

How to make your optimizer generalize better

Sharan Vaswani
Reza Babenzhad
Jose Gallego
Aaron Mishkin
Simon Lacoste-Julien
Nicolas Le Roux

UNIVERSITY OF ALBERTA
SAIT AI LAB, MONTREAL
MILA, UNIVERSITÉ DE MONTRÉAL
STANFORD UNIVERSITY
MILA, UNIVERSITÉ DE MONTRÉAL
MILA, GOOGLE BRAIN

Abstract

We study the implicit regularization of optimization methods for linear models *interpolating* the training data in the under-parametrized and over-parametrized regimes. For over-parameterized linear regression, where there are infinitely many interpolating solutions, different optimization methods can converge to solutions with varying generalization performance. In this setting, we show that projections onto linear spans can be used to move between solutions. Furthermore, via a simple reparameterization, we can ensure that an arbitrary optimizer converges to the minimum ℓ_2 -norm solution with favourable generalization properties. For under-parameterized linear classification, optimizers can converge to different decision boundaries separating the data. We prove that for any such classifier, there exists a family of quadratic norms $\|\cdot\|_{\mathbf{P}}$ such that the classifier's direction is the same as that of the maximum \mathbf{P} -margin solution. We argue that analyzing convergence to the standard maximum ℓ_2 -margin is arbitrary and show that minimizing the norm induced by the data can result in better generalization. We validate our theoretical results via experiments on synthetic and real datasets.

1. Introduction

Modern machine learning has seen the rise of large over-parameterized models such as deep neural networks [16]. These models are highly expressive and are able to fit or *interpolate* all the training data [7, 8, 41]. Since the number of parameters is much larger than the size of the training dataset, there are infinitely many solutions that can fit the data. These solutions can have vastly different generalization performance and the optimization method employed to minimize the training loss also influences the test performance [3, 17, 24, 34, 40]. This is in contrast to classical regularized, under-parameterized models where there is a unique solution and the optimization method is solely responsible for converging to this solution at an appropriate rate.

A recent line of work [3, 17–19, 22, 30, 31, 39] studies the *implicit regularization* of optimization methods in simplified settings. The implicit regularization of an optimizer biases it towards specific types of minimizers that are preferred amongst the infinite set of solutions. Two such simplified settings studied in this literature are over-parameterized linear regression with the squared

loss [17] and linear classification on separable data using losses with an exponential tail [39]. In each of these settings, recent works studies numerous optimization methods characterizing their implicit bias towards the minimum ℓ_2 norm or maximum ℓ_2 margin solutions respectively. Even in these simple scenarios, it is difficult to analyze the solutions of common optimization methods such as Adagrad [14]. Moreover, recent works [3, 35] have shown that it is not possible to derive closed-form expressions in more challenging settings like matrix factorization.

Perhaps more importantly, these works on implicit regularization do not give any practical guidance when choosing an optimizer even for over-parameterized linear problems. On the other hand, from an optimization perspective, generalization performance is rarely considered when designing a new optimization method. Consequently, numerous methods (for example [9, 32]) that are robust and have good training performance do not generalize well and are not used for training over-parameterized models. Only few recent works [11, 24, 27] consider the generalization performance when designing an optimization method, however, these methods are heuristic and do not show a provable improvement in the generalization. We address the following questions:

When does an optimizer converge to solutions that generalize well for over-parameterized linear regression and classification? Can we design techniques that ensure that optimization methods converge to interpolating solutions with good generalization performance.

1.1. Background and Contributions

We consider linear models interpolating the training data in the under-parametrized and over-parametrized regimes. In particular, we study the implicit bias of optimizers in the linear regression (Section 2) and classification (Sections 3) settings.

Over-parametrized linear regression: In [17], the authors study the implicit bias of gradient descent and its accelerated variants and show that it converges to the minimum ℓ_2 -norm solution. More generally, they characterize the implicit regularization for the steepest descent and mirror descent algorithms. In Section 2, we prove that every interpolating solution has a corresponding quadratic norm that it minimizes. This result enables the use of projections onto linear spans to move between interpolating solutions. This further implies that for any interpolating solution found by an optimizer, a projection onto the data-span recovers the min-norm solution and can potentially improve the optimizer’s generalization performance. Since such projections can be computationally expensive, we show that a simple reparameterization can ensure that an arbitrary optimization method converges to the minimum ℓ_2 norm solution. We investigate, both theoretically and empirically, whether it is possible to find a norm that generalizes better than the ℓ_2 norm.

Linear classification: The implicit bias of gradient descent minimizing losses with an exponential tail has been studied in [30, 31, 39]. In these works, it was shown that the direction of the gradient descent (GD) solution converges to the max-margin solution at a $1/\log(T)$ rate, where T is the number of GD iterations. In [17], the authors also outline the implicit regularization properties of steepest descent, whereas the implicit bias of Adagrad is studied in [34]. In Section 3, we first show that for any linear classifier perfectly separating the data, there exists a family of corresponding quadratic norms $\|\cdot\|_{\mathbf{P}}$ such that the classifiers direction is the same as that of the maximum

P-margin solution. We then show, both theoretically and empirically, that the max-margin solution in the norm induced by the data results in better generalization than the well-studied ℓ_2 -norm. This result implies that it is important to consider the properties of the data when reasoning about the implicit bias and generalization performance of methods. For the exponential loss, we propose heuristics that can bias optimization methods towards solutions with good generalization properties and can potentially improve their test performance. In Section 4, we validate our theoretical results on synthetic and real datasets.

2. Over-parametrized linear regression

We first consider over-parameterized linear regression with n points $\{\mathbf{x}_i, y_i\}$ consisting of d -dimensional feature vectors \mathbf{x}_i s.t. $d > n$ and the corresponding labels/measurements y_i . We make the standard assumption [6] that the true labels are corrupted with noise $\epsilon \sim \mathcal{N}(0, \sigma^2 \mathbf{I}_d)$, implying that $\mathbf{y} = \mathbf{X}\mathbf{w}^* + \epsilon$ where $\mathbf{X} \in \mathbb{R}^{n \times d}$ is the matrix of features, $\mathbf{y} \in \mathbb{R}^n$ is the vector of labels and $\mathbf{w}^* \in \mathbb{R}^d$ is the “true“ data-generating vector. We seek to minimize the squared loss, $\min_{\mathbf{w}} f(\mathbf{w}) := \|\mathbf{X}\mathbf{w} - \mathbf{y}\|_2^2$. If the matrix $\mathbf{X}\mathbf{X}^\top$ is full rank, there are infinitely many solutions that can *interpolate* or exactly fit the training dataset. Consequently, optimization methods achieving zero training loss converge to different solutions that can have vastly different generalization properties. We study the generalization performance of such interpolating solutions \mathbf{w}_{opt} , that is $\mathbf{X}\mathbf{w}_{\text{opt}} = \mathbf{y}$.

2.1. Convergence to the minimum norm solution

In the over-parameterized regime, given that $\mathbf{X}\mathbf{X}^\top$ is invertible, it is known that gradient descent (GD) initialized at the origin converges to the minimum ℓ_2 norm solution \mathbf{w}_{mn} , henceforth referred to as the *min-norm solution*.

$$\mathbf{w}_{\text{mn}} = \arg \min_{\mathbf{w}} \frac{1}{2} \|\mathbf{w}\|_2^2 \quad \text{s.t.} \quad \mathbf{X}\mathbf{w} = \mathbf{y} \quad \implies \quad \mathbf{w}_{\text{mn}} = \mathbf{X}^\top (\mathbf{X}\mathbf{X}^\top)^{-1} \mathbf{y} \quad (1)$$

The generalization properties of the min-norm solution have been thoroughly studied in the under-parameterized [5] and more recently the over-parameterized interpolation regime [6, 20, 29]. The min-norm solution is the unique point that interpolates the data and lies in the span of the feature vectors. This property has been used to analyze the implicit regularization of common optimization methods [17]. Note that this property is unaffected by using stochastic gradients of a finite-sum, implying that mini-batch variants of optimizers have the same implicit regularization. In Appendix A.7, we prove that iterates of the Newton method with Levenberg-Marquardt regularization [26, 28] lie in the span of the training data, implying that it converges to the min-norm solution. Similarly, we use this property to prove that full-matrix variants [1] (without the diagonal approximation) of adaptive gradient methods like Adagrad [14], Adam [25] also converge to the min-norm solution (Appendix A.8). These methods are more robust to the step-size than GD and have better empirical convergence, implying that their full-matrix variants converge faster but generalize as well as GD. However, with the commonly used diagonal approximation, these methods

result in iterates that do not lie in the data-span and are consequently not guaranteed to converge to the min-norm solution [17].

Preconditioned gradient descent (PGD) is the simplest method whose iterates do not lie in the span of the training data, but whose implicit regularization can be analyzed. The iterates corresponding to PGD with a constant positive-definite preconditioner \mathbf{P} and constant step-size η can be written as: $\mathbf{w}_{k+1} = \mathbf{w}_k - \eta \mathbf{P} \nabla f(\mathbf{w})$. In Appendix A.5, we show that PGD converges to the solution $\mathbf{w}_{\text{PGD}} := \lim_{k \rightarrow \infty} \mathbf{w}_k = \mathbf{w}_0 + \mathbf{P} \mathbf{X}^\top (\mathbf{X} \mathbf{P} \mathbf{X}^\top)^{-1} [\mathbf{y} - \mathbf{X} \mathbf{w}_0]$. Furthermore, when $\mathbf{w}_0 = 0$, \mathbf{w}_{PGD} is the minimum \mathbf{P}^{-1} -norm solution, meaning that $\mathbf{w}_{\text{PGD}} = \arg \min \frac{1}{2} \|\mathbf{w}\|_{\mathbf{P}^{-1}}^2 = \frac{1}{2} \mathbf{w}^\top \mathbf{P}^{-1} \mathbf{w}$ such that $\mathbf{X} \mathbf{w} = \mathbf{y}$. We note that this result can also be seen as a consequence of Theorem 1 of [17]. In Lemma 6 of Appendix A.6, we prove that the PGD solution is the unique point that interpolates the data and lies in the transformed data-span, that is $\mathbf{w}_{\text{PGD}} \in \text{span}(\mathbf{P} \mathbf{X}^\top)$. These results imply that the properties of PGD are equivalent to that of GD, differing only in the norm. In fact, we can show that a similar equivalence holds for general optimization methods. This enables us to express the solution of an arbitrary optimization method in terms of a PGD solution for a family of constant preconditioners. It implies that an interpolating solution found by an optimization method is the unique minimum-norm solution in the \mathbf{P}_{opt} norm. We now show that such an equivalence enables the use of projection operators to move between interpolating solutions.

Proposition 1 *Consider two optimization methods, their respective interpolating solutions \mathbf{w}_1 and \mathbf{w}_2 and equivalent preconditioners \mathbf{P}_1 and \mathbf{P}_2 , constructed according to Lemma 4 in Appendix A.1. Projecting \mathbf{w}_1 onto the span($\mathbf{P}_2 \mathbf{X}^\top$) using the operator $\pi = \mathbf{P}_2 \mathbf{X}^\top (\mathbf{X} \mathbf{P}_2 \mathbf{X}^\top)^{-1} \mathbf{X}$ recovers \mathbf{w}_2 .*

The above result implies that projecting any optimizer’s interpolating solution onto the data-span recovers the min-norm solution and can potentially improve its generalization performance. Proposition 1 can be generalized to an arbitrary iterate, in that the same projection operator π moves \mathbf{w}_1 to the span($\mathbf{P}_2 \mathbf{X}^\top$) without changing its training loss. The computational cost of such a projection is equal to that for solving the normal equations and approximate projection methods are required for making this a practical approach.

Instead of an expensive projection onto the data-span, we can reparameterize the problem to ensure that any optimization method will converge to the min-norm solution. Specifically, since \mathbf{w}_{mn} is the unique interpolating solution in the data-span, let us seek an interpolating solution of the form $\mathbf{w} = \mathbf{X}^\top \alpha$ where $\alpha \in \mathbb{R}^n$. The corresponding minimization problem is $\min_{\alpha} \frac{1}{2} \|\mathbf{X} \mathbf{X}^\top \alpha - \mathbf{y}\|^2 = \min_{\alpha} \frac{1}{2} \|\mathbf{K} \alpha - \mathbf{y}\|^2$ where $\mathbf{K} = \mathbf{X} \mathbf{X}^\top$ is an $n \times n$ full-rank matrix. Note that \mathbf{K} corresponds to the *kernel matrix* under a linear kernel. This reparameterized objective is strictly convex w.r.t. α and thus admits a unique solution $\alpha^* = \mathbf{K}^{-1} \mathbf{y}$ regardless of the optimization method. In the original \mathbf{w} -space, this solution corresponds to $\mathbf{X}^\top \alpha^* = \mathbf{X}^\top (\mathbf{X} \mathbf{X}^\top)^{-1} \mathbf{y} = \mathbf{w}_{\text{mn}}$. This implies that any optimization method solving the reparameterized problem is ensured to converge to the min-norm solution with favourable generalization properties. In fact, the above reasoning is true for any generalized linear model with a loss function of the form $f(\mathbf{X} \mathbf{w})$ including the logistic/exponential and squared-hinge losses we consider in Section 3.

Since different norms have different generalization properties, we attempt to find a norm that results in better generalization than the minimum ℓ_2 norm solution. In Appendix A.10, we gen-

eralize the excess risk bounds in [6] to analyze solutions found by PGD. In Appendix A.10.1, we optimize an upper bound on the risk w.r.t \mathbf{P} . Even assuming full knowledge of the covariance matrix of the data, we show that it is not possible to uniformly improve over the min-norm solution in the noiseless case ($\sigma = 0$). We can obtain a better upper bound on the risk in the noisy case, however, we empirically demonstrate that the looseness of this bound prevents obtaining better generalizing solutions in practice.

3. Under-parameterized linear classification

In this section, we consider binary classification with a training dataset $\{\mathbf{x}_i, y_i\}_{i=1}^n$ of d -dimensional feature vectors \mathbf{x}_i and labels $y_i \in \{-1, 1\}$, with $d \leq n$. We seek to find a hyperplane \mathbf{w}^* that minimizes the 0-1 loss, $\mathbf{w}^* = \arg \min_{\mathbf{w}} \sum_{i=1}^n \mathbb{I}\{y_i \langle \mathbf{w}, \mathbf{x}_i \rangle \geq 0\}$ where \mathbb{I} is the indicator function equal to 1 when true and 0 otherwise¹. Unlike under-parameterized regression that has a unique minimizer, there can be infinitely many linear classifiers or hyperplanes that separate the data. We study the interpolation setting where the data is linearly separable by a non-zero margin, implying that there exist linear classifiers with *zero training error* or zero 0-1 loss. Similar to over-parameterized regression, optimization methods achieving zero training error are biased towards certain solutions and can converge to hyperplanes that have different generalization properties.

For a general positive definite matrix \mathbf{P} , if the data is separable by a margin (in the \mathbf{P} norm) equal to γ , the maximum \mathbf{P} -margin solution $\mathbf{w}_{\text{mm},\mathbf{P}}$ has the following equivalent forms²:

$$\mathbf{w}_{\text{mm},\mathbf{P}} := \arg \max_{\|\mathbf{w}\|_{\mathbf{P}} \leq 1/\gamma} \min_{i \in [n]} \langle \mathbf{w}, \mathbf{x}_i \rangle = \arg \min_{\mathbf{w}} \|\mathbf{w}\|_{\mathbf{P}^{-1}} \quad \text{s.t. for all } i, \quad \langle \mathbf{w}, \mathbf{x}_i \rangle \geq \gamma \quad (2)$$

When $\mathbf{P} = I_d$, the corresponding maximum margin solution \mathbf{w}_{mm} is the standard max ℓ_2 -margin solution, henceforth referred to as the *max-margin* solution. In this case, the quantity $\max_{\|\mathbf{w}\|_2 \leq 1/\gamma} \min_{i \in [n]} \langle \mathbf{w}, \mathbf{x}_i \rangle$ is the ℓ_2 margin and data points corresponding to the equality $\langle \mathbf{w}_{\text{mm}}, \mathbf{x}_i \rangle = \gamma$ are the support vectors for \mathbf{w}_{mm} . The max-margin solution is shown to have good generalization performance for under-parameterized models [23] and more recently in the over-parameterized setting [10]. We first show that the direction of any linear classifier separating the data is the same as that of a maximum \mathbf{P} -margin solution for an appropriately constructed \mathbf{P} . This equivalence can be used to get a handle on the generalization performance of \mathbf{w} . In particular, we first show that the generalization performance of the maximum \mathbf{P} -margin solution depends on the induced norm it minimizes. We then investigate whether it is possible to construct norms that generalize better than the ℓ_2 max-margin solution. From relation 2, observe that the maximum P -margin solution minimizes $\|\cdot\|_{\mathbf{P}^{-1}}$. Let us consider an equivalent hypothesis class that has a (small) bounded \mathbf{P}^{-1} norm and is given by $\mathcal{F}(\mathbf{P}) = \{\mathbf{x} \rightarrow y\mathbf{w}^\top \mathbf{x} \mid \frac{1}{2} \mathbf{w}^\top \mathbf{P}^{-1} \mathbf{w} \leq E\}$. We measure the complexity of this hypothesis class in terms of its Rademacher complexity or VC-dimension and obtain bounds on its generalization performance [36].

1. We only consider homogeneous linear classifiers without a bias term

2. For notational convenience, from now on, we absorb the label y_i into the feature \mathbf{x}_i .

Lemma 2 *The Rademacher complexity of the model family $\mathcal{F} = \{y\mathbf{w}^\top \mathbf{x} | \frac{1}{2}\mathbf{w}^\top \mathbf{P}^{-1}\mathbf{w} \leq E\}$ is upper bounded by*

$$\hat{\mathcal{R}}(\mathcal{F}) \leq \frac{2\sqrt{2E}}{n} \sqrt{\text{tr}(\mathbf{P}\hat{\Sigma})} \quad (3)$$


where $\hat{\Sigma}$ is the scaled covariance matrix of the data i.e. $\hat{\Sigma} = \mathbf{X}^\top \mathbf{X}$. By constraining \mathbf{P} to be symmetric positive definite, and using a log-det regularization to avoid degenerate solutions, we solve the following regularized problem:

$$\min_{\mathbf{P}} \frac{2\sqrt{2E}}{n} \sqrt{\text{tr}(\mathbf{P}\hat{\Sigma})} - \lambda \log |\mathbf{P}| \quad (4)$$

with the minimizer $\mathbf{P}^* = \lambda \hat{\Sigma}^{-1}$.

If the two classes have different covariances, it can be easily shown that $\mathbf{P}^* = (\hat{\Sigma}_+ + \hat{\Sigma}_-)^{-1}$ where $\hat{\Sigma}_+$ and $\hat{\Sigma}_-$ is the scaled covariance matrix for the positive and negative classes respectively. Furthermore under a mild assumption, we show that \mathbf{P}^* also minimizes an upper-bound on the VC-dimension. The above lemma shows that maximum margin solutions in the $\hat{\Sigma}$ -norm can generalize better than the ℓ_2 max-margin solution. Finally, we note that the margin in the $\hat{\Sigma}$ -norm was defined as the

relative margin in [38] where the authors modified the standard SVM formulation to maximize the relative margin [37]. We present a simple example [38] (details in Appendix B.3) to empirically validate the above result. In Figure 3, we show the effect of the proposed preconditioning for gradient descent: incorporating the covariance of the data maximizes the relative margin resulting in a solution which is better aligned with the Bayes optimal classifier. The above results show that measuring convergence w.r.t to the ℓ_2 max-margin is arbitrary, and norms incorporating the structure of the data can generalize better.



Since the 0-1 loss is non-convex and difficult to minimize, we consider optimizers minimizing the exponential loss. The exponential loss is a smooth, convex loss minimized at infinity for separable data and given by: $f(\mathbf{w}) := \frac{1}{n} \sum_{i=1}^n \exp(-\langle \mathbf{w}, \mathbf{x}_i \rangle)$. Previous works [17, 30, 39] show that (stochastic) gradient descent converges to the max-margin solution and the resulting empirical margin converges to the true margin at an $O(1/\log(T))$ rate. This result can be extended to general losses with an exponential tail, including the logistic loss [39]. Similar to the regression setting, we first analyze the implicit regularization of PGD. In particular, we use the result in [39] and state the following lemma (for completeness, we provide the proof in Appendix A.4),

Lemma 3 *The empirical margin for PGD with preconditioner P and constant step-size $\eta < 1/f(\mathbf{w}_0)$ satisfies:*

$$\frac{\min_j \langle \mathbf{w}_{k+1}, \mathbf{x}_j \rangle}{\|\mathbf{w}_{k+1}\|_{\mathbf{P}^{-1}}} \geq \gamma - \left[\frac{\gamma \|\mathbf{w}_0\|_{\mathbf{P}^{-1}} + \eta f(\mathbf{w}_0) + \log(f(\mathbf{w}_0))}{\|\mathbf{w}_0\|_{\mathbf{P}^{-1}} + \log(\eta \gamma^2 (k+1))} \right]$$

The empirical margin converges to the true margin γ (in the \mathbf{P}^{-1} -norm) at an $O(1/\log(k))$ rate.

