

# ON OPTIMAL EARLY STOPPING: OVERPARAMETRIZATION VERSUS UNDERPARAMETRIZATION

**Anonymous authors**

Paper under double-blind review

## ABSTRACT

Early stopping is a simple and widely used method to prevent over-training neural networks. We develop theoretical results to reveal the relationship between optimal early stopping time and model dimension as well as sample size of the dataset for certain linear regression models. Our results demonstrate two very different behaviors when the model dimension exceeds the number of features versus the opposite scenario. While most previous works on linear models focus on the latter setting, we observe that in common deep learning tasks, the dimension of the model often exceeds the number of features arising from data. We demonstrate experimentally that our theoretical results on optimal early stopping time corresponds to the training process of deep neural network. Moreover, we study the effect of early stopping on generalization and demonstrate that optimal early stopping can help mitigate “double descent” in various settings.

## 1 INTRODUCTION

Generalization, accuracy, and computation are three of the major aspects in deploying large scale machine learning models. They are often brought together into an optimization framework, where the first two aspects concern stationary solution and the last aspect concerns convergence properties of the algorithm. As a result, most previous works have analyzed them separately, seeking to strike a balance between generalization and accuracy first then understand computational complexity (Zhang, 2002). These approaches often design a regularizer in the form of a penalty term added to the objective function (c.f., Nakkiran et al., 2020b, and references therein). While this method is effective in many settings, it can prove challenging to determine the regularization that guarantees appealing behaviors. Adding to the complication is the fact that they are often designed to avoid overfitting when the optimization algorithm converges. It is not obvious how the behavior will change if the algorithm stops before the convergence of the optimization procedure.

Recent success of deep learning models combined with gradient based optimization has prompted increasingly many works to focus on blending the computation aspect into the picture of generalization–accuracy tradeoff (Du et al., 2018). Such approaches often leverage implicit properties of the optimization algorithms in conjunction of the model and the data structures and are thus referred to as implicit regularization methods. In training of deep neural networks, techniques such as stochastic gradient descent, batch normalization (Ioffe & Szegedy, 2015), and dropouts (Srivastava et al., 2014) are widely used to grant generalization properties implicitly.

In this paper, we study regularization via early stopping (Morgan & Bourlard, 1989; Zhang & Yu, 2005; Yao et al., 2007). While in practice, early stopping is often forced by computational budget constraints, we focus on its efficacy in avoiding overfitting towards the training dataset. In the presence of label noise, in particular, we study the common strategy of concluding the training process at the optimal early stopping time, which achieves the lowest test risk before it increases again. One crucial question is how this optimal early stopping time relates to the model size and sample size. Answering this question not only provides guidance for the model training process in practice, but also contributes to the understanding of the generalization property of different models.

**A Conundrum on Optimal Early Stopping** (Ali et al., 2019) studied the regularization effect of early stopping on least square regression, but did not provide explicit characterization of the optimal early stopping time with respect to the model and data complexity. As a first step, we

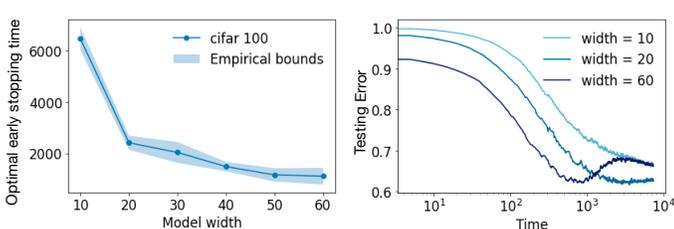


Figure 1: Left: Optimal early stopping time vs. model width  $d$  for CIFAR-100 with 20% label noise trained using ResNet networks with convolutional layers of widths  $[d, 2d, 4d, 8d]$ . Right: Testing error vs. time for CIFAR-100 with 20% label noise trained using ResNet networks using different model widths.

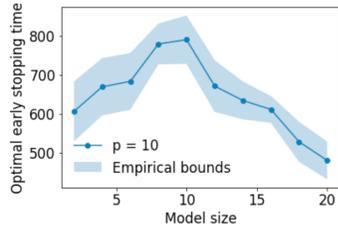


Figure 2: Optimal early stopping time vs. hidden layer width trained using two-layer ReLU networks on randomly generated points of  $p = 10$  clusters with 20% label noise.

compute the optimal early stopping time of the model studied in (Ali et al., 2019) and discover that under the setting adopted therein, the optimal early stopping time increases along with the model size. However, this trend is contrary to what we observed in training neural networks. We observed that when training popular datasets such as MNIST, CIFAR-10/100 using deep neural networks, the optimal early stopping time decreases as the model size increases (see Figure 1). These two contradicting trends indicate that previous theoretical models are not sufficient to explain what happens in practice.

**Overparametrization vs. Underparametrization** We noticed that despite the huge models used to train those datasets, their labels are actually generated by low dimensional representations or features (See Section 4.2). Previous theoretical works have primarily focused on the case where the model dimension is approximately the same or less than the number of features instead of the opposite scenario (Ali et al., 2019; Belkin et al., 2020; Nakkiran et al., 2020b). To complete the missing pieces in previous analyses, we propose a new model where the label is generated by a low-dimensional transformation of the input data. We define the case where the model size is larger than the number of features as being overparametrized. We verify experimentally that this overparametrized case better reflects what happens in training neural networks. This new overparametrization setting demonstrates very different behavior from the previous theoretical model, which resolves the contradiction we previously observed.

We illustrate both theoretically and experimentally that the optimal stopping time behaves very differently when the number of features is smaller (overparametrization) or larger (underparametrization) than the model size. In the overparametrization setting, the optimal early stopping time decreases as the model dimension increases or as the sample size decreases, similar to what we observe when training deep neural networks for image classification tasks. On the other hand, in the underparametrization setting, the optimal early stopping time increases as the model dimension increases or as the number of samples increases. We corroborate experimentally that these trends persist when training various linear models and neural networks (e.g., Figure 2).

**Generalization of the early stopped-models** Even for the same model, stopping at the optimal early stopping time versus after the training converges can lead to solutions with very different generalization properties. Another problem we study is how the test risk of models changes as a function of other training parameters when stopped at the optimal early stopping time. In particular, it has been shown that without regularization, the test risk can exhibit the so-called “double descent” phenomenon (Belkin et al., 2018; Nakkiran et al., 2020a), where the risk as a function of model size or sample size experience two distinct phases of descent. This phenomenon can make deciding the best model parameters challenging. We demonstrate that early stopping helps mitigate double descent in multiple settings.

### 1.1 RELATED WORKS

The test risk of simplified machine learning models has been studied in a long line of works. A partial list of papers that studied linear models similar to ours includes (Bartlett et al., 2020; Chen et al.,

2020a; Dobriban et al., 2018; Hastie et al., 2019; Montanari et al., 2019; Mitra, 2019; Muthukumar et al., 2020)

Regularization is widely used in training machine learning models to prevent overfitting. (Dobriban et al., 2018; Dobriban & Sheng, 2020; Hastie et al., 2019; Kobak et al., 2020; Mei & Montanari, 2019; Nakkiran et al., 2020b) studied the test risk of ridge regression. In addition, a lot of recent works studied the theoretical guarantee of implicit regularization in model training (Dereziński et al., 2019; Vaskevicius et al., 2019; HaoChen et al., 2020; Blanc et al., 2020; Razin & Cohen, 2020). In our work, we focus on regularization using early stopping. (Li et al., 2020) proved that early stopping is robust to label noise in certain settings. (Ali et al., 2019; 2020) related the risk of ridge regression to that of early stopping. (Vaškevičius et al., 2020) studied early stopped mirror descent. Our work characterizes the time that achieves the optimal stopping in various settings and demonstrates how optimal early stopping affects generalization.

The double descent phenomenon has a long history of being studied in the literature. (Loog et al., 2020) gave a brief history of the early works studying this topic. A partial list of works that tackled related problem in the setting of least square regression includes (Belkin et al., 2018; 2020; Bartlett et al., 2020; Deng et al., 2019; Hastie et al., 2019; Mei & Montanari, 2019; Muthukumar et al., 2020; d’Ascoli et al., 2020; Chen et al., 2020a; Yang et al., 2020). (Nakkiran et al., 2020a) demonstrated double descent can also occur as a function of number of samples. (Nakkiran et al., 2020b) shows that optimal regularization can mitigate double descent. Our work differs from previous works by studying the effect of early stopping on double descent.

Finally, our work studies the low rank structure of the feature representations. (Li et al., 2018; Dusenberry et al., 2020) have also empirically explored and leveraged this property.

## 2 OVERPARAMETRIZATION

In this section, we introduce and analyze the overparametrization setting, where the model size exceeds the number of features. This setting completes the missing piece of previous theoretical analysis on large models with excessive parameters. For data  $(x, y)$ , the input  $x \in \mathbb{R}^d$  is from a high dimensional space. The label  $y \in \mathbb{R}$  is generated from a low-dimensional feature representation  $z \in \mathbb{R}^p$ ,  $p \leq d$ , which is mapped from  $x$  by a low-dimensional transformation. A lot of common deep learning applications falls under this setting. Despite the huge model size used, the underlying feature space dimension can be low. For example, we observe that the labels  $y$  in CIFAR10, a widely used image classification problem with 10 classes, can be generated by a feature space of dimension a little more than 10 (See Section 4.2).

Formally, we consider the following linear regression problem with  $d > p$ . Let the covariate/input matrix  $X \in \mathbb{R}^{n \times d}$  be a random matrix such that each entry of it is generated from  $\mathcal{N}(0, 1)$  independently. Let  $x_1, \dots, x_n$  be the row vectors of the matrix  $X$ . For a semi-orthogonal matrix  $P \in \mathbb{R}^{p \times d}$  with  $p < d$  and some unknown parameter  $\theta^*$ , the response is generated by

$$y_i = \langle Px_i, \theta^* \rangle + \epsilon_i, \tag{1}$$

where the label noise  $\epsilon_i \sim \mathcal{N}(0, \sigma^2)$  is independent of  $x_i$  and is generated independently for each  $(x_i, y_i)$ . Let  $\epsilon \in \mathbb{R}^n$  be the vector with entries given by  $\epsilon_i$ .

We let  $\mathcal{S}_{\text{over}}$  denote the overparametrization setting studied in this section. This setting is inspired by the overparametrization property of many real datasets and models. For example, to classify images using deep neural network, the labels can usually be determined by features extracted from certain important parts of the images. The other parts of the image, though do not contribute to the label much, are still input to the model. When the neural network is wide, the network has the ability to extract far more features than what contribute to the label. This character of the deep neural network is what the setting  $\mathcal{S}_{\text{over}}$  attempts to model.

Our data assumption has been studied in many related literature (Belkin et al., 2020; Nakkiran et al., 2020b). Our analysis only relies on the concentration property of singular values of the data matrix  $X$ , so it is possible relax the data assumptions on  $X$  to bounds on the minimum and maximum singular values of the data matrix (Vaswani & Nayer, 2016). The linear model has been widely used as the first step to study more complicated models as well due to its simple form (Bartlett et al., 2020; Chen et al., 2020a; Dobriban et al., 2018; Hastie et al., 2019; Montanari et al., 2019; Mitra,

2019; Muthukumar et al., 2020). It is possible to extend our results to non-linear settings such as the random feature setting. In this paper, we stick to the simplest model so that we can get a simple and explicit characterization of the optimal early stopping time.

For fixed semi-orthogonal matrix  $P$  and parameter  $\theta^*$ , let  $\mathcal{D}_{P,\theta^*}$  be the joint distribution of a single data point  $(x, y)$  given  $P$  and  $\theta^*$ . For an estimator  $\beta$ , define the population risk

$$R(\beta) = \mathbb{E}_{(\tilde{x}, \tilde{y}) \sim \mathcal{D}_{P,\theta^*}} \left[ (\langle \tilde{x}, \beta \rangle - \tilde{y})^2 \right].$$

We consider the standard linear least squares problem

$$\min_{\beta \in \mathbb{R}^d} \frac{1}{2n} \|y - X\beta\|_2^2. \quad (2)$$

Applying gradient flow starting from 0 on equation 2 gives a continuous differential equation over time  $t$ ,

$$\frac{d}{dt} \beta_{X,y}(t) = \frac{X^\top}{n} (y - X\beta_{X,y}(t)). \quad (3)$$

with an initial condition  $\beta_{X,y}(0) = 0$ . We can solve the differential equation exactly and obtain

$$\beta(t) = (X^\top X)^+ (I - \exp(-tX^\top X/n)) X^\top y, \quad (4)$$

for all  $t \geq 0$  (Ali et al., 2019). Here, for any matrix  $A$ ,  $\exp(A)$  is defined as  $\exp(A) := \sum_{n=0}^{\infty} \frac{A^n}{n!}$ . In practice, one would discretize the gradient flow and perform gradient descent. We discuss the error due to such discretization in Appendix E.

For any fixed dataset  $y \in \mathbb{R}^n$  and  $X \in \mathbb{R}^{n \times d}$  given by fixed  $P$  and  $\epsilon$ , let  $\beta_{X,y}(t)$  be the gradient flow solution of standard linear regression problem at time  $t$ . We further assume that  $P \in \mathbb{R}^{p \times d}$  is a uniformly random semi-orthogonal matrix.

Our goal is to study the expected risk of the estimator  $\beta_{X,y}(t)$  over  $P$  and  $\epsilon$ . We denote the expected risk at time  $t$  for dataset size  $n$  as

$$\bar{R}_{X,\theta^*}(t) = \mathbb{E}_{P,\epsilon} [R(\beta_{X,y}(t))].$$

We study the optimal early stopping time, the time  $t$  that achieves the first local minimum of the expected risk. We characterize the optimal early stopping time for different parameter choices. We let  $t_{\text{opt}}$  denote the optimal stopping time, omitting  $(X, \theta^*)$  when clear from context,

$$t_{\text{opt}}(X, \theta^*) = \min t \text{ s.t. } \frac{d}{dt} \bar{R}_{X,\theta^*}(t') < 0 \text{ for } 0 < t' < t, \text{ and } \frac{d}{dt} \bar{R}_{X,\theta^*}(t) = 0. \quad (5)$$

## 2.1 OPTIMAL EARLY STOPPING TIME

With the gradient flow solution  $\beta_{X,y}$ , we can derive a high probability upper and lower bound on  $t_{\text{opt}}$ .

**Theorem 1** (Optimal early stopping time in overparametrization setting). *In setting  $\mathcal{S}_{\text{over}}$ , for a fixed parameter  $\theta^*$  and noise variance  $\sigma^2$ , when  $n \leq d$ , let  $\gamma = \left( \frac{\sqrt{n} + \sqrt{2 \log n}}{\sqrt{d}} \right)^2$ . For  $\gamma \leq 1$ , with probability at least  $1 - \frac{2}{n}$  over the randomness of  $X$ , the optimal early stopping time  $t_{\text{opt}}$  satisfies*

$$\frac{n}{(1 + \sqrt{\gamma})^2 d} \log \left( 1 + \frac{(1 - \sqrt{\gamma})^2 \|\theta^*\|_2^2}{\sigma^2} \right) \leq t_{\text{opt}} \leq \frac{n}{(1 - \sqrt{\gamma})^2 d} \log \left( 1 + \frac{(1 + \sqrt{\gamma})^2 \|\theta^*\|_2^2}{\sigma^2} \right).$$

*When  $n > d$ , let  $\gamma = \left( \frac{\sqrt{n}}{\sqrt{d} + \sqrt{2 \log d}} \right)^2$ . For  $\gamma \geq 1$ , with probability at least  $1 - \frac{2}{d}$  over the randomness of  $X$ , the optimal early stopping time  $t_{\text{opt}}$  satisfies*

$$\frac{1}{\left(1 + \frac{1}{\sqrt{\gamma}}\right)^2} \log \left( 1 + \left(1 - \frac{1}{\sqrt{\gamma}}\right)^2 \frac{n \|\theta^*\|_2^2}{d\sigma^2} \right) \leq t_{\text{opt}} \leq \frac{1}{\left(1 - \frac{1}{\sqrt{\gamma}}\right)^2} \log \left( 1 + \left(1 + \frac{1}{\sqrt{\gamma}}\right)^2 \frac{n \|\theta^*\|_2^2}{d\sigma^2} \right).$$

Note that our definitions of  $\gamma$  differ slightly in the case  $n \leq d$  and  $n > d$ , but they are both of order  $\tilde{\Theta}\left(\frac{n}{d}\right)$ . Theorem 1 gives an approximation of the optimal stopping time up to a constant as long as  $\gamma$  is constantly bounded away from 1. When  $\gamma \leq \frac{1}{4}$  (roughly  $n \leq 4d$ ), omitting the logarithmic terms,  $\frac{9n}{4d} \leq t_{\text{opt}} \leq \frac{2n}{d}$ . When  $\gamma \geq 4$  (roughly  $n \geq 4d$ ), **omitting the logarithmic terms**,  $\frac{n}{3(d+n)} \log\left(1 + \frac{n\|\theta^*\|}{4d\sigma^2}\right) \leq t_{\text{opt}} \leq \frac{5n}{d+n} \log\left(1 + 3\frac{n\|\theta^*\|}{d\sigma^2}\right)$ . Combining these two cases shows that when  $\gamma$  is bounded away from 1,  $t_{\text{opt}}$  is about  $\Theta\left(\frac{n}{d+n} \log \frac{n}{d}\right)$ . When  $\gamma$  is very close to 1, the upper bound on  $t_{\text{opt}}$  in Theorem 1 can be loose. We show in Appendix C, that in the asymptotic case where  $n$  and  $d$  are large and  $\frac{\|\theta^*\|}{\sigma^2}$  is lower bounded, the approximation  $t_{\text{opt}} = \tilde{\Theta}\left(\frac{n}{d+n}\right)$  still holds when  $\frac{1}{4} < \gamma < 4$ . In the asymptotic analysis we also achieve tighter constants in the bounds of the optimal early stopping time. Furthermore, in Appendix F, using numerical experiments, we show that our theoretical result can be a good approximation of the optimal stopping time for all  $\gamma$ .

