *Lecture Notes in Computer Science*, pp. 630–645. Springer, 2016a. doi: 10.1007/978-3-319-46493-0_38. URL https://doi.org/10.1007/978-3-319-46493-0_38.

Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016b.

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

Arthur Jacot, Clément Hongler, and Franck Gabriel. Neural tangent kernel: Convergence and generalization in neural networks. pp. 8580–8589, 2018.

Aref Jafari, Mehdi Rezagholizadeh, Pranav Sharma, and Ali Ghodsi. Annealing knowledge distillation. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pp. 2493–2504, Online, April 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.eacl-main.212. URL https://aclanthology.org/2021.eacl-main.212.

Nandan Kumar Jha, Rajat Saini, and Sparsh Mittal. On the demystification of knowledge distillation: A residual network perspective. 2020.

Guangda Ji and Zhanxing Zhu. Knowledge distillation in wide neural networks: Risk bound, data efficiency and imperfect teacher. In Hugo Larochelle, Marc'Aurelio Ranzato, Raia Hadsell, Maria-Florina Balcan, and Hsuan-Tien Lin (eds.), *Advances in Neural Information Processing Systems*, 2020.

Gal Kaplun, Eran Malach, Preetum Nakkiran, and Shai Shalev-Shwartz. Knowledge distillation: Bad models can be good role models. *CoRR*, 2022.

Jaehoon Lee, Lechao Xiao, Samuel S. Schoenholz, Yasaman Bahri, Roman Novak, Jascha Sohl-Dickstein, and Jeffrey Pennington. Wide neural networks of any depth evolve as linear models under gradient descent. pp. 8570–8581, 2019.

Mingchen Li, Mahdi Soltanolkotabi, and Samet Oymak. Gradient descent with early stopping is provably robust to label noise for overparameterized neural networks. In *The 23rd International Conference on Artificial Intelligence and Statistics, AISTATS 2020*, Proceedings of Machine Learning Research. PMLR, 2020.

Sheng Liu, Jonathan Niles-Weed, Narges Razavian, and Carlos Fernandez-Granda. Early-learning regularization prevents memorization of noisy labels. In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020*, 2020.

Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. Roberta: A robustly optimized BERT pretraining approach. *CoRR*, 2019. URL http://arxiv.org/abs/1907.11692.

D. Lopez-Paz, B. Schölkopf, L. Bottou, and V. Vapnik. Unifying distillation and privileged information. In *International Conference on Learning Representations (ICLR)*, November 2016.

Michal Lukasik, Srinadh Bhojanapalli, Aditya Krishna Menon, and Sanjiv Kumar. Teacher's pet: understanding and mitigating biases in distillation. *CoRR*, abs/2106.10494, 2021. URL https://arxiv.org/abs/2106.10494.

Aditya Krishna Menon, Ankit Singh Rawat, Sashank J. Reddi, Seungyeon Kim, and Sanjiv Kumar. A statistical perspective on distillation. In *Proceedings of the 38th International Conference on Machine Learning, ICML 2021*, volume 139 of *Proceedings of Machine Learning Research*, pp. 7632–7642. PMLR, 2021.

Hossein Mobahi, Mehrdad Farajtabar, and Peter L. Bartlett. Self-distillation amplifies regularization in hilbert space. In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020*, 2020.

Rafael Müller, Simon Kornblith, and Geoffrey E. Hinton. When does label smoothing help? In *Advances in Neural Information Processing Systems 32*, pp. 4696–4705, 2019.

Utkarsh Ojha, Yuheng Li, and Yong Jae Lee. What knowledge gets distilled in knowledge distillation? 2022. doi: 10.48550/arXiv.2205.16004. URL https://doi.org/10.48550/arXiv.2205.16004.

Wonpyo Park, Dongju Kim, Yan Lu, and Minsu Cho. Relational knowledge distillation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019.

Mary Phuong and Christoph Lampert. Towards understanding knowledge distillation. In *Proceedings of the 36th International Conference on Machine Learning*, pp. 5142–5151, 2019.

Ilija Radosavovic, Piotr Dollár, Ross B. Girshick, Georgia Gkioxari, and Kaiming He. Data distillation: Towards omni-supervised learning. In *2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018*, pp. 4119–4128, 2018.

Arman Rahbar, Ashkan Panahi, Chiranjib Bhattacharyya, Devdatt Dubhashi, and Morteza Haghir Chehreghani. On the unreasonable effectiveness of knowledge distillation: Analysis in the kernel regime, 2020.

Yi Ren, Shangmin Guo, and Danica J. Sutherland. Better supervisory signals by observing learning paths. In *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022*. OpenReview.net, 2022.

Adriana Romero, Nicolas Ballas, Samira Ebrahimi Kahou, Antoine Chassang, Carlo Gatta, and Yoshua Bengio. Fitnets: Hints for thin deep nets. In Yoshua Bengio and Yann LeCun (eds.), *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015.

Mark Sandler, Andrew Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. Mobilenetv2: Inverted residuals and linear bottlenecks. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2018.

Samuel Stanton, Pavel Izmailov, Polina Kirichenko, Alexander A. Alemi, and Andrew Gordon Wilson. Does knowledge distillation really work? In *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021*, 2021.

Jiaxi Tang, Rakesh Shivanna, Zhe Zhao, Dong Lin, Anima Singh, Ed H. Chi, and Sagar Jain. Understanding and improving knowledge distillation. *CoRR*, abs/2002.03532, 2020.

Adina Williams, Nikita Nangia, and Samuel Bowman. A broad-coverage challenge corpus for sentence understanding through inference. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pp. 1112–1122. Association for Computational Linguistics, 2018. URL http://aclweb.org/anthology/N18-1101.

Qizhe Xie, Eduard Hovy, Minh-Thang Luong, and Quoc V. Le. Self-training with noisy student improves imagenet classification, 2019.

Li Yuan, Francis E. H. Tay, Guilin Li, Tao Wang, and Jiashi Feng. Revisiting knowledge distillation via label smoothing regularization. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 3902–3910. IEEE, 2020.

Linfeng Zhang, Jiebo Song, Anni Gao, Jingwei Chen, Chenglong Bao, and Kaisheng Ma. Be your own teacher: Improve the performance of convolutional neural networks via self distillation. In *2019 IEEE/CVF International Conference on Computer Vision, ICCV 2019*, 2019.

Xiang Zhang, Junbo Jake Zhao, and Yann LeCun. Character-level convolutional networks for text classification. In *NIPS*, 2015.

Zhilu Zhang and Mert R. Sabuncu. Self-distillation as instance-specific label smoothing. In Hugo Larochelle, Marc'Aurelio Ranzato, Raia Hadsell, Maria-Florina Balcan, and Hsuan-Tien Lin (eds.), *Advances in Neural Information Processing Systems*, 2020.

Helong Zhou, Liangchen Song, Jiajie Chen, Ye Zhou, Guoli Wang, Junsong Yuan, and Qian Zhang. Rethinking soft labels for knowledge distillation: A bias-variance tradeoff perspective. In *9th International Conference on Learning Representations, ICLR 2021*, 2021.

Zaida Zhou, Chaoran Zhuge, Xinwei Guan, and Wen Liu. Channel distillation: Channel-wise attention for knowledge distillation. *CoRR*, abs/2006.01683, 2020. URL https://arxiv.org/abs/2006.01683.

A LIMITATIONS

We highlight a few key limitations to our results that may be relevant for future work to look at:

1. While we show underfitting across many datasets, we do not establish when this form of underfitting can help in situations besides label noise. Indeed, we find that, in our language datasets, underfitting can occur even in the absence of any improvement in generalization. Future work may explore more sophisticated measures of student-teacher deviation that correlates with the gains of distillation. We highlight further caveats of our empirical results in Sec C.2.1.
2. Both our formalizations are based on linear models. On the one hand, this implies that our results are general and fundamental in that they are not unique to neural network training. On the other, this means our result may not capture other interesting effects in the context of distillation (such as that of feature-learing as in Allen-Zhu & Li (2020)). Indeed, we hypothesize that the eigenspace regularization effect may compound across multiple layers of a neural network. This is an open direction for future work.
3. While our gradient-space view suggests why switching to one-hot can hurt, it does not explain why switching to one-hot for a very short amount of time increases the performance — even beyond distillation — in the case of ResNet50 self-distillation (cf. Sec C.2.3 Fig 16 and Fig 15). Analyzing this nuanced effect of distillation-to-one-hot switching is an interesting question for future work.

B EIGENSPACE VIEW

Below, we provide the proof for the sparsity levels of an early-stopped teacher and early-stopped student.