Note that the empirical margin is measured in the \mathbf{P}^{-1} norm and PGD with a constant step-size converges to the maximum \mathbf{P} -margin solution. For a general optimizer, consider an iterate \mathbf{w}_T obtained after T iterations. Furthermore, using the above result, this direction corresponds to the direction of a PGD solution with a preconditioner \mathbf{P} . This result implies that, as in linear regression, an interpolating solution is equivalent (in direction) to a PGD solution.

However, unlike regression, it is not possible to provably recover the maximum ℓ_2 -margin solution from a general interpolating solution found by an optimizer. Instead, we consider an empirical strategy that *switches* from the original optimization method and runs “some” iterations of PGD to bias the resulting solution towards the corresponding max-margin direction. For example, switching to GD will result in a bias to the max- ℓ_2 -margin, and we can obtain better generalization by using $\mathbf{P} = \Sigma^{-1}$ to bias the optimizer to the maximum relative-margin solution. Since Lemma 3 holds for an arbitrary initialization, we invoke it with \mathbf{w}_0 equal to the iterate obtained by the original optimization method. Since convergence to the max-margin solution depends on the loss $f(\mathbf{w}_0)$, if the original optimization method is effective in minimizing the loss, the resulting $f(\mathbf{w}_0)$ is small and $\|\mathbf{w}_0\|$ is large, making it possible to approach the corresponding max-margin solution in fewer iterations. In Section 4, we empirically demonstrate that switching to GD for only a few iterations can significantly improve the generalization performance of the original optimizer. We note that such a strategy of switching to GD (from Adam) has been explored in [24] in the context of deep networks and found to improve the generalization performance of Adam. Our reasoning using Lemma 3 gives further intuition on why such a strategy is reasonable. Using projections onto the data-span or reparameterization allows us to use these heuristics in the over-parameterized linear classification setting as well.

4. Experimental Evaluation

We verify the theoretical results for both the regression and classification settings, and evaluate the proposed techniques for improving the test performance of optimizers.

Regression: We consider over-parameterized regression with one-layer neural networks. We generate data-points from a Gaussian distribution and use the resulting neural tangent kernel [21] (NTK) features as input to our problem. The targets are generated to ensure that the problem is realizable. For this problem, we study the generalization of SGD and Adagrad. From Figure 1, we observe that (i) Adagrad is robust to the choice of step-size and converges quickly, (ii) hand-tuned SGD converges slowly due to the problem’s ill-conditioning, (iii) Adagrad’s test performance highly depends on the step-size, and (iv) the min-norm solution found by SGD consistently generalizes well. We can project the Adagrad solution onto the data span as in Proposition 1 and recover the min-norm solution. This suggests we can benefit from the robustness of Adagrad, while also ensuring good generalization performance.

In Appendix C.1.1, we present additional results on the `wine` and `mushroom` datasets [13]. In Appendix C.1.2, we validate that full-matrix Adagrad and the Newton methods converge to




Figure 1: Performance of SGD and Adagrad on a synthetic regression problem using the NTK of one-layer networks with 50 and 100 hidden units. Tuned SGD converges slowly to the min-norm solution. In contrast, the convergence of Adagrad is more robust to the step-size, however its generalization depends on the step-size.

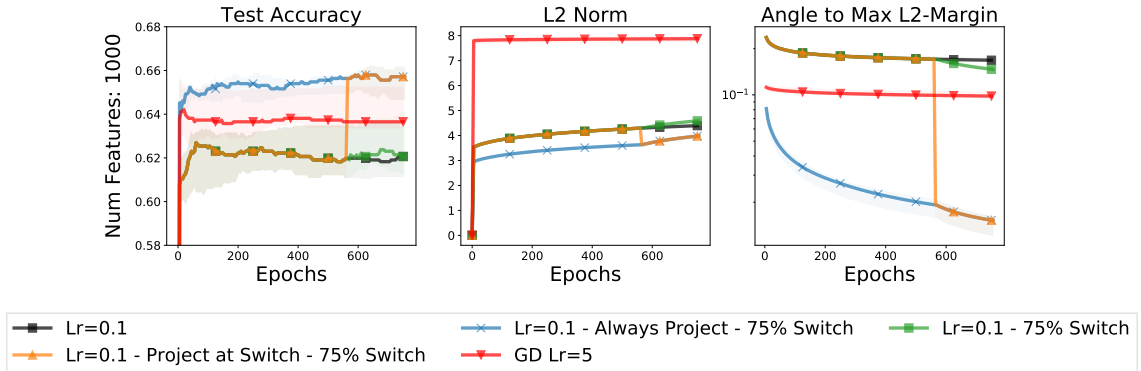


Figure 2: Performance of GD and Adagrad on a synthetic overparameterized classification problem with random Gaussian features. Projecting onto the data-span improves the test accuracy, while decreasing the solution’s norm and angle to the max-margin solution.

the min-norm solution and verify the construction of the equivalent preconditioner in Lemma 4 in Appendix A.1. In Appendix C.1.3, we investigate whether it is possible to obtain solutions that generalize better than the min-norm solution by minimizing upper bounds on the excess risk.

Classification: We evaluate effectiveness of switching to GD and projection in improving the generalization for over-parameterized linear classification. In particular, we present the results for Adagrad and GD minimizing the logistic loss on a synthetic dataset with random Gaussian features. We ensure that both GD and Adagrad converge to solutions which interpolate the training data. We consider four variants: (i) standard Adagrad (black), (ii) switch to GD after 75% of the iterations (green) (iii) project onto the data-span after every iteration (blue) and (iv) project the final solution

onto the data-span (orange). For the latter two variants, we also switch to GD with a fixed step-size after 75% of the iterations. From Figure 2, we observe that (i) switching to GD results in a small improvement in the test accuracy. (ii) projections after every iteration result in a smaller norm and angle to the max-margin solution as well as higher test accuracy, (iii) projecting only the final solution is sufficient to improve the test performance (from an accuracy of 64% to 66%) (iv) compared to projection, switching to GD has a small effect on the test accuracy. Our results indicate that being in the correct subspace can improve the generalization performance.

In Appendix B.4, we presents additional results for this setting. In Appendix B.1, we verify the construction of the equivalent preconditioner in Lemma 14.

5. Conclusion

For both linear regression and classification, we saw that an interpolating solution found by an optimizer minimizes an equivalent quadratic norm. This reasoning enabled us to use projections to move between norms (and therefore solutions) for over-parameterized settings. For classification, we showed that it important to consider the geometry induced by the data to measure generalization. We also proposed techniques to improve the generalization of optimization methods. We consider extending our techniques to non-linear models including deep networks as important future work. Finally, we hope to use our insights to develop optimizers that are guaranteed to generalize well.

References

- [1] Naman Agarwal, Brian Bullins, Xinyi Chen, Elad Hazan, Karan Singh, Cyril Zhang, and Yi Zhang. Efficient full-matrix adaptive regularization. In *International Conference on Machine Learning*, 2019.
- [2] Shun-Ichi Amari. Natural gradient works efficiently in learning. *Neural computation*, 10(2): 251–276, 1998.
- [3] Sanjeev Arora, Nadav Cohen, Wei Hu, and Yuping Luo. Implicit regularization in deep matrix factorization. In *Advances in Neural Information Processing Systems*, 2019.
- [4] Sanjeev Arora, Simon S Du, Zhiyuan Li, Ruslan Salakhutdinov, Ruosong Wang, and Dingli Yu. Harnessing the power of infinitely wide deep nets on small-data tasks. *arXiv preprint arXiv:1910.01663*, 2019.
- [5] Peter L Bartlett and Shahar Mendelson. Rademacher and gaussian complexities: Risk bounds and structural results. *Journal of Machine Learning Research*, 3(Nov), 2002.
- [6] Peter L Bartlett, Philip M Long, Gábor Lugosi, and Alexander Tsigler. Benign overfitting in linear regression. *arXiv preprint arXiv:1906.11300*, 2019.
- [7] Mikhail Belkin, Siyuan Ma, and Soumik Mandal. To understand deep learning we need to understand kernel learning. *arXiv preprint arXiv:1802.01396*, 2018.

- [8] Mikhail Belkin, Daniel Hsu, Siyuan Ma, and Soumik Mandal. Reconciling modern machine-learning practice and the classical bias–variance trade-off. *Proceedings of the National Academy of Sciences*, 116(32), 2019.
- [9] Raghu Bollapragada, Dheevatsa Mudigere, Jorge Nocedal, Hao-Jun Michael Shi, and Ping Tak Peter Tang. A progressive batching l-bfgs method for machine learning. *arXiv preprint arXiv:1802.05374*, 2018.
- [10] Niladri S Chatterji and Philip M Long. Finite-sample analysis of interpolating linear classifiers in the overparameterized regime. *arXiv preprint arXiv:2004.12019*, 2020.
- [11] Jinghui Chen and Quanquan Gu. Closing the generalization gap of adaptive gradient methods in training deep neural networks. *arXiv preprint arXiv:1806.06763*, 2018.
- [12] Felix Dangel, Frederik Kunstner, and Philipp Hennig. BackPACK: Packing more into backprop. In *International Conference on Learning Representations*, 2020. URL <https://openreview.net/forum?id=BJlrf24twB>.
- [13] Dheeru Dua and Casey Graff. UCI machine learning repository, 2017. URL <http://archive.ics.uci.edu/ml>.
- [14] John Duchi, Elad Hazan, and Yoram Singer. Adaptive subgradient methods for online learning and stochastic optimization. *JMLR*, 2011.
- [15] Manuel Fernández-Delgado, Eva Cernadas, Senén Barro, and Dinani Amorim. Do we need hundreds of classifiers to solve real world classification problems? *The journal of machine learning research*, 15(1):3133–3181, 2014.
- [16] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep learning*. MIT press, 2016.
- [17] Suriya Gunasekar. Characterizing implicit bias in terms of optimization geometry. In *In International Conference on Machine Learning*, 2018.
- [18] Suriya Gunasekar, Blake E Woodworth, Srinadh Bhojanapalli, Behnam Neyshabur, and Nati Srebro. Implicit regularization in matrix factorization. In *Advances in Neural Information Processing Systems*, 2017.
- [19] Suriya Gunasekar, Jason D Lee, Daniel Soudry, and Nati Srebro. Implicit bias of gradient descent on linear convolutional networks. In *Advances in Neural Information Processing Systems*, 2018.
- [20] Trevor Hastie, Andrea Montanari, Saharon Rosset, and Ryan J Tibshirani. Surprises in high-dimensional ridgeless least squares interpolation. *arXiv preprint arXiv:1903.08560*, 2019.

- [21] Arthur Jacot, Franck Gabriel, and Clément Hongler. Neural tangent kernel: Convergence and generalization in neural networks. In *Advances in neural information processing systems*, 2018.
- [22] Ziwei Ji and Matus Telgarsky. Gradient descent aligns the layers of deep linear networks. *arXiv preprint arXiv:1810.02032*, 2018.
- [23] Sham M Kakade, Karthik Sridharan, and Ambuj Tewari. On the complexity of linear prediction: Risk bounds, margin bounds, and regularization. In *Advances in neural information processing systems*, 2009.
- [24] Nitish Shirish Keskar and Richard Socher. Improving generalization performance by switching from adam to sgd. *arXiv preprint arXiv:1712.07628*, 2017.
- [25] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [26] Kenneth Levenberg. A method for the solution of certain non-linear problems in least squares. *Quarterly of applied mathematics*, 2(2), 1944.
- [27] Liangchen Luo, Yuanhao Xiong, Yan Liu, and Xu Sun. Adaptive gradient methods with dynamic bound of learning rate. In *ICLR*, 2019.
- [28] Donald W Marquardt. An algorithm for least-squares estimation of nonlinear parameters. *Journal of the society for Industrial and Applied Mathematics*, 11(2), 1963.
- [29] Vidya Muthukumar, Kailas Vodrahalli, Vignesh Subramanian, and Anant Sahai. Harmless interpolation of noisy data in regression. *IEEE Journal on Selected Areas in Information Theory*, 2020.
- [30] Mor Shpigel Nacson, Jason Lee, Suriya Gunasekar, Pedro HP Savarese, Nathan Srebro, and Daniel Soudry. Convergence of gradient descent on separable data. *arXiv preprint arXiv:1803.01905*, 2018.
- [31] Mor Shpigel Nacson, Nathan Srebro, and Daniel Soudry. Stochastic gradient descent on separable data: Exact convergence with a fixed learning rate. In *The 22nd International Conference on Artificial Intelligence and Statistics*, 2019.
- [32] Francesco Orabona and Tatiana Tommasi. Training deep networks without learning rates through coin betting. In *Advances in Neural Information Processing Systems*, 2017.
- [33] Daniel S Park, Jascha Sohl-Dickstein, Quoc V Le, and Samuel L Smith. The effect of network width on stochastic gradient descent and generalization: an empirical study. *arXiv preprint arXiv:1905.03776*, 2019.

- [34] Qian Qian and Xiaoyuan Qian. The implicit bias of adagrad on separable data. In *Advances in Neural Information Processing Systems*, 2019.
- [35] Noam Razin and Nadav Cohen. Implicit regularization in deep learning may not be explainable by norms. *arXiv preprint arXiv:2005.06398*, 2020.
- [36] Shai Shalev-Shwartz and Shai Ben-David. *Understanding machine learning: From theory to algorithms*. Cambridge university press, 2014.
- [37] Pannagadatta K Shivaswamy and Tony Jebara. Ellipsoidal machines. In *Artificial Intelligence and Statistics*, 2007.
- [38] Pannagadatta K Shivaswamy and Tony Jebara. Maximum relative margin and data-dependent regularization. *Journal of Machine Learning Research*, 11(Feb), 2010.
- [39] Daniel Soudry, Elad Hoffer, Mor Shpigel Nacson, Suriya Gunasekar, and Nathan Srebro. The implicit bias of gradient descent on separable data. *The Journal of Machine Learning Research*, 19(1), 2018.
- [40] Ashia C Wilson, Rebecca Roelofs, Mitchell Stern, Nati Srebro, and Benjamin Recht. The marginal value of adaptive gradient methods in machine learning. In *Advances in Neural Information Processing Systems*, 2017.
- [41] Chiyuan Zhang, Samy Bengio, Moritz Hardt, Benjamin Recht, and Oriol Vinyals. Understanding deep learning requires rethinking generalization. *arXiv preprint arXiv:1611.03530*, 2016.

Appendix A. Appendix

A.1. Optimizer-based Preconditioner

Lemma 4 For the solution \mathbf{w}_{opt} obtained by a generic optimizer on the squared loss, it is possible to construct a family of constant preconditioners \mathbf{P}_{opt} such that PGD with a fixed preconditioner in \mathbf{P}_{opt} converges to \mathbf{w}_{opt} . One such family can be given as:

$$\mathbf{P}_{opt} = \left[\|\mathbf{w}_{opt}\|^2 I_d - \mathbf{w}_{opt} \mathbf{w}_{opt}^\top + \frac{\nu \nu^\top}{\langle \mathbf{w}_{opt}, \nu \rangle} \right]^{-1}$$

Here $\nu = \mathbf{X}^\top \alpha$ where α is a random vector such that $\langle \mathbf{w}_{opt}, \nu \rangle > 0$, and I_d is the $d \times d$ identity matrix. Note that for any ν , either ν or $-\nu$ satisfies this constraint.

Proof Let \mathbf{w}_{opt} be the solution to which a given optimizer converges to. Since \mathbf{w}_{opt} interpolates the data, $\mathbf{X} \mathbf{w}_{opt} = \mathbf{y}$. This solution also corresponds to the min-norm solution measured in the \mathbf{M} norm, implying we want to find a positive definite matrix \mathbf{M} s.t.

$$\mathbf{w}_{opt} = \arg \min_z \frac{1}{2} z^\top \mathbf{M} z \quad \text{s.t.} \quad \mathbf{X} z = \mathbf{y}$$

The Lagrangian for the optimization on the RHS can be written as follows. Here $\lambda \in \mathbb{R}^{n \times 1}$

$$\begin{aligned} L(z, \lambda) &= \frac{1}{2} z^\top \mathbf{M} z + \lambda^\top (\mathbf{X} z - \mathbf{y}) \\ \frac{\partial L(z, \lambda)}{\partial z} &= \mathbf{M} + \mathbf{X}^\top \lambda \end{aligned}$$

Since \mathbf{w}_{opt} is the solution of this optimization problem,

$$\mathbf{M} \mathbf{w}_{opt} = -\mathbf{X}^\top \lambda$$

implying $\mathbf{M} \mathbf{w}_{opt} \in \text{span}\{\mathbf{X}^\top\}$. We choose \mathbf{M} to be the following matrix,

$$\mathbf{M} = \|\mathbf{w}_{opt}\|^2 \mathbf{I}_d - \mathbf{w}_{opt} \mathbf{w}_{opt}^\top + \frac{\nu \nu^\top}{\langle \mathbf{w}_{opt}, \nu \rangle}$$

Here $\nu = \mathbf{X}^\top \alpha$ where α is a random vector such that $\langle \mathbf{w}_{opt}, \nu \rangle > 0$. Note that for any ν , either ν or $-\nu$ satisfies this constraint. We now verify that $\mathbf{M} \mathbf{w}_{opt} = \nu = \mathbf{X}^\top \alpha$.

$$\mathbf{M} \mathbf{w}_{opt} = \|\mathbf{w}_{opt}\|^2 \mathbf{w}_{opt} - (\mathbf{w}_{opt} \mathbf{w}_{opt}^\top) \mathbf{w}_{opt} + \frac{\nu \nu^\top \mathbf{w}_{opt}}{\langle \mathbf{w}_{opt}, \nu \rangle} = \|\mathbf{w}_{opt}\|^2 \mathbf{w}_{opt} - \|\mathbf{w}_{opt}\|^2 \mathbf{w}_{opt} + \nu = \nu$$

We now compute $\mathbf{a}^\top \mathbf{M} \mathbf{a}$ to verify that \mathbf{M} is positive definite.

$$\begin{aligned} \mathbf{a}^\top \mathbf{M} \mathbf{a} &= \|\mathbf{w}_{opt}\|^2 \|\mathbf{a}\|^2 - (\mathbf{a}^\top \mathbf{w}_{opt}) (\mathbf{a}^\top \mathbf{w}_{opt}) + \frac{(\mathbf{a}^\top \nu) (\mathbf{a}^\top \nu)}{\langle \mathbf{w}_{opt}, \nu \rangle} = \|\mathbf{w}_{opt}\|^2 \|\mathbf{a}\|^2 - \|\mathbf{a} \mathbf{w}_{opt}\|^2 + \frac{\|\mathbf{a}^\top \nu\|^2}{\langle \mathbf{w}_{opt}, \nu \rangle} \\ &> \|\mathbf{w}_{opt}\|^2 \|\mathbf{a}\|^2 - \|\mathbf{w}_{opt}\|^2 \|\mathbf{a}\|^2 + \frac{\|\mathbf{a}^\top \nu\|^2}{\langle \mathbf{w}_{opt}, \nu \rangle} = \frac{\|\mathbf{a}^\top \nu\|^2}{\langle \mathbf{w}_{opt}, \nu \rangle} \end{aligned}$$

Since $\langle \mathbf{w}_{opt}, \nu \rangle > 0$ by construction, $\mathbf{a}^\top \mathbf{M} \mathbf{a} > 0$ for all \mathbf{a} , implying \mathbf{M} is positive definite.