Now, we discuss some implications of Theorem 1. The following lemma computes  $\bar{R}_{X,\theta^*}(t_{\text{opt}})$ , the expected risk at time  $t_{\text{opt}}$ , and holds for all input matrix  $X$  regardless of how it is generated.

**Lemma 1.** *For all  $n, p, d \in \mathbb{N}$  such that  $p \leq d$ , let  $X \in \mathbb{R}^{n \times d}$  and  $\theta^* \in \mathbb{R}^p$  be fixed. Let  $\lambda_1 \geq \dots \geq \lambda_d$  be the eigenvalues of the matrix  $\frac{1}{n}X^\top X$ . Then,*

$$\bar{R}_{X,\theta^*}(t) = \sigma^2 + \sum_{i=1}^d \exp(-2t\lambda_i) \frac{\|\theta^*\|_2^2}{d} + \frac{\sigma^2}{n} \sum_{i=1}^d \mathbb{1}\{\lambda_i \neq 0\} \frac{1}{\lambda_i} (1 - \exp(-t\lambda_i))^2. \quad (6)$$

In equation 6, the second term is the approximation error due to early stopping. The third term is due to overfitting to the label noise, which increases as the training time increases. In presence of label noise, the optimal early stopping finds the balance point of these two terms.

Theorem 1 shows that the optimal early stopping time is roughly  $\tilde{\Theta}\left(\frac{n}{n+d}\right)$ , **when  $n$  and  $d$  are around the same order. When  $n \gg d$ , the optimal early stopping time is roughly  $\Theta(\log \frac{n}{d})$ .** It implies that **when the number of features is smaller than the model dimension, optimal early stopping time increases as  $n$  increases or  $d$  decreases.** Section F.1 corroborates this trend using experiments. In addition, we show in Section 4.2 and Section 4.3 that training deep neural networks on large real datasets can also follow these two trends. An intuitive explanation of such phenomenon is when the sample size is large, the label noise has less effect on model training, so the model can be trained longer before overfitting starts. On the contrary, when  $d$  increases, the number of parameters is large so that even small updates of each parameter can result in overfitting.

## 2.2 RISK MONOTONICITY

It has been observed empirically that without early stopping or other types of regularization, when the number of samples  $n$  increases, the test error can experience double descent in presence of label noise. However, optimal early stopping can possibly mitigate the double descent phenomenon (Nakkiran et al., 2020a). In Proposition 1, we bound the expected risk at the optimal stopping time for varying data size  $n$  and attempt to explain this phenomenon.

**Proposition 1.** *In setting  $\mathcal{S}_{\text{over}}$ , for a fixed parameter  $\theta^*$  and noise variance  $\sigma^2$ . Let  $\bar{t}_{\text{opt}} = \alpha n / (n + d)$  for some constant  $\alpha$ . Let  $\lambda_1 \geq \dots \geq \lambda_d$  be the eigenvalues of the matrix  $\frac{1}{n}X^\top X$ . Then,*

$$\bar{R}_1 \leq \mathbb{E}_X[\bar{R}_{X,\theta^*}(\bar{t}_{\text{opt}})] \leq \bar{R}_2. \quad (7)$$

where

$$\bar{R}_1 = \sigma^2 + \mathbb{E}_X \left[ \sum_{i=1}^d \exp(-2\alpha n \lambda_i / d) \right] \frac{\|\theta^*\|_2^2}{d}, \quad \text{and} \quad \bar{R}_2 = \bar{R}_1 + \frac{1}{2} \alpha \sigma^2.$$

$\bar{R}_1$  and  $\bar{R}_2$  decrease monotonically as  $n$  increases.

Theorem 1 shows that the optimal early stopping time is roughly  $t_{\text{opt}} \approx \frac{\alpha n}{n+d}$  for a small constant  $\alpha$  up the logarithmic factors. When  $\frac{\|\theta^*\|_2^2}{\sigma^2}$  is not too small, equation 7 gives a small region that bounds the expected risk at  $\bar{t}_{\text{opt}}$ . We are able to show that both the upper and the lower bounds of the

region decrease as  $n$  increases. This implies that **optimal early stopping can mitigate sample-wise double descent**. We use experiments to corroborate our computation and give more details on this observation in Section 4 and Appendix F.

A related question is whether stopping at optimal early stopping time can mitigate the double descent due to increasing model size  $d$ . (Nakkiran et al., 2020a) observed that double descent can still exist at optimal early stopping. An theoretical analysis of this question can be an interesting future direction.

### 3 UNDERPARAMETRIZATION

In this section, we study the underparametrization setting, where the number of features exceeds the model size. This setting shows very different behavior from the overparametrization setting. In this setting, the label  $y$  is generated by some covariate  $z$ , which is in a high-dimensional space, but the model only has access to a low-dimension projection of  $z$ .

Formally, we consider the following linear regression problem with  $d \leq p$ . Let the covariate matrix  $Z \in \mathbb{R}^{n \times p}$  be a random matrix such that each entry of it is generated from  $\mathcal{N}(0, 1)$  independently. Let  $z_1, \dots, z_n$  be the row vectors of the matrix  $Z$ . For some unknown parameter  $\theta^*$ , the response is generated by

$$y_i = \langle z_i, \theta^* \rangle + \epsilon_i. \quad (8)$$

where the label noise  $\epsilon_i \sim \mathcal{N}(0, \sigma^2)$  is independent of  $z_i$  and  $y_i$ . For a semi-orthogonal matrix  $P \in \mathbb{R}^{p \times d}$  with  $d \leq p$ , let  $X = ZP$ . we apply gradient flow on  $(X, y)$ . For fixed semi-orthogonal matrix  $P$  and parameter  $\theta^*$ , let  $\mathcal{D}_{P, \theta^*}$  be the joint distribution of a single data point  $(x, y)$ . For an estimator  $\beta$ , define the population risk

$$R(\beta) = \mathbb{E}_{(\tilde{x}, \tilde{y}) \sim \mathcal{D}_{P, \theta^*}} \left[ (\langle \tilde{x}, \beta \rangle - \tilde{y})^2 \right].$$

Similar as in the overparametrizations setting, we consider the gradient flow solution (equation 4) of the standard linear least square problem (equation 2). For any fixed dataset  $y \in \mathbb{R}^n$  and  $X \in \mathbb{R}^{n \times d}$  given by fixed  $Z, P$  and  $\epsilon$ , let  $\beta_{X, y}(t)$  be the solution of the gradient flow given in equation 4. We further assume that  $P \in \mathbb{R}^{p \times d}$  is a uniformly random semi-orthogonal matrix.

We let  $\mathcal{S}_{\text{under}}$  denote the above setting.  $\mathcal{S}_{\text{under}}$  studies the case where we only have partial access to the features that determine the label  $y$  and train the model on what is available.

Our goal is to study the expected risk of the estimator  $\beta_{X, y}(t)$  over  $Z, P$  and  $\epsilon$ . We denote the expected risk over  $Z, P$  and  $\epsilon$  at time  $t$  as

$$\bar{R}_{\theta^*}(t) = \mathbb{E}_{Z, P, \epsilon} [R(\beta_{X, y}(t))].$$

Note that here in  $\mathcal{S}_{\text{under}}$ , we take expectation on  $Z$  when computing the expected risk. In  $\mathcal{S}_{\text{over}}$ , we do not take expectation on  $X$  and instead get a high probability bound on  $X$ . This slight difference is due to the different methods we use in analyzing these two settings. We study the optimal early stopping time, the time  $t$  that achieves the first local minimum of the expected risk. We characterize the optimal early stopping time for different choices of  $p, n$  and  $d$ . We let  $t_{\text{opt}}$  denote the optimal stopping time, omitting  $\theta^*$  when clear from context,

$$t_{\text{opt}}(\theta^*) = \min t \text{ s.t. } \frac{d}{dt} \bar{R}_{\theta^*}(t') < 0 \text{ for } 0 < t' < t, \text{ and } \frac{d}{dt} \bar{R}_{\theta^*}(t) = 0. \quad (9)$$

We are able to give a linear approximation of the optimal stopping time,  $t_{\text{opt}}$ .

**Theorem 2** (Optimal early stopping time in underparametrization setting). *In setting  $\mathcal{S}_{\text{under}}$ , for a fixed parameter  $\theta^*$  and noise variance  $\sigma^2$ , if  $(8n + 9d + 16) \|\theta^*\|_2^2 \leq p\sigma^2 + p \|\theta^*\|_2^2$ , the optimal early stopping time  $t_{\text{opt}}$  satisfies*

$$\frac{n \|\theta^*\|_2^2}{2(p\sigma^2 + (p-d) \|\theta^*\|_2^2)} \leq t_{\text{opt}} \leq \frac{2n \|\theta^*\|_2^2}{p\sigma^2 + (p-d) \|\theta^*\|_2^2}.$$

Theorem 2 approximates the early optimal stopping  $t_{\text{opt}}$  up to a constant when  $p$  is larger than  $n$  and  $d$  by a constant factor. For  $d$  that is close to  $p$ , numerical experiments show that the approximation

$\tilde{\Theta} \left( \frac{n \|\theta^*\|_2^2}{p\sigma^2 + (p-d)\|\theta^*\|_2^2} \right)$  still holds. Theorem 2 implies that **when the number of features is larger than the model dimension, optimal early stopping time increases as  $n$  or  $d$  increases**. An intuitive explanation is when the model size and the sample size are small, the model does not have enough information to train so that the model tends to fit the noise easier.

**Proposition 2.** *In setting  $\mathcal{S}_{\text{under}}$ , for a fixed parameter  $\theta^*$  and noise variance  $\sigma^2$ , let  $\bar{t}_{\text{opt}} = \frac{n \|\theta^*\|_2^2}{p\sigma^2 + (p-d)\|\theta^*\|_2^2}$ . Then, when  $(8n + 9d + 16) \|\theta^*\|_2^2 \leq p\sigma^2 + p \|\theta^*\|_2^2$ , we have*

$$\bar{R}_1 \leq \bar{R}_{\theta^*}(\bar{t}_{\text{opt}}) \leq \bar{R}_2,$$

where  $\bar{R}_1 = \sigma^2 + \|\theta^*\|_2^2 - \frac{2nd\|\theta^*\|_2^4}{p^2\sigma^2 + p(p-d)\|\theta^*\|_2^2}$  and  $\bar{R}_2 = \sigma^2 + \|\theta^*\|_2^2 - \frac{3nd\|\theta^*\|_2^4}{4(p^2\sigma^2 + p(p-d)\|\theta^*\|_2^2)}$ .

Here,  $\bar{R}_1 \geq 0.8\bar{R}_2$  and both  $\bar{R}_1$  and  $\bar{R}_2$  decrease when  $d$  and  $n$  increase.

Proposition 2 gives an upper bound and a lower bound on the risk at the approximated optimal early stopping time. Under the assumption  $d$  and  $n$  are small,  $\bar{R}_1$  and  $\bar{R}_2$  give a tight region ( $\bar{R}_1 \geq 0.8\bar{R}_2$ ). Proposition 2 supports the observation that **risk at optimal early stopping time monotonically decreases as  $n$  and  $d$  increases**. Numerical experiments are presented in Section F.1.

## 4 EXPERIMENT

In this section, we run simulations for settings  $\mathcal{S}_{\text{over}}$  and  $\mathcal{S}_{\text{under}}$  and demonstrate that training deep neural networks also follows the trends our theoretical analysis identifies. We provide additional experimental results in Appendix F.

### 4.1 LINEAR REGRESSION

(We moved the experiments results on linear regression to the appendix so we can have more space to address the comments from the reviewers.)

### 4.2 DEEP LEARNING WITH NEURAL NETWORKS

We study the optimal early stopping time in neural network training. We demonstrate that datasets MNIST, CIFAR-10 and CIFAR-100 have low-rank representations. Such representations show that training common deep neural networks on these datasets falls under the overparametrization setting. Then, we demonstrate experimentally that the optimal early stopping time follows the pattern we identify for the overparametrization setting.

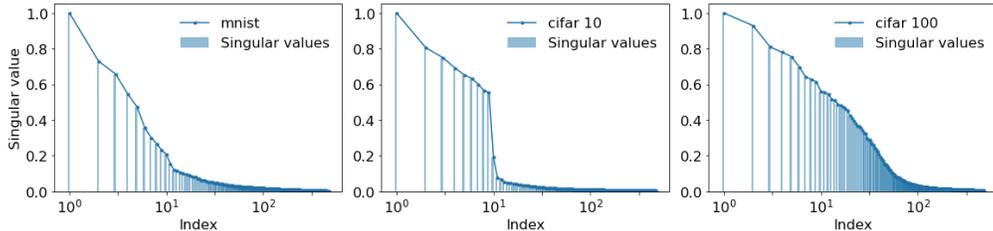


Figure 3: Scaled singular values of representations obtained from MNIST, CIFAR-10 and CIFAR-100.

**Existence of low-rank structure in image classification.** Many previous works observe that low-rank structures exist and can be leveraged in deep image classification tasks (Van der Maaten & Hinton, 2008; Dusenberry et al., 2020). This structure facilitates the self-supervised learning and transfer learning paradigm where image features are learned and fixed close to the input layer, reducing the complexity of image-to-label mapping. Then only the last layer is trained to adapt to and classify out-of-domain images (Chen et al., 2020b). We follow this approach and demonstrate the existence of the low-rank structures in Figure 3, where we plot the singular values of the representations of MNIST, CIFAR-10, and CIFAR-100. First, we train CNN/ResNet-60 (He et al., 2016a;b) to convergence with zero training loss (no label noise). Then, we obtain the representations produced by the trained neural networks on MNIST, CIFAR-10, and CIFAR-100. Even when trained with a very wide neural network, only the first 10/100 singular values of the representations are significant

compared to the others. This suggests that low-rank feature representations of the datasets exist. We believe this phenomenon may be of independent interest to understanding neural networks.

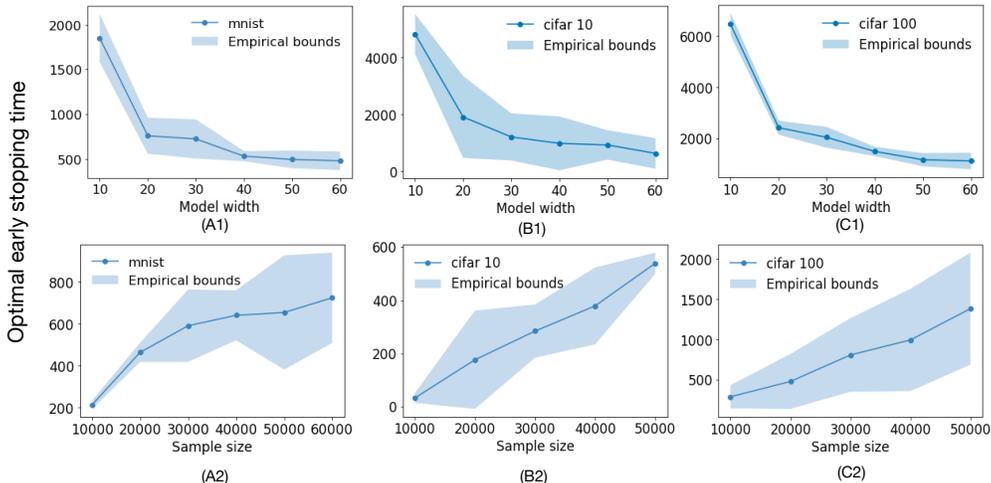


Figure 4: In (A1) (B1) (C1), optimal early stopping time decreases with increasing model sizes for MNIST, CIFAR-10, and CIFAR-100. In (A2) (B2) (C2), optimal early stopping time increases with increasing sample sizes for MNIST, CIFAR-10, and CIFAR-100.

**Overparameterization behavior in image classification using neural networks** In this section, we study the optimal early stopping time of MNIST, CIFAR-10, and CIFAR-100. We train MNIST using a fully connected network, and CIFAR-10/100 using ResNet. Part of our code in this part builds on open source code from (Nakkiran et al., 2020a). We add 20% of label noise to all datasets. For MNIST, a simple ReLU network with one hidden layer is used. We vary the hidden layer width over  $d = [10, 20, \dots, 60]$ . For CIFAR-10/100, we use a family of ResNet networks with convolutional layers of widths  $[d, 2d, 4d, 8d]$  for different layer depths. The sample size ranges from  $n = [1e^5, 2e^5, \dots, 6e^5]$  for MNIST, and  $n = [1e^5, 2e^5, \dots, 5e^5]$  for CIFAR-10/100. Inputs are normalized to  $[-1, 1]$  with standard data-augmentation. We use stochastic gradient descent with cross-entropy loss, learning rate  $\eta = 0.1$  for MNIST, CIFAR-10, and  $\eta = 0.05$  for CIFAR-100, and minibatch size  $B = 512$ . The optimal stopping time is computed as  $t_{opt} = \eta T \frac{n}{B}$ , where  $T$  is the epoch with the optimal early stopping performance. To select the optimal stopping time, we first apply a moving averaging window with length 5 to mitigate the effect of mini-batch noise. After smoothing, we select the point with the optimal stopping time with the first minimal testing error (the lowest point of the first ‘valley’, ignoring small local minima).