Proof. (of Theorem 4.1) For the early-stopped model, we need $1 - e^{-\lambda_{k_1} t} \geq \alpha^\star_{k_1}$. For any value of t that satisfies this, we can lower bound $\alpha_{k_2}(t)$ as:

$$\alpha_{k_2}(t) = 1 - e^{-\lambda_{k_2} t} \tag{7}$$

$$= 1 - \left(e^{-\lambda_{k_1} t}\right)^{\frac{\lambda_{k_2}}{\lambda_{k_1}}} \tag{8}$$

$$\geq 1 - \left(1 - \alpha^\star_{k_1}\right)^{\frac{\lambda_{k_2}}{\lambda_{k_1}}} . \tag{9}$$

Now, recall that we need to show that *there exists* a set of distillation hyperparameters $\tilde{t}, T^{\mathsf{te}}$ that satisfy our required bounds. We will show that such a set of hyperparameters exist under the simplification that $T^{\mathsf{te}} = \tilde{t}$. Thus, for the student model, we need $(1 - e^{-\lambda_{k_1} \tilde{t}})(1 - e^{-\lambda_{k_1} t}) \geq \alpha^\star_{k_1}$. To achieve the largest skew, we should choose the smallest $\tilde{t}$, which is achieved when $(1 - e^{-\lambda_{k_1} \tilde{t}})^2 = \alpha^\star_{k_1}$. Plugging this into the expression for $\tilde{\alpha}_{k_2}(\tilde{t})$, we get:

$$\tilde{\alpha}_{k_2}(\tilde{t}) = (1 - e^{-\lambda_{k_2} \tilde{t}})^2 \tag{10}$$

$$= \left(1 - \left(e^{-\lambda_{k_1} \tilde{t}}\right)^{\frac{\lambda_{k_2}}{\lambda_{k_1}}}\right)^2 \tag{11}$$

$$= \left(1 - (1 - \sqrt{\alpha^\star_{k_1}})^{\frac{\lambda_{k_2}}{\lambda_{k_1}}}\right)^2 . \tag{12}$$

For simplicity, let us define the fraction $\kappa := \frac{\lambda_{k_2}}{\lambda_{k_1}}$. Putting the above two sets of equation together, we have:

$$\alpha_{k_2}(t) - \tilde{\alpha}_{k_2}(\tilde{t}) \leq 1 - \left(1 - \alpha^\star_{k_1}\right)^f - \left(1 - (1 - \sqrt{\alpha^\star_{k_1}})^f\right)^2 \tag{13}$$

$$\leq 1 - \left(1 - \alpha^\star_{k_1}\right)^f - 1 - ((1 - \sqrt{\alpha^\star_{k_1}})^2)^f + 2(1 - \sqrt{\alpha^\star_{k_1}})^f \tag{14}$$

$$\leq 2(1 - \sqrt{\alpha^\star_{k_1}})^f - ((1 - \sqrt{\alpha^\star_{k_1}})^2)^f - \left(1 - \alpha^\star_{k_1}\right)^f \tag{15}$$

$$\leq 2 - (1 - \sqrt{\alpha^\star_{k_1}}) \left((1 - \sqrt{\alpha^\star_{k_1}}))^f - \left(1 - \sqrt{\alpha^\star_{k_1}}\right)^f\right) \qquad \leq 0 \tag{16}$$

The last step follows from the power means inequality: we have that since $f < 1$, $\left(\frac{(1-\sqrt{\alpha^\star_{k_1}}))^f + (1+\sqrt{\alpha^\star_{k_1}}))^f}{2}\right)^f$ is at most $\left(\frac{(1-\sqrt{\alpha^\star_{k_1}})) + (1+\sqrt{\alpha^\star_{k_1}}))}{2}\right) = 1$. □

B.1 Gradient-space view

In this section, we prove Theorem 4.2 where we analyzed the quality of one-hot and distillation gradients in a setting with class similarity and zero label noise. We first recall the setting below.

Setting. We consider a K-class classification dataset, where the ith datapoint's features can be written as a K-"channel" input $\mathbf{x}_i = (\mathbf{x}_1^{(i)}, \ldots, \mathbf{x}_K^{(i)})$ where each channel $\mathbf{x}_k^{(i)} \in \mathbb{R}^D$ is D-dimensional. Assuming a uniform distribution over K classes, given a label y, we generate $\mathbf{x}_y$ from a Gaussian $\mathcal{N}(0, \mathbf{I}/D)$ truncated to the support $\mathbf{x}_y \cdot \boldsymbol{\mu}^\star_{y_i} = \tau$ for $\tau > 0$ and for a pre-defined class vector $\boldsymbol{\mu}^\star_k$. We also assume that for point $\mathbf{x}_i$, we pick a non-target "similar" class z_i for which we pick $\mathbf{x}_{z_i}$ from a Gaussian $\mathcal{N}(0, \mathbf{I}/D)$ truncated to the support $\mathbf{x}_{z_i} \cdot \boldsymbol{\mu}^\star_{z_i} = \tau/2$. All other co-ordinates are set to zero. We consider a simplified linear architecture whose K-dimensional output is of the form $\mathbf{f}(\mathbf{x}) = (\mathbf{w}_1 \cdot \mathbf{x}_1, \ldots, \mathbf{w}_K \cdot \mathbf{x}_K)$ where the k'th node acts only on the k'th channel.

Notations: Throughout the rest of the discussion, we fix some node k, and define $S_+ = \{(\mathbf{x}, y) \in S | y = k\}$ and $S_- = \{(\mathbf{x}, y) \in S | y \neq k, z = k\}$. Intuitively, S_+ is the subset of samples that provide "constructive" gradients for node k under the one-hot loss, while S_- is the subset that provides "destructive" gradients under the one-hot loss (i.e., gradients in the opposite direction of optimality). With an abuse of notation, we will write $\mathbf{x} \in S$ to denote an $\mathbf{x}$ such that $(\mathbf{x}, y) \in S$.

We will use $\mathbf{p}^{\mathsf{te}}$ to denote a perfect teacher's probabilities, and $\mathbf{q}^{\mathsf{te}}$ to denote an imperfect teacher's probabilities.

Also recall that α^{te} denotes the "scale" of the teacher's logits — think of this as a proxy for how long the teacher has been run to minimize one-hot loss. Similarly, α denotes the scale of the student model's logits. We now state some simplifying assumptions.

Assumption B.1. *We make the following simplifying assumptions:*

1. *We assume that we have a training dataset S such that $S = mK$. The dataset is balanced in that for every k, $|S_+| = |S_-| = m$.*
2. *We assume $\|\boldsymbol{\mu}^\star_k\| = 1$.*
3. *We assume $\epsilon^{te} < 1$ and $\tau > 1$.*
4. *We assume $4\alpha^{te}\tau\epsilon^{te} \leq 1$.*

In Lemma B.1, we show that as long as both the teacher and the student are not trained for too long (and thus their scales α^{te} and α are low), the (perfect) teacher's supervision is "constructive" on node k for all points in S_- and S_+. This is crucial for us to show that distillation is able to denoise the one-hot gradients (which are destructive on S_-). In the subsequent Lemma B.2, we bound the difference in supervision of a perfect and imperfect teacher. Using this, we can show that despite imperfection, the supervision of the teacher can be constructive for node k on both S_- and S_+.

Lemma B.1. *Assuming $\alpha^{te}\tau \leq \log(K-2)$ and $\alpha^{te} - \alpha \geq \frac{2}{\tau}\log\frac{3}{1-\kappa}$, then for any $\mathbf{x} \in S_+ \cup S_-$, $\frac{p_k^{te}(\mathbf{x}) - p_k(\mathbf{x})}{p_k^{te}(\mathbf{x})} \geq \kappa$.*

Proof. If the point belongs to S_-, then:

$$p_k(\mathbf{x}) = \frac{e^{\alpha\tau/2}}{e^{\alpha\tau} + e^{\alpha\tau/2} + K - 2} \text{ and } \qquad p_k^{\mathsf{te}}(\mathbf{x}) = \frac{e^{\alpha^{\mathsf{te}}\tau/2}}{e^{\alpha^{\mathsf{te}}\tau} + e^{\alpha^{\mathsf{te}}\tau/2} + K - 2}. \tag{17}$$

Then, the ratio is :

$$\frac{p_k(\mathbf{x})}{p_k^{\mathsf{te}}(\mathbf{x})} = \frac{e^{\alpha^{\mathsf{te}}\tau/2} + 1 + (K-2)e^{-\alpha^{\mathsf{te}}\tau/2}}{e^{\alpha\tau/2} + 1 + (K-2)e^{-\alpha\tau/2}}. \tag{18}$$

Since $\alpha^{\mathsf{te}}\tau \leq \log(K-2)$, we can upper-bound the first two terms in the numerator by the last term. Hence,

$$\frac{p_k(\mathbf{x})}{p_k^{\mathsf{te}}(\mathbf{x})} \leq \frac{3(K-2)e^{-\alpha^{\mathsf{te}}\tau/2}}{(K-2)e^{-\alpha\tau/2}} \leq 3e^{(\alpha-\alpha^{\mathsf{te}})\frac{\tau}{2}} \leq 1 - \kappa.$$

If the point belongs to S_+, then:

$$p_k(\mathbf{x}) = \frac{e^{\alpha\tau}}{e^{\alpha\tau} + e^{\alpha\tau/2} + K - 2} \text{ and } \qquad p_k^{\mathsf{te}}(\mathbf{x}) = \frac{e^{\alpha^{\mathsf{te}}\tau}}{e^{\alpha^{\mathsf{te}}\tau} + e^{\alpha^{\mathsf{te}}\tau/2} + K - 2}. \tag{19}$$

The ratio is

$$\frac{p_k(\mathbf{x})}{p_k^{\mathsf{te}}(\mathbf{x})} = \frac{1 + e^{-\alpha^{\mathsf{te}}\tau/2} + (K-2)e^{-\alpha^{\mathsf{te}}\tau}}{1 + e^{-\alpha\tau/2} + (K-2)e^{-\alpha\tau}}$$