The preconditioner resulting in the min-norm solution in the \mathbf{M} norm is \mathbf{M}^{-1} which is a family of preconditioners that result in the same solution as \mathbf{w}_{opt} . ■

A.2. Proof of Proposition 1

Proof Recall that the solution found by an optimizer with equivalent preconditioner \mathbf{P}_1 can be written as

$$\mathbf{w}_\infty^{\mathbf{P}_1} = \mathbf{P}_1 \mathbf{X}^\top (\mathbf{X} \mathbf{P}_1 \mathbf{X}^\top)^{-1} \mathbf{y}$$

Projecting this solution using the projection matrix, $\pi_P = \mathbf{P}_2 \mathbf{X}^\top (\mathbf{X} \mathbf{P}_2 \mathbf{X}^\top)^{-1} \mathbf{X}$,

$$\begin{aligned} \pi_P[\mathbf{w}_{opt}] &= \mathbf{P}_2 \mathbf{X}^\top (\mathbf{X} \mathbf{P}_2 \mathbf{X}^\top)^{-1} \mathbf{X} \mathbf{w}_\infty^{\mathbf{P}_1} \\ &= \mathbf{P}_2 \mathbf{X}^\top (\mathbf{X} \mathbf{P}_2 \mathbf{X}^\top)^{-1} \mathbf{X} \mathbf{P}_1 \mathbf{X}^\top (\mathbf{X} \mathbf{P}_1 \mathbf{X}^\top)^{-1} \mathbf{y} \\ &= \mathbf{P}_2 \mathbf{X}^\top (\mathbf{X} \mathbf{P}_2 \mathbf{X}^\top)^{-1} \mathbf{y} = \mathbf{w}_\infty^{\mathbf{P}_2} \end{aligned}$$

which is the solution obtained by PGD with a preconditioner \mathbf{P}_2 and interpolates the data. ■

A.3. Proof of Lemma 2

Proof Based on the definition of Rademacher complexity we have

$$\hat{\mathcal{R}}(\mathcal{F}) = \mathbb{E}_\sigma \left[\sup_{f \in \mathcal{F}} \left| \frac{2}{n} \sum_{i=1}^n \sigma_i y_i \mathbf{w}^\top \mathbf{x}_i \right| \right] = \mathbb{E}_\sigma \left[\sup_{\mathbf{w}: 1/2 \mathbf{w}^\top \mathbf{P}^{-1} \mathbf{w} \leq E} \left| \frac{2}{n} \sum_{i=1}^n \sigma_i y_i \mathbf{w}^\top \mathbf{x}_i \right| \right] \quad (5)$$

$$\leq \frac{2\sqrt{2E}}{n} \mathbb{E}_\sigma \left[\left(\sum_{i=1}^n \sigma_i y_i \mathbf{x}_i^\top \mathbf{P} \sum_{j=1}^n \sigma_j y_j \mathbf{x}_j \right)^{1/2} \right] \quad (6)$$

$$\leq \frac{2\sqrt{2E}}{n} \left[\left(\mathbb{E}_\sigma \sum_{j,i=1}^n \sigma_i \sigma_j y_i y_j \mathbf{x}_i^\top \mathbf{P} \mathbf{x}_j \right)^{1/2} \right] = \frac{2\sqrt{2E}}{n} \sqrt{\text{tr}(\mathbf{P}\mathbf{K})} \quad (7)$$

To find the optimal \mathbf{P} , since the space of positive definite matrices is an open set and also $\text{tr}(\mathbf{P}\mathbf{K})$ is optimal over the boundary of that space, we add a regularizer term $-\lambda \ln(\det(\mathbf{P}))$ which prevents degenerate solutions. So the objective function is

$$\arg \min_{\mathbf{P}} \text{tr}(\mathbf{P}\mathbf{K}) - \lambda \ln(\det(\mathbf{P})) \quad (8)$$

Taking derivative w.r.t. \mathbf{P} and set it to zero we have

$$\mathbf{P}^* = \lambda \mathbf{K}^{-1}. \quad (9)$$

We can assume that \mathbf{K} is symmetric positive definite. Furthermore, if we assume \mathbf{K}_+ and \mathbf{K}_- are scaled covariance matrix for data belonging to $+1$ and -1 classes, we get the following upper bound for the Rademacher complexity

$$\hat{\mathcal{R}}(\mathcal{F}) \leq \frac{2\sqrt{2E}}{n} \sqrt{\text{tr}(\mathbf{P}(\mathbf{K}_+ + \mathbf{K}_-))}. \quad (10)$$

Using the above argument, then the optimal precondition is

$$\mathbf{P}^* = \lambda(\mathbf{K}_+ + \mathbf{K}_-)^{-1}. \quad (11)$$

Similarly we can find an upper bound for the VC-dimension of max-margin which has the same upper bound as we get for the Rademacher complexity. For this part we assume that there exist a data set \mathcal{D}_v such that max-margin model can shatter it and besides for any positive definite \mathbf{P} we have $\text{tr}(\mathbf{P}\mathbf{K}_v) \leq \text{tr}(\mathbf{P}\mathbf{K})$ where \mathbf{K}_v is $\mathbf{X}_v^\top \mathbf{X}_v$. To be more precise, let assume v is the VC-dimension of max-margin problem, therefore there exist an a set of size v i.e. $\mathcal{D}_v = \{(\mathbf{x}_j, y_j)\}_{j=1:v}$ such that for all j $y_j \mathbf{w}^\top \mathbf{x}_j \geq 1$. If we sum up the both side of this inequality for all j data point we have:

$$\begin{aligned} v &\leq \sum_{j=1}^v y_j \mathbf{w}^\top \mathbf{x}_j \\ &\leq \|\mathbf{w}\|_{\mathbf{P}^{-1}} \left[\left(\sum_{j=1}^v y_j \mathbf{x}_j^\top \mathbf{P} \sum_{k=1}^v y_k \mathbf{x}_k \right)^{1/2} \right] \\ &\leq \sqrt{2E} \left[\left(\sum_{j=1}^v y_j \mathbf{x}_j^\top \mathbf{P} \sum_{k=1}^v y_k \mathbf{x}_k \right)^{1/2} \right] \end{aligned}$$

where the second inequality is due to Cauchy-Schwarz. Since max-margin can shatter this set, we can assume that y_j 's are independent random variables, and take expectation from both sides of the above bound and then apply the Jensen inequality to get the following

$$v \leq \sqrt{2E} \left[\mathbb{E} \left(\sum_{j=1}^v y_j \mathbf{x}_j^\top \mathbf{P} \sum_{k=1}^v y_k \mathbf{x}_k \right) \right]^{1/2} \quad (12)$$

$$\leq \sqrt{2E} \left[\sum_{j=1}^v \mathbf{x}_j^\top \mathbf{P} \mathbf{x}_j \right]^{1/2} \quad (13)$$

$$\leq \sqrt{2E} \sqrt{\text{tr}(\mathbf{P} \mathbf{K}_v)} \leq \sqrt{2E} \sqrt{\text{tr}(\mathbf{P} \mathbf{K})} \quad (14)$$

where the second inequality due to the fact that $\mathbb{E}(y_j) = 0$ and y_j and y_k are independent. The last inequality is due to our assumption. Here we show that when $\mathcal{D}_v \subset \mathcal{D}$ which means there is a subset of training data with size v such that max-margin can shatter it, then the assumption $\text{tr}(\mathbf{P} \mathbf{K}_v) \leq \text{tr}(\mathbf{P} \mathbf{K})$. Let $\mathbf{S} \in \mathbb{R}^{n \times n}$ be a diagonal matrix with $\mathbf{S}_{i,i} = 1$ if $\mathbf{x}_i \in \mathcal{D}_v$ and 0 everywhere else. Therefore $\mathbf{S} \mathbf{X}$ represents data points in \mathcal{D}_v and we have $\mathbf{K}_v = \mathbf{X}^\top \mathbf{S} \mathbf{X}$. Let $\mathbf{M} = \mathbf{X} \mathbf{P} \mathbf{X}^\top$ and we have $\text{tr}(\mathbf{M}) = \text{tr}(\mathbf{P} \mathbf{K})$. We observe that $\text{tr}(\mathbf{P} \mathbf{K}_v) = \text{tr}(\mathbf{M} \mathbf{S})$. Since \mathbf{P} is positive definite, all the diagonal elements of \mathbf{M} are positive. Therefore multiplying \mathbf{M} by \mathbf{S} sets some of the diagonal elements to 0, therefore we have $\text{tr}(\mathbf{P} \mathbf{K}_v) = \text{tr}(\mathbf{M} \mathbf{S}) \leq \text{tr}(\mathbf{M}) = \text{tr}(\mathbf{P} \mathbf{K})$. \blacksquare

A.4. Proof of Lemma 3

We map preconditioned gradient descent to the notation of [39] and use Theorem 7 that considers steepest descent to prove the statement of the lemma. We first consider a constant step-size s.t. $\eta_k = \eta < 1/f(\mathbf{w}_0)$.

Proof PGD has the following update:

$$\mathbf{w}_{k+1} = \mathbf{w}_k - \eta \mathbf{P} \nabla f(\mathbf{w}_k)$$

The update considered in Appendix B.2. of [39] can be written as follows:

$$\mathbf{w}_{k+1} = \mathbf{w}_k - \eta \gamma_k \Delta \mathbf{w}_k$$

where γ_k and $\Delta \mathbf{w}_k$ are defined as:

$$\begin{aligned} \gamma_k &= \|\nabla f(\mathbf{w}_k)\|_* \quad ; \quad \langle \Delta \mathbf{w}_k, \nabla f(\mathbf{w}_k) \rangle = \|\nabla f(\mathbf{w}_k)\|_* \\ \|\Delta \mathbf{w}_k\| &= 1 \end{aligned}$$

Here, $\|\cdot\|_*$ is the dual norm.

Mapping PGD to this update, we get the following equivalence:

$$\gamma_k = \|\nabla f(\mathbf{w}_k)\|_{\mathbf{P}} \quad ; \quad \Delta \mathbf{w}_k = \frac{\mathbf{P} \nabla f(\mathbf{w}_k)}{\|\nabla f(\mathbf{w}_k)\|_{\mathbf{P}}}$$

It is easy to verify that $\|\cdot\| = \|\cdot\|_{\mathbf{P}^{-1}}$ and $\|\cdot\|_* = \|\cdot\|_{\mathbf{P}}$, and that γ_k and $\Delta \mathbf{w}_k$ satisfy the above relations.

Given this mapping, we follow the proof of Theorem 7. Using the descent lemma and the property of the exponential loss,

$$f(\mathbf{w}_{k+1}) \leq f(\mathbf{w}_k) \left(1 - \frac{\eta \gamma_k}{f(\mathbf{w}_k)} + \frac{\eta^2 \gamma_k^2}{2} \right)$$

Recurring, multiplying both sides by $-\log$ and using Jensen's inequality, we get a bound on the unnormalized margin

$$\min_j \langle \mathbf{w}_{k+1}, \mathbf{x}_j \rangle \geq \sum_{i=0}^k \frac{\eta \gamma_i^2}{f(\mathbf{w}_i)} - \frac{1}{2} \eta^2 \sum_{i=0}^k \gamma_i^2 - \log(f(\mathbf{w}_0))$$

We now upper bound the $\|\mathbf{w}_{k+1}\|$. Note for PGD, this norm is the induced norm wrt to the P^{-1} matrix.

$$\begin{aligned}\|\mathbf{w}_{k+1}\| &= \left\| \mathbf{w}_0 - \eta \sum_{i=0}^t \gamma_i \Delta \mathbf{w}_i \right\| \leq \|\mathbf{w}_0\| + \eta \left\| \sum_{i=0}^t \gamma_i \Delta \mathbf{w}_i \right\| \\ &\leq \|\mathbf{w}_0\| + \eta \sum_{i=0}^t \|\gamma_i \Delta \mathbf{w}_i\| \leq \|\mathbf{w}_0\| + \eta \sum_{i=0}^t \gamma_i\end{aligned}$$

Dividing the above inequalities, we get a bound on normalized margin,

$$\frac{\min_j \langle \mathbf{w}_{k+1}, \mathbf{x}_j \rangle}{\|\mathbf{w}_{k+1}\|} \geq \frac{\sum_{i=0}^k \frac{\eta \gamma_i^2}{f(\mathbf{w}_i)} - \frac{1}{2} \eta^2 \sum_{i=0}^k \gamma_i^2 - \log(f(\mathbf{w}_0))}{\|\mathbf{w}_0\| + \eta \sum_{i=0}^t \gamma_i}$$

Using the result in Lemma 2 of [39], $\gamma_i \geq \gamma f(\mathbf{w}_i)$,

$$\frac{\min_j \langle \mathbf{w}_{k+1}, \mathbf{x}_j \rangle}{\|\mathbf{w}_{k+1}\|} \geq \frac{\gamma \sum_{i=0}^k \eta \gamma_i - \frac{1}{2} \eta^2 \sum_{i=0}^k \gamma_i^2 - \log(f(\mathbf{w}_0))}{\|\mathbf{w}_0\| + \eta \sum_{i=0}^t \gamma_i}$$

Using the descent lemma,

$$f(\mathbf{w}_{k+1}) \leq f(\mathbf{w}_k) - \frac{\eta}{2} \gamma_k^2 \implies \sum_{i=0}^k \frac{\eta}{2} \gamma_i^2 \leq f(\mathbf{w}_0) \quad (\text{By telescoping from } i = 0 \text{ to } k)$$

Using this bound in the above inequality,

$$\frac{\min_j \langle \mathbf{w}_{k+1}, \mathbf{x}_j \rangle}{\|\mathbf{w}_{k+1}\|} \geq \frac{\gamma \sum_{i=0}^k \eta \gamma_i - \eta f(\mathbf{w}_0) - \log(f(\mathbf{w}_0))}{\|\mathbf{w}_0\| + \eta \sum_{i=0}^t \gamma_i}$$

Rearranging,

$$\geq \gamma - \left[\frac{\gamma \|\mathbf{w}_0\| + \eta f(\mathbf{w}_0) + \log(f(\mathbf{w}_0))}{\|\mathbf{w}_0\| + \eta \sum_{i=0}^t \gamma_i} \right]$$

Finally, we use Claim 1 in Theorem 7 of [39] to bound $\eta \sum_{i=0}^t \gamma_i$,

$$\eta \sum_{i=0}^t \gamma_i \geq \log(\eta \gamma^2 (k+1))$$

Using this bound in the above inequality,

$$\frac{\min_j \langle \mathbf{w}_{k+1}, \mathbf{x}_j \rangle}{\|\mathbf{w}_{k+1}\|} \geq \gamma - \left[\frac{\gamma \|\mathbf{w}_0\| + \eta f(\mathbf{w}_0) + \log(f(\mathbf{w}_0))}{\|\mathbf{w}_0\| + \log(\eta \gamma^2 (k+1))} \right]$$

If $\mathbf{w}_0 = 0$, $f(\mathbf{w}_0) = 1$, $\eta = 1/f(\mathbf{w}_0) = 1$

$$\implies \frac{\min_j \langle \mathbf{w}_{k+1}, \mathbf{x}_j \rangle}{\|\mathbf{w}_{k+1}\|} \geq \gamma - \frac{1}{\log(\gamma^2 (k+1))}$$

Making the dependence on P explicit,

$$\implies \frac{\min_j \langle \mathbf{w}_{k+1}, \mathbf{x}_j \rangle}{\|\mathbf{w}_{k+1}\|_{P^{-1}}} \geq \gamma - \frac{1}{\log(\gamma^2 (k+1))}$$

As $k \rightarrow \infty$,

$$\frac{\min_j \langle \mathbf{w}_{k+1}, \mathbf{x}_j \rangle}{\|\mathbf{w}_{k+1}\|_{P-1}} \rightarrow \gamma$$

We now show that using backtracking line-search procedure to set the step-size does not change the implicit regularization and results in a similar bound. The line-search procedure picks the largest step-size that satisfies the Armijo condition:

$$f(\mathbf{w}_{k+1}) \leq f(\mathbf{w}_k) - c\eta_k \|\nabla f(\mathbf{w}_k)\|_*^2$$

Here, c is a hyper-parameter. We assume that the resulting step-size $\eta_k \in [\eta_{\min}, \eta_{\max}]$. Using the PGD update:

$$\mathbf{w}_{k+1} = \mathbf{w}_k - \eta_k P \nabla f(\mathbf{w}_k)$$

Following the same analysis,

$$\begin{aligned} \frac{\min_j \langle \mathbf{w}_{k+1}, \mathbf{x}_j \rangle}{\|\mathbf{w}_{k+1}\|} &\geq \frac{\gamma \sum_{i=0}^k \eta_i \gamma_i - \frac{1}{2} \sum_{i=0}^k \eta_i^2 \gamma_i^2 - \log(f(\mathbf{w}_0))}{\|\mathbf{w}_0\| + \sum_{i=0}^t \eta_i \gamma_i} \\ &\geq \frac{\gamma \sum_{i=0}^k \eta_i \gamma_i - \frac{\eta_{\max}}{2} \sum_{i=0}^k \eta_i \gamma_i^2 - \log(f(\mathbf{w}_0))}{\|\mathbf{w}_0\| + \sum_{i=0}^t \eta_i \gamma_i} \end{aligned}$$

Using the line-search condition,

$$f(\mathbf{w}_{k+1}) \leq f(\mathbf{w}_k) - c\eta_k \gamma_k^2 \implies \sum_{i=0}^k \eta_i \gamma_i^2 \leq \frac{f(\mathbf{w}_0)}{c}$$

Using this bound in the above inequality,

$$\geq \frac{\gamma \sum_{i=0}^k \eta_i \gamma_i - \frac{\eta_{\max} f(\mathbf{w}_0)}{2c} - \log(f(\mathbf{w}_0))}{\|\mathbf{w}_0\| + \sum_{i=0}^t \eta_i \gamma_i}$$

Rearranging,

$$\geq \gamma - \left[\frac{\gamma \|\mathbf{w}_0\| + \frac{\eta_{\max} f(\mathbf{w}_0)}{2c} + \log(f(\mathbf{w}_0))}{\|\mathbf{w}_0\| + \sum_{i=0}^t \eta_i \gamma_i} \right] \geq \gamma - \left[\frac{\gamma \|\mathbf{w}_0\| + \frac{\eta_{\max} f(\mathbf{w}_0)}{2c} + \log(f(\mathbf{w}_0))}{\|\mathbf{w}_0\| + \eta_{\min} \sum_{i=0}^t \gamma_i} \right]$$

The line-search implies

$$f(\mathbf{w}_{k+1}) \leq f(\mathbf{w}_k) - c\eta_k \gamma_k^2 \leq f(\mathbf{w}_k) - c\eta_k \gamma f(\mathbf{w}_k) \leq f(\mathbf{w}_k) - c\eta_{\min} \gamma f(\mathbf{w}_k)$$

Using Claim 1 in Theorem 7 of [39],

$$c\eta_{\min} \sum_{i=0}^k \gamma_i \geq \log(c\eta_{\min} \gamma^2 (k+1))$$

Using this bound with the above inequality,

$$\frac{\min_j \langle \mathbf{w}_{k+1}, \mathbf{x}_j \rangle}{\|\mathbf{w}_{k+1}\|} \geq \gamma - \left[\frac{\gamma \|\mathbf{w}_0\| + \frac{\eta_{\max} f(\mathbf{w}_0)}{2c} + \log(f(\mathbf{w}_0))}{\|\mathbf{w}_0\| + \frac{\log(c\eta_{\min} \gamma^2 (k+1))}{c}} \right]$$

For line-search, $\eta_{\min} = \frac{2(1-c)}{L}$,

$$\geq \gamma - \left[\frac{\gamma \|\mathbf{w}_0\| + \frac{\eta_{\max} f(\mathbf{w}_0)}{2c} + \log(f(\mathbf{w}_0))}{\|\mathbf{w}_0\| + \frac{\log(2c(1-c)\gamma^2 (k+1)/L)}{c}} \right]$$

If $\mathbf{w}_0 = 0$, $f(\mathbf{w}_0) = 1$,

$$\frac{\min_j \langle \mathbf{w}_{k+1}, \mathbf{x}_j \rangle}{\|\mathbf{w}_{k+1}\|_{P^{-1}}} \geq \gamma - \left[\frac{\eta_{\max}}{2 \log(2c(1-c)\gamma^2(k+1)/L)} \right]$$

And as $k \rightarrow \infty$,

$$\frac{\min_j \langle \mathbf{w}_{k+1}, \mathbf{x}_j \rangle}{\|\mathbf{w}_{k+1}\|_{P^{-1}}} \rightarrow \gamma$$

■

A.5. Convergence of preconditioned gradient descent

Lemma 5 When optimizing the squared loss, the iterates of PGD with preconditioner \mathbf{P} and constant step-size $\eta \leq \frac{1}{\lambda_{\max}(\mathbf{X}\mathbf{P}\mathbf{X}^\top)}$ evolve as:

$$\mathbf{w}_{PGD} = \lim_{k \rightarrow \infty} \mathbf{w}_k = \mathbf{w}_0 + \mathbf{P}\mathbf{X}^\top(\mathbf{X}\mathbf{P}\mathbf{X}^\top)^{-1}[\mathbf{y} - \mathbf{X}\mathbf{w}_0]$$

Furthermore, \mathbf{w}_{PGD} is the solution to a constrained minimization problem,

$$\mathbf{w}_{PGD} = \arg \min \frac{1}{2} \|\mathbf{w} - \mathbf{w}_0\|_{\mathbf{P}^{-1}}^2 \quad s.t. \quad \mathbf{X}\mathbf{w} = \mathbf{y}$$

Proof The PGD update for linear-regression can be written as:

$$\mathbf{w}_{k+1} = \mathbf{w}_k - \eta \mathbf{P} \nabla f(\mathbf{w}_k) = \mathbf{w}_k - \eta \mathbf{P}\mathbf{X}^\top(\mathbf{X}\mathbf{w}_k - \mathbf{y})$$

Starting at \mathbf{w}_0 and defining $\mathbf{y}_0 = \mathbf{X}\mathbf{w}_0$.

$$\mathbf{w}_1 = \mathbf{w}_0 - \eta \mathbf{P}\mathbf{X}^\top(\mathbf{X}\mathbf{w}_0 - \mathbf{y}) = \mathbf{w}_0 - \eta \mathbf{P}\mathbf{X}^\top[\mathbf{y}_0 - \mathbf{y}]$$