We show in Figure 4 the optimal early stopping time for all three datasets with varying  $n$  and  $d$ . In (A1) (A2) (A3), we observe that the optimal early stopping time decreases with increasing model width. In (B1) (B2) (B3), we observe that the optimal early stopping time increases with larger sample sizes. Due to the existence of low-rank representations, Figure 4 falls under the overparameterization setting even for small  $d$ . The observations are consistent with our theoretical analysis and the experimental results on linear models in the overparameterization setting. When trained with ultra small  $d < 10$ , the models show great instability. We provide additional experimental results supporting the underparameterization setting in Section 4.3.

### 4.3 CLASSIFICATION USING TWO-LAYER NEURAL NETWORKS

Training datasets like CIFAR-10/100 using ultra small models is unstable, so in this section, we use simpler datasets to verify our theoretical results for the underparameterization setting. In this experiment, we use `make_classification` function in `sklearn` (Pedregosa et al., 2011) to generate clusters of points normally distributed around vertices of a high-dimensional hypercube and assigns an equal number of clusters to each class. We generate 2000 samples (50-50 train-test split) of 20 dimensions/features, but only  $p = 10$  dimensions/features of each sample are informative features. We add 20% label noise to the dataset. The optimal stopping time is selected similarly as in Section 4.2.

We train the data using two-layer ReLU networks with varying hidden layer width  $d$ . The blue area in Figure 5 gives 0.5 standard deviation of 6 runs and the blue line gives the mean. In Figure 5a, when  $d \leq p = 10$ , we observe the underparametrization behavior where the optimal early stopping time increases as the model size  $d$  increases. When  $d \geq 10 = p$ , the optimal stopping time decreases with  $d$ . We can further observe that the test loss decreases with increasing model sizes at the optimal early stopping time.

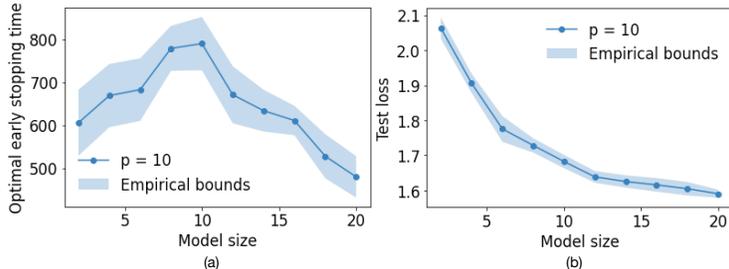


Figure 5: (a) Optimal stopping time as a function of model size. (b) Test loss as a function of model size.

## 5 CONCLUSION AND FUTURE WORK

In this paper, we propose a new model of overparametrization to study the optimal early stopping time. This model captures the phenomenon that the model size usually exceeds the number of features in practice. We believe this model of overparametrization can be of independent interest. One future direction is the performance of other optimizers or techniques, e.g., stochastic gradient descent or data augmentation, on this model. Moreover, we give an explicit characterization of the early stopping time and the resulting model. Future works can explore other regularization properties of this early stopped model, e.g., whether this early stopping helps when there is a domain shift in the testing dataset.

Another direction worth exploring is to extend our results to nonlinear settings. Extending our result to the lazy training regime of neural networks (e.g., random feature regime and neural tangent kernel regime) can be relatively straightforward. One approach is to first locate the features extracted at initialization and then show how gradient flow progresses on the features extracted. Theoretical analysis of the non-lazy training is still a largely open question, so extending this result to non-lazy setting can be more challenging. To extend this result to the non-lazy setting for a fully nonlinear neural network model, one needs to study how the features extracted change and how the weights on the features change at the same time.

## REFERENCES

- Alnur Ali, J Zico Kolter, and Ryan J Tibshirani. A continuous-time view of early stopping for least squares regression. In *The 22nd International Conference on Artificial Intelligence and Statistics* , pp. 1370–1378. PMLR, 2019.
- Alnur Ali, Edgar Dobriban, and Ryan Tibshirani. The implicit regularization of stochastic gradient flow for least squares. In *Proceedings of the 37th International Conference on Machine Learning* , pp. 233–244. PMLR, 2020.
- Peter L Bartlett, Philip M Long, Gábor Lugosi, and Alexander Tsigler. Benign overfitting in linear regression. *Proceedings of the National Academy of Sciences* , 117(48):30063–30070, 2020.
- Mikhail Belkin, Daniel Hsu, Siyuan Ma, and Soumik Mandal. Reconciling modern machine learning practice and the bias-variance trade-off. *arXiv preprint arXiv:1812.11118* , 2018.
- Mikhail Belkin, Daniel Hsu, and Ji Xu. Two models of double descent for weak features. *SIAM Journal on Mathematics of Data Science* , 2(4):1167–1180, 2020.
- Guy Blanc, Neha Gupta, Gregory Valiant, and Paul Valiant. Implicit regularization for deep neural networks driven by an ornstein-uhlenbeck like process. In *Conference on learning theory* , pp. 483–513. PMLR, 2020.
- Lin Chen, Yifei Min, Mikhail Belkin, and Amin Karbasi. Multiple descent: Design your own generalization curve. *arXiv preprint arXiv:2008.01036* , 2020a.
- Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *International conference on machine learning* , pp. 1597–1607. PMLR, 2020b.
- Stéphane d’Ascoli, Levent Sagun, and Giulio Biroli. Triple descent and the two kinds of overfitting: Where & why do they appear? *arXiv preprint arXiv:2006.03509* , 2020.
- Kenneth R Davidson and Stanislaw J Szarek. Local operator theory, random matrices and banach spaces. *Handbook of the geometry of Banach spaces* , 1(317-366):131, 2001.
- Zeyu Deng, Abla Kammoun, and Christos Thrampoulidis. A model of double descent for high-dimensional binary linear classification. *arXiv preprint arXiv:1911.05822* , 2019.
- Michał Dereziński, Feynman Liang, and Michael W Mahoney. Exact expressions for double descent and implicit regularization via surrogate random design. *arXiv preprint arXiv:1912.04533* , 2019.
- Edgar Dobriban and Yue Sheng. Wonder: Weighted one-shot distributed ridge regression in high dimensions. *Journal of Machine Learning Research* , 21(66):1–52, 2020.
- Edgar Dobriban, Stefan Wager, et al. High-dimensional asymptotics of prediction: Ridge regression and classification. *The Annals of Statistics* , 46(1):247–279, 2018.
- Simon S Du, Wei Hu, and Jason D Lee. Algorithmic regularization in learning deep homogeneous models: Layers are automatically balanced. *arXiv preprint arXiv:1806.00900* , 2018.
- Michael Dusenberry, Ghassen Jerfel, Yeming Wen, Yian Ma, Jasper Snoek, Katherine Heller, Balaji Lakshminarayanan, and Dustin Tran. Efficient and scalable bayesian neural nets with rank-1 factors. In *International conference on machine learning* , pp. 2782–2792. PMLR, 2020.
- Jeff Z HaoChen, Colin Wei, Jason D Lee, and Tengyu Ma. Shape matters: Understanding the implicit bias of the noise covariance. *arXiv preprint arXiv:2006.08680* , 2020.
- Trevor Hastie, Andrea Montanari, Saharon Rosset, and Ryan J Tibshirani. Surprises in high-dimensional ridgeless least squares interpolation. *arXiv preprint arXiv:1903.08560* , 2019.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition* , pp. 770–778, 2016a.

- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Identity mappings in deep residual networks. In *European conference on computer vision* , pp. 630–645. Springer, 2016b.
- Roger A Horn and Charles R Johnson. *Matrix analysis* . Cambridge university press, 2012.
- Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *International conference on machine learning* , pp. 448–456. PMLR, 2015.
- Dmitry Kobak, Jonathan Lomond, and Benoit Sanchez. The optimal ridge penalty for real-world high-dimensional data can be zero or negative due to the implicit ridge regularization. *Journal of Machine Learning Research* , 21(169):1–16, 2020.
- Chunyuan Li, Heerad Farkhoor, Rosanne Liu, and Jason Yosinski. Measuring the intrinsic dimension of objective landscapes. *arXiv preprint arXiv:1804.08838* , 2018.
- Mingchen Li, Mahdi Soltanolkotabi, and Samet Oymak. Gradient descent with early stopping is provably robust to label noise for overparameterized neural networks. In *International Conference on Artificial Intelligence and Statistics* , pp. 4313–4324. PMLR, 2020.
- Marco Loog, Tom Viering, Alexander Mey, Jesse H Krijthe, and David MJ Tax. A brief prehistory of double descent. *Proceedings of the National Academy of Sciences* , 117(20):10625–10626, 2020.
- Song Mei and Andrea Montanari. The generalization error of random features regression: Precise asymptotics and double descent curve. *arXiv preprint arXiv:1908.05355* , 2019.
- Partha P Mitra. Understanding overfitting peaks in generalization error: Analytical risk curves for  $l_2$  and  $l_1$  penalized interpolation. *arXiv preprint arXiv:1906.03667* , 2019.
- Andrea Montanari, Feng Ruan, Youngtak Sohn, and Jun Yan. The generalization error of max-margin linear classifiers: High-dimensional asymptotics in the overparametrized regime. *arXiv preprint arXiv:1911.01544* , 2019.
- Nelson Morgan and Hervé Bouchard. Generalization and parameter estimation in feedforward nets: Some experiments. *Advances in neural information processing systems* , 2:630–637, 1989.
- Vidya Muthukumar, Kailas Vodrahalli, Vignesh Subramanian, and Anant Sahai. Harmless interpolation of noisy data in regression. *IEEE Journal on Selected Areas in Information Theory* , 1(1): 67–83, 2020.
- Preetum Nakkiran, Gal Kaplun, Yamini Bansal, Tristan Yang, Boaz Barak, and Ilya Sutskever. Deep double descent: Where bigger models and more data hurt. In *International Conference on Learning Representations* , 2020a.
- Preetum Nakkiran, Prayaag Venkat, Sham Kakade, and Tengyu Ma. Optimal regularization can mitigate double descent. *arXiv preprint arXiv:2003.01897* , 2020b.
- F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research* , 12:2825–2830, 2011.
- Noam Razin and Nadav Cohen. Implicit regularization in deep learning may not be explainable by norms. *arXiv preprint arXiv:2005.06398* , 2020.
- Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: a simple way to prevent neural networks from overfitting. *The journal of machine learning research* , 15(1):1929–1958, 2014.
- Laurens Van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of machine learning research* , 9(11), 2008.
- Tomas Vaskevicius, Varun Kanade, and Patrick Rebeschini. Implicit regularization for optimal sparse recovery. *Advances in Neural Information Processing Systems* , 32:2972–2983, 2019.

- Tomas Vaškevičius, Varun Kanade, and Patrick Rebeschini. The statistical complexity of early-stopped mirror descent. arXiv preprint arXiv:2002.00189 , 2020.
- Namrata Vaswani and Seyedehsara Nayer. A simple generalization of a result for random matrices with independent sub-gaussian rows. arXiv preprint arXiv:1612.00127 , 2016.
- Zitong Yang, Yaodong Yu, Chong You, Jacob Steinhardt, and Yi Ma. Rethinking bias-variance trade-off for generalization of neural networks. In International Conference on Machine Learning , pp. 10767–10777. PMLR, 2020.
- Yuan Yao, Lorenzo Rosasco, and Andrea Caponnetto. On early stopping in gradient descent learning. *Constructive Approximation* , 26(2):289–315, 2007.
- Tong Zhang. Covering number bounds of certain regularized linear function classes. *Journal of Machine Learning Research* , 2(Mar):527–550, 2002.
- Tong Zhang and Bin Yu. Boosting with early stopping: Convergence and consistency. *The Annals of Statistics* , 33(4):1538–1579, 2005.

## A APPENDIX

## B PROOF OF MAIN RESULTS

## B.1 PROOF OF THEOREM 1

We first compute the expected risk in this setting.

**Lemma 1.** *For all  $n, p, d \in \mathbb{N}$  such that  $p \leq d$ , let  $X \in \mathbb{R}^{n \times d}$  and  $\theta^* \in \mathbb{R}^p$  be fixed. Let  $\lambda_1 \geq \dots \geq \lambda_d$  be the eigenvalues of the matrix  $\frac{1}{n}X^\top X$ . Then,*

$$\bar{R}_{X, \theta^*}(t) = \sigma^2 + \sum_{i=1}^d \exp(-2t\lambda_i) \frac{\|\theta^*\|_2^2}{d} + \frac{\sigma^2}{n} \sum_{i=1}^d \mathbb{1}\{\lambda_i \neq 0\} \frac{1}{\lambda_i} (1 - \exp(-t\lambda_i))^2. \quad (6)$$

*Proof.* We first compute  $R(\beta)$  for any  $\beta \in \mathbb{R}^d$  and a fixed  $\theta^*$ . Let  $\beta^* = P^\top \theta^*$ .

$$\begin{aligned} R(\beta) &= \mathbb{E}_{(\tilde{x}, \tilde{y}) \sim \mathcal{D}_{P, \theta^*}} \left[ (\langle \tilde{x}, \beta \rangle - \tilde{y})^2 \right] \\ &= \mathbb{E}_{(\tilde{x}, \tilde{y}) \sim \mathcal{D}_{P, \theta^*}} \left[ (\langle \tilde{x}, \beta \rangle - \langle P\tilde{x}, \theta^* \rangle - \tilde{\epsilon})^2 \right] \\ &= \mathbb{E}_{(\tilde{x}, \tilde{y}) \sim \mathcal{D}_{P, \theta^*}} \left[ (\langle \tilde{x}, \beta - \beta^* \rangle - \tilde{\epsilon})^2 \right] \\ &= \sigma^2 + \|\beta - \beta^*\|_2^2. \end{aligned} \quad (10)$$

The first step follows from the definition of  $R$ . The second step follows from equation 1. The third step follows from  $\langle P\tilde{x}, \theta^* \rangle = \langle \tilde{x}, P^\top \theta^* \rangle$ . The last step follows from  $\tilde{\epsilon} \sim \mathcal{N}(0, \sigma^2)$ ,  $\tilde{x} \sim \mathcal{N}(0, \mathbf{I}_d)$  and  $\tilde{\epsilon}$  is independent of  $\tilde{x}$ .

Next, we compute the expected risk  $\bar{R}_{X, \theta^*}(t)$ .

$$\begin{aligned} \bar{R}_{X, \theta^*}(t) &= \sigma^2 + \mathbb{E}_{P, \epsilon} [\|\beta_{X, y}(t) - \beta^*\|_2^2] \\ &= \sigma^2 + \mathbb{E}_{P, \epsilon} \left[ \left\| (X^\top X)^+ (I - \exp(-tX^\top X/n)) X^\top (X\beta^* + \epsilon) - \beta^* \right\|_2^2 \right] \\ &= \sigma^2 + \mathbb{E}_P \left[ \left\| (X^\top X)^+ (I - \exp(-tX^\top X/n)) X^\top X\beta^* - \beta^* \right\|_2^2 \right] \\ &\quad + \sigma^2 \left[ \left\| (X^\top X)^+ (I - \exp(-tX^\top X/n)) X^\top \right\|_F^2 \right] \end{aligned} \quad (11)$$

Here, we use  $\|\cdot\|_F$  to denote the Frobenius norm. The first step follows from equation 10. The second step follows from equation 4. The third step follows from  $\epsilon \sim \mathcal{N}(0, \mathbf{I}_n)$  is independent of  $X$  and  $\beta^*$ . Now let  $\frac{1}{n}X^\top X = U\Lambda U^\top$  be the eigenvalue decomposition of  $\frac{1}{n}X^\top X$ . Let  $\lambda_1 \geq \dots \geq \lambda_d$  be the eigenvalues of  $\frac{1}{n}X^\top X$ . Then,

$$\begin{aligned} &\mathbb{E}_P \left[ \left\| (X^\top X)^+ (I - \exp(-tX^\top X/n)) X^\top X\beta^* - \beta^* \right\|_2^2 \right] \\ &= \mathbb{E}_P \left[ \left\| (X^\top X)^+ (X^\top X) (I - \exp(-tX^\top X/n)) \beta^* - \beta^* \right\|_2^2 \right] \\ &= \mathbb{E}_P \left[ \left\| (I - \exp(-tX^\top X/n)) \beta^* - \beta^* \right\|_2^2 \right] \\ &= \mathbb{E}_P \left[ \left\| U \exp(-t\Lambda) U^\top P^\top \theta^* \right\|_2^2 \right] = \sum_{i=1}^d \exp(-2t\lambda_i) \frac{\|\theta^*\|_2^2}{d}. \end{aligned} \quad (12)$$

The first step follows from  $X^\top X$ ,  $(X^\top X)^+$  and  $I - \exp(-tX^\top X/n)$  are simultaneously diagonalizable. The second step follows from  $I - \exp(-tX^\top X/n)$  is already in row space of  $X^\top X$ . The third step follows from writing  $\frac{1}{n}X^\top X$  as  $U\Lambda U^\top$ . The last step follows from  $P$  is a random

semi-orthogonal matrix. Similarly, we have

$$\begin{aligned}
& \left\| (X^\top X)^\dagger (I - \exp(-tX^\top X/n)) X^\top \right\|_F^2 \\
&= \text{tr} \left( (X^\top X)^\dagger (I - \exp(-tX^\top X/n)) X^\top X (I - \exp(-tX^\top X/n)) (X^\top X)^\dagger \right) \\
&= \text{tr} \left( (I - \exp(-tX^\top X/n))^2 (X^\top X)^\dagger \right) \\
&= \frac{1}{n} \sum_{i=1}^d \mathbb{1}\{\lambda_i \neq 0\} \frac{1}{\lambda_i} (1 - \exp(-t\lambda_i))^2.
\end{aligned} \tag{13}$$

Combining equation 11, equation 12 and equation 13 completes the proof.  $\square$

We state a useful concentration result.