Again, we can upper-bound each term in the denominator by the last term.

$$\frac{p_k(\mathbf{x})}{p_k^{\mathsf{te}}(\mathbf{x})} \geq \frac{3(K-2)e^{-\alpha^{\mathsf{te}}\tau}}{(K-2)e^{-\alpha\tau}} \leq 1 - \kappa.$$

□

Lemma B.2. *Consider a perfect teacher model with weights $\alpha^{\mathsf{te}}\boldsymbol{\mu}_k^\star$ and corresponding probabilities $\mathbf{p}^{\mathsf{te}}(\cdot)$. Consider an ϵ^{te}-imperfect teacher model with weights $\mathbf{w}_k$ such that $\|\mathbf{w}_k\| = \alpha^{\mathsf{te}}$ and $\frac{\|\mathbf{w}_k^\perp\|^2}{\|\mathbf{w}_k^\parallel\|^2} \leq \epsilon^{\mathsf{te}}$. Let the corresponding probabilities be given by $\mathbf{q}^{\mathsf{te}}(\cdot)$.*

Assuming $4\alpha^{\mathsf{te}}\tau\epsilon^{\mathsf{te}} \leq 1$, on any $\mathbf{x} \in S_+ \cup S_-$,

$$|p_k^{\mathsf{te}}(\mathbf{x}) - q_k^{\mathsf{te}}(\mathbf{x})| \leq 12\alpha^{\mathsf{te}}\tau \cdot \sqrt{\epsilon^{\mathsf{te}}} \cdot p_k^{\mathsf{te}}(\mathbf{x}). \tag{20}$$

Proof. We have $\|\mathbf{w}^\perp\|^2 \leq (\alpha^{\mathsf{te}})^2 \frac{\epsilon^{\mathsf{te}}}{1+\epsilon^{\mathsf{te}}}$ and $\|\mathbf{w}^\parallel\|^2 \geq (\alpha^{\mathsf{te}})^2 \frac{1}{1+\epsilon^{\mathsf{te}}}$.

We have

$$\mathbf{w}_k \cdot \mathbf{x}_k = \|\mathbf{w}^\parallel\| \cdot (\mathbf{x}_k \cdot \boldsymbol{\mu}_k^\star) + \mathbf{w}^\perp \cdot \mathbf{x}_k^\perp. \tag{21}$$

We have $\|\mathbf{w}^\parallel\| \leq \alpha^{\mathsf{te}}$ and $\|\mathbf{w}^\parallel\| \geq \alpha^{\mathsf{te}} \frac{1}{\sqrt{1+\epsilon^{\mathsf{te}}}} \geq \alpha^{\mathsf{te}}(1-\epsilon^{\mathsf{te}})$. We also have $|\mathbf{x}_k \cdot \boldsymbol{\mu}_k^\star| \leq \tau$. Next, $\|\mathbf{w}^\perp\| \leq \alpha^{\mathsf{te}}\sqrt{\frac{\epsilon^{\mathsf{te}}}{1+\epsilon^{\mathsf{te}}}} \leq \alpha^{\mathsf{te}}\sqrt{\epsilon^{\mathsf{te}}}$. Finally, $\lim_{D\to\infty} \|\mathbf{x}_k^\perp\| \leq 1$. Putting these together, we get:

$$|\mathbf{w}_k \cdot \mathbf{x}_k - \alpha^{\mathsf{te}}(\mathbf{x}_k \cdot \mu_k^\star)| \leq \alpha^{\mathsf{te}}\tau\epsilon^{\mathsf{te}} + \alpha^{\mathsf{te}}\sqrt{\epsilon^{\mathsf{te}}} \leq 2\alpha^{\mathsf{te}}\tau\sqrt{\epsilon^{\mathsf{te}}}. \tag{22}$$

Here, we've assumed $\epsilon^{\mathsf{te}} < 1$ and $\tau > 1$.

The above gives us a bound on the differences in the logits. But to bound the probabilities themselves, note that due to the structure of the softmax, both the numerator and the denominator can either grow or diminish by a factor of $e^{2\alpha^{\mathsf{te}}\tau\sqrt{\epsilon^{\mathsf{te}}}}$. Thus, $e^{4\alpha^{\mathsf{te}}\tau\sqrt{\epsilon^{\mathsf{te}}}} \geq \frac{q_k^{\mathsf{te}}(\mathbf{x})}{p_k^{\mathsf{te}}(\mathbf{x})} \geq e^{-4\alpha^{\mathsf{te}}\tau\sqrt{\epsilon^{\mathsf{te}}}}$. We have $e^{-4\alpha^{\mathsf{te}}\tau\sqrt{\epsilon^{\mathsf{te}}}} \geq 1 - 4\alpha^{\mathsf{te}}\tau\sqrt{\epsilon^{\mathsf{te}}}$.

Also, given that $4\alpha^{\mathsf{te}}\tau\epsilon^{\mathsf{te}} \leq 1$, $e^{4\alpha^{\mathsf{te}}\tau\epsilon^{\mathsf{te}}} \leq 1 + e \cdot 4\alpha^{\mathsf{te}}\tau\epsilon^{\mathsf{te}}$. Rearranging this gives us our final inequality.

□

We are now ready to state and prove our main theorem:

Theorem B.3. *Consider a student model with weights $\mathbf{w}_k = \alpha \boldsymbol{\mu}_k^\star$ such that $\alpha \leq \frac{1}{\tau}\log(K-2)$. Let $\boldsymbol{\mu}_k$ denote the one-hot gradients on this model under the cross-entropy loss on the dataset S.*

Consider an imperfect teacher model with weights $\mathbf{w}_k$ such that $\|\mathbf{w}_k\| = \alpha^{te}$ and $\frac{\|\mathbf{w}_k^\perp\|^2}{\|\mathbf{w}_k^\parallel\|^2} \leq \epsilon^{te}$ where $\epsilon^{te} > \frac{1}{m\tau^2}$. Assume the teacher model is early-stopped in that $\alpha^{te} \leq \frac{1}{\tau}\log(K-2)$.

Assuming $\alpha^{te} \geq \frac{1}{\tau}\log 3$, let $\kappa > 0$ denote the level of early-stopping in the student, such that $\alpha^{te} - \alpha \geq \frac{2}{\tau}\log\frac{3}{1-\kappa}$. Let $\tilde{\mathbf{u}}_k$ denote the gradient under distillation loss with this teacher on the above student model.

Denote the components of $\mathbf{u}_k$ and $\tilde{\mathbf{u}}_k$ along $\mu_k^\star$ as $\mathbf{u}_k^\parallel$ and $\tilde{\mathbf{u}}_k^\parallel$ and orthogonal to $\mu_k^\star$ as $\mathbf{u}_k^\perp$ and $\tilde{\mathbf{u}}_k^\perp$ respectively.

Then, for a sufficiently large τ, there exists values of α, κ such that $\alpha \in [0, \alpha^{te}]$ such that $\kappa^3 \geq 192\alpha^{te}\tau\sqrt{m\epsilon^{te}}$ and so we have for any class k that:

$$\underbrace{\lim_{D\to\infty}\frac{\|\mathbf{u}_k^\parallel\|^2}{\|\mathbf{u}_k^\perp\|^2} \leq m\tau^2\left(1-\frac{e^{-\alpha\tau}}{9}\right)}_{\text{Upper-bound on quality of one-hot gradients}} \leq \underbrace{m\tau^2\left(1+\frac{\kappa^2}{2}\right) \leq \lim_{D\to\infty}\frac{\|\tilde{\mathbf{u}}_k^\parallel\|^2}{\|\tilde{\mathbf{u}}_k^\perp\|^2}}_{\text{Lower-bound on quality of distillation gradients}}. \tag{23}$$

Furthermore, the quality of the teacher weights can be upper bounded by the quality of distillation gradients as :

$$\frac{1}{\epsilon^{te}} \leq \underbrace{m\tau^2\left(1+\frac{\kappa^2}{2}\right) \leq \lim_{D\to\infty}\frac{\|\tilde{\mathbf{u}}_k^\parallel\|^2}{\|\tilde{\mathbf{u}}_k^\perp\|^2}}_{\text{Lower-bound on quality of distillation gradients}}. \tag{24}$$

Proof. We have assumed $|S_+| = |S_-| = m$. Note that for all $\mathbf{x}$ in S_+, $p_k(\mathbf{x}) = p_+$ for some constant $p_+ > 0$. Similarly, for all $\mathbf{x}$ in S_-, $p_k(\mathbf{x}) = p_-$ for some constant $p_+ > 0$.

Let the gradient descent update on node k be written as $\mathbf{u}_k = \mathbf{u}_k^\parallel + \mathbf{u}_k^\perp$, where $\mathbf{u}_k^\perp \cdot \boldsymbol{\mu}_k^\star = 0$.