Further unfolding the iterates,

$$\mathbf{w}_2 = \mathbf{w}_1 - \eta \mathbf{P}\mathbf{X}^\top(\mathbf{X}\mathbf{w}_1 - \mathbf{y}) = \mathbf{w}_0 - \eta \mathbf{P}\mathbf{X}^\top[\mathbf{y}_0 - \mathbf{y}] + \eta \mathbf{P}\mathbf{X}^\top \mathbf{y} - \eta \mathbf{P}\mathbf{X}^\top(\mathbf{X}\mathbf{w}_1)$$

Adding and subtracting $\eta \mathbf{P}\mathbf{X}^\top[\mathbf{y}_0 - \mathbf{y}]$

$$\begin{aligned} &= \mathbf{w}_0 - 2\eta \mathbf{P}\mathbf{X}^\top[\mathbf{y}_0 - \mathbf{y}] + \eta \mathbf{P}\mathbf{X}^\top[\mathbf{y}_0 - \mathbf{y}] + \eta \mathbf{P}\mathbf{X}^\top \mathbf{y} - \eta \mathbf{P}\mathbf{X}^\top(\mathbf{X}\mathbf{w}_1) \\ &= \mathbf{w}_0 - 2\eta \mathbf{P}\mathbf{X}^\top[\mathbf{y}_0 - \mathbf{y}] + \eta \mathbf{P}\mathbf{X}^\top \mathbf{y}_0 - \eta \mathbf{P}\mathbf{X}^\top(\mathbf{X}(\mathbf{w}_0 - \eta \mathbf{P}\mathbf{X}^\top[\mathbf{y}_0 - \mathbf{y}])) \\ &= \mathbf{w}_0 - 2\eta \mathbf{P}\mathbf{X}^\top[\mathbf{y}_0 - \mathbf{y}] + \eta \mathbf{P}\mathbf{X}^\top \mathbf{y}_0 - \eta \mathbf{P}\mathbf{X}^\top \mathbf{y}_0 + \eta^2 \mathbf{P}\mathbf{X}^\top \mathbf{X}\mathbf{P}\mathbf{X}^\top[\mathbf{y}_0 - \mathbf{y}] \\ &= \mathbf{w}_0 - 2\eta \mathbf{P}\mathbf{X}^\top[\mathbf{y}_0 - \mathbf{y}] + \eta^2 \mathbf{P}\mathbf{X}^\top \mathbf{X}\mathbf{P}\mathbf{X}^\top[\mathbf{y}_0 - \mathbf{y}] \\ &= \mathbf{w}_0 - \mathbf{P}\mathbf{X}^\top [2\eta - \eta^2(\mathbf{X}\mathbf{P}\mathbf{X}^\top)] [\mathbf{y}_0 - \mathbf{y}] \end{aligned}$$

Defining $\mathbf{K} = \mathbf{X}\mathbf{P}\mathbf{X}^\top$

$$\mathbf{w}_2 = \mathbf{w}_0 - \mathbf{P}\mathbf{X}^\top [2\eta - \eta^2 \mathbf{K}] [\mathbf{y}_0 - \mathbf{y}]$$

Similarly writing down \mathbf{w}_3 ,

$$\begin{aligned} \mathbf{w}_3 &= \mathbf{w}_2 - \eta \mathbf{P}\mathbf{X}^\top(\mathbf{X}\mathbf{w}_2 - \mathbf{y}) = \mathbf{w}_2 + \eta \mathbf{P}\mathbf{X}^\top \mathbf{y} - \eta \mathbf{P}\mathbf{X}^\top \mathbf{X}\mathbf{w}_2 \\ &= \mathbf{w}_0 - \mathbf{P}\mathbf{X}^\top [2\eta - \eta^2 \mathbf{K}] [\mathbf{y}_0 - \mathbf{y}] + \eta \mathbf{P}\mathbf{X}^\top \mathbf{y} - \eta \mathbf{P}\mathbf{X}^\top \mathbf{X}\mathbf{w}_2 \end{aligned}$$

By adding, subtracting $\eta \mathbf{P}\mathbf{X}^\top \mathbf{y}_0$,

$$\begin{aligned}
&= \mathbf{w}_0 - \mathbf{P}\mathbf{X}^\top [3\eta - \eta^2 \mathbf{K}] [\mathbf{y}_0 - \mathbf{y}] - \eta \mathbf{P}\mathbf{X}^\top \mathbf{y}_0 - \eta \mathbf{P}\mathbf{X}^\top \mathbf{X} \mathbf{w}_2 \\
&= \mathbf{w}_0 - \mathbf{P}\mathbf{X}^\top [3\eta - \eta^2 \mathbf{K}] [\mathbf{y}_0 - \mathbf{y}] + \eta \mathbf{P}\mathbf{X}^\top \mathbf{X} \mathbf{P}\mathbf{X}^\top [2\eta - \eta^2 \mathbf{K}] [\mathbf{y}_0 - \mathbf{y}] \\
&= \mathbf{w}_0 - \mathbf{P}\mathbf{X}^\top [3\eta - \eta^2 \mathbf{K}] [\mathbf{y}_0 - \mathbf{y}] + \eta \mathbf{P}\mathbf{X}^\top \mathbf{K} [2\eta - \eta^2 \mathbf{K}] [\mathbf{y}_0 - \mathbf{y}] \\
&= \mathbf{w}_0 - \mathbf{P}\mathbf{X}^\top [3\eta - \eta^2 \mathbf{K} + 2\eta^2 \mathbf{K} - \eta^3 \mathbf{K}^2] [\mathbf{y}_0 - \mathbf{y}] \\
\mathbf{w}_3 &= \mathbf{w}_0 - \mathbf{P}\mathbf{X}^\top [3\eta + \eta^2 \mathbf{K} - \eta^3 \mathbf{K}^2] [\mathbf{y}_0 - \mathbf{y}]
\end{aligned}$$

Writing down the general form,

$$\begin{aligned}
\mathbf{w}_1 &= \mathbf{w}_0 - \mathbf{P}\mathbf{X}^\top [K^0] [\mathbf{y}_0 - \mathbf{y}] \\
\mathbf{w}_2 &= \mathbf{w}_0 - \mathbf{P}\mathbf{X}^\top [2\eta \mathbf{K}^0 - \eta^2 \mathbf{K}^1] [\mathbf{y}_0 - \mathbf{y}] \\
\mathbf{w}_3 &= \mathbf{w}_0 - \mathbf{P}\mathbf{X}^\top [3\eta \mathbf{K}^0 + \eta^2 \mathbf{K}^1 - \eta^3 \mathbf{K}^2] [\mathbf{y}_0 - \mathbf{y}] \\
\implies \mathbf{w}_k &= \mathbf{w}_0 - \mathbf{P}\mathbf{X}^\top \left[\sum_{i=1}^k (-1)^{i-1} \binom{k}{i} \eta^i \mathbf{K}^{i-1} \right] [\mathbf{y}_0 - \mathbf{y}]
\end{aligned}$$

Using the fact that $\left[\sum_{i=1}^k (-1)^{i-1} \binom{k}{i} \eta^i \mathbf{K}^{i-1} \right] = -K^{-1} \left[(\mathbf{I}_n - \eta \mathbf{K})^k - \mathbf{I}_n \right]$,

$$\mathbf{w}_k = \mathbf{w}_0 + \mathbf{P}\mathbf{X}^\top \mathbf{K}^{-1} \left[(\mathbf{I}_n - \eta \mathbf{K})^k - \mathbf{I}_n \right] [\mathbf{y}_0 - \mathbf{y}]$$

If $\eta < \frac{1}{\lambda_{\max}(\mathbf{K})}$ and as $k \rightarrow \infty$,

$$\begin{aligned}
\lim_{k \rightarrow \infty} \mathbf{w}_k &= \mathbf{w}_0 - \mathbf{P}\mathbf{X}^\top \mathbf{K}^{-1} [\mathbf{y}_0 - \mathbf{y}] \\
\implies \mathbf{w}_\infty &= \mathbf{w}_0 + \mathbf{P}\mathbf{X}^\top (\mathbf{X} \mathbf{P}\mathbf{X}^\top)^{-1} [\mathbf{y} - \mathbf{X} \mathbf{w}_0]
\end{aligned}$$

■

A.6. Properties of PGD solution

Similar to GD, zero initialized PGD converges to a unique solution lying in the span of the transformed data. Specifically, we obtain the following lemma:

Lemma 6 *The solution found by PGD initialized the origin is the unique point satisfying the constraints: (i) lies in the span of the feature vectors, $\mathbf{w}_\infty = \arg \min_{\mathbf{z} \in \text{span}(\mathbf{P}\mathbf{X}^\top)} \|\mathbf{w}_\infty - \mathbf{z}\|_{\mathbf{P}^{-1}}^2$ and (ii) interpolates the data, implying that $\mathbf{X} \mathbf{w}_\infty = \mathbf{y}$.*

Proof

$$\mathbf{w}_\infty = \arg \min_{\mathbf{z} = \mathbf{P}\mathbf{X}^\top \alpha} \|\mathbf{w}_\infty - \mathbf{z}\|_{\mathbf{P}^{-1}}^2$$

Let $\mathbf{w}_\infty = \mathbf{P}\mathbf{X}^\top \alpha_\infty$.

$$\begin{aligned}
\implies \alpha_\infty &= \arg \min_{\alpha} \|\mathbf{w}_\infty - \mathbf{P}\mathbf{X}^\top \alpha\|_{\mathbf{P}^{-1}}^2 = \arg \min_{\alpha} [\mathbf{w}_\infty^\top \mathbf{P}^{-1} \mathbf{w}_\infty - 2\mathbf{w}_\infty^\top \mathbf{X}^\top \alpha + \alpha^\top \mathbf{X} \mathbf{P}^{-1} \mathbf{X}^\top \alpha] \\
\implies \mathbf{X} \mathbf{w}_\infty &= \alpha_\infty
\end{aligned}$$

Since \mathbf{w}_∞ interpolates the data, $\mathbf{X} \mathbf{w}_\infty = \mathbf{y}$,

$$\implies \alpha_\infty = (\mathbf{X} \mathbf{P}\mathbf{X}^\top)^{-1} \mathbf{y}$$

Since, $\mathbf{w}_\infty = \mathbf{P}\mathbf{X}^\top \alpha_\infty$,

$$\implies \mathbf{w}_\infty = \mathbf{P}\mathbf{X}^\top (\mathbf{X}\mathbf{P}\mathbf{X}^\top)^{-1} \mathbf{y}$$

■

The above lemma implies that the PGD solution interpolates the data and is a projection onto $\mathbf{P}\mathbf{X}^\top$ but with the distance measured in the \mathbf{P}^{-1} norm. When $\mathbf{P} = \mathbf{I}_d$, the solution is a projection onto the span of the data points measured in the l_2 norm, recovering the known result for GD. The above lemma show that the solutions of both GD and PGD are unique interpolating solutions in their respective subspaces.

Lemma 7 *The solution to the min \mathbf{P}^{-1} norm least square i.e. $w_{opt} = \arg \min_w \|w - w_0\|_{\mathbf{P}^{-1}}$ is equal to the solution of PGD when the model interpolate the data i.e. $\mathbf{X}\mathbf{w} = \mathbf{y}$.*

Proof To solve this we reformulate our objective in the Lagrangian form where $\lambda \in R^{n \times 1}$:

$$\mathcal{L} = 1/2 \|\mathbf{w} - \mathbf{w}_0\|_{\mathbf{P}^{-1}}^2 - \lambda^\top (\mathbf{y} - \mathbf{X}\mathbf{w})$$

$$\frac{\partial \mathcal{L}}{\partial \mathbf{w}} = 0 \implies \mathbf{P}^{-1}(\mathbf{w} - \mathbf{w}_0) + \mathbf{X}^\top \lambda = 0 \implies \lambda^\top = -(\mathbf{w} - \mathbf{w}_0)^\top \mathbf{X}^\top (\mathbf{X}\mathbf{P}\mathbf{X}^\top)^{-1}$$

$$\text{Let } \mathbf{A} = \mathbf{X}^\top (\mathbf{X}\mathbf{P}\mathbf{X}^\top)^{-1}$$

$$\implies \mathcal{L} = 1/2 \|\mathbf{w} - \mathbf{w}_0\|_{\mathbf{P}^{-1}}^2 + (\mathbf{w} - \mathbf{w}_0)^\top \mathbf{A}(\mathbf{y} - \mathbf{X}\mathbf{w})$$

$$\frac{\partial \mathcal{L}}{\partial \mathbf{w}} = 0 \implies \mathbf{P}^{-1}(\mathbf{w} - \mathbf{w}_0) + \mathbf{A}\mathbf{y} - (2\mathbf{A}\mathbf{X})\mathbf{w} + \mathbf{X}^\top \mathbf{A}^\top \mathbf{w}_0 = 0$$

Let \mathbf{w}_∞ be the solution of the above inequality.

$$\implies (\mathbf{w}_\infty - \mathbf{w}_0) + \mathbf{P}\mathbf{A}\mathbf{y} - (2\mathbf{P}\mathbf{A}\mathbf{X})\mathbf{w}_\infty + \mathbf{P}\mathbf{A}\mathbf{X}\mathbf{w}_0 = 0$$

$$\mathbf{X}\mathbf{w}_\infty = \mathbf{y} \text{ because of the constraint. Let } \mathbf{y}_0 = \mathbf{X}\mathbf{w}_0 \implies \mathbf{w}_\infty - \mathbf{w}_0 + \mathbf{P}\mathbf{A}\mathbf{y} - 2\mathbf{P}\mathbf{A}\mathbf{y} + \mathbf{P}\mathbf{A}\mathbf{y}_0 = 0$$

$$\implies \mathbf{w}_\infty = \mathbf{w}_0 + \mathbf{P}\mathbf{A}(\mathbf{y} - \mathbf{y}_0) = \mathbf{w}_0 + \mathbf{P}\mathbf{X}^\top (\mathbf{X}\mathbf{P}\mathbf{X}^\top)^{-1} [\mathbf{y} - \mathbf{X}\mathbf{w}_0]$$

$$\implies \mathbf{w}_\infty = [\mathbf{I}_d - \mathbf{P}\mathbf{X}^\top (\mathbf{X}\mathbf{P}\mathbf{X}^\top)^{-1} \mathbf{X}] \mathbf{w}_0 + \mathbf{P}\mathbf{X}^\top (\mathbf{X}\mathbf{P}\mathbf{X}^\top)^{-1} \mathbf{y}$$

■

A.7. Convergence of Newton method

For the Newton method, for which the iterates can be written as: $\mathbf{w}_{k+1} = \mathbf{w}_k - \eta [\nabla^2 f(\mathbf{w}_k) + \lambda \mathbf{I}_d]^{-1} \nabla f(\mathbf{w}_k)$. Here, $\nabla^2 f(\mathbf{w}_k)$ is the Hessian equal to $\mathbf{X}^\top \mathbf{X}$ for linear regression and λ is the LM regularization parameter.

Lemma 8 *The Newton method remains in the span of the data-points and hence converges to the min-norm solution.*

Proof Let assume $f(\mathbf{w}) = F(\mathbf{X}\mathbf{w})$. Therefore we have $\nabla f(\mathbf{w}) = \mathbf{X}^\top \nabla F(\mathbf{X}\mathbf{w})$ and $\nabla^2 f(\mathbf{w}) = \mathbf{X}^\top \nabla^2 F(\mathbf{X}\mathbf{w}) \mathbf{X}$. If we assume \mathbf{w}_k in the span, we just need to show that $\nu_k = \nabla^2 f(\mathbf{w}_k) + \lambda \mathbf{I}_d]^{-1} \nabla f(\mathbf{w}_k)$ lays in the span.

$$\nu_k = [\nabla^2 f(\mathbf{w}_k) + \lambda \mathbf{I}_d]^{-1} \nabla f(\mathbf{w}_k) \implies \nu_k = \frac{1}{\lambda} (\nabla f(\mathbf{w}_k) - \nabla^2 f(\mathbf{w}_k) \nu_k)$$

$$\implies \nu_k = \frac{\mathbf{X}^\top}{\lambda} (\nabla F(\mathbf{X}\mathbf{w}) - \nabla^2 F(\mathbf{X}\mathbf{w}) \mathbf{X} \nu_k)$$

$$\implies \nu_k = \mathbf{X}^\top \alpha_k$$

for some vector α_k . This shows that ν_k and therefore \mathbf{w}_{k+1} lay in the span.

■

A.8. Convergence of full-matrix Adagrad

For full matrix Adagrad [1], the iterates can be written as: $\mathbf{w}_{k+1} = \mathbf{w}_k - \eta_k \mathbf{G}_k \nabla f(\mathbf{w}_k)$ where $\mathbf{G}_k = \mathbf{S}_k^{-1/2}$ and $\mathbf{S}_k = \mathbf{S}_{k-1} + (\nabla f(\mathbf{w}_{k-1}) \nabla f(\mathbf{w}_{k-1})^T)$. The more commonly used diagonal version of Adagrad uses $\mathbf{G}_k = \text{diag}(\mathbf{S}_k)$. We obtain the following lemma, analyzing the convergence of these variants.

Lemma 9 *The iterates of full matrix Adagrad lie in the span of the data and converges to the min-norm solution. However, convergence to the min-norm solution is not ensured for the diagonal version of Adagrad.*

Proof The proof is similar to the proof of 8. ■

A.9. Generalization bounds

Lemma 10 *The excess risk for the solution \mathbf{w}_{PGD} of PGD with a preconditioner P and initialized at \mathbf{w}_0 can be bounded as:*

$$R(\mathbf{w}_{PGD}) := \mathbb{E}_{\mathbf{x}, \epsilon} [\mathbf{x}^T (\mathbf{w}_{PGD} - \mathbf{w}^*)]^2 \leq (\mathbf{w}_0 - \mathbf{w}^*)^T \mathbf{B}_P^T \Sigma \mathbf{B}_P (\mathbf{w}_0 - \mathbf{w}^*) + \sigma^2 \text{tr}(\mathbf{C}_P) \quad (15)$$

where $\Sigma = \mathbb{E}[\mathbf{x}\mathbf{x}^T]$ is the covariance matrix s.t. $\mathbf{x} \sim N(0, \Sigma)$ and $\mathbf{B}_P = \mathbf{I} - \mathbf{P}(\mathbf{X}^T(\mathbf{X}\mathbf{P}\mathbf{X}^T)^{-1})\mathbf{X}$ and $\mathbf{C}_P = (\mathbf{X}\mathbf{P}\mathbf{X}^T)^{-1}\mathbf{X}\mathbf{P}^T \Sigma \mathbf{P}\mathbf{X}^T(\mathbf{X}\mathbf{P}\mathbf{X}^T)^{-1}$.

Proof Recall that $\mathbf{w}_\infty = \mathbf{w}_0 + \mathbf{P}\mathbf{A}(\mathbf{y} - \mathbf{y}_0) = \mathbf{w}_0 + \mathbf{P}\mathbf{X}^T(\mathbf{X}\mathbf{P}\mathbf{X}^T)^{-1}[\mathbf{y} - \mathbf{X}\mathbf{w}_0]$. Define matrix $\mathbf{A} = \mathbf{X}^T(\mathbf{X}\mathbf{P}\mathbf{X}^T)^{-1}$, implying $\mathbf{w}_\infty = \mathbf{w}_0 + \mathbf{P}\mathbf{A}[\mathbf{y} - \mathbf{X}\mathbf{w}_0] = \mathbf{w}_0 + \mathbf{P}\mathbf{A}[\mathbf{X}\mathbf{w}^* + \epsilon - \mathbf{X}\mathbf{w}_0]$

$$\begin{aligned} R(\mathbf{w}_\infty) &= \mathbb{E}_x [(\mathbf{X}^T (\mathbf{w}_\infty - \mathbf{w}^*))^2] \\ &= \mathbb{E}_x [(\mathbf{X}^T (\mathbf{w}_0 + \mathbf{P}\mathbf{A}[\mathbf{X}\mathbf{w}^* + \epsilon - \mathbf{X}\mathbf{w}_0] - \mathbf{w}^*))^2] \\ &= \mathbb{E}_x [((\mathbf{X}^T ((\mathbf{I} - \mathbf{P}\mathbf{A}\mathbf{X})) (\mathbf{w}_0 - \mathbf{w}^*) + \mathbf{X}^T \mathbf{P}\mathbf{A}\epsilon)^2] \\ &\leq 2\mathbb{E}_x [(\mathbf{X}^T ((\mathbf{I} - \mathbf{P}\mathbf{A}\mathbf{X})) (\mathbf{w}_0 - \mathbf{w}^*))^2] + 2\mathbb{E}_x [(\mathbf{X}^T \mathbf{P}\mathbf{A}\epsilon)^2] \end{aligned}$$

Using the fact that $\Sigma = \mathbb{E}[\mathbf{x}\mathbf{x}^T]$, $\mathbf{B} = \mathbf{I} - \mathbf{P}\mathbf{A}\mathbf{X}$ and $\mathbf{C} = \mathbf{A}^T \mathbf{P}^T \Sigma \mathbf{P}\mathbf{A}$.

$$= (\mathbf{w}_0 - \mathbf{w}^*)^T \mathbf{B}^T \Sigma \mathbf{B} (\mathbf{w}_0 - \mathbf{w}^*) + 2\epsilon^T \mathbf{C} \epsilon$$

Take expectation w.r.t. the noise:

$$\mathbb{E}_\epsilon [R(\mathbf{w}_\infty)] \leq (\mathbf{w}_0 - \mathbf{w}^*)^T \mathbf{B}^T \Sigma \mathbf{B} (\mathbf{w}_0 - \mathbf{w}^*) + \mathbb{E}_\epsilon [\epsilon^T \mathbf{C} \epsilon] = (\mathbf{w}_0 - \mathbf{w}^*)^T \mathbf{B}^T \Sigma \mathbf{B} (\mathbf{w}_0 - \mathbf{w}^*) + \sigma^2 \text{tr}(\mathbf{C})$$

since the noise has mean zero and variance σ^2 . ■

A.10. Generalization bounds

Lemma 11 *The excess risk for the solution \mathbf{w}_{PGD} of PGD with a preconditioner P and initialized at \mathbf{w}_0 can be bounded as:*

$$R(\mathbf{w}_{PGD}) := \mathbb{E}_{\mathbf{x}, \epsilon} [\mathbf{x}^T (\mathbf{w}_{PGD} - \mathbf{w}^*)]^2 \leq (\mathbf{w}_0 - \mathbf{w}^*)^T \mathbf{B}_P^T \Sigma \mathbf{B}_P (\mathbf{w}_0 - \mathbf{w}^*) + \sigma^2 \text{tr}(\mathbf{C}_P) \quad (16)$$

where $\Sigma = \mathbb{E}[\mathbf{x}\mathbf{x}^T]$ is the covariance matrix s.t. $\mathbf{x} \sim N(0, \Sigma)$ and $\mathbf{B}_P = \mathbf{I} - \mathbf{P}(\mathbf{X}^T(\mathbf{X}\mathbf{P}\mathbf{X}^T)^{-1})\mathbf{X}$ and $\mathbf{C}_P = (\mathbf{X}\mathbf{P}\mathbf{X}^T)^{-1}\mathbf{X}\mathbf{P}^T \Sigma \mathbf{P}\mathbf{X}^T(\mathbf{X}\mathbf{P}\mathbf{X}^T)^{-1}$.