**Lemma 2.** Let  $X \in \mathbb{R}^{n \times d}$ ,  $d \geq n$  be a matrix whose entries are independent Gaussian random variables following  $\mathcal{N}(0, 1)$ . Let  $\lambda_1 \geq \dots \geq \lambda_d$  be the  $d$  eigenvalues of the matrix  $\frac{1}{n} X^\top X$ . Let

$\gamma = \left( \frac{\sqrt{n} + \sqrt{2 \log n}}{\sqrt{d}} \right)^2$ . For  $\gamma \geq 1$ ,  $\lambda_{n+1} = \dots = \lambda_d = 0$ , and with probability at least  $1 - \frac{2}{n}$ ,

$$\frac{d}{n} (1 - \sqrt{\gamma})^2 \leq \lambda_n \leq \lambda_1 \leq \frac{d}{n} (1 + \sqrt{\gamma})^2.$$

*Proof.*  $\lambda_{n+1} = \dots = \lambda_d = 0$  follows from the matrix  $\frac{1}{n} X^\top X$  is at most rank  $n$ . By Lemma 6, with probability at least  $1 - \frac{2}{n}$ ,

$$\frac{1}{n} \left( \sqrt{d} - \sqrt{n} - \sqrt{2 \log n} \right)^2 \leq \lambda_n \leq \lambda_1 \leq \frac{1}{n} \left( \sqrt{d} + \sqrt{n} + \sqrt{2 \log n} \right)^2.$$

For  $\gamma = \left( \frac{\sqrt{n} + \sqrt{2 \log n}}{\sqrt{d}} \right)^2$ ,  $\frac{1}{n} \left( \sqrt{d} - \sqrt{n} - \sqrt{2 \log n} \right)^2 = \frac{d}{n} (1 - \sqrt{\gamma})^2$  and  $\frac{1}{n} \left( \sqrt{d} + \sqrt{n} + \sqrt{2 \log n} \right)^2 = \frac{d}{n} (1 + \sqrt{\gamma})^2$ , which completes the proof.  $\square$

**Lemma 3.** Let  $X \in \mathbb{R}^{n \times d}$ ,  $d \leq n$  be a matrix whose entries are independent Gaussian random variables following  $\mathcal{N}(0, 1)$ . Let  $\lambda_1 \geq \dots \geq \lambda_d$  be the  $d$  eigenvalues of the matrix  $\frac{1}{n} X^\top X$ . Let

$\gamma = \left( \frac{\sqrt{n}}{\sqrt{d} + \sqrt{2 \log d}} \right)^2$ . For  $\gamma \geq 1$ , and with probability at least  $1 - \frac{2}{d}$ ,

$$\left( 1 - \frac{1}{\sqrt{\gamma}} \right)^2 \leq \lambda_d \leq \lambda_1 \leq \left( 1 + \frac{1}{\sqrt{\gamma}} \right)^2.$$

*Proof.* By Lemma 6, with probability at least  $1 - \frac{2}{d}$ ,

$$\frac{1}{n} \left( \sqrt{n} - \sqrt{d} - \sqrt{2 \log d} \right)^2 \leq \lambda_d \leq \lambda_1 \leq \frac{1}{n} \left( \sqrt{n} + \sqrt{d} + \sqrt{2 \log d} \right)^2.$$

For  $\gamma = \left( \frac{\sqrt{n}}{\sqrt{d} + \sqrt{2 \log d}} \right)^2$ ,  $\frac{1}{n} \left( \sqrt{n} - \sqrt{d} - \sqrt{2 \log d} \right)^2 = \left( 1 - \frac{1}{\sqrt{\gamma}} \right)^2$  and  $\frac{1}{n} \left( \sqrt{n} + \sqrt{d} + \sqrt{2 \log d} \right)^2 = \left( 1 + \frac{1}{\sqrt{\gamma}} \right)^2$ .  $\square$

With Lemma 2, we are able to bound the optimal stopping time.

**Theorem 1** (Optimal early stopping time in overparametrization setting). *In setting  $\mathcal{S}_{over}$ , for a fixed parameter  $\theta^*$  and noise variance  $\sigma^2$ , when  $n \leq d$ , let  $\gamma = \left( \frac{\sqrt{n} + \sqrt{2 \log n}}{\sqrt{d}} \right)^2$ . For  $\gamma \leq 1$ , with probability at least  $1 - \frac{2}{n}$  over the randomness of  $X$ , the optimal early stopping time  $t_{\text{opt}}$  satisfies*

$$\frac{n}{(1 + \sqrt{\gamma})^2 d} \log \left( 1 + \frac{(1 - \sqrt{\gamma})^2 \|\theta^*\|_2^2}{\sigma^2} \right) \leq t_{\text{opt}} \leq \frac{n}{(1 - \sqrt{\gamma})^2 d} \log \left( 1 + \frac{(1 + \sqrt{\gamma})^2 \|\theta^*\|_2^2}{\sigma^2} \right).$$

When  $n > d$ , let  $\gamma = \left(\frac{\sqrt{n}}{\sqrt{d} + \sqrt{2 \log d}}\right)^2$ . For  $\gamma \geq 1$ , with probability at least  $1 - \frac{2}{d}$  over the randomness of  $X$ , the optimal early stopping time  $t_{\text{opt}}$  satisfies

$$\frac{1}{\left(1 + \frac{1}{\sqrt{\gamma}}\right)^2} \log \left(1 + \left(1 - \frac{1}{\sqrt{\gamma}}\right)^2 \frac{n \|\theta^*\|_2^2}{d\sigma^2}\right) \leq t_{\text{opt}} \leq \frac{1}{\left(1 - \frac{1}{\sqrt{\gamma}}\right)^2} \log \left(1 + \left(1 + \frac{1}{\sqrt{\gamma}}\right)^2 \frac{n \|\theta^*\|_2^2}{d\sigma^2}\right).$$

*Proof of Theorem 1.* We first compute the derivative of the expected risk. By Lemma 1,

$$\begin{aligned} \frac{d}{dt} \bar{R}_{X, \theta^*}(t) &= \frac{d}{dt} \left[ \sigma^2 + \sum_{i=1}^d \exp(-2t\lambda_i) \frac{\|\theta^*\|_2^2}{d} + \frac{\sigma^2}{n} \sum_{i=1}^d \mathbb{1}\{\lambda_i \neq 0\} \frac{1}{\lambda_i} (1 - \exp(-t\lambda_i))^2 \right] \\ &= - \sum_{i=1}^d 2\lambda_i \exp(-2t\lambda_i) \frac{\|\theta^*\|_2^2}{d} + \frac{2\sigma^2}{n} \sum_{i=1}^d (\exp(-t\lambda_i) - \exp(-2t\lambda_i)). \end{aligned}$$

Let

$$d\bar{R}_{X, \theta^*}^{(i)}(t) = -2\lambda_i \exp(-2t\lambda_i) \frac{\|\theta^*\|_2^2}{d} + \frac{2\sigma^2}{n} (\exp(-t\lambda_i) - \exp(-2t\lambda_i)).$$

We first consider the case  $d > n$ . For  $i$  such that  $\lambda_i = 0$ ,  $d\bar{R}_{X, \theta^*}^{(i)}(t) = 0$  for all  $t$ . For  $i$  such that  $\lambda_i \neq 0$ , we can show that if  $t < \frac{n}{(1+\sqrt{\gamma})^2 d} \log \left(1 + \frac{(1-\sqrt{\gamma})^2 \|\theta\|_2^2}{\sigma^2}\right)$ ,  $d\bar{R}_{X, \theta^*}^{(i)}(t) < 0$ . Let  $\alpha = \exp(-t\lambda_i)$ , then

$$d\bar{R}_{X, \theta^*}^{(i)}(t) = - \left( \frac{2\|\theta^*\|_2^2}{d} \lambda_i + \frac{2\sigma^2}{n} \right) \alpha^2 + \frac{2\sigma^2}{n} \alpha.$$

When  $\alpha > \frac{\sigma^2 d}{\sigma^2 d + \|\theta^*\|_2^2 n \lambda_i}$ ,  $d\bar{R}_{X, \theta^*}^{(i)}(t) < 0$ . When  $t < \log \left( \frac{\sigma^2 d + \|\theta^*\|_2^2 n \lambda_i}{\sigma^2 d} \right) / \lambda_i$ ,  $\alpha > \frac{\sigma^2 d}{\sigma^2 d + \|\theta^*\|_2^2 n \lambda_i}$ . By Lemma 2, with probability at least  $1 - \frac{2}{n}$ ,

$$\frac{d}{n} (1 - \sqrt{\gamma})^2 \leq \lambda_n \leq \frac{d}{n} (1 + \sqrt{\gamma})^2,$$

where  $\gamma = \left(\frac{\sqrt{n} + \sqrt{2 \log n}}{\sqrt{d}}\right)^2$ . Then, when  $t < \frac{n}{(1+\sqrt{\gamma})^2 d} \log \left(1 + \frac{(1-\sqrt{\gamma})^2 \|\theta\|_2^2}{\sigma^2}\right)$ ,  $d\bar{R}_{X, \theta^*}^{(i)}(t) < 0$  for all  $i$  such that  $\lambda_i \neq 0$ , which shows  $t_{\text{opt}} \geq \frac{n}{(1+\sqrt{\gamma})^2 d} \log \left(1 + \frac{(1-\sqrt{\gamma})^2 \|\theta\|_2^2}{\sigma^2}\right)$ . Similarly, we can upper bound  $t_{\text{opt}}$ . For  $i$  such that  $\lambda_i \neq 0$ , when  $0 < \alpha < \frac{\sigma^2 d}{\sigma^2 d + \|\theta^*\|_2^2 n \lambda_i}$ ,  $d\bar{R}_{X, \theta^*}^{(i)}(t) > 0$ . When  $t > \log \left( \frac{\sigma^2 d + \|\theta^*\|_2^2 n \lambda_i}{\sigma^2 d} \right) / \lambda_i$ ,  $\alpha < \frac{\sigma^2 d}{\sigma^2 d + \|\theta^*\|_2^2 n \lambda_i}$ . Then, when  $t > \frac{n}{(1-\sqrt{\gamma})^2 d} \log \left(1 + \frac{(1+\sqrt{\gamma})^2 \|\theta\|_2^2}{\sigma^2}\right)$ ,  $d\bar{R}_{X, \theta^*}^{(i)}(t) > 0$  for all  $i$  such that  $\lambda_i \neq 1$ , which shows  $t_{\text{opt}} \leq \frac{n}{(1-\sqrt{\gamma})^2 d} \log \left(1 + \frac{(1+\sqrt{\gamma})^2 \|\theta\|_2^2}{\sigma^2}\right)$ .

Next, we consider the case when  $n > d$ . For  $i$  such that  $\lambda_i = 0$ ,  $d\bar{R}_{X, \theta^*}^{(i)}(t) = 0$  for all  $t$ . For  $i$  such that  $\lambda_i \neq 0$ , we can show that if  $t < \frac{1}{\left(1 + \frac{1}{\sqrt{\gamma}}\right)^2} \log \left(1 + \left(1 - \frac{1}{\sqrt{\gamma}}\right)^2 \frac{n \|\theta\|_2^2}{d\sigma^2}\right)$ ,  $d\bar{R}_{X, \theta^*}^{(i)}(t) < 0$ . Let  $\alpha = \exp(-t\lambda_i)$ , then

$$d\bar{R}_{X, \theta^*}^{(i)}(t) = - \left( \frac{2\|\theta^*\|_2^2}{d} \lambda_i + \frac{2\sigma^2}{n} \right) \alpha^2 + \frac{2\sigma^2}{n} \alpha.$$

When  $\alpha > \frac{\sigma^2 d}{\sigma^2 d + \|\theta^*\|_2^2 n \lambda_i}$ ,  $d\bar{R}_{X, \theta^*}^{(i)}(t) < 0$ . When  $t < \log \left( \frac{\sigma^2 d + \|\theta^*\|_2^2 n \lambda_i}{\sigma^2 d} \right) / \lambda_i$ ,  $\alpha > \frac{\sigma^2 d}{\sigma^2 d + \|\theta^*\|_2^2 n \lambda_i}$ . By Lemma 3, with probability at least  $1 - \frac{2}{d}$ ,

$$\left(1 - \frac{1}{\sqrt{\gamma}}\right)^2 \leq \lambda_d \leq \lambda_1 \leq \left(1 + \frac{1}{\sqrt{\gamma}}\right)^2,$$

where  $\gamma = \left(\frac{\sqrt{n}}{\sqrt{d} + \sqrt{2 \log d}}\right)^2$ . Then, when  $t < \frac{1}{\left(1 + \frac{1}{\sqrt{\gamma}}\right)^2} \log \left(1 + \left(1 - \frac{1}{\sqrt{\gamma}}\right)^2 \frac{n \|\theta\|_2^2}{d \sigma^2}\right)$ ,  $d \bar{R}_{X, \theta^*}^{(i)}(t) < 0$  for all  $i$  such that  $\lambda_i \neq 0$ , which shows  $t_{\text{opt}} \geq \frac{1}{\left(1 + \frac{1}{\sqrt{\gamma}}\right)^2} \log \left(1 + \left(1 - \frac{1}{\sqrt{\gamma}}\right)^2 \frac{n \|\theta\|_2^2}{d \sigma^2}\right)$ . Similarly, we can upper bound  $t_{\text{opt}}$ . For  $i$  such that  $\lambda_i \neq 0$ , when  $0 < \alpha < \frac{\sigma^2 d}{\sigma^2 d + \|\theta^*\|^2 n \lambda_i}$ ,  $d \bar{R}_{X, \theta^*}^{(i)}(t) > 0$ . When  $t > \log \left(\frac{\sigma^2 d + \|\theta^*\|^2 n \lambda_i}{\sigma^2 d}\right) / \lambda_i$ ,  $\alpha < \frac{\sigma^2 d}{\sigma^2 d + \|\theta^*\|^2 n \lambda_i}$ . Then, when  $t > \frac{1}{\left(1 - \frac{1}{\sqrt{\gamma}}\right)^2} \log \left(1 + \left(1 + \frac{1}{\sqrt{\gamma}}\right)^2 \frac{n \|\theta\|_2^2}{d \sigma^2}\right)$ ,  $d \bar{R}_{X, \theta^*}^{(i)}(t) > 0$  for all  $i$  such that  $\lambda_i \neq 0$ , which shows  $t_{\text{opt}} \leq \frac{1}{\left(1 - \frac{1}{\sqrt{\gamma}}\right)^2} \log \left(1 + \left(1 + \frac{1}{\sqrt{\gamma}}\right)^2 \frac{n \|\theta\|_2^2}{d \sigma^2}\right)$ .  $\square$

## B.2 PROOF OF PROPOSITION 1

**Proposition 1.** *In setting  $S_{\text{over}}$ , for a fixed parameter  $\theta^*$  and noise variance  $\sigma^2$ . Let  $\bar{t}_{\text{opt}} = \alpha n / (n + d)$  for some constant  $\alpha$ . Let  $\lambda_1 \geq \dots \geq \lambda_d$  be the eigenvalues of the matrix  $\frac{1}{n} X^\top X$ . Then,*

$$\bar{R}_1 \leq \mathbb{E}_X [\bar{R}_{X, \theta^*}(\bar{t}_{\text{opt}})] \leq \bar{R}_2. \quad (7)$$

where

$$\bar{R}_1 = \sigma^2 + \mathbb{E}_X \left[ \sum_{i=1}^d \exp(-2\alpha n \lambda_i / d) \right] \frac{\|\theta^*\|^2}{d}, \quad \text{and} \quad \bar{R}_2 = \bar{R}_1 + \frac{1}{2} \alpha \sigma^2.$$

$\bar{R}_1$  and  $\bar{R}_2$  decrease monotonically as  $n$  increases.

*Proof.* By Lemma 1,

$$\mathbb{E}_X [\bar{R}_{X, \theta^*}(\bar{t}_{\text{opt}})] = \sigma^2 + \mathbb{E}_X \left[ \sum_{i=1}^d \exp(-2\bar{t}_{\text{opt}} \lambda_i) \frac{\|\theta^*\|^2}{d} + \frac{\sigma^2}{n} \sum_{i=1}^d \mathbb{1}\{\lambda_i \neq 0\} \frac{1}{\lambda_i} (1 - \exp(-\bar{t}_{\text{opt}} \lambda_i))^2 \right].$$

where  $\lambda_i$ 's are the eigenvalues of the matrix  $\frac{1}{n} X^\top X$ . Next, we can bound

$$\begin{aligned} \mathbb{E}_X \left[ \frac{\sigma^2}{n} \sum_{i=1}^d \mathbb{1}\{\lambda_i \neq 0\} \frac{1}{\lambda_i} (1 - \exp(-\bar{t}_{\text{opt}} \lambda_i))^2 \right] &\leq \mathbb{E}_X \left[ \frac{\sigma^2}{n} \sum_{i=1}^d \mathbb{1}\{\lambda_i \neq 0\} \frac{1}{\lambda_i} (1 - \exp(-\bar{t}_{\text{opt}} \lambda_i)) \right] \\ &\leq \mathbb{E}_X \left[ \frac{\sigma^2}{n} \sum_{i=1}^d \mathbb{1}\{\lambda_i \neq 0\} \frac{1}{\lambda_i} \cdot \bar{t}_{\text{opt}} \lambda_i \right] \\ &= \frac{\alpha \sigma^2 \min\{n, d\}}{n + d} \leq \frac{1}{2} \alpha \sigma^2. \end{aligned}$$

We used  $0 \leq 1 - \exp(-x) \leq 1$  and  $\exp(-x) \geq 1 - x$  for  $x \geq 0$ . It is trivial to show  $\mathbb{E}_X \left[ \frac{\sigma^2}{n} \sum_{i=1}^d \frac{1}{\lambda_i} (1 - \exp(-\bar{t}_{\text{opt}} \lambda_i))^2 \right] \geq 0$ .