The one-hot gradient at node k can be written as:

$$\mathbf{u}_k = \sum_{\mathbf{x}\in S_+}(1-p_k(\mathbf{x}))\mathbf{x}_k - \underbrace{\sum_{\mathbf{x}\in S_-} p_k(\mathbf{x})\mathbf{x}_k}_{\text{destructive gradients}} \tag{25}$$

$$= (1-p_+)\sum_{\mathbf{x}\in S_+}\mathbf{x}_k - p_-\left(\sum_{\mathbf{x}\in S_-}\mathbf{x}_k\right). \tag{26}$$

We project the sum total of all these gradients along $\boldsymbol{\mu}_k^\star$ to get

$$\|\mathbf{u}_k^\parallel\| = \mathbf{u}_k \cdot \boldsymbol{\mu}_k^\star = m\tau\left((1-p_+) - p_-\right).$$

Next, we want to compute the total magnitude of all the gradients orthogonal to $\boldsymbol{\mu}_k^\star$, namely $\|\mathbf{u}_k^\perp\|$. Note that the projection of any sample $\mathbf{x}$ orthogonal to $\boldsymbol{\mu}_k^\star$ is essentially a $D-1$ multivariate zero-mean Gaussian with variance $1/D$ along each direction. Thus, the projection of $\mathbf{u}_k^\perp$ along each of these $D-1$ directions is a zero-mean Gaussian with variance equal to $\frac{m}{D}\left((1-p_+)^2 + p_-^2\right)$. Thus $\|\mathbf{u}_k^\perp\|^2$ is the sum of squared of $D-1$ of these random variables. Note that:

$$\lim_{D\to\infty}\|\mathbf{u}_k^\perp\|1 = m\left((1-p_+)^2 + p_-^2\right).$$

Thus, the ratio can be written as:

$$\lim_{D\to\infty} \frac{\|\mathbf{u}_k^{\|}\|^2}{\|\mathbf{u}_k^{\perp}\|^2} = \frac{m^2\tau^2\left((1-p_+)-p_-\right)^2}{m\left((1-p_+)^2+p_-^2\right)} \leq m\tau^2\left(1-\frac{2(1-p_+)p_-}{(1-p_+)^2+p_-^2}\right) \tag{27}$$

$$\leq m\tau^2\left(1-\left(\frac{p_-}{p_+}\right)^2\right). \tag{28}$$

We can lower bound p_-/p_+ as:

$$\frac{p_+}{p_-} \geq \frac{1+(K-2)e^{-\alpha\tau}+e^{-\alpha\tau/2}}{1+(K-2)e^{-\alpha\tau/2}+e^{\alpha\tau/2}} \geq \frac{(K-2)e^{-\alpha\tau}}{3(K-2)e^{-\alpha\tau/2}} \geq \frac{e^{-\alpha\tau/2}}{3}, \tag{29}$$

where we use $\alpha\tau \leq \log(K-2)$ to bound each of the denominator terms.

Plugging this back, we get the following bound on the quality of one-hot gradients:

$$\lim_{D\to\infty} \frac{\|\mathbf{u}_k^{\|}\|^2}{\|\mathbf{u}_k^{\perp}\|^2} \leq m\tau^2\left(1-\frac{e^{-\alpha\tau}}{9}\right). \tag{30}$$

Now let p_+^{te} and p_-^{te} denote the probabilities of a *perfect* teacher. We can apply an identical argument for distillation to bound the ratio of the gradients as:

$$\lim_{D\to\infty} \frac{\|\tilde{\mathbf{u}}^{\|}\|^2}{\|\tilde{\mathbf{u}}^{\perp}\|^2} := \frac{m^2\tau^2((p_+^{\mathsf{te}}-p_+)+(p_-^{\mathsf{te}}-p_-))}{m((p_+^{\mathsf{te}}-p_+)^2+(p_-^{\mathsf{te}}-p_-)^2)} \geq m\tau^2\left(1+2\frac{(p_+^{\mathsf{te}}-p_+)(p_-^{\mathsf{te}}-p_-)}{(p_+^{\mathsf{te}}-p_+)^2+(p_-^{\mathsf{te}}-p_-)^2}\right) \tag{31}$$

$$\geq m\tau^2\left(1+2\frac{\kappa^2 p_+^{\mathsf{te}}p_-^{\mathsf{te}}}{(p_+^{\mathsf{te}})^2+(p_-^{\mathsf{te}})^2}\right) \geq m\tau^2(1+\kappa^2), \tag{32}$$

where in the second line we've used Lemma B.1 to lower-bound the teacher-student probability gap, subsequently followed by the AM-GM inequality. Here, we have used the fact that $p_-^{\mathsf{te}} - p_- > 0$, or in other words, that the gradients from S_- are constructive.

The next step in the proof is to extend this lower-bound for an *imperfect* teacher with probabilities $\mathbf{q}^{\mathsf{te}}(\mathbf{x})$. For convenience, let us assume that for $\mathbf{x} \in S_+$ $|p_k^{\mathsf{te}}(\mathbf{x}) - q_k^{\mathsf{te}}(\mathbf{x})| \leq \gamma_+$ and $\mathbf{x} \in S_-$, $|p_k^{\mathsf{te}}(\mathbf{x}) - q_k^{\mathsf{te}}(\mathbf{x})| \leq \gamma_-$ for some constants γ_+ and γ_-. We can say that:

$$\|\tilde{\mathbf{u}}^{\|}\|^2 = \tau^2\left(\sum_{\mathbf{x}\in S_-\cup S_+} q_k^{\mathsf{te}}(\mathbf{x}) - p_k(\mathbf{x})\right)^2 \geq m^2\tau^2\left((p_+^{\mathsf{te}}-p_+)+(p_-^{\mathsf{te}}-p_-)-\gamma_+-\gamma_-\right)^2. \tag{33}$$

If we have $\gamma_- \leq \frac{\epsilon'}{2}\kappa p_-^{\mathsf{te}}$ for some ϵ', from Lemma B.1, we have $\gamma_- \leq \frac{\epsilon'}{2}\kappa(p_-^{\mathsf{te}} - p_-)$. We can make a similar claim for γ_+. Thus,

$$\|\tilde{\mathbf{u}}^{\|}\|^2 \geq m^2\tau^2\left((p_+^{\mathsf{te}}-p_+)+(p_-^{\mathsf{te}}-p_-)\right)^2\left(1-\frac{\epsilon'}{2}\right)^2 \tag{34}$$

$$\geq m^2\tau^2\left((p_+^{\mathsf{te}}-p_+)+(p_-^{\mathsf{te}}-p_-)\right)^2(1-\epsilon') \tag{35}$$

$$\tag{36}$$

As for the orthogonal gradient term, recall that $\|\mathbf{u}_k^\perp\|^2$ is essentially the summation of $D-1$ squared random variables. With a perfect teacher, we could write each of the $D-1$ variables as a sum of m Gaussians each scaled by a constant probability $(p_+^{\mathsf{te}} - p_+)$ and $(p_-^{\mathsf{te}} - p_-)$ independent of the sample $\mathbf{x}$. However, with the imperfect teacher, the probabilities are themselves random variables dependent on the draw of $\mathbf{x}$. So we first bound the gap between this norm and the norm under the perfect teacher.

Without loss of generality, let us denote the $D-1$ dimensions of $\mathbf{x}_k$ that is orthogonal to $\boldsymbol{\mu}_k{}^\star$ as $x_{k,1}, x_{k,2}, \ldots, x_{k,D-1}$. Then,

$$\lim_{D\to\infty} \|\tilde{\mathbf{u}}^\perp\|^2 \leq \lim_{D\to\infty} \sum_{j=1}^{D-1} \left(\sum_{\mathbf{x}\in S_+\cup S_-} (p_k^{\mathsf{te}}(\mathbf{x}) - p_k(\mathbf{x}))x_{k,j} \right)^2 \tag{37}$$

$$\leq \lim_{D\to\infty} \sum_{j=1}^{D-1} \left(\sum_{\mathbf{x}\in S_+\cup S_-} (q_k^{\mathsf{te}}(\mathbf{x}) - p_k(\mathbf{x}))x_{k,j} + x_{k,j} \underbrace{(p_k^{\mathsf{te}}(\mathbf{x}) - q_k^{\mathsf{te}}(\mathbf{x}))}_{\leq\gamma_+ \text{ or } \leq\gamma_-} \right)^2 \tag{38}$$

$$\leq \underbrace{\lim_{D\to\infty} \sum_{j=1}^{D-1} \left(\sum_{\mathbf{x}\in S_+\cup S_-} (p_k^{\mathsf{te}}(\mathbf{x}) - p_k(\mathbf{x}))x_{k,j} \right)^2}_{\leq m((p_+^{\mathsf{te}}-p_+)^2+(p_-^{\mathsf{te}}-p_-)^2)} \tag{39}$$

$$+ \lim_{D\to\infty} \sum_{j=1}^{D-1} \underbrace{\left(\sum_{\mathbf{x}\in S_+} \gamma_+|x_{k,j}| + \sum_{\mathbf{x}\cup S_-} \gamma_-|x_{k,j}| \right)^2}_{A_j} \tag{40}$$

$$+ \lim_{D\to\infty} 2\sum_{j=1}^{D-1} \underbrace{\sum_{\mathbf{x}\in S_+\cup S_-} (p_k^{\mathsf{te}}(\mathbf{x}) - p_k(\mathbf{x}))|x_{k,j}| \left(\sum_{\mathbf{x}\in S_+} \gamma_+|x_{k,j}| + \sum_{\mathbf{x}\in S_-} \gamma_-|x_{k,j}| \right)}_{B_j} \tag{41}$$

$$\tag{42}$$

For each j, A_j is the squared of the sum of the absolute value m Gaussians with variance $\frac{\gamma_+^2}{D}$ and another m with variance $\frac{\gamma_-^2}{D}$. The means of these variables are $\gamma_+\sqrt{\frac{1}{D}\frac{2}{\pi}}$ and $\gamma_-\sqrt{\frac{}{D}\frac{2}{\pi}}$ respectively. The expectation of A_j is the squared sum of the means and the sum of the variances of all these variables, equalling $\frac{m(\gamma_+^2+\gamma_-^2)}{D} + m^2\left(\gamma_+\sqrt{\frac{}{D}\frac{2}{\pi}} + \gamma_-\sqrt{\frac{1}{D}\frac{2}{\pi}}\right)^2$. The second term can be upper bounded by $m^2\left(\frac{\gamma_+^2}{D} + \frac{\gamma_-^2}{D}\right)$. Therefore, we can conclude that $\lim_{D\to\infty}\sum_{j=1}^{D-1} A_j \leq 3m^2(\gamma_+^2+\gamma_-^2)$.