Proof Recall that $\mathbf{w}_\infty = \mathbf{w}_0 + \mathbf{P}\mathbf{A}(\mathbf{y} - \mathbf{y}_0) = \mathbf{w}_0 + \mathbf{P}\mathbf{X}^T(\mathbf{X}\mathbf{P}\mathbf{X}^T)^{-1}[\mathbf{y} - \mathbf{X}\mathbf{w}_0]$. Define matrix $\mathbf{A} = \mathbf{X}^T(\mathbf{X}\mathbf{P}\mathbf{X}^T)^{-1}$, implying $\mathbf{w}_\infty = \mathbf{w}_0 + \mathbf{P}\mathbf{A}[\mathbf{y} - \mathbf{X}\mathbf{w}_0] = \mathbf{w}_0 + \mathbf{P}\mathbf{A}[\mathbf{X}\mathbf{w}^* + \epsilon - \mathbf{X}\mathbf{w}_0]$

$$\begin{aligned} R(\mathbf{w}_\infty) &= \mathbb{E}_x [(\mathbf{X}^T (\mathbf{w}_\infty - \mathbf{w}^*))^2] \\ &= \mathbb{E}_x [(\mathbf{X}^T (\mathbf{w}_0 + \mathbf{P}\mathbf{A}[\mathbf{X}\mathbf{w}^* + \epsilon - \mathbf{X}\mathbf{w}_0] - \mathbf{w}^*))^2] \\ &= \mathbb{E}_x [((\mathbf{X}^T ((\mathbf{I} - \mathbf{P}\mathbf{A}\mathbf{X})) (\mathbf{w}_0 - \mathbf{w}^*) + \mathbf{X}^T \mathbf{P}\mathbf{A}\epsilon)^2] \\ &\leq 2\mathbb{E}_x [(\mathbf{X}^T ((\mathbf{I} - \mathbf{P}\mathbf{A}\mathbf{X})) (\mathbf{w}_0 - \mathbf{w}^*))^2] + 2\mathbb{E}_x [(\mathbf{X}^T \mathbf{P}\mathbf{A}\epsilon)^2] \end{aligned}$$

Using the fact that $\Sigma = \mathbb{E}[\mathbf{x}\mathbf{x}^\top]$, $\mathbf{B} = \mathbf{I} - \mathbf{P}\mathbf{A}\mathbf{X}$ and $\mathbf{C} = \mathbf{A}^\top \mathbf{P}^\top \Sigma \mathbf{P}\mathbf{A}$.

$$= (\mathbf{w}_0 - \mathbf{w}^*)^\top \mathbf{B}^\top \Sigma \mathbf{B} (\mathbf{w}_0 - \mathbf{w}^*) + 2\epsilon^\top \mathbf{C} \epsilon$$

Take expectation w.r.t. the noise:

$$\mathbb{E}_\epsilon[R(\mathbf{w}_\infty)] \leq (\mathbf{w}_0 - \mathbf{w}^*)^\top \mathbf{B}^\top \Sigma \mathbf{B} (\mathbf{w}_0 - \mathbf{w}^*) + \mathbb{E}_\epsilon[\epsilon^\top \mathbf{C} \epsilon] = (\mathbf{w}_0 - \mathbf{w}^*)^\top \mathbf{B}^\top \Sigma \mathbf{B} (\mathbf{w}_0 - \mathbf{w}^*) + \sigma^2 \text{tr}(\mathbf{C})$$

since the noise has mean zero and variance σ^2 . ■

A.10.1. FINDING AN OPTIMAL PRECONDITIONER

It is clear that if we do not have additional information in the form of a validation set, remaining in the span of the points is the optimal strategy for an optimizer and the min-norm solution results in the best generalization for a general \mathbf{w}^* and Σ .

Consequently, we consider a semi-supervised setting and investigate whether we can exploit unlabelled data and obtain better generalization. We consider an idealized case where we have infinite unlabelled data that enables us to get an accurate estimate of the true covariance matrix Σ . We prove the following theorem,

Lemma 12 *Assume we have sufficient unlabeled data to estimate Σ accurately. In the regression with noisy data we can leverage the unlabeled data to design an optimal preconditioner so as to reduce generalization risk. In the noiseless case, the unlabeled data doesn't help.*

This lemma shows that in the noiseless setting, having a preconditioner that keeps the update inside the data span can only help improving the convergence speed of the optimization regardless of the size of unlabeled data or the exact value of covariance matrix. However in the noisy case, we can find a preconditioner that has the minimum generalization error among all preconditioners including the identity matrix which gives us GD. Our experimental results for this section confirm our theoretical result.

Proof We can decompose Σ as follows:

$$\Sigma = \mathbf{X}^\top \mathbf{S}_1 \mathbf{X} + \mathbf{X}^\top \mathbf{S}_2 \mathbf{X}_\perp + \mathbf{X}_\perp^\top \mathbf{S}_2^\top \mathbf{X} + \mathbf{X}_\perp^\top \mathbf{S}_3 \mathbf{X}_\perp, \quad (17)$$

and similarly we can decompose preconditioner:

$$\mathbf{P} = \mathbf{X}^\top \mathbf{P}_1 \mathbf{X} + \mathbf{X}^\top \mathbf{P}_2 \mathbf{X}_\perp + \mathbf{X}_\perp^\top \mathbf{P}_2^\top \mathbf{X} + \mathbf{X}_\perp^\top \mathbf{P}_3 \mathbf{X}_\perp, \quad (18)$$

where $\Sigma, \mathbf{P} \in \mathbb{R}^{d \times d}$, $\mathbf{S}_1, \mathbf{P}_1 \in \mathbb{R}^{n \times n}$, $\mathbf{S}_2, \mathbf{P}_2 \in \mathbb{R}^{n \times d-n}$, and $\mathbf{S}_3, \mathbf{P}_3 \in \mathbb{R}^{d-n \times d-n}$. Recall the excess risk for linear regression

$$R(\mathbf{w}_\infty^{opt}) = \mathbf{w}^{*\top} (\mathbf{I} - \mathbf{P}\mathbf{X}^\top (\mathbf{X}\mathbf{P}\mathbf{X}^\top)^{-1} \mathbf{X})^\top \Sigma (\mathbf{I} - \mathbf{P}\mathbf{X}^\top (\mathbf{X}\mathbf{P}\mathbf{X}^\top)^{-1} \mathbf{X}) \mathbf{w}^*.$$

To apply the above decompositions, note that we have

$$\mathbf{P}\mathbf{X}^\top = \mathbf{X}^\top\mathbf{P}_1\mathbf{X}\mathbf{X}^\top + \mathbf{X}_\perp^\top\mathbf{P}_2^\top\mathbf{X}\mathbf{X}^\top \quad (19)$$

$$(\mathbf{X}\mathbf{P}\mathbf{X}^\top)^{-1} = (\mathbf{X}\mathbf{X}^\top)^{-1}\mathbf{P}_1^{-1}(\mathbf{X}\mathbf{X}^\top)^{-1} \quad (20)$$

$$\mathbf{P}\mathbf{X}^\top(\mathbf{X}\mathbf{P}\mathbf{X}^\top)^{-1} = \mathbf{X}^\top(\mathbf{X}\mathbf{X}^\top)^{-1} + \mathbf{X}_\perp^\top\mathbf{P}_2^\top\mathbf{P}_1^{-1}(\mathbf{X}\mathbf{X}^\top)^{-1} \quad (21)$$

$$\mathbf{B} = \mathbf{I} - \mathbf{P}\mathbf{X}^\top(\mathbf{X}\mathbf{P}\mathbf{X}^\top)^{-1}\mathbf{X} = \mathbf{I} - \mathbf{X}^\top(\mathbf{X}\mathbf{X}^\top)^{-1}\mathbf{X} - \mathbf{X}_\perp^\top\mathbf{P}_2^\top\mathbf{P}_1^{-1}(\mathbf{X}\mathbf{X}^\top)^{-1}\mathbf{X} \quad (22)$$

$$\mathbf{Q} = \mathbf{P}_2^\top\mathbf{P}_1^{-1} \quad (23)$$

$$\mathbf{F} = \mathbf{X}_\perp^\top\mathbf{P}_2^\top\mathbf{P}_1^{-1}(\mathbf{X}\mathbf{X}^\top)^{-1}\mathbf{X} \quad (24)$$

$$\Sigma\mathbf{B} = \Sigma - \mathbf{X}^\top\mathbf{S}_1\mathbf{X} - \mathbf{X}_\perp^\top\mathbf{S}_2^\top\mathbf{X} - \mathbf{X}^\top\mathbf{S}_2\mathbf{X}_\perp\mathbf{F} - \mathbf{X}_\perp^\top\mathbf{S}_3\mathbf{X}_\perp^\top\mathbf{F} \quad (25)$$

$$= \mathbf{X}^\top\mathbf{S}_2\mathbf{X}_\perp + \mathbf{X}_\perp^\top\mathbf{S}_3\mathbf{X}_\perp - \mathbf{X}^\top\mathbf{S}_2\mathbf{F} - \mathbf{X}_\perp^\top\mathbf{S}_3\mathbf{X}_\perp\mathbf{F} \quad (26)$$

$$\mathbf{B}^\top\Sigma\mathbf{B} = (\mathbf{I} - \mathbf{X}^\top(\mathbf{X}\mathbf{X}^\top)^{-1}\mathbf{X} - \mathbf{F}^\top)\Sigma\mathbf{B} \quad (27)$$

$$= \mathbf{X}^\top\mathbf{S}_2\mathbf{X}_\perp + \mathbf{X}_\perp^\top\mathbf{S}_3\mathbf{X}_\perp - \mathbf{X}^\top\mathbf{S}_2\mathbf{F} - \mathbf{X}_\perp^\top\mathbf{S}_3\mathbf{X}_\perp\mathbf{F} \quad (28)$$

$$- \mathbf{X}^\top\mathbf{S}_2\mathbf{X}_\perp + \mathbf{X}^\top\mathbf{S}_2\mathbf{F} \quad (29)$$

$$- \mathbf{F}^\top\mathbf{X}^\top\mathbf{S}_2\mathbf{X}_\perp - \mathbf{F}^\top\mathbf{X}_\perp^\top\mathbf{S}_3\mathbf{X}_\perp + \mathbf{F}^\top\mathbf{X}^\top\mathbf{S}_2\mathbf{X}_\perp\mathbf{F} + \mathbf{F}^\top\mathbf{X}_\perp^\top\mathbf{S}_3\mathbf{X}_\perp\mathbf{F} \quad (30)$$

$$= \mathbf{X}_\perp^\top\mathbf{S}_3\mathbf{X}_\perp - \mathbf{X}_\perp^\top\mathbf{S}_3\mathbf{X}_\perp\mathbf{F} - \mathbf{F}^\top\mathbf{X}_\perp^\top\mathbf{S}_3\mathbf{X}_\perp + \mathbf{F}^\top\mathbf{X}_\perp^\top\mathbf{S}_3\mathbf{X}_\perp\mathbf{F} \quad (31)$$

$$\text{tr}(\mathbf{F}^\top\mathbf{X}_\perp^\top\mathbf{S}_3\mathbf{X}_\perp) = 0 \quad (32)$$

$$\text{tr}(\mathbf{B}^\top\Sigma\mathbf{B}) = \text{tr}(\mathbf{X}_\perp^\top\mathbf{S}_3\mathbf{X}_\perp) + \text{tr}(\mathbf{F}^\top\mathbf{X}_\perp^\top\mathbf{S}_3\mathbf{X}_\perp\mathbf{F}) = \text{const} + \text{tr}(\mathbf{F}^\top\mathbf{X}_\perp^\top\mathbf{S}_3\mathbf{X}_\perp\mathbf{F}) \quad (33)$$

In this case we see that the best trace is achieved when we set $\mathbf{P}_2 = 0$ which is independent of the information of Σ . Now assume noisy case whose excess risk has extra term which is:

$$\text{tr}(\mathbf{C}) = \sigma^2\text{tr}((\mathbf{X}\mathbf{P}\mathbf{X}^\top)^{-1}\mathbf{X}\mathbf{P}\Sigma\mathbf{P}\mathbf{X}^\top(\mathbf{X}\mathbf{P}\mathbf{X}^\top)^{-1}). \quad (34)$$

Now we expand this extra term based on the aforementioned decompositions

$$\Sigma\mathbf{P}\mathbf{X}^\top(\mathbf{X}\mathbf{P}\mathbf{X}^\top)^{-1} = (\mathbf{X}^\top\mathbf{S}_1\mathbf{X} + \mathbf{X}^\top\mathbf{S}_2\mathbf{X}_\perp + \mathbf{X}_\perp^\top\mathbf{S}_2^\top\mathbf{X} + \mathbf{X}_\perp^\top\mathbf{S}_3\mathbf{X}_\perp)(\mathbf{X}^\top(\mathbf{X}\mathbf{X}^\top)^{-1} + \mathbf{X}_\perp^\top\mathbf{P}_2^\top\mathbf{P}_1^{-1}(\mathbf{X}\mathbf{X}^\top)^{-1}) \quad (35)$$

$$= \mathbf{X}^\top\mathbf{S}_1 + \mathbf{X}^\top\mathbf{S}_2\mathbf{X}_\perp\mathbf{X}_\perp^\top\mathbf{Q}(\mathbf{X}\mathbf{X}^\top)^{-1} + \mathbf{X}_\perp^\top\mathbf{S}_2^\top + \mathbf{X}_\perp^\top\mathbf{S}_3\mathbf{X}_\perp\mathbf{X}_\perp^\top\mathbf{Q}(\mathbf{X}\mathbf{X}^\top)^{-1} \quad (36)$$

$$(\mathbf{X}\mathbf{P}\mathbf{X}^\top)^{-1}\mathbf{X}\mathbf{P}\Sigma\mathbf{P}\mathbf{X}^\top(\mathbf{X}\mathbf{P}\mathbf{X}^\top)^{-1} = \mathbf{S}_1 + \mathbf{S}_2\mathbf{X}_\perp\mathbf{X}_\perp^\top\mathbf{Q}(\mathbf{X}\mathbf{X}^\top)^{-1} \quad (37)$$

$$+ (\mathbf{X}\mathbf{X}^\top)^{-1}\mathbf{Q}^\top\mathbf{X}_\perp\mathbf{X}_\perp^\top\mathbf{S}_2^\top + (\mathbf{X}\mathbf{X}^\top)^{-1}\mathbf{Q}^\top\mathbf{X}_\perp\mathbf{X}_\perp^\top\mathbf{S}_3\mathbf{X}_\perp\mathbf{X}_\perp^\top\mathbf{Q}(\mathbf{X}\mathbf{X}^\top)^{-1} \quad (38)$$

So the the excess risk is upperbounded by

$$\text{tr}(\mathbf{B}^\top\Sigma\mathbf{B}) + \sigma^2\text{tr}(\mathbf{C}) = \text{tr}(\mathbf{X}_\perp^\top\mathbf{S}_3\mathbf{X}_\perp) + \text{tr}(\mathbf{F}^\top\mathbf{X}_\perp^\top\mathbf{S}_3\mathbf{X}_\perp\mathbf{F}) \quad (39)$$

$$+ \sigma^2(\mathbf{S}_1 + \mathbf{S}_2\mathbf{X}_\perp\mathbf{X}_\perp^\top\mathbf{Q}(\mathbf{X}\mathbf{X}^\top)^{-1} + (\mathbf{X}\mathbf{X}^\top)^{-1}\mathbf{Q}^\top\mathbf{X}_\perp\mathbf{X}_\perp^\top\mathbf{S}_2^\top + (\mathbf{X}\mathbf{X}^\top)^{-1}\mathbf{Q}^\top\mathbf{X}_\perp\mathbf{X}_\perp^\top\mathbf{S}_3\mathbf{X}_\perp\mathbf{X}_\perp^\top\mathbf{Q}(\mathbf{X}\mathbf{X}^\top)^{-1}) \quad (40)$$

Now if we take derivative w.r.t. \mathbf{Q} and set it to zero we get

$$\mathbf{Q}^* = -(\mathbf{K}_\perp\mathbf{S}_3\mathbf{K}_\perp)^{-1}(\sigma^2\mathbf{K}_\perp\mathbf{S}_2^\top\mathbf{K}^{-1})(\mathbf{K}^{-1} + \mathbf{K}^{-2})^{-1} \quad (41)$$

where $\mathbf{K} = \mathbf{X}\mathbf{X}^\top$ and $\mathbf{K}_\perp = \mathbf{X}_\perp\mathbf{X}_\perp^\top$. By setting $\mathbf{P}_1 = \mathbf{I}$, we have $\mathbf{P}_2 = \mathbf{Q}^{*\top}$. Therefore an optimal \mathbf{P} is

$$\mathbf{P} = \mathbf{X}^\top\mathbf{X} + \mathbf{X}^\top\mathbf{Q}^{*\top}\mathbf{X}_\perp + \mathbf{X}_\perp^\top\mathbf{Q}^*\mathbf{P} \quad (42)$$

Note that here we assume we have enough unlabeled data which among them we can pick $n - d$ of them which are orthogonal to the training data i.e.

$$\mathbf{X}^\top\mathbf{X}_\perp = \mathbf{X}_\perp^\top\mathbf{X} = 0 \quad (43)$$

■

A.11. Finding an Optimal Preconditioner

It is clear that if we do not have additional information in the form of a validation set, remaining in the span of the points is the optimal strategy for an optimizer and the min-norm solution results in the best generalization for a general \mathbf{w}^* and Σ .

Consequently, we consider a semi-supervised setting and investigate whether we can exploit unlabelled data and obtain better generalization. We consider an idealized case where we have infinite unlabelled data that enables us to get an accurate estimate of the true covariance matrix Σ . We prove the following theorem,

Lemma 13 *Assume we have sufficient unlabeled data to estimate Σ accurately. In the regression with noisy data we can leverage the unlabeled data to design an optimal preconditioner so as to reduce generalization risk. In the noiseless case, the unlabeled data doesn't help.*

This lemma shows that in the noiseless setting, having a preconditioner that keeps the update inside the data span can only help improving the convergence speed of the optimization regardless of the size of unlabeled data or the exact value of covariance matrix. However in the noisy case, we can find a preconditioner that has the minimum generalization error among all preconditioners including the identity matrix which gives us GD. Our experimental results for this section confirm our theoretical result.