Finally, we show that  $\bar{R}$  decreases as  $n$  increases. Let  $X' \in \mathbb{R}^{(n+1) \times d}$  be a random matrix with the first  $n$  rows the same as  $X$  and the last row follow  $\mathcal{N}(0, \mathbf{I})$ . Let  $\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_d$  be the eigenvalues of the matrix  $\frac{1}{n} X^\top X$  and  $\lambda'_1 \geq \lambda'_2 \geq \dots \geq \lambda'_d$  be the eigenvalues of the matrix  $\frac{1}{n} X'^\top X'$ . Then, by Lemma 8, we have  $\lambda'_i \geq \lambda_i$  for all  $i = 1, \dots, d$ . The derivative

$$\frac{d}{d\lambda_i} \left( \sum_{i=1}^d \exp(-2\bar{t}_{\text{opt}} \lambda_i) \right) = -2\bar{t}_{\text{opt}} \exp(-2\bar{t}_{\text{opt}} \lambda_i) \leq 0.$$

By coupling  $X$  and  $X'$ , we have

$$\mathbb{E}_X \left[ \sum_{i=1}^d \exp \left( -2 \frac{\alpha n}{n + d} \lambda_i \right) \right] \geq \mathbb{E}_{X'} \left[ \sum_{i=1}^d \exp \left( -2 \frac{\alpha n}{n + d} \lambda'_i \right) \right] \geq \mathbb{E}_{X'} \left[ \sum_{i=1}^d \exp \left( -2 \frac{\alpha(n+1)}{n+1+d} \lambda'_i \right) \right],$$

which shows  $\bar{R}$  decreases as  $n$  increases.  $\square$

## B.3 PROOF OF THEOREM 2

We first compute the expected risk  $\bar{R}_{\theta^*}(t)$ .

**Lemma 4.** For all  $n, p, d \in \mathbb{N}$  such that  $p \geq d$ , let  $\theta^* \in \mathbb{R}^p$  be fixed. Let  $\mathcal{D}_\Lambda$  be the distribution of eigenvalues of the matrix  $\frac{1}{n} \tilde{X}^\top \tilde{X}$ , where each entry of the matrix  $\tilde{X} \in \mathbb{R}^{n \times d}$  are i.i.d standard normal random variable. Then,

$$\begin{aligned} \bar{R}_{\theta^*}(t) &= \sigma^2 + \left(1 - \frac{d}{p}\right) \|\theta^*\|_2^2 + \mathbb{E}_{\Lambda \sim \mathcal{D}_\Lambda} \left[ \sum_{i=1}^d \exp(-2t\lambda_i) \frac{\|\theta^*\|_2^2}{p} \right] + \\ &\mathbb{E}_{\Lambda \sim \mathcal{D}_\Lambda} \left[ \frac{1}{n} \sum_{i=1}^d \mathbb{1}\{\lambda_i \neq 0\} \frac{1}{\lambda_i} (1 - \exp(-t\lambda_i))^2 \left( \sigma^2 + \frac{p-d}{p} \|\theta^*\|_2^2 \right) \right]. \end{aligned}$$

*Proof.* We first compute  $R(\beta)$  for any  $\beta \in \mathbb{R}^d$  and a fixed  $\theta$ . Let  $\beta^* = P^\top \theta^*$ , which minimizes the population risk  $R(\beta)$ .

$$\begin{aligned} R(\beta) &= \mathbb{E}_{(\tilde{x}, \tilde{y}) \sim \mathcal{D}_{P, \theta^*}} \left[ (\langle \tilde{x}, \beta \rangle - \tilde{y})^2 \right] \\ &= \mathbb{E}_{(\tilde{x}, \tilde{y}) \sim \mathcal{D}_{P, \theta^*}} \left[ (\langle \tilde{z}, P\beta \rangle - \langle \tilde{z}, \theta^* \rangle - \tilde{\epsilon})^2 \right] \\ &= \sigma^2 + \|P\beta - \theta^*\|_2^2 \\ &= \sigma^2 + \|\theta^* - P\beta^*\|_2^2 + \|P\beta - P\beta^*\|_2^2 + 2 \langle P(\beta - \beta^*), \theta^* - P\beta^* \rangle \\ &= \sigma^2 + \|\theta^* - P\beta^*\|_2^2 + \|\beta - \beta^*\|_2^2 \end{aligned} \tag{14}$$

The first step follows from the definition of  $R(t)$ . The second step follows from equation 8. The third step follows from  $\epsilon \sim \mathcal{N}(0, \sigma^2)$ ,  $\tilde{z} \sim \mathcal{N}(0, \mathbf{I}_d)$  and  $\tilde{\epsilon}$  is independent of  $\tilde{x}$ . The last two steps follows from expanding and  $P^\top \theta^* = \beta^* = P^\top P\beta^*$ .

Next, we compute the expected risk  $\bar{R}_{\theta^*}(t)$ . Since  $P$  is a random semi-orthogonal matrix,

$$\mathbb{E}_P \left[ \|\theta^* - P\beta^*\|_2^2 \right] = \mathbb{E}_P \left[ \|\theta^* - PP^\top \theta^*\|_2^2 \right] = \left(1 - \frac{d}{p}\right) \|\theta^*\|_2^2.$$

For the term  $\|\beta - \beta^*\|_2^2$ ,

$$\begin{aligned} &\mathbb{E}_{Z, P, \epsilon} \left[ \|\beta - \beta^*\|_2^2 \right] \\ &= \mathbb{E}_{Z, P, \epsilon} \left[ \left\| (X^\top X)^\dagger (I - \exp(-tX^\top X/n)) X^\top (X\beta^* + Z(I - PP^\top)\theta^* + \epsilon) - \beta^* \right\|_2^2 \right] \\ &= \mathbb{E}_{Z, P} \left[ \left\| (X^\top X)^\dagger (I - \exp(-tX^\top X/n)) X^\top X\beta^* - \beta^* \right\|_2^2 \right] \\ &+ \mathbb{E}_{Z, P, \epsilon} \left[ \left\| (X^\top X)^\dagger (I - \exp(-tX^\top X/n)) X^\top Z(I - PP^\top)\theta^* \right\|_2^2 \right] \\ &+ \sigma^2 \left\| (X^\top X)^\dagger (I - \exp(-tX^\top X/n)) X^\top \right\|_F^2 \end{aligned} \tag{15}$$

The first step follows from equation 4. The second step follows from  $\epsilon \sim \mathcal{N}(0, \mathbf{I}_n)$  is independent of  $X$  and  $\beta^*$  and  $I - PP^\top$  projects to orthogonal complement of the column space of  $P$ .

Now let  $\frac{1}{n}X^\top X = U\Lambda U^\top$  be the eigenvalue decomposition of  $\frac{1}{n}X^\top X$ . We have

$$\begin{aligned}
& \mathbb{E}_{Z,P} \left[ \left\| (X^\top X)^\dagger (I - \exp(-tX^\top X/n)) X^\top X \beta^* - \beta^* \right\|_2^2 \right] \\
&= \mathbb{E}_{Z,P} \left[ \left\| (X^\top X)^\dagger (X^\top X) (I - \exp(-tX^\top X/n)) \beta^* - \beta^* \right\|_2^2 \right] \\
&= \mathbb{E}_{Z,P} \left[ \left\| (I - \exp(-tX^\top X/n)) \beta^* - \beta^* \right\|_2^2 \right] \\
&= \mathbb{E}_{Z,P} \left[ \left\| U \exp(-t\Lambda) U^\top P^\top \theta^* \right\|_2^2 \right] \\
&= \mathbb{E}_{\Lambda,P} \left[ \sum_{i=1}^d \mathbb{1}\{\lambda_i \neq 0\} \exp(-2t\lambda_i) \frac{\|P^\top \theta^*\|_2^2}{d} \right] \\
&= \mathbb{E}_{\Lambda} \left[ \sum_{i=1}^d \mathbb{1}\{\lambda_i \neq 0\} \exp(-2t\lambda_i) \frac{\|\theta^*\|_2^2}{p} \right]
\end{aligned} \tag{16}$$

The first step follows from  $X^\top X$ ,  $(X^\top X)^\dagger$  and  $I - \exp(-tX^\top X/n)$  are simultaneously diagonalizable. The second step follows from  $I - \exp(-tX^\top X/n)$  is already in row space of  $X^\top X$ . The third step follows from writing  $\frac{1}{n}X^\top X$  as  $U\Lambda U^\top$ . The last two step follows from  $X$  and  $P$  is independent by rotational invariance of  $X$ . Similarly, we have

$$\begin{aligned}
& \left\| (X^\top X)^\dagger (I - \exp(-tX^\top X/n)) X^\top \right\|_F^2 \\
&= \text{tr} \left( (X^\top X)^\dagger (I - \exp(-tX^\top X/n)) X^\top X (I - \exp(-tX^\top X/n)) (X^\top X)^\dagger \right) \\
&= \text{tr} \left( (I - \exp(-tX^\top X/n))^2 (X^\top X)^\dagger \right) \\
&= \frac{1}{n} \sum_{i=1}^d \mathbb{1}\{\lambda_i \neq 0\} \frac{1}{\lambda_i} (1 - \exp(-t\lambda_i))^2.
\end{aligned} \tag{17}$$

For the second term,

$$\begin{aligned}
& \mathbb{E}_{Z,P} \left[ \left\| (X^\top X)^\dagger (I - \exp(-tX^\top X/n)) X^\top Z (I - PP^\top) \theta^* \right\|_2^2 \right] \\
&= \mathbb{E}_{Z,P} \left[ \left\| (X^\top X)^\dagger (I - \exp(-tX^\top X/n)) X^\top \right\|_F^2 \left\| Z (I - PP^\top) \theta^* \right\|_2^2 \right] \\
&= \mathbb{E}_{\Lambda} \left[ \frac{1}{n} \sum_{i=1}^d \mathbb{1}\{\lambda_i \neq 0\} \frac{1}{\lambda_i} (1 - \exp(-t\lambda_i))^2 \cdot \frac{p-d}{p} \|\theta^*\|_2^2 \right]
\end{aligned} \tag{18}$$

The first step follows from  $X$  and  $Z(I - PP^\top)$  are independent. Finally, we note that since  $P$  is a semi-orthogonal matrix, the eigenvalues of  $\frac{1}{n}X^\top X$  follows the distribution  $\mathcal{D}_\Lambda$ . Combining equation 15, equation 16, equation 17 and equation 18 completes the proof.  $\square$

Now, we are ready to give an approximation of the optimal stopping time  $t_{\text{opt}}$ .

**Theorem 2** (Optimal early stopping time in underparametrization setting). *In setting  $\mathcal{S}_{\text{under}}$ , for a fixed parameter  $\theta^*$  and noise variance  $\sigma^2$ , if  $(8n + 9d + 16) \|\theta^*\|_2^2 \leq p\sigma^2 + p\|\theta^*\|_2^2$ , the optimal early stopping time  $t_{\text{opt}}$  satisfies*

$$\frac{n \|\theta^*\|_2^2}{2(p\sigma^2 + (p-d)\|\theta^*\|_2^2)} \leq t_{\text{opt}} \leq \frac{2n\|\theta^*\|_2^2}{p\sigma^2 + (p-d)\|\theta^*\|_2^2}.$$

*Proof of Theorem 2.* In this proof, all expectations are taken over the distribution  $\mathcal{D}_\Lambda$ , so we omit them. We first compute the derivative of  $\bar{R}_{\theta^*}$  with respect to time  $t$ . By Lemma 4,

$$\frac{d}{dt} \bar{R}_{\theta^*}(t) = \mathbb{E} \left[ -\frac{2\|\theta^*\|_2^2}{p} \sum_{i=1}^d \lambda_i \exp(-2t\lambda_i) + \frac{2}{n} \left( \sigma^2 + \frac{p-d}{p} \|\theta^*\|_2^2 \right) \sum_{i=1}^d (\exp(-t\lambda_i) - \exp(-2t\lambda_i)) \right].$$

We first derive an upper bound on  $\frac{d}{dt}\bar{R}_{\theta^*}(t)$ ,

$$\begin{aligned}\frac{d}{dt}\bar{R}_{\theta^*}(t) &\leq \mathbb{E} \left[ -\frac{2\|\theta^*\|_2^2}{p} \sum_{i=1}^d (\lambda_i - 2t\lambda_i^2) + \frac{2}{n} \left( \sigma^2 + \frac{p-d}{p} \|\theta^*\|_2^2 \right) \sum_{i=1}^d t\lambda_i \right] \\ &= \mathbb{E} \left[ -\frac{2\|\theta^*\|_2^2}{p} \sum_{i=1}^d \lambda_i + \left( \frac{4\|\theta^*\|_2^2}{p} \sum_{i=1}^d \lambda_i^2 + \frac{2}{n} \left( \sigma^2 + \frac{p-d}{p} \|\theta^*\|_2^2 \right) \sum_{i=1}^d \lambda_i \right) t \right].\end{aligned}$$

We used  $\exp(-x) \geq 1-x$  and  $\exp(-x) - \exp(-2x) \leq x$  for all  $x \geq 0$ . Then, we have  $\frac{d}{dt}\bar{R}_{\theta^*}(t) < 0$  at time  $t < t_1$ , where

$$\begin{aligned}t_1 &= \frac{\frac{2\|\theta^*\|_2^2}{p} \mathbb{E} \left[ \sum_{i=1}^d \lambda_i \right]}{\frac{4\|\theta^*\|_2^2}{p} \mathbb{E} \left[ \sum_{i=1}^d \lambda_i^2 \right] + \frac{2}{n} \left( \sigma^2 + \frac{p-d}{p} \|\theta^*\|_2^2 \right) \mathbb{E} \left[ \sum_{i=1}^d \lambda_i \right]} \\ &= \frac{n\|\theta^*\|_2^2}{2(n+d+2)\|\theta^*\|_2^2 + p\sigma^2 + (p-d)\|\theta^*\|_2^2} \geq \frac{n\|\theta^*\|_2^2}{p\sigma^2 + 2(p-d)\|\theta^*\|_2^2}.\end{aligned}$$

The second equation follows from Lemma 7. The last step follows from the assumption  $(8n+9d+16)\|\theta^*\|_2^2 \leq p\sigma^2 + p\|\theta^*\|_2^2$ . Next, we derive a lower bound on  $\frac{d}{dt}\bar{R}_{\theta^*}(t)$  for  $t < T = \frac{n}{4(n+d+2)}$ ,

$$\frac{d}{dt}\bar{R}_{\theta^*}(t) \geq \mathbb{E} \left[ -\frac{2\|\theta^*\|_2^2}{p} \sum_{i=1}^d \lambda_i + \frac{2}{n} \left( \sigma^2 + \frac{p-d}{p} \|\theta^*\|_2^2 \right) \sum_{i=1}^d (t\lambda_i - 2T\lambda_i^2) \right].$$

We used  $\exp(-2x) \leq 1-2x+2x^2$  for all  $x \geq 0$ . Then, we have  $\frac{d}{dt}\bar{R}_{\theta^*}(t) \geq 0$  at time

$$t_2 = \frac{\frac{2\|\theta^*\|_2^2}{p} \mathbb{E} \left[ \sum_{i=1}^d \lambda_i \right]}{\frac{2}{n} \left( \sigma^2 + \frac{p-d}{p} \|\theta^*\|_2^2 \right) \mathbb{E} \left[ \sum_{i=1}^d \lambda_i - 2T\lambda_i^2 \right]} \leq \frac{2n\|\theta^*\|_2^2}{p\sigma^2 + (p-d)\|\theta^*\|_2^2}$$

which follows from Lemma 7. Under assumption  $(8n+9d+16)\|\theta^*\|_2^2 \leq p\sigma^2 + p\|\theta^*\|_2^2$ ,  $t_2 < T$ . Thus, we have

$$\frac{n\|\theta^*\|_2^2}{2(p\sigma^2 + (p-d)\|\theta^*\|_2^2)} \leq t_{\text{opt}} \leq \frac{2n\|\theta^*\|_2^2}{p\sigma^2 + (p-d)\|\theta^*\|_2^2}$$

□

**Proposition 2.** In setting  $\mathcal{S}_{\text{under}}$ , for a fixed parameter  $\theta^*$  and noise variance  $\sigma^2$ , let  $\bar{t}_{\text{opt}} = \frac{n\|\theta^*\|_2^2}{p\sigma^2 + (p-d)\|\theta^*\|_2^2}$ . Then, when  $(8n+9d+16)\|\theta^*\|_2^2 \leq p\sigma^2 + p\|\theta^*\|_2^2$ , we have

$$\bar{R}_1 \leq \bar{R}_{\theta^*}(\bar{t}_{\text{opt}}) \leq \bar{R}_2,$$

where  $\bar{R}_1 = \sigma^2 + \|\theta^*\|_2^2 - \frac{2nd\|\theta^*\|_2^4}{p^2\sigma^2 + p(p-d)\|\theta^*\|_2^2}$  and  $\bar{R}_2 = \sigma^2 + \|\theta^*\|_2^2 - \frac{3nd\|\theta^*\|_2^4}{4(p^2\sigma^2 + p(p-d)\|\theta^*\|_2^2)}$ .

Here,  $\bar{R}_1 \geq 0.8\bar{R}_2$  and both  $\bar{R}_1$  and  $\bar{R}_2$  decrease when  $d$  and  $n$  increase.