For each j, B_j can be expanded into $(2m)^2$ terms. The expected value of each of these terms can be bounded in the form of $(p_k^{\mathsf{te}}(\mathbf{x}) - p_k(\mathbf{x}))\left(\frac{\gamma_+}{D}\right)$ or $(p_k^{\mathsf{te}}(\mathbf{x}) - p_k(\mathbf{x}))\left(\frac{\gamma_-}{D}\right)$. Thus, the expected value of B_j is $m\left(\frac{\gamma_-}{D} + \frac{\gamma_+}{D}\right)\sum_{\mathbf{x}}(p_k^{\mathsf{te}}(\mathbf{x}) - p_k(\mathbf{x}))$. Let $\lim_{D\to\infty} B_j = m\left(\gamma_- + \gamma_+\right)\sum_{\mathbf{x}}(p_k^{\mathsf{te}}(\mathbf{x}) - p_k(\mathbf{x}))$

Now, if we have that $\sqrt{m}\gamma_- \leq \frac{\epsilon'}{4}\kappa p_-^{\mathsf{te}}$ and $\sqrt{m}\gamma_+ \leq \frac{\epsilon'}{4}\kappa p_+^{\mathsf{te}}$, for some ϵ' we have:

$$\lim_{D\to\infty} \|\tilde{\mathbf{u}}^\perp\|^2 \leq m((p_+^{\mathsf{te}} - p_+)^2 + (p_-^{\mathsf{te}} - p_-)^2) \tag{43}$$

$$+ 3m\frac{(\epsilon')^2}{16}\left((\kappa p_+^{\mathsf{te}})^2 + (\kappa p_-^{\mathsf{te}})^2\right) + m\frac{\epsilon'}{4}(\kappa p_+^{\mathsf{te}} + \kappa p_-^{\mathsf{te}})(p_+^{\mathsf{te}} - p_+) + (p_-^{\mathsf{te}} - p_-)) \tag{44}$$

$$\tag{45}$$

In the second term, we can upper bound $(\kappa p_+^{\mathsf{te}})^2 + (\kappa p_-^{\mathsf{te}})^2$ by $((p_+^{\mathsf{te}} - p_+)^2 + (p_-^{\mathsf{te}} - p_-)^2)$ using Lemma B.1. In the third term, we can similarly upper bound κp_+^{te} and κp_-^{te} followed by an AM-GM inequality to upper-bound by $2((p_+^{\mathsf{te}} - p_+)^2 + (p_-^{\mathsf{te}} - p_-)^2)$. Therefore,

$$\lim_{D\to\infty} \|\tilde{\mathbf{u}}^{\perp}\|^2 \leq m((p_+^{\mathsf{te}} - p_+)^2 + (p_-^{\mathsf{te}} - p_-)^2)\left(1 + 3\frac{(\epsilon')^2}{16} + \frac{\epsilon'}{4}\right) \tag{46}$$

$$\leq m((p_+^{\mathsf{te}} - p_+)^2 + (p_-^{\mathsf{te}} - p_-)^2)\,(1+\epsilon')\,. \tag{47}$$

$$\tag{48}$$

Then, we have that for our imperfect teacher:

$$\lim_{D\to\infty} \frac{\|\tilde{\mathbf{u}}^{\|}\|^2}{\|\tilde{\mathbf{u}}^{\perp}\|^2} \geq m\tau^2(1+\kappa^2)\frac{(1-\epsilon')}{(1+\epsilon')} \geq m(1+\kappa^2)(1-\epsilon')^2$$

$$\geq m\tau^2(1+\kappa^2)(1-2\epsilon') \geq m\tau^2(1+\kappa^2-4\epsilon') \geq m\left(1+\frac{\kappa^2}{2}\right).$$

The last step follows if we have $\epsilon' \leq \frac{\kappa^2}{8}$. Plugging this back in our required inequality of the form $\sqrt{m}\gamma_- \leq \frac{\epsilon'}{4}\kappa p_-^{\mathsf{te}}$, we need $\sqrt{m}\gamma_- \leq \frac{\kappa^3}{16}p_-^{\mathsf{te}}$ and similarly for γ_+. From Lemma B.2, this means we need $\kappa^3 \geq 192\alpha^{\mathsf{te}}\tau\sqrt{m\epsilon^{\mathsf{te}}}$.

Now, we need to make sure that there are possible values of α for which the above lower-bound on κ can be realized. Note that κ can only attain values of $1 - 3e^{-(\alpha^{\mathsf{te}}-\alpha)\frac{\tau}{2}}$ for various values of α. This follows from the statement of Lemma B.1. Thus, we want $1 - 3e^{-(\alpha^{\mathsf{te}}-\alpha)\frac{\tau}{2}} > \left(192\alpha^{\mathsf{te}}\tau\sqrt{m\epsilon^{\mathsf{te}}}\right)$ for there to be any values of α to achieve our lower-bound on κ. Furthermore, we have that $\epsilon = \frac{c}{m\tau^2}$ for some $c > 1$. Thus, we need $1 - 3e^{-(\alpha^{\mathsf{te}}-\alpha)\frac{\tau}{2}} > \left(192\alpha^{\mathsf{te}}\sqrt{c}\right)$ to have feasible solutions for α. This is possible as long as τ is sufficiently large. □

C Experiments and plots

C.1 Details of experimental setup

We present details on relevant hyper-parameters for our experiments.

Model architectures. For all image datasets (CIFAR10, CIFAR100, Tiny-ImageNet), we use ResNet-v2 (He et al., 2016a) and MobileNet-v2 (Sandler et al., 2018), models. Specifically, for CIFAR, we consider the CIFAR ResNet-$\{56, 20\}$ family and MobileNet-v2 architectures; for Tiny-ImageNet, we consider the ResNet-$\{50, 18\}$ family and MobileNet-v2 architectures; For all ResNet models, we employ standard augmentations as per He et al. (2016b).

For all text datasets (MNLI, AGNews, QQP, IMDB), we fine-tune a pre-trained RoBERTa (Liu et al., 2019) model. We consider a combinations of cross-architecture- and self-distillation with RoBERTa -Base, -Medium and -Small architectures.

Training settings. We train using minibatch SGD applied to the softmax cross-entropy loss. For all image datasets, we follow the settings in Table 1.

For all text datasets, we use a batch size of 64, and train for 25000 steps. We use a peak learning rate of 10^{-5}, with 1000 warmup steps, decayed linearly. For the distillation experiments, we use a distillation weight of 1.0. We use temperature $\tau = 2.0$ for MNLI, $\tau = 16.0$ for IMDB, $\tau = 1.0$ for QQP, and $\tau = 1.0$ for AGNews.

Parameter	**CIFAR10***	**Tiny-ImageNet**
Weight decay	10^{-4}	$5 \cdot 10^{-4}$
Batch size	1024	128
Epochs	450	200
Peak learning rate	1.0	0.1
Learning rate warmup epochs	15	5
Learning rate decay factor	0.1	0.1
Learning rate decay epochs	$200, 300, 400$	$75, 135$
Nesterov momentum	0.9	0.9
Distillation weight	1.0	1.0
Distillation temperature	4.0	4.0
Gradual loss switch window	$1k$ steps	$10k$ steps

Table 1: Summary of training settings.