Proof We can decompose Σ as follows:

$$\Sigma = \mathbf{X}^\top \mathbf{S}_1 \mathbf{X} + \mathbf{X}^\top \mathbf{S}_2 \mathbf{X}_\perp + \mathbf{X}_\perp^\top \mathbf{S}_2^\top \mathbf{X} + \mathbf{X}_\perp^\top \mathbf{S}_3 \mathbf{X}_\perp, \quad (44)$$

and similarly we can decompose preconditioner:

$$\mathbf{P} = \mathbf{X}^\top \mathbf{P}_1 \mathbf{X} + \mathbf{X}^\top \mathbf{P}_2 \mathbf{X}_\perp + \mathbf{X}_\perp^\top \mathbf{P}_2^\top \mathbf{X} + \mathbf{X}_\perp^\top \mathbf{P}_3 \mathbf{X}_\perp, \quad (45)$$

where $\Sigma, \mathbf{P} \in \mathbb{R}^{d \times d}$, $\mathbf{S}_1, \mathbf{P}_1 \in \mathbb{R}^{n \times n}$, $\mathbf{S}_2, \mathbf{P}_2 \in \mathbb{R}^{n \times d-n}$, and $\mathbf{S}_3, \mathbf{P}_3 \in \mathbb{R}^{d-n \times d-n}$. Recall the excess risk for linear regression

$$R(\mathbf{w}_\infty^{opt}) = \mathbf{w}^{*\top} (\mathbf{I} - \mathbf{P} \mathbf{X}^\top (\mathbf{X} \mathbf{P} \mathbf{X}^\top)^{-1} \mathbf{X})^\top \Sigma (\mathbf{I} - \mathbf{P} \mathbf{X}^\top (\mathbf{X} \mathbf{P} \mathbf{X}^\top)^{-1} \mathbf{X}) \mathbf{w}^*.$$

To apply the above decompositions, note that we have

$$\mathbf{P} \mathbf{X}^\top = \mathbf{X}^\top \mathbf{P}_1 \mathbf{X} \mathbf{X}^\top + \mathbf{X}_\perp^\top \mathbf{P}_2^\top \mathbf{X} \mathbf{X}^\top \quad (46)$$

$$(\mathbf{X} \mathbf{P} \mathbf{X}^\top)^{-1} = (\mathbf{X} \mathbf{X}^\top)^{-1} \mathbf{P}_1^{-1} (\mathbf{X} \mathbf{X}^\top)^{-1} \quad (47)$$

$$\mathbf{P} \mathbf{X}^\top (\mathbf{X} \mathbf{P} \mathbf{X}^\top)^{-1} = \mathbf{X}^\top (\mathbf{X} \mathbf{X}^\top)^{-1} + \mathbf{X}_\perp^\top \mathbf{P}_2^\top \mathbf{P}_1^{-1} (\mathbf{X} \mathbf{X}^\top)^{-1} \quad (48)$$

$$\mathbf{B} = \mathbf{I} - \mathbf{P} \mathbf{X}^\top (\mathbf{X} \mathbf{P} \mathbf{X}^\top)^{-1} \mathbf{X} = \mathbf{I} - \mathbf{X}^\top (\mathbf{X} \mathbf{X}^\top)^{-1} \mathbf{X} - \mathbf{X}_\perp^\top \mathbf{P}_2^\top \mathbf{P}_1^{-1} (\mathbf{X} \mathbf{X}^\top)^{-1} \mathbf{X} \quad (49)$$

$$\mathbf{Q} = \mathbf{P}_2^\top \mathbf{P}_1^{-1} \quad (50)$$

$$\mathbf{F} = \mathbf{X}_\perp^\top \mathbf{P}_2^\top \mathbf{P}_1^{-1} (\mathbf{X} \mathbf{X}^\top)^{-1} \mathbf{X} \quad (51)$$

$$\Sigma \mathbf{B} = \Sigma - \mathbf{X}^\top \mathbf{S}_1 \mathbf{X} - \mathbf{X}_\perp^\top \mathbf{S}_2^\top \mathbf{X} - \mathbf{X}^\top \mathbf{S}_2 \mathbf{X}_\perp \mathbf{F} - \mathbf{X}_\perp^\top \mathbf{S}_3 \mathbf{X}_\perp^\top \mathbf{F} \quad (52)$$

$$= \mathbf{X}^\top \mathbf{S}_2 \mathbf{X}_\perp + \mathbf{X}_\perp^\top \mathbf{S}_3 \mathbf{X}_\perp - \mathbf{X}^\top \mathbf{S}_2 \mathbf{F} - \mathbf{X}_\perp^\top \mathbf{S}_3 \mathbf{X}_\perp \mathbf{F} \quad (53)$$

$$\mathbf{B}^\top \Sigma \mathbf{B} = (\mathbf{I} - \mathbf{X}^\top (\mathbf{X} \mathbf{X}^\top)^{-1} \mathbf{X} - \mathbf{F}^\top) \Sigma \mathbf{B} \quad (54)$$

$$= \mathbf{X}^\top \mathbf{S}_2 \mathbf{X}_\perp + \mathbf{X}_\perp^\top \mathbf{S}_3 \mathbf{X}_\perp - \mathbf{X}^\top \mathbf{S}_2 \mathbf{F} - \mathbf{X}_\perp^\top \mathbf{S}_3 \mathbf{X}_\perp \mathbf{F} \quad (55)$$

$$- \mathbf{X}^\top \mathbf{S}_2 \mathbf{X}_\perp + \mathbf{X}^\top \mathbf{S}_2 \mathbf{F} \quad (56)$$

$$- \mathbf{F}^\top \mathbf{X}^\top \mathbf{S}_2 \mathbf{X}_\perp - \mathbf{F}^\top \mathbf{X}_\perp^\top \mathbf{S}_3 \mathbf{X}_\perp + \mathbf{F}^\top \mathbf{X}^\top \mathbf{S}_2 \mathbf{X}_\perp \mathbf{F} + \mathbf{F}^\top \mathbf{X}_\perp^\top \mathbf{S}_3 \mathbf{X}_\perp \mathbf{F} \quad (57)$$

$$= \mathbf{X}_\perp^\top \mathbf{S}_3 \mathbf{X}_\perp - \mathbf{X}_\perp^\top \mathbf{S}_3 \mathbf{X}_\perp \mathbf{F} - \mathbf{F}^\top \mathbf{X}_\perp^\top \mathbf{S}_3 \mathbf{X}_\perp + \mathbf{F}^\top \mathbf{X}_\perp^\top \mathbf{S}_3 \mathbf{X}_\perp \mathbf{F} \quad (58)$$

$$\text{tr}(\mathbf{F}^\top \mathbf{X}_\perp^\top \mathbf{S}_3 \mathbf{X}_\perp) = 0 \quad (59)$$

$$\text{tr}(\mathbf{B}^\top \Sigma \mathbf{B}) = \text{tr}(\mathbf{X}_\perp^\top \mathbf{S}_3 \mathbf{X}_\perp) + \text{tr}(\mathbf{F}^\top \mathbf{X}_\perp^\top \mathbf{S}_3 \mathbf{X}_\perp \mathbf{F}) = \text{cnst} + \text{tr}(\mathbf{F}^\top \mathbf{X}_\perp^\top \mathbf{S}_3 \mathbf{X}_\perp \mathbf{F}) \quad (60)$$

In this case we see that the best trace is achieved when we set $\mathbf{P}_2 = 0$ which is independent of the information of Σ . Now assume noisy case whose excess risk has extra term which is:

$$\text{tr}(\mathbf{C}) = \sigma^2 \text{tr}((\mathbf{X} \mathbf{P} \mathbf{X}^\top)^{-1} \mathbf{X} \mathbf{P} \Sigma \mathbf{P} \mathbf{X}^\top (\mathbf{X} \mathbf{P} \mathbf{X}^\top)^{-1}). \quad (61)$$

Now we expand this extra term based on the aforementioned decompositions

$$\Sigma \mathbf{P} \mathbf{X}^\top (\mathbf{X} \mathbf{P} \mathbf{X}^\top)^{-1} = (\mathbf{X}^\top \mathbf{S}_1 \mathbf{X} + \mathbf{X}^\top \mathbf{S}_2 \mathbf{X}_\perp + \mathbf{X}_\perp^\top \mathbf{S}_2^\top \mathbf{X} + \mathbf{X}_\perp^\top \mathbf{S}_3 \mathbf{X}_\perp) (\mathbf{X}^\top (\mathbf{X} \mathbf{X}^\top)^{-1} + \mathbf{X}_\perp^\top \mathbf{P}_2^\top \mathbf{P}_1^{-1} (\mathbf{X} \mathbf{X}^\top)^{-1}) \quad (62)$$

$$= \mathbf{X}^\top \mathbf{S}_1 + \mathbf{X}^\top \mathbf{S}_2 \mathbf{X}_\perp \mathbf{X}_\perp^\top \mathbf{Q} (\mathbf{X} \mathbf{X}^\top)^{-1} + \mathbf{X}_\perp^\top \mathbf{S}_2^\top + \mathbf{X}_\perp^\top \mathbf{S}_3 \mathbf{X}_\perp \mathbf{X}_\perp^\top \mathbf{Q} (\mathbf{X} \mathbf{X}^\top)^{-1} \quad (63)$$

$$(\mathbf{X} \mathbf{P} \mathbf{X}^\top)^{-1} \mathbf{X} \mathbf{P} \Sigma \mathbf{P} \mathbf{X}^\top (\mathbf{X} \mathbf{P} \mathbf{X}^\top)^{-1} = \mathbf{S}_1 + \mathbf{S}_2 \mathbf{X}_\perp \mathbf{X}_\perp^\top \mathbf{Q} (\mathbf{X} \mathbf{X}^\top)^{-1} \quad (64)$$

$$+ (\mathbf{X} \mathbf{X}^\top)^{-1} \mathbf{Q}^\top \mathbf{X}_\perp \mathbf{X}_\perp^\top \mathbf{S}_2^\top + (\mathbf{X} \mathbf{X}^\top)^{-1} \mathbf{Q}^\top \mathbf{X}_\perp \mathbf{X}_\perp^\top \mathbf{S}_3 \mathbf{X}_\perp \mathbf{X}_\perp^\top \mathbf{Q} (\mathbf{X} \mathbf{X}^\top)^{-1} \quad (65)$$

So the the excess risk is upperbounded by

$$tr(\mathbf{B}^\top \Sigma \mathbf{B}) + \sigma^2 tr(\mathbf{C}) = tr(\mathbf{X}_\perp^\top \mathbf{S}_3 \mathbf{X}_\perp) + tr(\mathbf{F}^\top \mathbf{X}_\perp^\top \mathbf{S}_3 \mathbf{X}_\perp \mathbf{F}) \quad (66)$$

$$+ \sigma^2 (\mathbf{S}_1 + \mathbf{S}_2 \mathbf{X}_\perp \mathbf{X}_\perp^\top \mathbf{Q} (\mathbf{X} \mathbf{X}^\top)^{-1} + (\mathbf{X} \mathbf{X}^\top)^{-1} \mathbf{Q}^\top \mathbf{X}_\perp \mathbf{X}_\perp^\top \mathbf{S}_2^\top + (\mathbf{X} \mathbf{X}^\top)^{-1} \mathbf{Q}^\top \mathbf{X}_\perp \mathbf{X}_\perp^\top \mathbf{S}_3 \mathbf{X}_\perp \mathbf{X}_\perp^\top \mathbf{Q} (\mathbf{X} \mathbf{X}^\top)^{-1}) \quad (67)$$

Now if we take derivative w.r.t. \mathbf{Q} and set it to zero we get

$$\mathbf{Q}^* = -(\mathbf{K}_\perp \mathbf{S}_3 \mathbf{K}_\perp)^{-1} (\sigma^2 \mathbf{K}_\perp \mathbf{S}_2^\top \mathbf{K}^{-1}) (\mathbf{K}^{-1} + \mathbf{K}^{-2})^{-1} \quad (68)$$

where $\mathbf{K} = \mathbf{X} \mathbf{X}^\top$ and $\mathbf{K}_\perp = \mathbf{X}_\perp \mathbf{X}_\perp^\top$. By setting $\mathbf{P}_1 = \mathbf{I}$, we have $\mathbf{P}_2 = \mathbf{Q}^{*\top}$. Therefore an optimal \mathbf{P} is

$$\mathbf{P} = \mathbf{X}^\top \mathbf{X} + \mathbf{X}^\top \mathbf{Q}^{*\top} \mathbf{X}_\perp + \mathbf{X}_\perp^\top \mathbf{Q}^* \mathbf{P} \quad (69)$$

Note that here we assume we have enough unlabeled data which among them we can pick $n - d$ of them which are orthogonal to the training data i.e.

$$\mathbf{X}^\top \mathbf{X}_\perp = \mathbf{X}_\perp^\top \mathbf{X} = 0 \quad (70)$$

■

A.12. Optimization based Preconditioner for Linear Classification

Lemma 14 For an interpolating linear classifier \mathbf{w} , it is possible to construct a family of quadratic norms $\|\cdot\|_{\mathbf{P}}$ such that the direction of the classifier is equivalent to the direction corresponding to the max \mathbf{P} -margin solution where

$$\mathbf{P} = \left[\|\mathbf{w}\|^2 I_d - \mathbf{w} \mathbf{w}^\top + \nu \nu^\top \right]^{-1} \quad \text{s.t.} \quad \langle \mathbf{w}, \nu \rangle = 1.$$

Here, $\nu = \mathbf{V}^\top \alpha$ where $\mathbf{V} \in \mathbb{R}^{|S| \times d}$ is the feature-matrix corresponding to the set S of support vectors for \mathbf{w} , α is a random vector satisfying the above constraints.

Proof Given \mathbf{w}_{opt} , the iterate obtained after running an arbitrary optimizer for T iterations, we want to find matrix \mathbf{M} such that

$$\frac{\mathbf{w}_{opt}}{\|\mathbf{w}_{opt}\|} = \arg \max_{\|\mathbf{z}\|_{\mathbf{M}} \leq 1} \min_i \langle \mathbf{z}, \mathbf{x}_i \rangle$$

We constrain \mathbf{M} such that $\|\mathbf{w}_{opt}\|_{\mathbf{M}} = 1$. If \mathbf{M} satisfies this constraint, then

$$\mathbf{w}_{opt} = \arg \max_{\|\mathbf{z}\|_{\mathbf{M}} \leq 1} \min_i \langle \mathbf{z}, \mathbf{x}_i \rangle$$

For the given \mathbf{w}_{opt} , there exist a unique set of support vectors S s.t for $\mathbf{s} \in S$, $\min_i \langle \mathbf{w}_{opt}, \mathbf{s} \rangle = \min_i \langle \mathbf{w}_{opt}, \mathbf{x}_i \rangle$. Simplifying, matrix \mathbf{M} should satisfy the following equality for $\mathbf{s} \in S$,

$$\mathbf{w}_{opt} = \arg \max_{\|\mathbf{z}\|_{\mathbf{M}} \leq 1} \langle \mathbf{z}, \mathbf{s} \rangle \implies \mathbf{w}_{opt} = \arg \max_{\|\mathbf{z}\|_{\mathbf{M}} \leq 1} \sum_{\mathbf{s} \in S} \langle \mathbf{z}, \mathbf{s} \rangle$$

The Lagrangian for the RHS can be written as:

$$L(\mathbf{z}, \lambda) = \sum_{\mathbf{s} \in S} \langle \mathbf{z}, \mathbf{s} \rangle + \lambda \left(\|\mathbf{z}\|_{\mathbf{M}}^2 - 1 \right)$$

$$\frac{\partial L(\mathbf{z}, \lambda)}{\partial \mathbf{z}} = 0 \implies \mathbf{M}\mathbf{w}_{opt} = -\frac{\sum_{\mathbf{s} \in S} \mathbf{s}}{\lambda}$$

Implying that vector $\mathbf{M}\mathbf{w}_{opt}$ lies in the span of the support vectors. Let $\mathbf{V} \in \mathbb{R}^{|S| \times d}$ be the matrix of support vectors. And let α be an $|S|$ -dimensional vector of coefficients.

$$\implies \mathbf{M}\mathbf{w}_{opt} = \mathbf{V}^T \alpha$$

We now use the norm constraint $\|\mathbf{w}_{opt}\|_{\mathbf{M}} = 1$ to constrain the α . Specifically, we want, $\|\mathbf{w}_{opt}\|_{\mathbf{M}}^2 = \mathbf{w}_{opt}^T \mathbf{V}^T \alpha = 1$. This implies that α should satisfy the following equality that ensures $\|\mathbf{w}_{opt}\|_{\mathbf{M}} = 1$,

$$\langle \mathbf{V}\mathbf{w}_{opt}, \alpha \rangle = 1$$

Since \mathbf{M} is the inverse of a preconditioner matrix, it needs to be positive definite. Using the same construction as in Lemma 4 with the additional constraint,

$$\mathbf{M} = \|\mathbf{w}_{opt}\|^2 \mathbf{I}_d - \mathbf{w}_{opt} \mathbf{w}_{opt}^T + \frac{(\mathbf{V}^T \alpha)(\mathbf{V}^T \alpha)^T}{\langle \mathbf{w}_{opt}, (\mathbf{V}^T \alpha) \rangle} \quad (71)$$

$$\text{where } \mathbf{V}^T \alpha > 0 \quad ; \quad \langle \mathbf{V}\mathbf{w}_{opt}, \alpha \rangle = 1. \quad (72)$$

Similar to Lemma 4, we can verify that \mathbf{M} is positive definite if $\mathbf{V}^T \alpha > 0$ and satisfies the equality $\mathbf{M}\mathbf{w}_{opt} = \mathbf{V}^T \alpha$. ■

A.13. Counter-examples for squared-hinge loss

Example 1 Consider two points $\mathbf{x}_1 = (-1, 0)$ and $\mathbf{x}_2 = (a, b = \sqrt{1-a^2})$ where $0 < a < 1$ with $y_1 = -1$ and $y_2 = 1$ as labels respectively. It can be shown that $\mathbf{w}^* = (1, \frac{b}{1+a})$ is the min-norm solution and we have $y_1 \langle \mathbf{w}^*, \mathbf{x}_1 \rangle = y_2 \langle \mathbf{w}^*, \mathbf{x}_2 \rangle = 1$. In the following we assume three different cases where gradient descent starting from \mathbf{w}_0 over squared hinge loss $\mathcal{L}(\mathbf{w}) = \frac{1}{4} \sum_{i=1}^2 (\max\{0, 1 - y_i \langle \mathbf{w}, \mathbf{x}_i \rangle\})^2$ with any fixed step size $\eta > 0$ won't converge to \mathbf{w}^* .

1. let $\mathbf{w}^0 = \mathbf{w}^* + (\alpha, \beta)$ where $\alpha, \beta \geq 0$ and $\alpha + \beta > 0$. Since $y_1 \langle \mathbf{w}_0, \mathbf{x}_1 \rangle \geq 1$ and $y_2 \langle \mathbf{w}_0, \mathbf{x}_2 \rangle \geq 1$, therefore the squared hinge loss is zero and so is its gradient. Therefore GD would not progress and \mathbf{w}_0 is an answer which is not min-norm solution.
2. Now assume \mathbf{w}_0 classifies \mathbf{x}_2 with margin bigger than one and \mathbf{x}_1 with margin smaller than one i.e. $y_2 \langle \mathbf{w}_0, \mathbf{x}_2 \rangle \geq 1$ and $y_1 \langle \mathbf{w}_0, \mathbf{x}_1 \rangle < 1$. In this situation the gradient of loss function at \mathbf{w}_0

$$\nabla \mathcal{L}(\mathbf{w}_0) = -0.5 y_1 \mathbf{x}_1 (1 - y_1 \mathbf{x}_1^T \mathbf{w}_0) = -0.5 \alpha_0 y_1 \mathbf{x}_1.$$

Note that $\nabla \mathcal{L}(\mathbf{w}_0) = (-0.5 \alpha_0, 0)$ and $\alpha_0 > 0$. Now if we run GD for one step we have

$$\mathbf{w}_1 = \mathbf{w}_0 + (0.5 \eta \alpha_0, 0).$$

Observe that $y_2 \langle \mathbf{w}_1, \mathbf{x}_2 \rangle = y_2 \langle \mathbf{w}_0, \mathbf{x}_2 \rangle + 0.5 \eta \alpha_0 a \geq 1$ since $0.5 \eta \alpha_0 a \geq 0$. Therefore the loss value at \mathbf{w}_2 for \mathbf{x}_2 is also zero. Therefore for all \mathbf{w}_t we have $y_2 \langle \mathbf{w}_t, \mathbf{x}_2 \rangle \geq 1$ and loss functions gradient at any \mathbf{w}_t just add some positive value to first component of \mathbf{w}_t . Therefore if we assume GD converges to $\mathbf{w}_\infty = (\mathbf{w}_\infty^1, \mathbf{w}_\infty^2)$, we have $\mathbf{w}_\infty^2 = \mathbf{w}_0^2 \neq \mathbf{w}^{*2}$.