*Proof of Proposition 2.* By Lemma 4,

$$\begin{aligned}\bar{R}_{\theta^*}(\bar{t}_{\text{opt}}) &= \sigma^2 + \left(1 - \frac{d}{p}\right) \|\theta^*\|_2^2 + \mathbb{E}_{\Lambda \sim \mathcal{D}_\Lambda} \left[ \sum_{i=1}^d \exp(-2\bar{t}_{\text{opt}}\lambda_i) \frac{\|\theta^*\|_2^2}{p} \right] + \\ &\quad \mathbb{E}_{\Lambda \sim \mathcal{D}_\Lambda} \left[ \frac{1}{n} \sum_{i=1}^d \frac{1}{\lambda_i} (1 - \exp(-\bar{t}_{\text{opt}}\lambda_i))^2 \left( \sigma^2 + \frac{p-d}{p} \|\theta^*\|_2^2 \right) \right]\end{aligned}$$

$\mathcal{D}_\Lambda$  is the distribution of eigenvalues of a matrix  $\frac{1}{n} \tilde{X}^\top \tilde{X}$  where each entry of  $\tilde{X}$  is i.i.d standard normal random variable. We first derive an upper bound on  $\bar{R}_{\theta^*}(\bar{t}_{\text{opt}})$ . For  $\bar{t}_{\text{opt}} = \frac{n \|\theta^*\|_2^2}{p\sigma^2 + (p-d)\|\theta^*\|_2^2}$ ,

$$\begin{aligned} & \bar{R}_{\theta^*}(\bar{t}_{\text{opt}}) \\ & \leq \sigma^2 + \|\theta^*\|_2^2 + \frac{\|\theta^*\|_2^2}{p} \mathbb{E} \left[ \sum_{i=1}^d -2\bar{t}_{\text{opt}}\lambda_i + 2\bar{t}_{\text{opt}}^2\lambda_i^2 \right] + \frac{1}{n} \left( \sigma^2 + \frac{p-d}{p} \|\theta^*\|_2^2 \right) \mathbb{E} \left[ \sum_{i=1}^d \bar{t}_{\text{opt}}^2 \lambda_i \right] \\ & = \sigma^2 + \|\theta^*\|_2^2 - \frac{2\bar{t}_{\text{opt}} \|\theta^*\|_2^2 d}{p} + \frac{2\bar{t}_{\text{opt}}^2 \|\theta^*\|_2^2 d(d+n+2)}{pn} + \frac{d}{n} \left( \sigma^2 + \frac{p-d}{p} \|\theta^*\|_2^2 \right) \bar{t}_{\text{opt}}^2 \\ & = \sigma^2 + \|\theta^*\|_2^2 - \frac{nd \|\theta^*\|_2^4}{p^2\sigma^2 + p(p-d)\|\theta^*\|_2^2} + \frac{2\|\theta^*\|_2^6 nd(d+n+2)}{p(p\sigma^2 + (p-d)\|\theta^*\|_2^2)^2}. \end{aligned}$$

The inequality follows from  $\exp(-x) \leq 1 - x + x^2$  for  $x \geq 0$ . The second step follows from Lemma 7. When  $8(n+d+2)\|\theta^*\|_2^2 \leq p\sigma^2 + (p-d)\|\theta^*\|_2^2$ , we have

$$\bar{R}_{\theta^*}(\bar{t}_{\text{opt}}) \leq \sigma^2 + \|\theta^*\|_2^2 - \frac{3nd \|\theta^*\|_2^4}{4p^2\sigma^2 + p(p-d)\|\theta^*\|_2^2} = \bar{R}_2.$$

Next, we derive a lower bound on  $\bar{R}_{\theta^*}(\alpha\bar{t}_{\text{opt}})$ .

$$\bar{R}_{\theta^*}(\bar{t}_{\text{opt}}) \geq \sigma^2 + \|\theta^*\|_2^2 + \frac{\|\theta^*\|_2^2}{p} \mathbb{E} \left[ \sum_{i=1}^d -2\bar{t}_{\text{opt}}\lambda_i \right] = \sigma^2 + \|\theta^*\|_2^2 - \frac{2nd \|\theta^*\|_2^4}{p^2\sigma^2 + p(p-d)\|\theta^*\|_2^2} = \bar{R}_1.$$

The inequality follow from  $\exp(-x) \geq 1 - x$ . The second step follows from Lemma 7. Since  $d \leq p$  and  $(8n+9d+16)\|\theta^*\|_2^2 \leq p\sigma^2 + p\|\theta^*\|_2^2$ ,  $\frac{nd\|\theta^*\|_2^4}{p^2\sigma^2 + p(p-d)\|\theta^*\|_2^2} \leq \frac{1}{8}\|\theta^*\|_2^2$ , which shows

$$\frac{\bar{R}_1}{\bar{R}_2} \geq \frac{1 - \frac{1}{8} \cdot 2}{1 - \frac{1}{8} \cdot \frac{3}{4}} \geq 0.8.$$

Both  $\bar{R}_1$  and  $\bar{R}_2$  decreases when  $d$  and  $n$  increases.  $\square$

## C ASYMPTOTIC REGIME

In this section, we examine the behavior of the optimal stopping time in the asymptotic regime where  $n, d \rightarrow \infty$ . We consider the setting  $\mathcal{S}_{\text{over}}$ . We define  $\gamma = n/d$ .

By Lemma 1, the derivative of the expected risk is given by

$$\frac{d}{dt} \bar{R}_{X, \theta^*}(t) = - \sum_{i=1}^d 2\lambda_i \exp(-2t\lambda_i) \frac{\|\theta^*\|_2^2}{d} + \frac{2\sigma^2}{n} \sum_{i=1}^d (\exp(-t\lambda_i) - \exp(-2t\lambda_i))$$

$\lambda_{n+1} = \dots = \lambda_d = 0$  if  $d > n$ .

Using the Marchenko–Pastur distribution, we can take the limit of  $d, n \rightarrow \infty$  and get

$$\begin{aligned} \frac{d}{dt} \bar{R}_{X_\infty, \theta^*}(t) &= \int_{(\sqrt{1/\gamma-1})^2}^{(\sqrt{1/\gamma+1})^2} \left( -2\|\theta^*\|_2^2 \lambda \exp(-2t\lambda) + \frac{2}{\gamma} \sigma^2 (\exp(-t\lambda) - \exp(-2t\lambda)) \right) d\mu(\lambda) \\ &= \int_{(\sqrt{1/\gamma-1})^2}^{(\sqrt{1/\gamma+1})^2} \left( -2\|\theta^*\|_2^2 \lambda \exp(-2t\lambda) + \frac{2}{\gamma} \sigma^2 (\exp(-t\lambda) - \exp(-2t\lambda)) \right) dF(\lambda), \end{aligned}$$

where

$$d\mu(\lambda) = \mathbf{1} \left\{ \lambda \in [(\sqrt{1/\gamma-1})^2, (\sqrt{1/\gamma+1})^2] \right\} dF(\lambda) + \max\{0, 1-\gamma\} \mathbf{1} \{ \lambda = 0 \} d\lambda,$$

and

$$dF(\lambda) = \frac{\gamma}{2\pi} \frac{\sqrt{((\sqrt{1/\gamma} + 1)^2 - \lambda)(\lambda - (\sqrt{1/\gamma} - 1)^2)}}{\lambda} d\lambda.$$

The second line of equation 19 holds because  $-2\|\theta^*\|_2^2 \gamma \lambda \exp(-2t\lambda) + 2\sigma^2 (\exp(-t\lambda) - \exp(-2t\lambda)) = 0$  when  $\lambda = 0$ .

We first consider the case where  $1/\gamma \geq 4$  ( $d \geq 4n$ ). Then on the support of  $\mu(\lambda)$ ,  $\lambda \in [\frac{1}{4\gamma}, \frac{9}{4\gamma}]$ . Therefore, when  $t < \frac{4}{9}\gamma \log\left(1 + \frac{\|\theta\|_2^2}{4\sigma^2}\right)$ ,  $\frac{d}{dt}\bar{R}_{X_\infty, \theta^*}(t) > 0$ ; when  $t > 4\gamma \log\left(1 + \frac{9\|\theta\|_2^2}{4\sigma^2}\right)$ ,  $\frac{d}{dt}\bar{R}_{X_\infty, \theta^*}(t) < 0$ .

We then consider the case where  $\gamma \geq 4$  ( $d < \frac{n}{4}$ ). Then on the support of  $\mu(\lambda)$ ,  $\lambda \in [\frac{1}{4}, \frac{9}{4}]$ . Therefore, when  $t < \frac{4}{9} \log\left(1 + \frac{\|\theta\|_2^2 \gamma}{4\sigma^2}\right)$ ,  $\frac{d}{dt}\bar{R}_{X_\infty, \theta^*}(t) > 0$ ; when  $t > 4 \log\left(1 + \frac{9\|\theta\|_2^2 \gamma}{4\sigma^2}\right)$ ,  $\frac{d}{dt}\bar{R}_{X_\infty, \theta^*}(t) < 0$ .

We now study the intermediate regime where  $\frac{1}{4} < \gamma < 4$ . We need the following lemma.

**Lemma 5.** For any  $t\lambda \geq 0$ , let

$$\begin{aligned} G_1(t, \lambda) &= -2\|\theta^*\|_2^2 \gamma(1 - t\lambda) + 2\sigma^2 \left(t - \frac{3}{2}t^2\lambda\right), \\ G_2(t, \lambda) &= -2\|\theta^*\|_2^2 \gamma + 2\sigma^2 \left(t - \frac{1}{2}t^2\lambda\right), \\ G_3(t, \lambda) &= -2\|\theta^*\|_2^2 \gamma(1 - t\lambda) + 2\sigma^2 t; \\ \Gamma_1(t) &= 1 - t(1 - \sqrt{1/\gamma})^2, \\ \Gamma_2(t) &= 1 - t(1 + \sqrt{1/\gamma})^2. \end{aligned}$$

Then,

$$\begin{aligned} -2\|\theta^*\|_2^2 \gamma \exp(-2t\lambda) + 2\sigma^2 \frac{\exp(-t\lambda) - \exp(-2t\lambda)}{\lambda} \\ \geq \min\{G_1(t, \lambda), G_2(t, \lambda), \Gamma_1(t)G_3(t, \lambda), \Gamma_2(t)G_3(t, \lambda)\}, \end{aligned} \quad (19)$$

and

$$\begin{aligned} -2\|\theta^*\|_2^2 \gamma \exp(-2t\lambda) + 2\sigma^2 \frac{\exp(-t\lambda) - \exp(-2t\lambda)}{\lambda} \\ \leq \max\{G_3(t, \lambda), \Gamma_1(t)G_2(t, \lambda), \Gamma_2(t)G_2(t, \lambda), \Gamma_1(t)G_3(t, \lambda), \Gamma_2(t)G_3(t, \lambda)\}. \end{aligned} \quad (20)$$

**Theorem 3.** When  $\frac{1}{4} < \gamma < 4$  and when  $\rho = (\gamma + 1) \frac{\|\theta^*\|_2^2}{\sigma^2} < 2 - \sqrt{3}$ , the risk  $\bar{R}_{X_\infty, \theta^*}(t)$  is decreasing in the interval

$$t < \frac{\gamma}{\gamma + 1} \cdot \frac{\rho}{\rho + 1}, \quad (21)$$

and is increasing in the interval

$$\frac{2\gamma}{\gamma + 1} \cdot \frac{\rho}{1 + \rho + \sqrt{(1 + \rho)^2 - 6\rho}} < t < \frac{1}{(1 + \sqrt{1/\gamma})^2}. \quad (22)$$

Therefore, the optimal early stopping time  $t^*$  satisfy:

$$\frac{\gamma}{\gamma + 1} \frac{\rho}{\rho + 1} \leq t^* \leq \frac{2\gamma}{\gamma + 1} \frac{\rho}{1 + \rho + \sqrt{(1 + \rho)^2 - 6\rho}}.$$

Theorem 3 shows that when  $\frac{1}{4} < \gamma < 4$ ,  $t = \Theta\left(\frac{\gamma}{\gamma + 1}\right) = \Theta\left(\frac{n}{n + d}\right)$  still holds, which completes the missing piece of Theorem 1.

*Proof.* By equation 19, it suffices to find the positive and negative parts of

$$\begin{aligned} & \int_{(\sqrt{1/\gamma-1})^2}^{(\sqrt{1/\gamma+1})^2} \left( -2 \|\theta^*\|_2^2 \gamma \lambda \exp(-2t\lambda) + 2\sigma^2 (\exp(-t\lambda) - \exp(-2t\lambda)) \right) dF(\lambda) \\ &= \int_{(\sqrt{1/\gamma-1})^2}^{(\sqrt{1/\gamma+1})^2} \left( -2 \|\theta^*\|_2^2 \gamma \exp(-2t\lambda) + 2\sigma^2 \frac{\exp(-t\lambda) - \exp(-2t\lambda)}{\lambda} \right) \lambda dF(\lambda), \end{aligned}$$

where  $\lambda dF(\lambda) = \frac{\gamma}{2\pi} \sqrt{((\sqrt{\gamma} + 1)^2 - \lambda)(\lambda - (\sqrt{\gamma} - 1)^2)} d\lambda$ .

We note that

$$\int_{(\sqrt{1/\gamma-1})^2}^{(\sqrt{1/\gamma+1})^2} \lambda dF(\lambda) = \lim_{d, n \rightarrow \infty} \frac{1}{d} \text{Tr} \left( \mathbb{E} \left[ \frac{1}{n} X_n^\top X_n \right] \right) = 1,$$

and that

$$\int_{(\sqrt{1/\gamma-1})^2}^{(\sqrt{1/\gamma+1})^2} \lambda^2 dF(\lambda) = \lim_{d, n \rightarrow \infty} \frac{1}{d} \mathbb{E} \left\| \frac{1}{n} (X_n^\top X_n) \right\|_F^2 = 1 + \frac{1}{\gamma}.$$

Applying these results on  $G_1(t, \lambda)$ ,  $G_2(t, \lambda)$ , and  $G_3(t, \lambda)$  (defined in Lemma 5), we obtain that

$$\begin{aligned} E_1(t) &= \int_{(\sqrt{1/\gamma-1})^2}^{(\sqrt{1/\gamma+1})^2} G_1(t, \lambda) \lambda dF(\lambda) = -2 \|\theta^*\|_2^2 \gamma (1 - t(1 + 1/\gamma)) + 2\sigma^2 \left( t - \frac{3}{2} t^2 (1 + 1/\gamma) \right), \\ E_2(t) &= \int_{(\sqrt{1/\gamma-1})^2}^{(\sqrt{1/\gamma+1})^2} G_2(t, \lambda) \lambda dF(\lambda) = -2 \|\theta^*\|_2^2 \gamma + 2\sigma^2 \left( t - \frac{1}{2} t^2 (1 + 1/\gamma) \right), \\ E_3(t) &= \int_{(\sqrt{1/\gamma-1})^2}^{(\sqrt{1/\gamma+1})^2} G_3(t, \lambda) \lambda dF(\lambda) = -2 \|\theta^*\|_2^2 \gamma (1 - t(1 + 1/\gamma)) + 2\sigma^2 t. \end{aligned}$$

From inequalities equation 19 and equation 20 in Lemma 5, we know that  $\frac{d}{dt} \bar{R}_{X_\infty, \theta^*}(t) > 0$  if

$$\min \{E_1(t), E_2(t), \Gamma_1(t)E_3(t), \Gamma_2(t)E_3(t)\} > 0,$$

and  $\frac{d}{dt} \bar{R}_{X_\infty, \theta^*}(t) < 0$  if

$$\max \{E_3(t), \Gamma_1(t)E_2(t), \Gamma_2(t)E_2(t), \Gamma_1(t)E_3(t), \Gamma_2(t)E_3(t)\} < 0.$$

**Increasing interval:** We can develop this condition further to be:  $\frac{d}{dt} \bar{R}_{X_\infty, \theta^*}(t) > 0$  if the following event happens:

$$\begin{aligned} & \{E_1(t) > 0\} \cap \{E_2(t) > 0\} \cap \left( \left( \{\Gamma_1(t) > 0\} \cap \{\Gamma_2(t) > 0\} \cap \{E_3(t) > 0\} \right) \right. \\ & \quad \left. \cup \left( \{\Gamma_1(t) < 0\} \cap \{\Gamma_2(t) < 0\} \cap \{E_3(t) < 0\} \right) \right). \end{aligned}$$

$E_1(t) > 0$  is equivalent to the condition that  $1 + \gamma > (2 + \sqrt{3}) \frac{\sigma^2}{\|\theta^*\|_2^2}$  or  $1 + \gamma < (2 - \sqrt{3}) \frac{\sigma^2}{\|\theta^*\|_2^2}$ , and that

$$\begin{aligned} & \frac{1 + \frac{\|\theta^*\|_2^2}{\sigma^2} (\gamma + 1) - \sqrt{\left(1 + \frac{\|\theta^*\|_2^2}{\sigma^2} (\gamma + 1)\right)^2 - 6 \frac{\|\theta^*\|_2^2}{\sigma^2} (\gamma + 1)}}{3(1/\gamma + 1)} \\ & < t < \frac{1 + \frac{\|\theta^*\|_2^2}{\sigma^2} (\gamma + 1) + \sqrt{\left(1 + \frac{\|\theta^*\|_2^2}{\sigma^2} (\gamma + 1)\right)^2 - 6 \frac{\|\theta^*\|_2^2}{\sigma^2} (\gamma + 1)}}{3(1/\gamma + 1)}. \end{aligned}$$

$E_2(t) > 0$  is equivalent to the condition that  $1 + \frac{1}{\gamma} < \frac{1}{2} \frac{\sigma^2}{\|\theta^*\|_2^2}$  and that

$$\frac{\gamma - \gamma \sqrt{1 - 2(\gamma + 1) \frac{\|\theta^*\|_2^2}{\sigma^2}}}{\gamma + 1} < t < \frac{\gamma + \gamma \sqrt{1 - 2(\gamma + 1) \frac{\|\theta^*\|_2^2}{\sigma^2}}}{\gamma + 1}.$$

Then the event  $\{E_1(t) > 0\} \cap \{E_2(t) > 0\}$  leads to  $\rho < 2 - \sqrt{3}$  and that  $t$  needs to satisfy the following condition for both events to happen:

$$\frac{2\gamma}{\gamma + 1} \cdot \frac{\rho}{1 + \rho + \sqrt{(1 + \rho)^2 - 6\rho}} < t < \frac{\gamma}{3(\gamma + 1)} \left(1 + \rho + \sqrt{(1 + \rho)^2 - 6\rho}\right).$$

Since

$$t > \frac{2\gamma}{\gamma + 1} \cdot \frac{\rho}{1 + \rho + \sqrt{(1 + \rho)^2 - 6\rho}} > \frac{\gamma}{\gamma + 1} \cdot \frac{\rho}{\rho + 1} = \frac{\gamma}{\rho + 1} \frac{\|\theta^*\|_2^2}{\sigma^2},$$

$E_3(t) > 0$ . Then  $\{\Gamma_1(t) > 0\} \cap \{\Gamma_2(t) > 0\}$  leads to  $t < \frac{1}{(1 + \sqrt{1/\gamma})^2}$ .