Dataset	**Teacher**	**Student**	**Test accuracy**		
			Teacher	**Student (OH)**	**Student (DIST)**
CIFAR10	ResNet-56	ResNet-56	93.72	93.72	93.9
	ResNet-56	ResNet-20	93.72	91.83	92.94
	ResNet-56	MobileNet-v2-1.0	93.72	85.11	87.81
CIFAR100	ResNet-56	ResNet-56	72.52	72.52	74.55
	ResNet-56	ResNet-20	72.52	67.52	70.87
	MobileNet-v2-1.0	MobileNet-v2-1.0	54.32	54.32	56.32
	ResNet-56	MobileNet-v2-1.0	72.52	54.32	62.4
CIFAR100 noise	ResNet-56	ResNet-56	69.8	69.8	72.7
	ResNet-56	ResNet-20	69.8	64.9	69.2
Tiny-ImageNet	ResNet-50	ResNet-50	66	66	66.44
	ResNet-50	ResNet-18	66	62.78	63.98
	ResNet-50	MobileNet-v2-1.0	66	62.75	63.97
	MobileNet-v2-1.0	MobileNet-v2-1.0	62.75	62.75	63.28
MNLI	RoBERTa-Base	RoBERTa-Small	87.4	69.9	70.3
MNLI	RoBERTa-Small	RoBERTa-Small	69.9	69.9	69.9
MNLI	RoBERTa-Medium	RoBERTa-Medium	83.8	83.8	83.5
IMDB	RoBERTa-Small	RoBERTa-Small	90.4	90.4	91.0
QQP	RoBERTa-Small	RoBERTa-Small	83.5	83.5	82.5
QQP	RoBERTa-Medium	RoBERTa-Medium	89.7	89.7	89.0
AGNews	RoBERTa-Small	RoBERTa-Small	93.8	93.8	93.9
AGNews	RoBERTa-Medium	RoBERTa-Medium	94.6	94.6	94.6

Table 2: Summary of train and test performance of various distillation settings.

C.2 Additional results

C.2.1 Scatter plots of probabilities

In this section, we present additional scatter plots of the teacher-student logit-transformed probabilities for the class corresponding to the teacher's top prediction: Fig 5 (for CIFAR100), Fig 6 (for

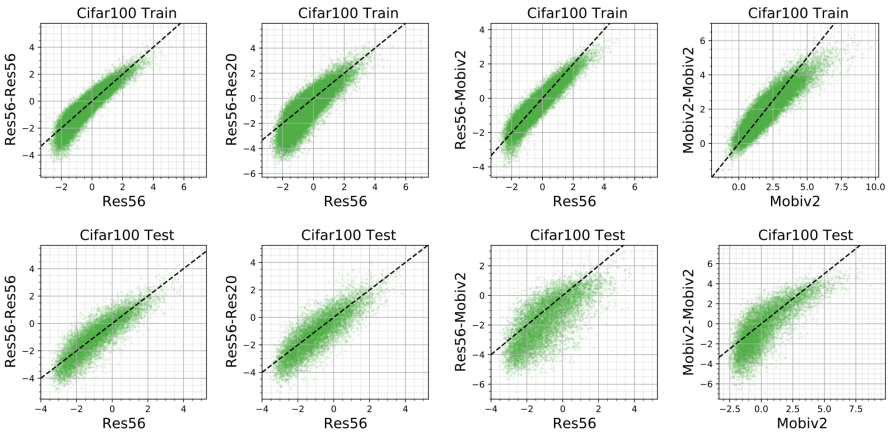

(a)

Figure 5: **Teacher-student logit plots for CIFAR100 experiments:** We report plots for various distillation settings involving ResNet56, ResNet20 and MobileNet-v2. We find underfitting of the hard points in the training set in all but the MobileNet self-distillation setting. Nevertheless, even in the MobileNet self-distillation setting, we find significant underfitting in the *test* dataset.

TinyImagenet), Fig 7 (for CIFAR10), Fig 8 (for MNLI and AGNews), Fig 9 (for self-distillation on QQP, IMDB and AGNews) and Fig 10 (for cross-architecture distillation on language datasets). We find underfitting of the hard points in a majority of the cases, with a few **notable exceptions and caveats**:

1. For MobileNet self-distillation on CIFAR100, and for a majority of the CIFAR10 experiments, we find no underfitting of the harder points *on the training dataset*. However, we do find underfitting of the harder points *on the test dataset*. Interestingly, we also find an underfitting of easier points in the training dataset.
2. In the language datasets, we generally find the plots to be different in pattern from the image datasets. In particular we find that, for harder points, there is both significant underfitting and overfitting. For easier points, there is less deviation, and if any, the deviation is from overfitting. We interpret this as the regularization from distillation inducing *less precision* on the harder points, rather than an underfitting per se.
3. Our patterns generally break down in the *cross-architecture* settings of language datasets (with the exception of a couple of settings). We suspect that this may be because certain cross-architecture effects dominate over the underfitting, which is a much more subtle effect.
4. The underfitting we observe in the language datasets are not always associated with an improvement in the student's generalization. e.g., we find no improvement in AGNews, and a decrease in generalization in QQP (see Table 2). (Note that we didn't finetune hyperparameters in these settings to make distillation work.)

C.2.2 Teacher's predicted class vs. ground truth class

Recall that in all our scatter plots we have looked at the probabilities of the teacher and the student on the teacher's predicted class i.e., $(p^{\mathsf{te}}_{y^{\mathsf{te}}}(x), p^{\mathsf{st}}_{y^{\mathsf{te}}}(x))$ where $y^{\mathsf{te}} \doteq \mathrm{argmax}_{y' \in [K]}\ p^{\mathsf{te}}_{y'}(x)$. Another natural alternative would have been to look at the probabilities for the ground truth class, $(p^{\mathsf{te}}_{y^\star}(x), p^{\mathsf{st}}_{y^\star}(x))$ where $y^\star$ is the ground truth label. We chose to look at y^{te} however, because we are interested in the "shortcomings" of the distillation procedure where the student only has access to teacher probabilities and not ground truth labels.

Nevertheless, one may still be curious as to what the probabilities for the ground truth class look like. First, we note that the plots look almost identical for the *training dataset* owing to the fact that the teacher model typically fits the data to zero training error (we skip these plots to avoid redundancy).

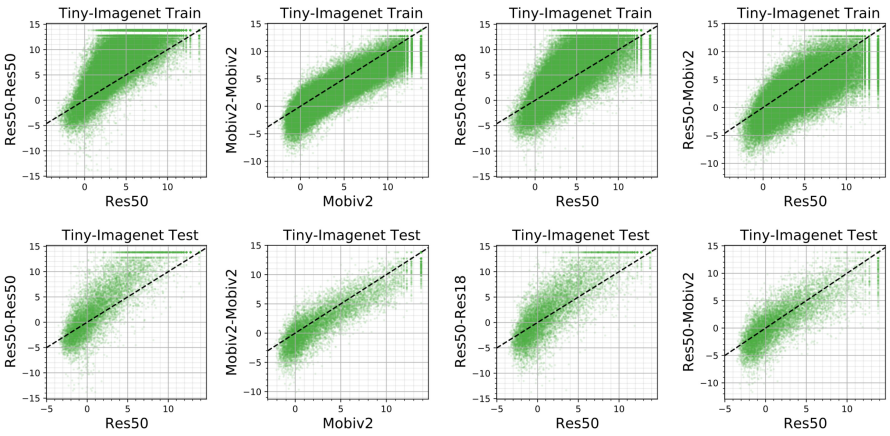

(a)

Figure 6: **Teacher-student logit plots for Tiny-Imagenet experiments:** We report plots for various distillation settings involving ResNet50, ResNet18 and MobileNet-v2. We find underfitting of the hard points in all the settings. We also find *overfitting of the easier points* when the student is a ResNet.

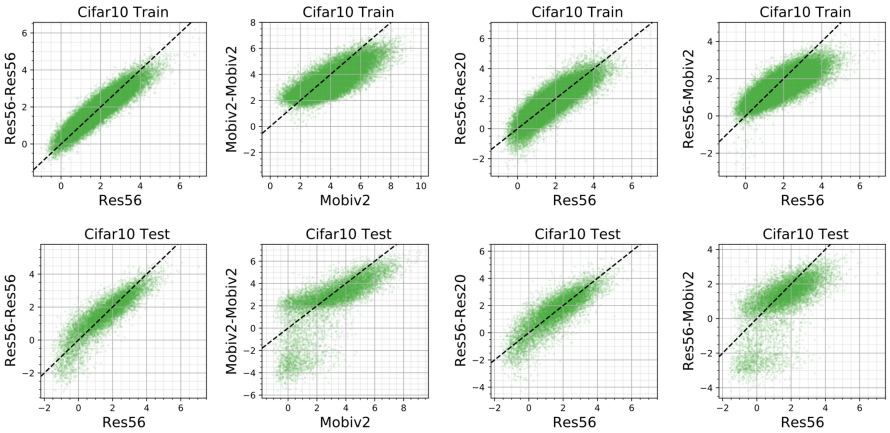

(a)

Figure 7: **Teacher-student logit plots for CIFAR10 experiments:** We report plots for various distillation settings involving ResNet56, ResNet20 and MobileNet-v2. We find that the underfitting phenomenon is almost non-existent in the training set (except for ResNet50 to ResNet20 distillation). However the phenomenon is prominent in the test dataset.

However, we find stark differences in the test dataset as shown in Fig 11. In particular, we see that the underfitting phenomenon is no longer prominent, and almost non-existent in many of our settings.

This is surprising as this suggests that the student somehow matches the probabilities on the ground truth class despite not knowing what the ground truth class is. This also tells us the mechanism by which the student manages to correct some of the teacher's labels: by reducing the probability of non-ground-truth classes rather than increasing the probability of the ground truth class. Indeed, in Fig 12, we verify this to be the case. We dissect the test data plots into four parts depending on which of the teacher and student model gets the label right. We consistently find that the points that the student gets right but the teacher not, fall in the underfit set of points. Interestingly, the underfit set of points is roughly *equivalent* to the set of all points where *at least one of the models is incorrect*. This suggests that in its attempt to underfit some of the points, the student can get some incorrect.