3. Here we consider the reverse of the above scenario i.e. $y_2 \langle \mathbf{w}_0, \mathbf{x}_2 \rangle < 1$ and $y_1 \langle \mathbf{w}_0, \mathbf{x}_1 \rangle = \mathbf{w}_0^1 \geq 1$.

$$\nabla \mathcal{L}(\mathbf{w}_0) = -0.5 \mathbf{x}_2 (1 - y_2 \mathbf{x}_2^T \mathbf{w}_0) = -0.5 \alpha_0 \mathbf{x}_2$$

$$\mathbf{w}_1 = \mathbf{w}_0 + 0.5 \eta \alpha_0 \mathbf{x}_2.$$

We can check that $y_1 \langle \mathbf{w}_1, \mathbf{x}_1 \rangle = y_1 \langle \mathbf{w}_0, \mathbf{x}_1 \rangle + y_1 \langle \mathbf{x}_2, \mathbf{x}_1 \rangle \geq 1 + 0.5 \eta \alpha_0 a \geq 1$ since $a, \alpha_0 > 0$. Therefore similar to above scenario, for all t we have $y_1 \langle \mathbf{w}_t, \mathbf{x}_1 \rangle \geq 1$. We can observe that $\nabla \mathcal{L}(\mathbf{w}_t) = \gamma \mathbf{x}_2$ for some $\gamma > 0$. Therefore $\mathbf{w}_\infty = \mathbf{w}_0 + \beta \mathbf{x}_2 = (\mathbf{w}_0^1 + \beta a, \mathbf{w}_0^2 + \beta b)$ for some $\beta > 0$ that means $\mathbf{w}_\infty^1 = \mathbf{w}_0^1 + \beta a \neq 1 = \mathbf{w}^*$.

Example 2 Consider the same dataset as above, and consider GD initialized with $\mathbf{w}_0 = 0$ vector. The GD update at step t is

$$\mathbf{w}_{t+1} = \mathbf{w}_t - \eta \nabla \mathcal{L}(\mathbf{w}_t).$$

Based on case 2 and case 3 of Example 1, if for any t we have $y_1 \langle \mathbf{w}_t, \mathbf{x}_1 \rangle \geq 1$ or $y_2 \langle \mathbf{w}_t, \mathbf{x}_2 \rangle \geq 1$ then we know that GD won't converge to \mathbf{w}^* . Now assume that for all $t < \infty$ we have $y_1 \langle \mathbf{w}_t, \mathbf{x}_1 \rangle < 1$ or $y_2 \langle \mathbf{w}_t, \mathbf{x}_2 \rangle < 1$. To make notation simpler let assume $\mathbf{x}_1 = y_1 \mathbf{x}_1$ and $\mathbf{x}_2 = y_2 \mathbf{x}_2$. In this case the update rule for GD is

$$\begin{aligned} \mathbf{w}_{t+1} &= \mathbf{w}_t + \frac{\eta}{2}(\mathbf{x}_1 + \mathbf{x}_2) - \frac{\eta}{2}(\mathbf{x}_1 \mathbf{x}_1^\top + \mathbf{x}_2 \mathbf{x}_2^\top) \mathbf{w}_t = (\mathbf{I} - \frac{\eta}{2} \mathbf{X}^\top \mathbf{X}) \mathbf{w}_t + \frac{\eta}{2}(\mathbf{x}_1 + \mathbf{x}_2) \\ &= \mathbf{A}^t \mathbf{w}_0 + \frac{\eta}{2}(\mathbf{x}_1 + \mathbf{x}_2) \sum_{i=0}^t \mathbf{A}^i = \frac{\eta}{2}(\mathbf{x}_1 + \mathbf{x}_2) \sum_{i=0}^t \mathbf{A}^i. \end{aligned}$$

It can be seen that the eigenvalues of $\mathbf{X}^\top \mathbf{X}$ are $\lambda_1 = 1 + a$ and $\lambda_2 = 1 - a$. To get convergence at $t \rightarrow \infty$, we need $\sum_{i=0}^{\infty} \mathbf{A}^i$ to be Neumann series. To get that, we need to $\eta \leq \frac{\alpha}{\lambda_1}$ where $\alpha < 1$. The largest eigenvalue of \mathbf{A} is $1 - \alpha \frac{\lambda_2}{\lambda_1}$. Hence we have

$$\mathbf{w}_\infty = \frac{\eta}{2}(\mathbf{x}_1 + \mathbf{x}_2) \sum_{i=0}^{\infty} \mathbf{A}^i = \frac{\eta}{2}(\mathbf{x}_1 + \mathbf{x}_2) \alpha \frac{\lambda_1}{\lambda_2} = \frac{\alpha^2}{1-a}(\mathbf{x}_1 + \mathbf{x}_2) = \frac{\alpha^2}{2(1-a)}(1+a, b).$$

To get convergence to \mathbf{w}^* i.e. $\mathbf{w}_\infty = \mathbf{w}^*$, we need that $\frac{\alpha^2(1+a)}{2(1-a)} = 1$ and $\frac{\alpha^2 b}{2(1-a)} = \frac{b}{1+a}$. Therefore we have to set $\alpha = \sqrt{\frac{2(1-a)}{1+a}}$. However if we pick $a \leq \frac{1}{3}$ then $\alpha \geq 1$ which is invalid value for α .

Appendix B. Experiments for linear classification

This section presents further experimental results for under and over-parameterized linear classification problems. In Appendix B.1, we verify the construction of the equivalent preconditioner in Lemma 14. In Appendix B.2, we verify that when minimizing the squared-hinge loss for datasets with a known margin, projections onto the data-span and the ℓ_2 ball ensure convergence to the max-margin solution. Finally, in Appendix B.3 we provide experimental details for preconditioned gradient descent converging to the maximum relative margin solution in Figure 3, and in Appendix B.4, we present additional results for over-parameterized linear classification with an exponential-tailed loss.

B.1. Verification of Lemma 14

In this section we empirically validate the result of Lemma 14: every interpolating solution \mathbf{w}_{opt} with zero training error has a corresponding preconditioner \mathbf{P} for which PGD with \mathbf{P} converges to \mathbf{w}_{opt} . We run the Adam, Adagrad, and Coin Betting optimizers initialized at the origin, and then construct the corresponding preconditioners. We use the exponential loss on a synthetic dataset of 500 training points using 50 features. Figure 3 shows the optimization performance of both the original optimizers and their associated PGD methods (dashed lines). The PGD methods converge to models with full training accuracy and the same direction as the original optimizers (right-most plot). Notice that some of the original optimizers (and their preconditioned equivalent) converge to solutions which do not align with the max ℓ_2 margin solution.

B.2. ℓ_2 ball projection for the squared hinge loss

In Section 4, we explored the effect of projecting onto an appropriate space in order to improve the generalization performance of a linear classifier under exponential-tailed losses. In this section we explore an analogous technique for the case of the squared hinge loss: projecting the iterates of an optimizer onto the ℓ_2 ball of radius $1/\gamma$, where the margin γ is assumed to be known, implies convergence to the ℓ_2 max-margin solution.

We construct synthetic classification datasets in which the samples from each class follow Gaussian distribution with different parameters for each class. We ensure that the training set is linearly separable and compute the maximum ℓ_2 margin attainable on it in order to perform the projection. We use a training set of 500 points and explore under- and over-parameterized settings with 100 and 1000-dimensional features, respectively.




Figure 3: Experimental validation of Lemma 14. For each optimizer, the corresponding dashed line shows the convergence of PGD with a preconditioner constructed from the interpolating solution found by that optimizer as in Lemma 14. The three PGD methods converge in *direction* to the same solutions as the original optimizers, as predicted.

In Figure 4 we present several learning rate configurations for each optimizer, as well as the effect that the projection onto the ℓ_2 ball (and data span for the over-parameterized case) have in relation to the convergence to the ℓ_2 max-margin solution.

We observe that even though the projection might slow down the convergence in terms of the training loss for large step-sizes, it improves the speed for the accuracy on the training set. Moreover, for every optimizer and step-size configuration, the projection onto the ball improves the convergence towards the max-margin solution compared to the un-projected optimizer. In the right-most plots, all the lines corresponding to projected optimizers overlap at a norm of $1/\gamma$.

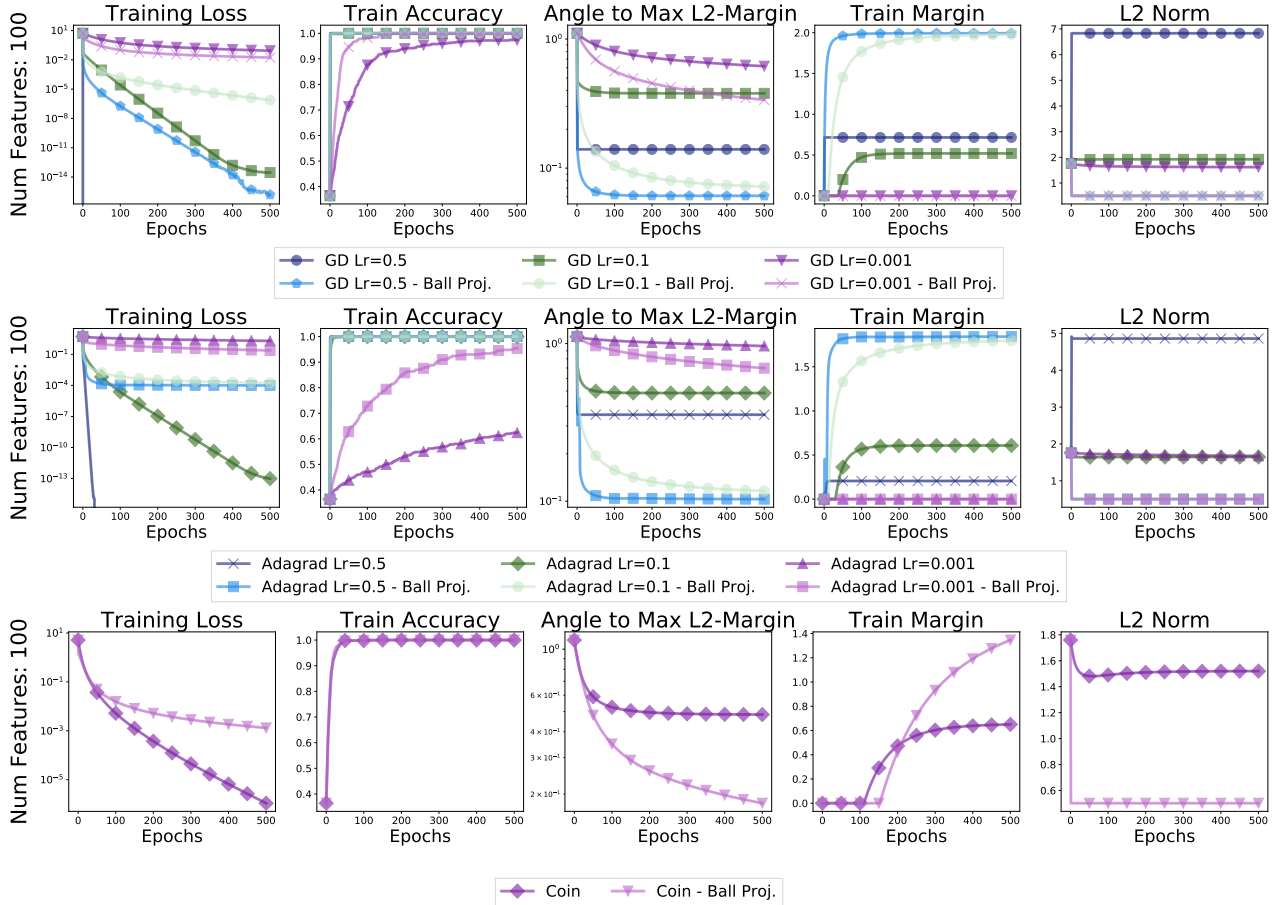


Figure 4: Projected and unprojected variants of several optimizers onto the ℓ_2 ball of (assumed known) radius $1/\gamma$.




Figure 5: Projected and unprojected variants of Adagrad optimizers onto the data-span and ℓ_2 ball of (assumed known) radius $1/\gamma$.

Finally, in the over-parameterized setting, in Figure 5, projection onto the span of the data and subsequent ball projection provides further improvements in terms of angle to the max-margin solution and the margin measured on the training points.

B.3. Maximum relative margin

We replicate the synthetic setting of [38] to study the convergence of preconditioned gradient descent to the maximum relative margin solution. We make minor modifications to the parameters of the originally proposed distribution to allow for separability under a homogeneous linear model. We consider a 2-dimensional classification dataset in which the distributions of each of the classes are Gaussian with matching covariance but different means.

$$\mathbf{x} | y = 1 \sim \mathcal{N}(\mu_+, \Sigma) \quad \mathbf{x} | y = -1 \sim \mathcal{N}(\mu_-, \Sigma) \quad \mu_+ = -\mu_- = -\frac{2}{5}[[19, 13]]^\top \quad \Sigma = \frac{1}{4} \begin{bmatrix} 17 & 16.9 \\ 16.9 & 17 \end{bmatrix}$$

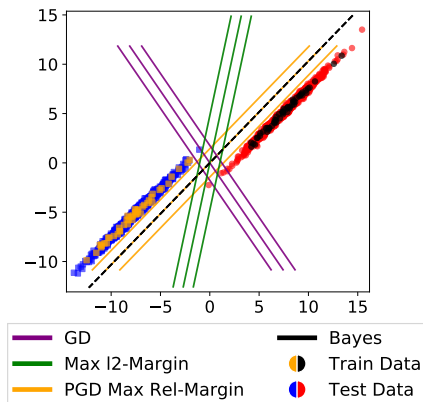


Figure 6: Comparing solutions with maximum ℓ_2 and relative margin performance on a synthetic mixture of Gaussians dataset. Incorporating the covariance of the data in the form of a preconditioning of the optimization, maximizes the relative margin and results in a solution which is better aligned with the Bayes optimal classifier. Note how, even though the maximum margin classifier is reasonable for the training set, it is “agnostic” to directions in which the data tends to spread.

Figure 6 displays again the training and test datasets as well as the classifiers presented in the main paper. The training set contains 100 points (orange and black), while the test set is formed of 1900 instances (blue and red). Training statistics for preconditioned gradient descent under the exponential loss are presented in Figure 7. The solution found by gradient descent stagnates after a few iterations due to the nature of the loss. However, the misalignment of the gradient descent solution causes it to mis-classify one of the test points. This, coupled with the norm of the iterates of gradient descent, induces a large test loss. Note that the mis-classified test point which represents an outlier for the GD and max ℓ_2 -margin solutions appears precisely in the directions “corrected” by the preconditioning via the covariance matrix.




Figure 7: Performance of gradient descent with preconditioner Σ^{-1} on a synthetic experiment with Gaussian class-conditionals.

B.4. Over-parameterized linear classification

The features and labels for this experiment are generated as the synthetic regression dataset presented in Appendix C.2.1 using 1-dimensional targets. The targets are then binarized depending on their sign. We explore different levels of over-parametrization by sampling 300 training points and 600 test points with 500 and 1000 random Gaussian features.

Figure 8 presents the impact of projection onto the data span, as well as switching to gradient descent with a step-size $\eta = 10$ (found to perform well via a grid-search). In the legend, the key-word “Always Project” indicates that the weights are projected onto the data span at every iteration, while “Project at Switch” indicates that only one projection onto the data span is performed and it takes place at the moment at which we switch from Adagrad to gradient descent. We experimented switching at different points in training (50%, 75% and 90%) and obtained qualitatively similar performance as those results presented here. The results are aggregated over 10 data samples and all instances are initialized at the origin.

As expected, we see that the projection results in a smaller ℓ_2 norm and angle with respect to the ℓ_2 maximum margin solution. More importantly, these results demonstrate that the generalization performance can be impacted by the “choice” of subspace in which the iterates of the optimizer lie.

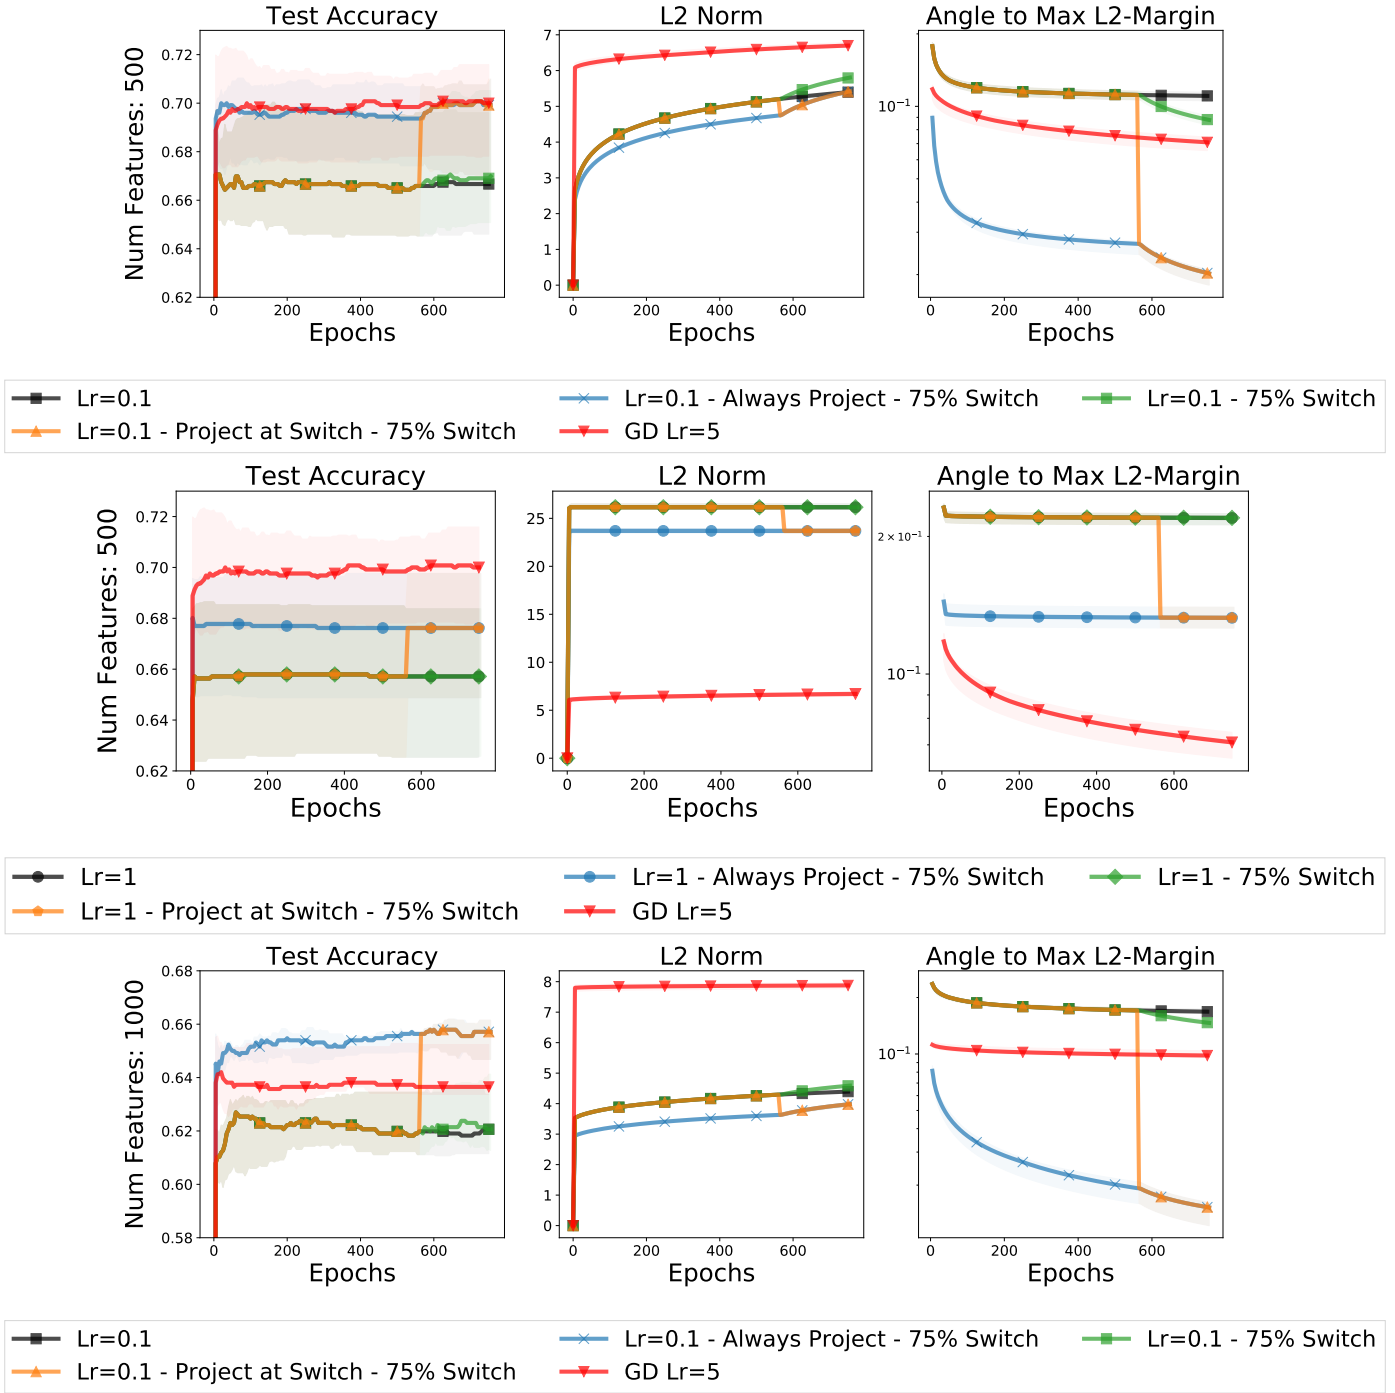


Figure 8: Performance of GD and Adagrad on a synthetic overparameterized classification problem with random Gaussian features. Projecting onto the data-span improves the test accuracy, while decreasing the solution’s norm and angle to the max-margin solution.