Joining the intervals, we obtain the sufficient condition for  $\frac{d}{dt} \bar{R}_{X_\infty, \theta^*}(t) > 0$  to be  $\rho = \frac{\gamma + 1}{\gamma} \frac{\|\theta^*\|_2^2}{\sigma^2} < 2 - \sqrt{3}$  and

$$\frac{2\gamma}{\gamma + 1} \cdot \frac{\rho}{1 + \rho + \sqrt{(1 + \rho)^2 - 6\rho}} < t < \frac{1}{(1 + \sqrt{1/\gamma})^2}. \quad (23)$$

**Decreasing interval:** We assume that  $\rho < 2 - \sqrt{3}$ .

$E_3(t) < 0$  leads to

$$t < \frac{\gamma}{\gamma + 1} \cdot \frac{\rho}{\rho + 1}.$$

Under this condition,  $\Gamma_1(t) > \Gamma_2(t) > 0$ . We therefore need  $E_2(t) < 0$  and  $E_3(t) < 0$ .

$E_2(t) < 0$  is equivalent to

$$t < \frac{\gamma}{\gamma + 1} \cdot \frac{2\rho}{1 + \sqrt{1 - 2\rho}} \quad \text{or} \quad t > \frac{\gamma}{\gamma + 1} \left(1 + \sqrt{1 - 2\rho}\right).$$

$E_1(t) < 0$  is equivalent to

$$t < \frac{2\gamma}{\gamma + 1} \cdot \frac{\rho}{1 + \rho + \sqrt{(1 + \rho)^2 - 6\rho}} \quad \text{or} \quad t > \frac{\gamma}{3(\gamma + 1)} \left(1 + \rho + \sqrt{(1 + \rho)^2 - 6\rho}\right).$$

Joining the intervals under the condition that  $\rho < 2 - \sqrt{3}$ , we obtain that the sufficient condition for  $\frac{d}{dt} \bar{R}_{X_\infty, \theta^*}(t) < 0$  is

$$t < \frac{\gamma}{\gamma + 1} \cdot \frac{\rho}{\rho + 1}. \quad (24)$$

□

Now, we prove Lemma 5.

*Proof of Lemma 5.* We first state an expansion result about the exponential function:

$$\begin{aligned} \exp(-t\lambda) &= 1 - t\lambda \exp(-t\lambda) \\ &= 1 - t\lambda + \frac{1}{2} t^2 \lambda^2 \exp(-t\lambda) \\ &= 1 - t\lambda + \frac{1}{2} t^2 \lambda^2 - \frac{1}{6} t^3 \lambda^3 \exp(-t\lambda). \end{aligned} \quad (25)$$

We then decompose our argument into two terms:

$$\begin{aligned} & -2 \|\theta^*\|_2^2 \gamma \exp(-2t\lambda) + 2\sigma^2 \frac{\exp(-t\lambda) - \exp(-2t\lambda)}{\lambda} \\ & = \exp(-t\lambda) \cdot \left( -2 \|\theta^*\|_2^2 \gamma \exp(-t\lambda) + 2\sigma^2 \frac{1 - \exp(-t\lambda)}{\lambda} \right). \end{aligned}$$

We use equation equation 25 to bound the following two terms that appear in the above multiplication:

$$1 - t\lambda \leq \exp(-t\lambda) \leq 1, \quad (26)$$

and

$$\begin{aligned} -2 \|\theta^*\|_2^2 \gamma + 2\sigma^2 \left( t - \frac{1}{2} t^2 \lambda \right) & \leq -2 \|\theta^*\|_2^2 \gamma \exp(-t\lambda) + 2\sigma^2 \frac{1 - \exp(-t\lambda)}{\lambda} \\ & \leq -2 \|\theta^*\|_2^2 \gamma (1 - t\lambda) + 2\sigma^2 t. \end{aligned}$$

Combining equations equation 26 and equation C, we obtain that

$$\begin{aligned} & -2 \|\theta^*\|_2^2 \gamma \exp(-2t\lambda) + 2\sigma^2 \frac{\exp(-t\lambda) - \exp(-2t\lambda)}{\lambda} \\ & \geq \min \left\{ (1 - t\lambda) \left( -2 \|\theta^*\|_2^2 \gamma + 2\sigma^2 \left( t - \frac{1}{2} t^2 \lambda \right) \right), -2 \|\theta^*\|_2^2 \gamma + 2\sigma^2 \left( t - \frac{1}{2} t^2 \lambda \right), \right. \\ & \quad \left. (1 - t\lambda) \left( -2 \|\theta^*\|_2^2 \gamma (1 - t\lambda) + 2\sigma^2 t \right) \right\} \\ & \geq \min \left\{ -2 \|\theta^*\|_2^2 \gamma (1 - t\lambda) + 2\sigma^2 \left( t - \frac{3}{2} t^2 \lambda \right), -2 \|\theta^*\|_2^2 \gamma + 2\sigma^2 \left( t - \frac{1}{2} t^2 \lambda \right), \right. \\ & \quad \left. (1 - t(1 + \sqrt{\gamma})) \left( -2 \|\theta^*\|_2^2 \gamma (1 - t\lambda) + 2\sigma^2 t \right), (1 - t(1 - \sqrt{\gamma})) \left( -2 \|\theta^*\|_2^2 \gamma (1 - t\lambda) + 2\sigma^2 t \right) \right\}, \end{aligned}$$

where we have used the fact that  $\lambda \in [(1 - \sqrt{1/\gamma})^2, (1 + \sqrt{1/\gamma})^2]$ . On the flip side, we obtain that

$$\begin{aligned} & -2 \|\theta^*\|_2^2 \gamma \exp(-2t\lambda) + 2\sigma^2 \frac{\exp(-t\lambda) - \exp(-2t\lambda)}{\lambda} \\ & \leq \max \left\{ -2 \|\theta^*\|_2^2 \gamma (1 - t\lambda) + 2\sigma^2 t, (1 - t\lambda) \left( -2 \|\theta^*\|_2^2 \gamma + 2\sigma^2 \left( t - \frac{1}{2} t^2 \lambda \right) \right), \right. \\ & \quad \left. (1 - t\lambda) \left( -2 \|\theta^*\|_2^2 \gamma (1 - t\lambda) + 2\sigma^2 t \right) \right\} \\ & \leq \max \left\{ -2 \|\theta^*\|_2^2 \gamma (1 - t\lambda) + 2\sigma^2 t, (1 - t(1 + \sqrt{1/\gamma})) \left( -2 \|\theta^*\|_2^2 \gamma + 2\sigma^2 \left( t - \frac{1}{2} t^2 \lambda \right) \right), \right. \\ & \quad \left. (1 - t(1 - \sqrt{1/\gamma})) \left( -2 \|\theta^*\|_2^2 \gamma + 2\sigma^2 \left( t - \frac{1}{2} t^2 \lambda \right) \right), (1 - t(1 + \sqrt{1/\gamma})) \left( -2 \|\theta^*\|_2^2 \gamma (1 - t\lambda) + 2\sigma^2 t \right), \right. \\ & \quad \left. (1 - t(1 - \sqrt{1/\gamma})) \left( -2 \|\theta^*\|_2^2 \gamma (1 - t\lambda) + 2\sigma^2 t \right) \right\}. \end{aligned}$$

For

$$G_1(t, \lambda) = -2 \|\theta^*\|_2^2 \gamma (1 - t\lambda) + 2\sigma^2 \left( t - \frac{3}{2} t^2 \lambda \right),$$

$$G_2(t, \lambda) = -2 \|\theta^*\|_2^2 \gamma + 2\sigma^2 \left( t - \frac{1}{2} t^2 \lambda \right),$$

$$G_3(t, \lambda) = -2 \|\theta^*\|_2^2 \gamma (1 - t\lambda) + 2\sigma^2 t;$$

$$\begin{aligned}\Gamma_1(t) &= 1 - t(1 - \sqrt{1/\gamma})^2, \\ \Gamma_2(t) &= 1 - t(1 + \sqrt{1/\gamma})^2,\end{aligned}$$

we have

$$\begin{aligned}-2 \|\theta^*\|_2^2 \gamma \exp(-2t\lambda) + 2\sigma^2 \frac{\exp(-t\lambda) - \exp(-2t\lambda)}{\lambda} \\ \geq \min \{G_1(t, \lambda), G_2(t, \lambda), \Gamma_1(t)G_3(t, \lambda), \Gamma_2(t)G_3(t, \lambda)\},\end{aligned}$$

and

$$\begin{aligned}-2 \|\theta^*\|_2^2 \gamma \exp(-2t\lambda) + 2\sigma^2 \frac{\exp(-t\lambda) - \exp(-2t\lambda)}{\lambda} \\ \leq \max \{G_3(t, \lambda), \Gamma_1(t)G_2(t, \lambda), \Gamma_2(t)G_2(t, \lambda), \Gamma_1(t)G_3(t, \lambda), \Gamma_2(t)G_3(t, \lambda)\}.\end{aligned}$$

□

## D TECHNICAL LEMMAS

**Lemma 6** (Davidson & Szarek (2001)). *Let  $A$  be a  $m \times N$  matrix with  $N \geq m$  whose entries are real, independent Gaussian random variables following  $\mathcal{N}(0, 1)$ . Let  $\sigma_1(A) \geq \sigma_2(A) \geq \dots \geq \sigma_m(A)$  be the singular values of  $A$ . Then,*

$$\Pr \left[ \sqrt{N} - \sqrt{m} - t \leq \sigma_1(A) \leq \sigma_m(A) \leq \sqrt{N} + \sqrt{m} + t \right] \leq 1 - 2 \exp(-t^2/2).$$

**Lemma 7.** *Let  $A$  be a  $m \times N$  matrix whose entries are real, independent Gaussian random variables following  $\mathcal{N}(0, 1)$ . Let  $\lambda_1(A) \geq \lambda_2(A) \geq \dots \geq \lambda_N(A)$  be the eigenvalues of  $\frac{1}{m} A^\top A$ . Then,*

$$\mathbb{E} \left[ \sum_{i=1}^N \lambda_i \right] = N$$

and

$$\mathbb{E} \left[ \sum_{i=1}^N \lambda_i^2 \right] = \frac{N}{m} (m + N + 2).$$

*Proof.* We first note that

$$m \mathbb{E} \left[ \sum_{i=1}^d \lambda_i \right] = \mathbb{E} \left[ \text{Tr} \left( \tilde{A}^\top \tilde{A} \right) \right] = \mathbb{E} \left[ \left\| \tilde{A} \right\|_F^2 \right] = mN.$$

We then note that

$$m^2 \mathbb{E} \left[ \sum_{i=1}^d \lambda_i^2 \right] = \mathbb{E} \left[ \text{Tr} \left( \left( \tilde{A}^\top \tilde{A} \right)^2 \right) \right] = \mathbb{E} \left[ \left\| \tilde{A}^\top \tilde{A} \right\|_F^2 \right],$$

where

$$\left( \tilde{A}^\top \tilde{A} \right)_{i,j} = \sum_{k=1}^n \tilde{a}_{k,i} \tilde{a}_{k,j}.$$

Therefore,

$$\mathbb{E} \left[ \left\| \tilde{A}^\top \tilde{A} \right\|_F^2 \right] = \sum_{i=1}^d \sum_{i=1}^d \mathbb{E} \left[ \left( \sum_{k=1}^n \tilde{a}_{k,i} \tilde{a}_{k,j} \right)^2 \right] = \sum_{i=1}^d \sum_{i=1}^d \sum_{k=1}^n \sum_{l=1}^n \mathbb{E} [\tilde{a}_{k,i} \tilde{a}_{k,j} \tilde{a}_{l,i} \tilde{a}_{l,j}].$$

Since all the odd moments of  $x$  are zero, the nonzero elements in the above equation appear when  $k = l$  or when  $i = j$ . Hence the number of nonzero elements is  $mN(m + N)$ . When either  $k = l$  or  $i = j$  happens, the second order moment of  $\tilde{a}$  is 1. When both  $k = l$  and  $i = j$  happen, the fourth order moment of  $\tilde{a}$  is 3. Therefore,

$$\mathbb{E} \left[ \sum_{i=1}^d \lambda_i^2 \right] = \frac{1}{m^2} \sum_{i=1}^d \sum_{i=1}^d \sum_{k=1}^n \sum_{l=1}^n \mathbb{E} [\tilde{a}_{k,i} \tilde{a}_{k,j} \tilde{a}_{l,i} \tilde{a}_{l,j}] = \frac{N}{m} (m + N + 2).$$

□

**Lemma 8** (Cauchy Interlacing Theorem (Corollary 4.3.9, Horn & Johnson (2012))). *Let  $A \in \mathbb{R}^{m \times m}$  be a symmetric matrix. Let  $\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_n$  be the eigenvalues of the matrix  $A$  and let  $z \in \mathbb{R}^m$  be a nonzero vector. Then,*

$$\begin{aligned}\lambda_i(A) &\leq \lambda_i(A + zz^\top) \leq \lambda_{i+1}(A), \quad i = 1, \dots, n-1 \\ \lambda_n(A) &\leq \lambda_n(A + zz^\top).\end{aligned}$$

## E DISCRETIZATION ERRORS

In practice, one would discretize the algorithm and perform gradient descent, instead of using the gradient flow, to optimize over empirical loss. This can be formulated in the following expression:

$$\beta_{k+1} = \beta_k - \frac{h}{n} X^\top X \beta_k + \frac{h}{n} X^\top y,$$

where  $h$  is the step size of the algorithm. A solution to the above iterative equation, starting at  $\beta_0$  is:

$$\beta_k = \left( \mathbf{I} - \frac{h}{n} X^\top X \right)^k \beta_0 + \frac{h}{n} \left( \sum_{i=0}^{k-1} \left( \mathbf{I} - \frac{h}{n} X^\top X \right)^i \right) X^\top y \quad (27)$$

$$= \left( \mathbf{I} - \frac{h}{n} X^\top X \right)^k \beta_0 + (X^\top X)^\dagger \left( \mathbf{I} - \left( \mathbf{I} - \frac{h}{n} X^\top X \right)^k \right) X^\top y. \quad (28)$$

Taking  $\beta_0 = 0$ , we obtain that for any  $h < 1/s_{\max}$ , where  $s_{\max}$  is the largest eigenvalue of  $X^\top X/n$ ,

$$\sup_k \left\| \beta_k - \hat{\beta}(kh) \right\| = (X^\top X)^\dagger \left( \exp\left(-\frac{kh}{n} X^\top X\right) - \left( \mathbf{I} - \frac{h}{n} X^\top X \right)^k \right) X^\top y \quad (29)$$

$$\leq \frac{h}{2n} \|X^\top y\| (\exp(hs_{\max}) - 1), \quad (30)$$

since  $\sup_{\tau > 0} |e^{-l\tau} - (1-l)^\tau| \leq \frac{l}{2} (e^l - 1)$ ,  $\forall l < 1$ .

## F ADDITIONAL DETAILS OF EXPERIMENTS

### F.1 LINEAR REGRESSION

(This was previously in the main paper.)

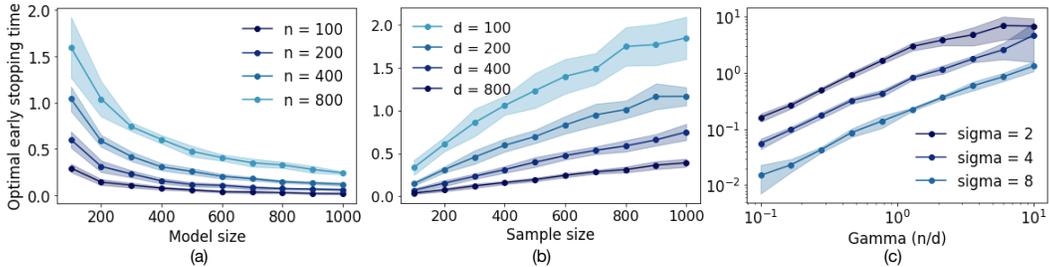


Figure 6: (a) Optimal early stopping time decreases with increasing model sizes. (b) Optimal early stopping time increases with increasing sample sizes. (c) Optimal early stopping time increases with increasing  $\gamma$ . Shaded areas give the empirical bounds generated by 10 runs of the experiments.