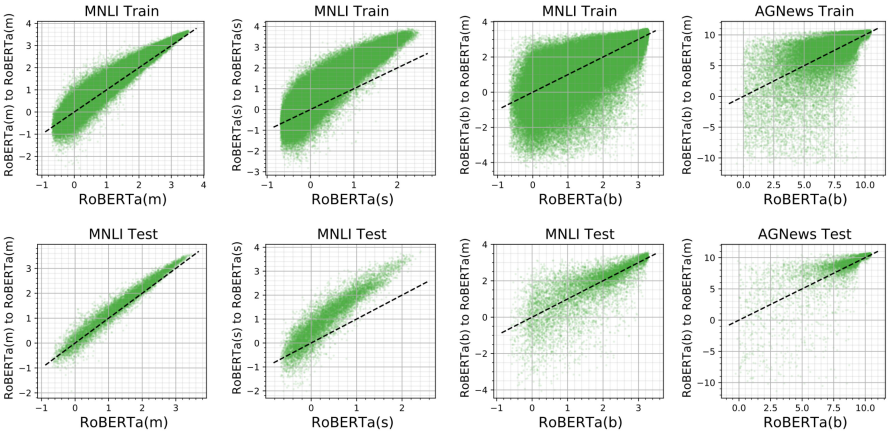

(a) Self-distillation in MNLI

(b) Cross-architecture distillation in MNLI and AGNews

Figure 8: **Teacher-student logit plots for MNLI and AGNews experiments:** We report plots for various distillation settings involving RoBERTa models. On the **left**, in the self-distillation settings on MNLI, we find significant underfitting of hard points (and also overfitting), while easy points are completely overfit. On the **right**, we report cross-architecture (Base to Medium) distillation for MNLI and AGNews. Here the plots are not as typical of our other plots. Neverthless, we observe significant underfitting *and* overfitting of hard points, and overfitting of the extremely easy points. We interpret this as distillation reducing its "precision" on the harder points (perhaps by ignoring lower eigenvectors that provide finer precision).

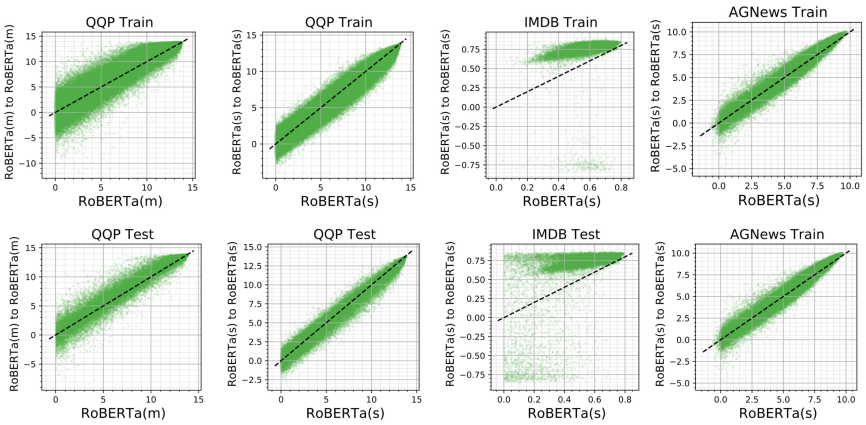

(a)

Figure 9: **Teacher-student logit plots for self-distillation in language datasets (QQP, IMDB, AGNews):** We report plots for various self-distillation settings involving RoBERTa models. As in the other language dataset settings, we find both significant underfitting and overfitting for harder points (indicating lack of precision), and with more precision for easier points (typically with more overfitting).

Finally, we note that previous work (Lukasik et al., 2021) has examined deviations on ground truth class probabilities albeit in an aggregated sense (at a class-level rather than at a sample-level). While they find that the student tends to have lower ground truth probability than the teacher on problems with label imbalance, they do *not* find any such difference on standard datasets without imbalance. This is in alignment with what we find above.

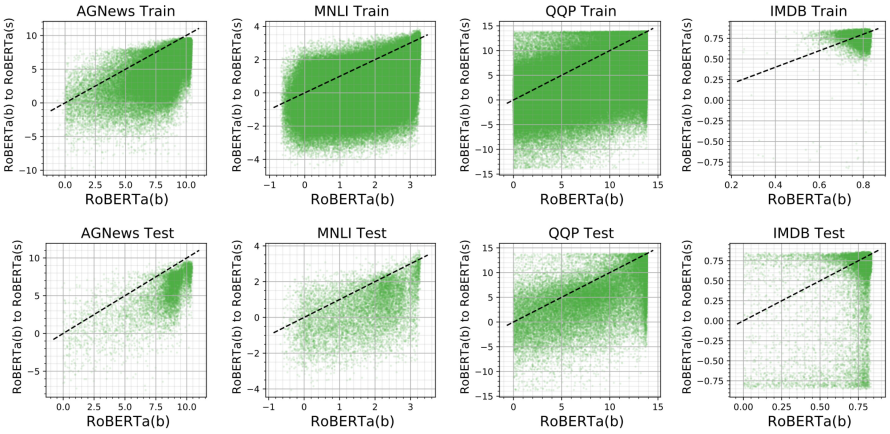

(a)

Figure 10: **Teacher-student logit plots for cross-architecture distillation in language datasets (AGNews, QQP, IMDB, MNLI):** We report plots for various cross-architecture distillation settings involving RoBERTa models. While we find significant student-teacher deviations in these settings, our typical patterns do not apply here. We believe that effects due to "cross-architecture gaps" may have likely drowned out the underfitting patterns, which is a more subtle phenomenon that shines in self-distillation settings.

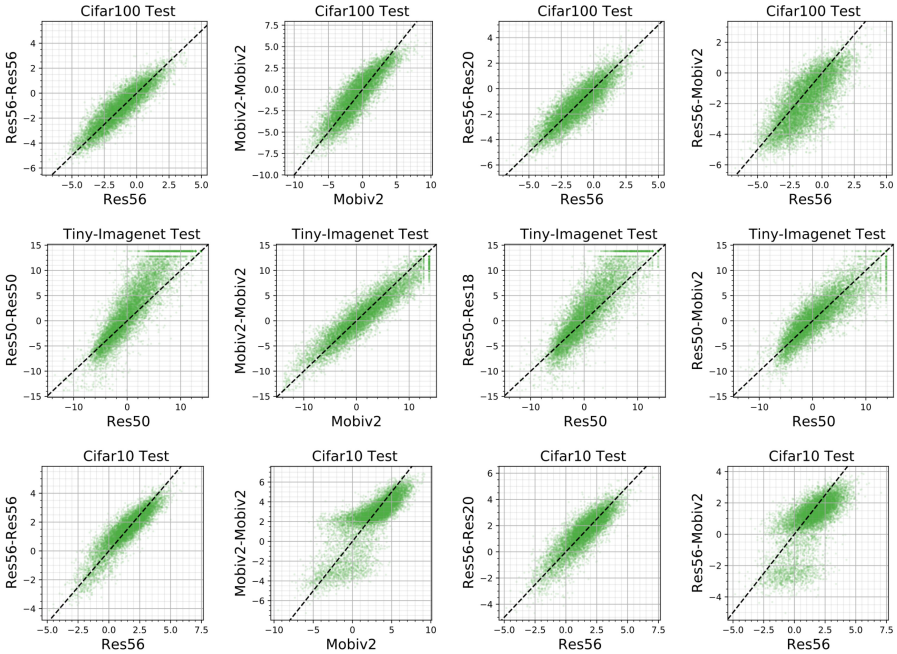

(a)

Figure 11: **Scatter plots for ground truth class:** Unlike in other plots where we report the probabilities for the class predicted by the teacher, here we focus on the ground truth class. Recall that the X-axis corresponds to the teacher, the Y-axis to the student, and all the probabilities are log-transformed. Surprisingly, we observe a much more subdued underfitting here, with the phenomenon completely disappearing in many situations. This suggests that the student magically *preserves* the ground-truth probabilities *despite no knowledge of what the ground-truth class is*, while underfitting on the teacher's predicted class.

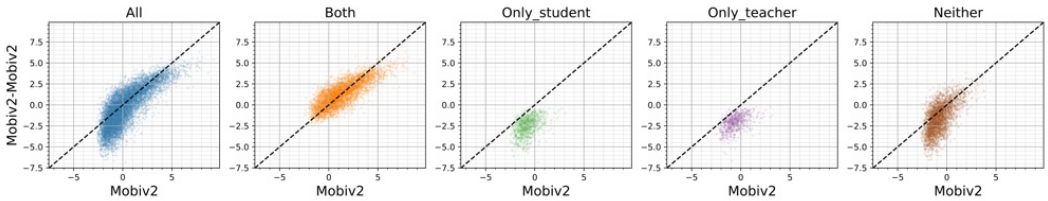

(a) CIFAR100 MobileNet-v2 self-distillation

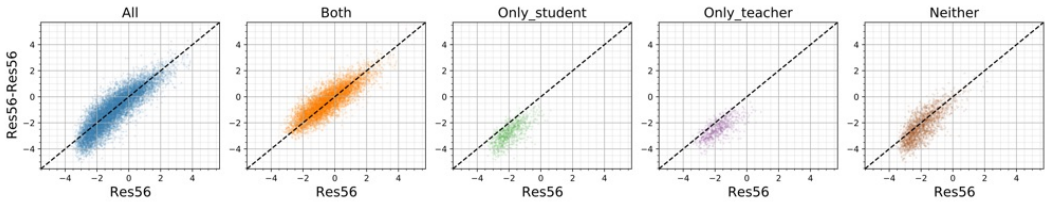

(b) CIFAR100 ResNet56 self-distillation

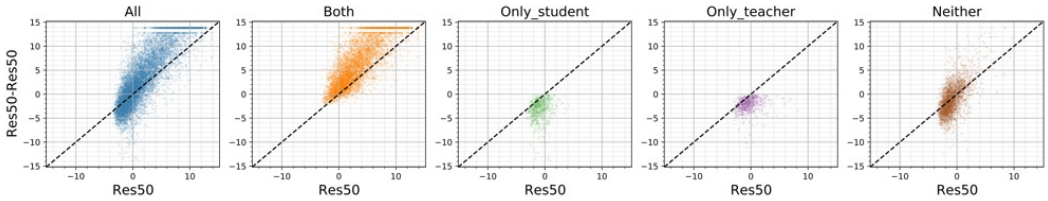

(c) TinyImageNet ResNet50 self-distillation

Figure 12: **Dissecting the underfit points:** Across a few settings on TinyImagenet and CIFAR100, we separate the teacher-student scatter plots of logit-transformed probabilities into four subsets: subsets where both models's top prediction is correct (titled as Both), where only the student gets correct (Only_student), where only the teacher gets correct (Only_teacher), where neither get correct (Neither). We consistently find that the student's "underfit" points are points where at least one of the models go wrong.

C.2.3 THE EFFECT OF LOSS SWITCHING

We provide additional results on loss-switching in this section. Fig 14 presents the accuracy values under a one-hot-to-distillation switch. Fig 15 presents the accuracy values under a distillation-to-one-hot switch. Fig 13 presents student-teacher scatter plots over the course of training under all the four methods (with and without loss switches). Fig 16 presents the trajectory of accuracies for ResNet50 on TinyImagenet under all the four settings.

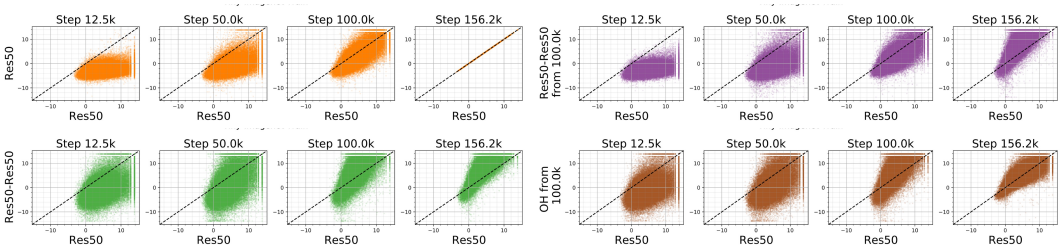

(a) One-hot and self-distillation. (b) Loss-switching to distillation/one-hot at $100k$ steps.

Figure 13: **Evolution of logit-logit plots over various steps of training for TinyImageNet ResNet50 self-distillation setup:** On the **left**, we present plots for one-hot training (**top**) and distillation (**bottom**). Like in the case of CIFAR100, we again see a significant deviation between the standalone and student plots early on during training. At step 12.5k, we see that the standalone model has prioritized fitting some of the harder points close to $x = y$ line. The student has however fit easier points more substantially. On the **right**, we present similar plots for experiments from Section 3.2, with the loss switched to distillation (**top**) and one-hot (**bottom**) at $100k$ steps. From the last two visualized plots in each, observe that switching to distillation introduces (a) underfitting of hard points and overfitting of easier points, (b) while switching to one-hot curiously undoes both of this.

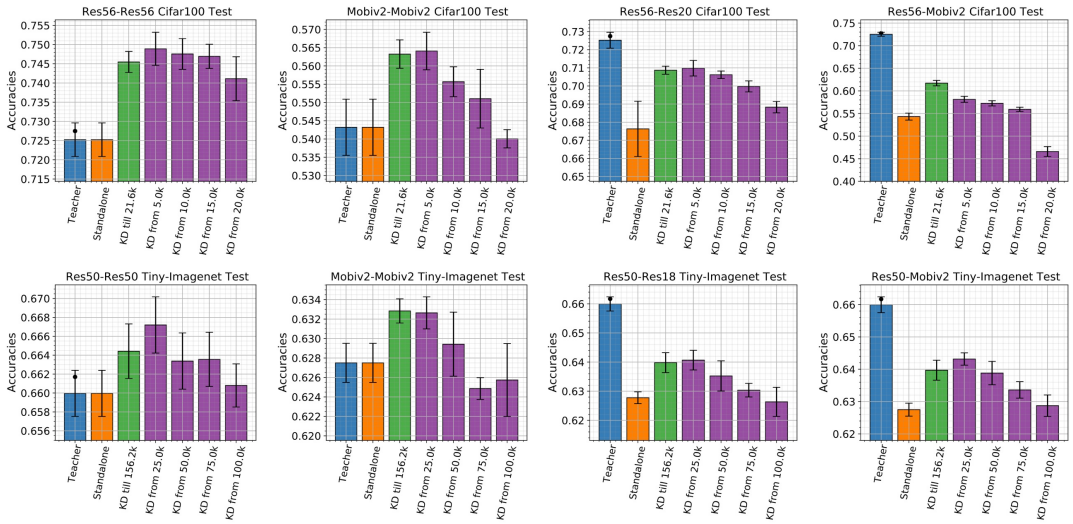

(a)

Figure 14: **Accuracies under late distillation for CIFAR100 and TinyImageNet:** We report the final accuracies for various distillation settings under ResNet56, ResNet20 and MobileNet-v2, with the loss switched to distillation in the middle of training. Observe that replacing the initial fourth/fifth of the training has zero effect on the distillation loss suggesting that deviations in the initial phase of training is not a crucial factor behind the success of distillation. Note that for MobileNet architectures, the loss-switch caused instability, which explains the strong dip in accuracy, even going below the one-hot model (top right).

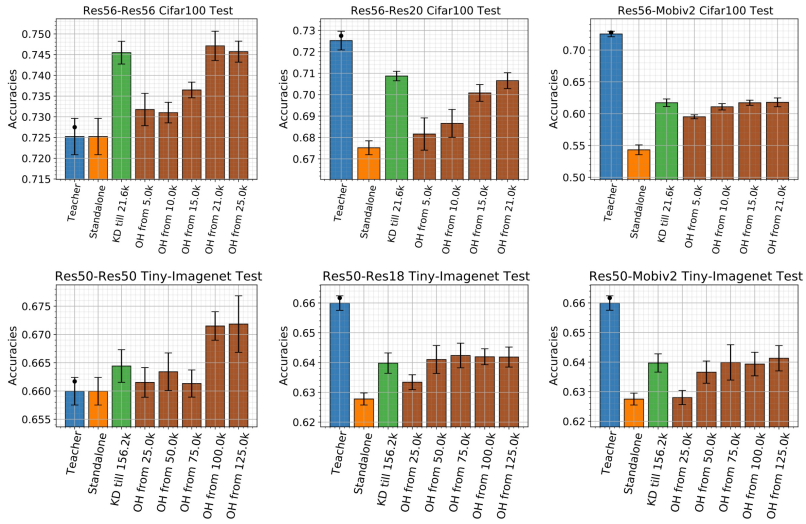

(a)

Figure 15: **Accuracies under late one-hot for CIFAR100 and TinyImageNet:** We report the final accuracies for various distillation-to-one-hot settings under ResNet56, ResNet20 and MobileNet-v2, with the loss switched to one-hot in the middle of training. Here, surprisingly, we find that in a majority of cases, switching to one-hot preserves the gains of distillation, and may even result an increase in accuracy, echoing the findings of (Zhang & Sabuncu, 2020; Yuan et al., 2020; Tang et al., 2020). However, in the ResNet/CIFAR100 experiments, we do find that one-hot training hurts the model.

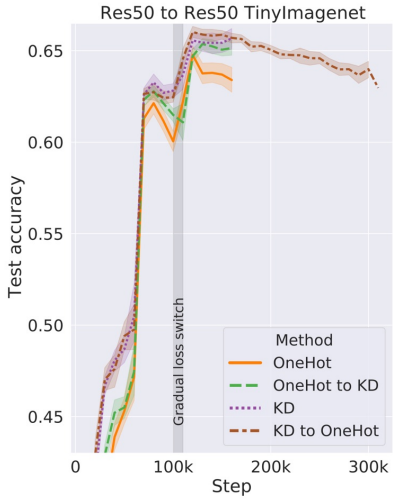

(a)

Figure 16: **Trajectory of test accuracy for loss-switching under ResNet50 self-distillation Tiny-ImageNet :** We gradually change the loss for a ResNet50 model between $100k - 110k$ steps. (Note that we report the test accuracy on a small random subset of the data.) We find that both switching to distillation and switching to one-hot result in an increase in accuracy. While the switch to one-hot attains the highest accuracy, if we were to train the model sufficiently long, the model does deterioriate, which we believe is due to the "destructive gradients" it incurs, as discussed in Section 4.2.