**Overparameterized optimal early stopping time.** In this experiment (Figure 6), we follow the setting  $S_{\text{over}}$ . We use  $p = 10$  for all  $n$  and  $d$ . **The optimal early stopping time is selected as the time with the first local minimal loss.** We observe from these three plots that the optimal early stopping time decreases with larger model size  $d$ , and increases with more samples  $n$ . Define  $\gamma = n/d$  to be the sample-model size ratio with various choices of  $d, n$ .  $t_{\text{opt}}$  increases as  $\gamma$  increases and decreases with a larger level of noise.

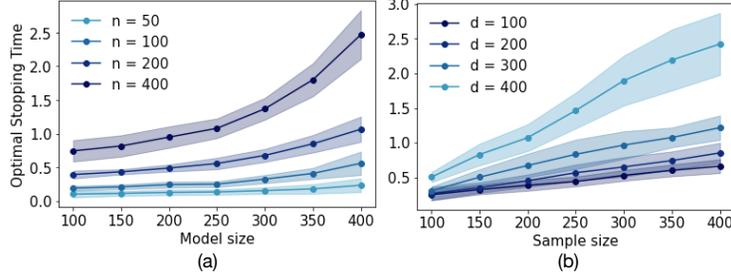


Figure 7: (a) Underparameterization optimal early stopping time increases with increasing model sizes. (b) Underparameterization optimal early stopping time increases with increasing sample sizes. Shaded areas give the empirical bounds generated by 10 runs of the experiments.

**Underparameterized optimal early stopping time.** In this experiment (Figure 7), we follow the setting  $\mathcal{S}_{\text{under}}$ . We use  $p = 400$  for all  $n$  and  $d$ . The optimal early stopping time increases with increasing model sizes  $d$ , and increasing sample sizes  $n$ .

**Risk monotonicity of underparameterized and overparameterized optimal early stopping.**

In Figure 8, for overparameterized models, we observe that the optimal early stopping risk decreases with increasing sample sizes  $n$ . For underparameterized models, the optimal early stopping risk decreases with both increasing  $d$  and  $n$ . This is consistent with the risk monotonicity analysis.

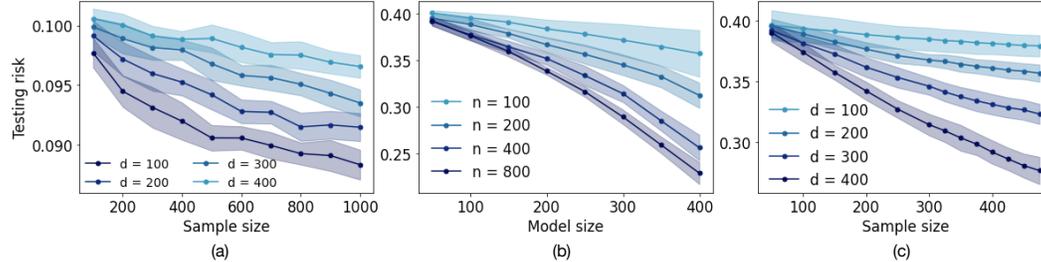


Figure 8: (a) Overparameterization testing risk at optimal early stopping time decreases with sample sizes. (b) (c) Underparameterization testing risk at optimal early stopping time decreases with model sizes and sample sizes.

**Testing loss vs time** Figure 9 and Figure 10 plot the testing loss as a function of time in the overparameterized and the underparameterized setting respectively. The solid dots on the curves represent the points with lowest testing loss. We use  $\theta^* \sim \mathcal{N}(0, I_d)$  and a randomly sample semi-orthogonal matrix  $P \in \mathbb{R}^{p \times d}$ . We choose  $p = 10, \sigma = 6$  and  $p = 500, \sigma = 2$  in Figure 9 and Figure 10. We can observe from Figure 9 the optimal early stopping time decreases with increasing  $d$  and increases with increasing  $n$ . Moreover, testing loss decreases with increasing  $n$ . We can observe from Figure 10 the optimal early stopping time increases with both  $d$  and  $n$ .

**Theoretical bounds on early optimal stopping time** We compare our theoretical bounds with the empirical bounds on the linear regression problem. In the overparameterization setting, we show in Theorem 1 that the optimal early stopping time scales as  $\tilde{\Theta}(\frac{n}{n+d})$ . In Figure 11, the yellow area shows the region  $(\frac{an}{n+d}, \frac{bn}{n+d})$  for some constant  $a$  and  $b$ . The blue area is the standard deviation of 20 runs of the experiments. We set  $p = 10, \sigma = 2, n = 400, d = [200, 300, \dots, 1000], a = 1.2, b = 2.7$  in Figure 11a,  $p = 10, \sigma = 2, n = [200, 300, \dots, 1000], d = 400, a = 2.5, b = 7$  in Figure 11b,  $p = 10, \sigma = 2 \gamma \in [0.1, 10], a = 0.5, b = 4.5$  in log-scale in Figure 11c.

In Figure 12, we plot the theory bounds of underparameterization setting. We have shown in Theorem 2 that the optimal early stopping time scales with  $\tilde{\Theta}(\frac{n\|\theta^*\|_2^2}{p\sigma^2 + (p-d)\|\theta^*\|_2^2})$ . The yellow area shows the theory bounds  $(\frac{an\|\theta^*\|_2^2}{p\sigma^2 + (p-d)\|\theta^*\|_2^2}, \frac{bn\|\theta^*\|_2^2}{p\sigma^2 + (p-d)\|\theta^*\|_2^2})$  for some constant  $a$  and  $b$ . The blue area is the the standard deviation of 20 runs of the experiments. We set  $p = 500, \sigma = 5, n =$

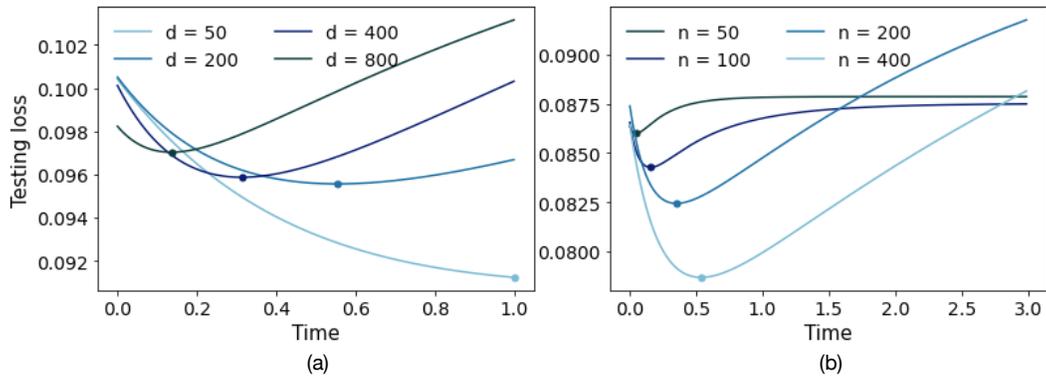


Figure 9: (a) Testing loss behaviour for increasing training time with various model sizes. (b) Testing loss behaviour for increasing training time with various number of samples.

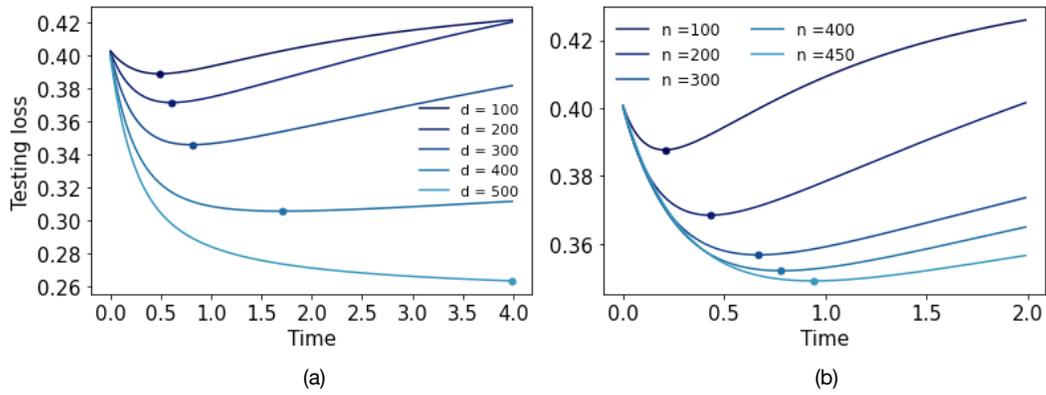


Figure 10: (a) Testing loss behaviour for increasing training time with various model sizes. (b) Testing loss behaviour for increasing training time with various sample sizes.

1000,  $d = [50, 100, 150, \dots, 400]$ ,  $a = 0.25, b = 0.85$  for Figure 12a, and  $p = 500, \sigma = 5, n = [50, 100, 150, \dots, 1000]$ ,  $d = 250, a = 0.3, b = 1.1$  for Figure 12b.

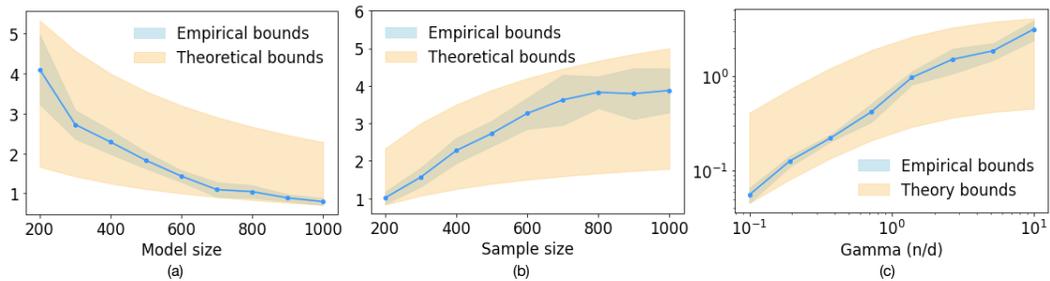


Figure 11: (a) Optimal early stopping time decreases with increasing model sizes. (b) Optimal early stopping time increases with increasing sample sizes. (c) Optimal early stopping time increases with  $\frac{n}{d}$ .

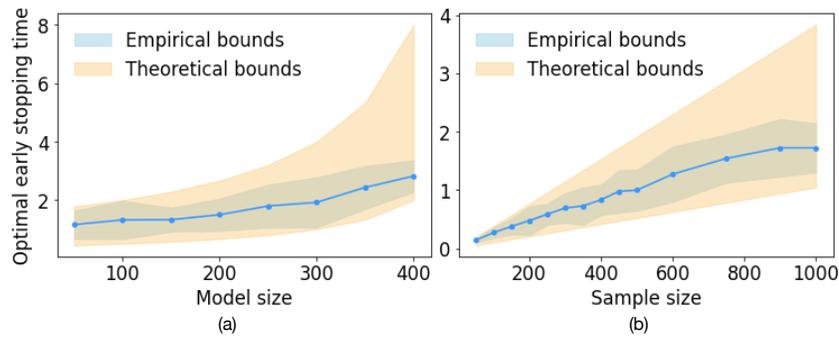


Figure 12: (a) Optimal early stopping time increases with increasing model sizes with bounds. (b) Optimal early stopping time increases with increasing sample sizes with bounds.

## F.2 ADDITIONAL PLOTS IMAGE CLASSIFICATION USING DEEP NEURAL NETWORKS

In Figure 13, we show the testing error as a function of epoch for training MNIST, CIFAR-10 and CIFAR-100 using models of different width. We can observe the double descent phenomenon and the varying optimal early stopping time for different models. In Figure 14, we show the testing error as a function of sample size for training MNIST, CIFAR-10 and CIFAR-100. We observe that the test error decreases for increasing sample size.

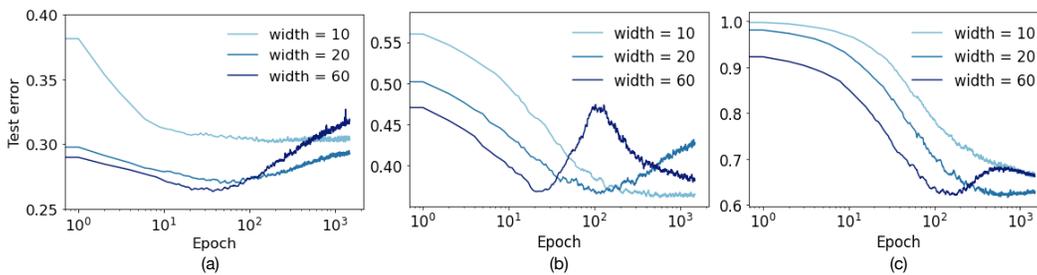


Figure 13: Epoch-wise double descent and optimal early stopping time for (a) MNIST (b) CIFAR-10 (c) CIFAR-100.

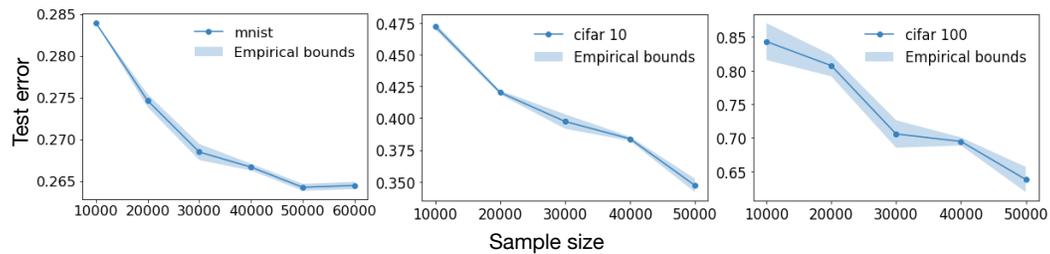


Figure 14: Test error decreases with increasing sample size.

(We added some additional experimental results below. Due to time constraints, we only run experiments for once. We will add more experimental results in the final version.)

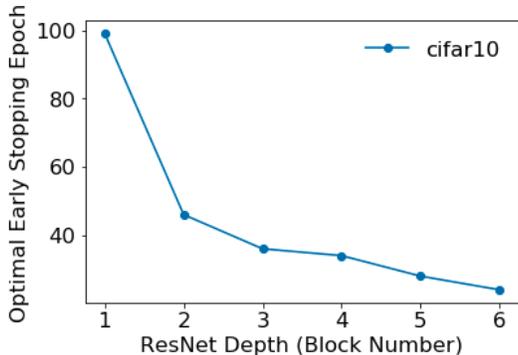


Figure 15: Optimal Early Stopping time decreases with increasing depth.

**ResNet Depths** In Figure 15, we show that the optimal stopping time decreases with increasing ResNet depth. We train CIFAR-10 using ResNet. We add 20% of label noise to all datasets. We use a family of ResNet networks with ResNet blocks of widths  $d$ . Suppose we have a ResNet with depth  $k$ , we will have  $k$  sequential ResNet blocks with width  $[2^0 d, 2^1 d, \dots, 2^k d]$  for different depth layers. The number of parameters grow from 100K ( $k = 1$ ) to 300M ( $k = 6$ ). Inputs are normalized to  $[-1, 1]$  with standard data-augmentation. We use stochastic gradient descent with cross-entropy loss, learning rate  $\eta = 0.1$ , and minibatch size  $B = 512$ . We choose the optimal stopping time similarly as in Section 4.2. Increasing the depth of ResNet would similarly enlarge the network parameter size, resulting an observable decrease of optimal early stopping time.

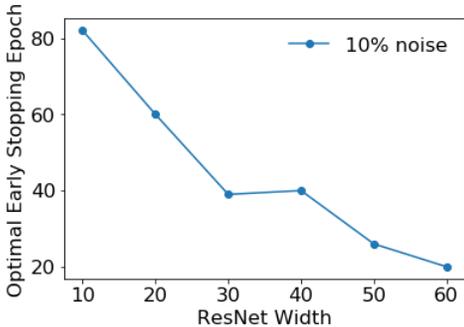


Figure 16: Optimal stopping time vs. model width with 10% label noise.

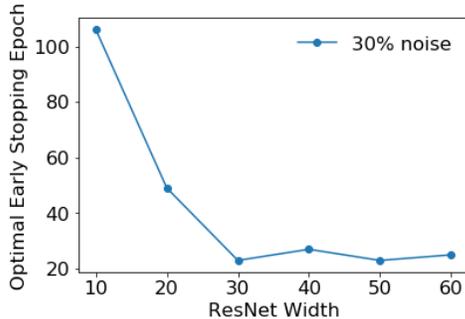


Figure 17: Optimal stopping time vs. model width with 30% label noise.

**Label Noise Levels** In Figure 16 and Figure 17, we show that the optimal stopping time decreases with increasing model width holds with 10% and 30 % label noise when training on CIFAR-10 using ResNet. The variance of each independent experimental trail slightly breaks the monotonicity, yet the overall trend is clear. We use a family of ResNet networks with convolutional layers of widths  $[d, 2d, 4d, 8d]$  for different layer, and we set ResNet width from 10 to 60. Inputs are normalized to  $[-1, 1]$  with standard data-augmentation. We use stochastic gradient descent with cross-entropy loss, learning rate  $\eta = 0.1$ , and minibatch size  $B = 512$ . We choose the optimal stopping time similarly as in Section 4.2.

## G CODE FOR REPRODUCIBILITY

All code and experiments in our paper are available: <https://anonymous.4open.science/r/On-Optimal-Early-Stopping-0E84>. Part of our code builds on open source code from (Nakkiran et al., 2020a).