Appendix C. Experiments for over-parameterized linear regression

C.1. Additional Results

This section presents further experimental results for over-parameterized linear regression problems. Appendix C.1.1 extends our investigation of regression with neural tangent kernels to several real-world datasets from the UCI repository [13]. For completeness, we also repeat the experiment from Figure 1 with batch (deterministic) gradients. Then, in Appendix C.1.2, we experimentally verify the theoretical results discussed in Section 2, including Lemmas 6, 9, 8, and 4.

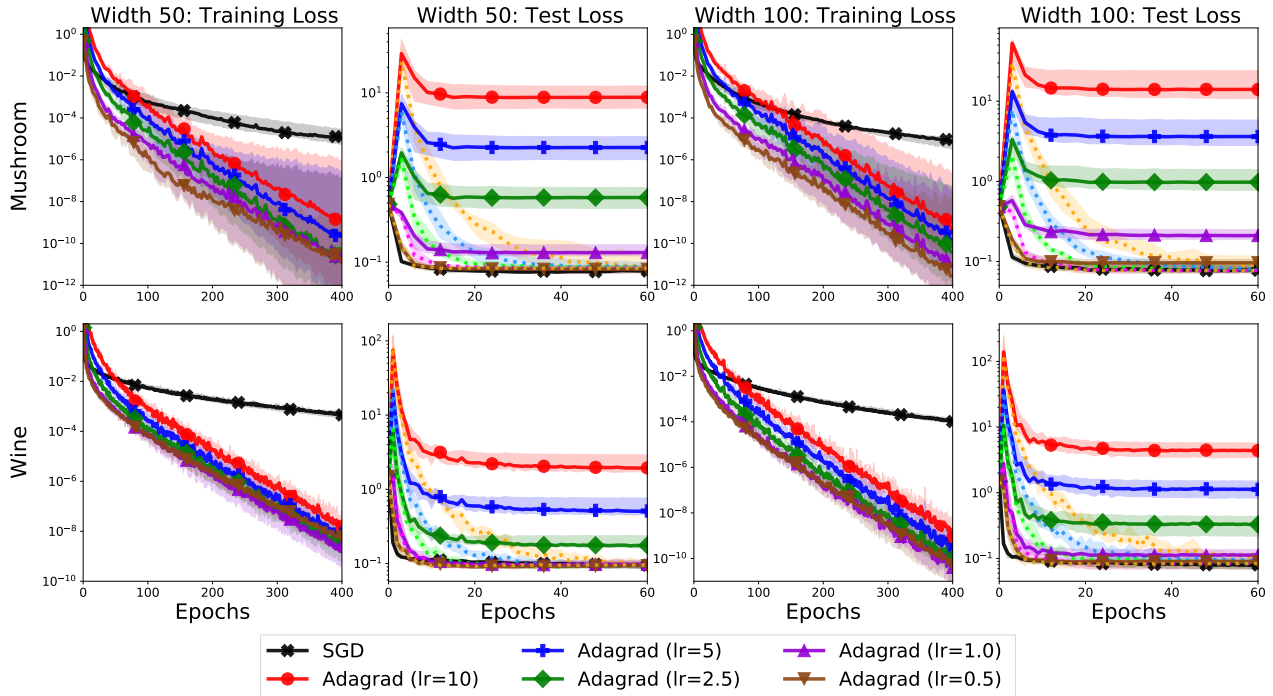


Figure 9: Performance of SGD and Adagrad for regression on the `mushroom` and `wine` datasets using squared-loss and the NTK of one-layer networks with 50 and 100 hidden units. We use the largest step-size for which SGD converges and consider a range of step-sizes for Adagrad. For each step-size, we plot as a dotted line the test loss (training loss is unchanged) for a variant of Adagrad where the model parameters are projected onto the span of the training data after every iteration. Tuned SGD stalls on the training loss but generalizes well, while Adagrad’s generalization depends strongly on step-size. Projecting onto the data span corrects this behavior.

Finally, Appendix C.1.3 investigates improving the generalization of preconditioned gradient-descent by optimizing over the space of preconditioners. In particular, we consider optimizing the excess-risk bound in Lemma 11 as well as several simplified upper-bounds based on this quantity.

C.1.1. REGRESSION WITH NEURAL TANGENT KERNELS

We investigate the implicit regularization of Adagrad and SGD for kernel regression on the `mushroom` and `wine` datasets. As in our synthetic experiments, we fit the model via the squared loss and use features from the neural tangent kernel of single-layer, feed-forward networks with 50 and 100 hidden units, respectively. Unlike Figure 1, we also consider variants of Adagrad where the model is projected on the data span after each iteration. We show only test loss for these projected optimizers, since training performance is unaffected by projecting onto the data span³.

Results are shown in Figure 9. The generalization performance for Adagrad shows a striking dependence on step-size — large step-sizes obtain test loss approximately two orders of magnitude larger than the smallest considered — while training loss is largely unaffected. Similarly to our synthetic experiments, SGD stalls on the training loss, but still generalizes well. Of particular interest are the projected versions of Adagrad, which completely correct for the poor generalization performance of the “vanilla” algorithm and obtain a test loss comparable to SGD.

As an ablation, we repeat the synthetic regression NTK experiment from the main paper (Figure 1) with batch gradients. We also include projected variants of Adagrad following the protocol above. Figure 10 shows that the trends from the stochastic setting also hold for deterministic optimization; Adagrad converges quickly in comparison to tuned gradient descent, which stalls on the ill-conditioned problem. Yet, Adagrad’s test performance depends strongly on the step-size chosen and never out-performs the min-norm solution. The projected variants of Adagrad correct for this poor generalization and also converge faster than gradient descent.

3. Projection operators with a large condition number do introduce a precision floor on the training loss.




Figure 10: Ablation for synthetic regression problem using NTK features. Unlike in Figure 1, we consider batch optimization. Moreover, for each step-size, we plot as a dotted line the test loss (training loss is unchanged) for a variant of Adagrad where the model parameters are projected onto the span of the training data after every iteration. We see that deterministic optimization does not change Adagrad’s varied generalization performance. Projecting onto the data span ensures comparable generalization to GD.

C.1.2. VERIFICATION OF THEORETICAL RESULTS

Now we verify our theoretical results for over-parameterized regression with several experiments on synthetic problems.

Lemma 6: Figure 11 examines the convergence of PGD with randomly generated preconditioners in the mini-batch and batch settings. For each generated preconditioner \mathbf{P} , we compute the solution to the normal equations

$$w_p^* = \mathbf{P}\mathbf{X}^\top(\mathbf{X}\mathbf{P}\mathbf{X}^\top)^{-1}\mathbf{y},$$

and plot convergence of the iterates generated by PGD (with \mathbf{P}) to the solution w_p^* . The convergence of gradient descent to the min-norm solution is shown as a baseline. We clearly see that PGD converges to the \mathbf{P}^{-1} -norm least squares solution as established by Lemma 6.

Lemmas 8 and 9: Figure 12 shows convergence of Newton’s method and full-matrix Adagrad for a synthetic regression problem. Unlike all other experiments with synthetic regression data, we generate a well-conditioned dataset to avoid complications with evaluating the Hessian. We see that Newton’s method and full-matrix Adagrad remain in the span of the data and converge to the min-norm solution as predicted.

Lemma 4: Now we experimentally confirm that every interpolating solution w_{opt} has a corresponding preconditioner \mathbf{P} for which PGD with \mathbf{P} converges to w_{opt} . We run the Adam [25], Adagrad [14], and Coin Betting [32] optimizers until convergence and then construct the corresponding preconditioners. Figure 13 shows the optimization and generalization performance of both the original optimizers and their associated PGD methods (dashed lines). The PGD methods converge to interpolating models with the same ℓ_2 norm and true risk as the original optimizers and show similar final training loss.

Proposition 1: Figure 14 explores the effects of projecting onto the span of the training data during optimization. We use the Adagrad, Adam, and Coin Betting optimizers and compare projecting the model parameters onto $\text{span}(\mathbf{X}^\top)$ at every iteration with the “default” algorithms. We make several observations: (i) the default algorithms show varied generalization performance and are all out-performed by GD, (ii) the projected variants converge to the min-norm solution and obtain test loss comparable to GD, and (iii) the improved generalization of the projected methods is consistent across the stochastic and deterministic cases.

C.1.3. IMPROVED GENERALIZATION VIA BETTER PRECONDITIONERS

In this section, we investigate choosing a preconditioner \mathbf{P} to minimize the bound on the excess risk given in Lemma 11. We consider minimizing the exact bound as well as two upper-bounds on the excess-risk bound that do not require full

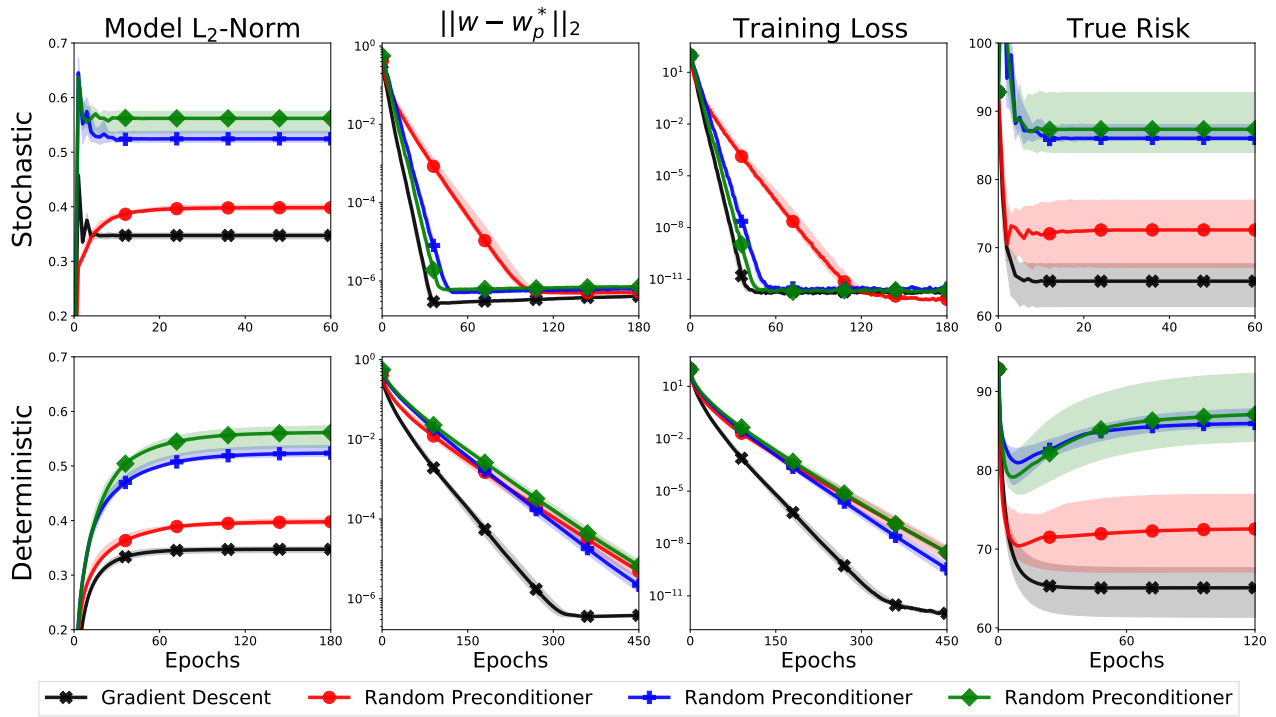


Figure 11: Experimental validation of Lemma 6. Each PGD method uses a randomly generated diagonal preconditioner \mathbf{P} . The metric $\|w - w_P\|_2$ is the ℓ_2 distance of the current weight vector w to w_P – the interpolating solution with minimum $\|\cdot\|_{\mathbf{P}^{-1}}$ norm. Each PGD method converges to the min-norm solution in its preconditioner \mathbf{P}^{-1} , as predicted.

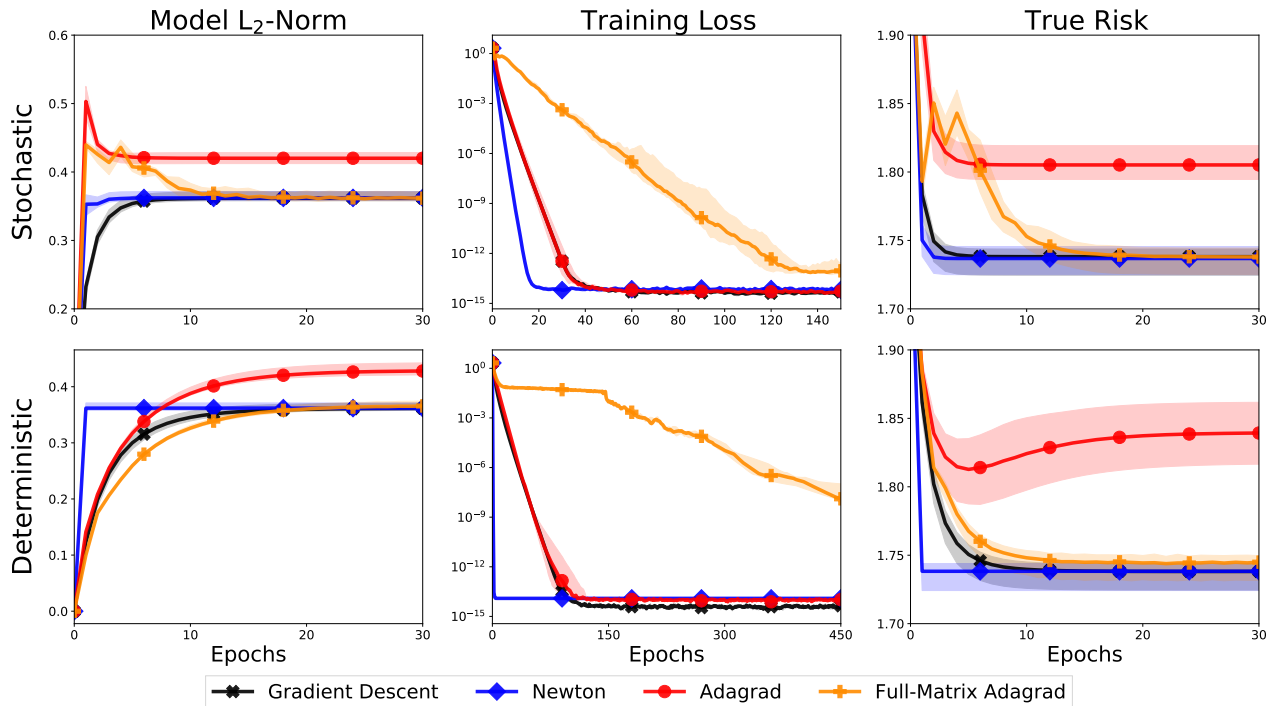


Figure 12: Experimental validation of Lemmas 8 and 9 with batch and mini-batch gradients. Newton’s method and full-matrix Adagrad converge to the min-norm solution, while diagonal Adagrad does not.

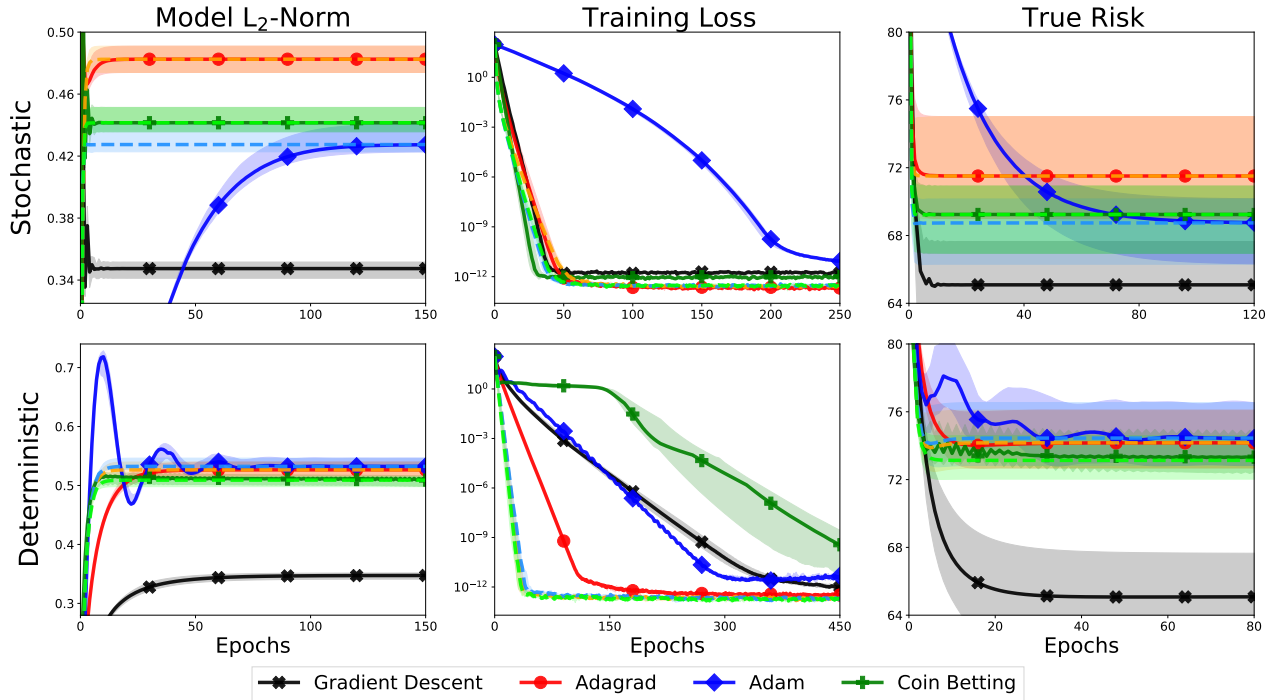


Figure 13: Experimental validation of Lemma 4. For each optimizer, the corresponding dashed line shows the convergence of PGD with a preconditioner constructed from the interpolating solution found by that optimizer as in Lemma 4. The three PGD methods converge to the same solutions as the original optimizers as predicted.

knowledge of the true model w^* . Starting from Lemma 11, we have

$$\begin{aligned}
 R(\mathbf{w}_{PGD}) &:= \mathbb{E}_{\mathbf{x}, \epsilon} [\mathbf{x}^\top (\mathbf{w}_{PGD} - \mathbf{w}^*)]^2 \leq (\mathbf{w}_0 - \mathbf{w}^*)^\top \mathbf{B}_P^\top \Sigma \mathbf{B}_P (\mathbf{w}_0 - \mathbf{w}^*) + \sigma^2 \text{tr}(\mathbf{C}_P) && \text{(Exact)} \\
 &\leq \|\mathbf{w}_0 - \mathbf{w}^*\| \|\mathbf{B}_P^\top \Sigma \mathbf{B}_P\|_2 + \sigma^2 \text{tr}(\mathbf{C}_P) && \text{(Operator)} \\
 &\leq \|\mathbf{w}_0 - \mathbf{w}^*\| \|\mathbf{B}_P^\top \Sigma \mathbf{B}_P\|_F + \sigma^2 \text{tr}(\mathbf{C}_P). && \text{(Frobenius)}
 \end{aligned}$$

When $\mathbf{w}_0 = 0$, these upper-bounds only require knowledge of the norm of the true model, rather than the true model itself. They are particularly attractive in the noiseless case ($\sigma^2 = 0$), where no knowledge of w^* is required and we need only optimize the two-operator norm or Frobenius norm of the matrix $\mathbf{B}_P^\top \Sigma \mathbf{B}_P$.

We optimize the exact excess risk bound (“Exact”) and upper-bound using the Frobenius norm (“Frobenius”) with respect to a diagonal preconditioner $\text{diag}(\mathbf{p})$ as well as a full-matrix preconditioner \mathbf{P} . The upper-bound using the 2-operator norm (“Operator”) is used to learn a diagonal preconditioner only, since evaluating the maximum singular value is computationally expensive. We also consider using PGD with Σ^{-1} as a preconditioner (“Sigma”), which corresponds to natural gradient descent [2]. Lemma 2 shows that the empirical precision $\hat{\Sigma}^{-1}$ minimizes an upper bound on the Rademacher complexity for the family of linear classifiers with bounded \mathbf{P}^{-1} norm, which suggests that PGD with Σ^{-1} may generalize well. For completeness, we also compare with tuned stochastic gradient descent and Adagrad.

Figure 15 shows the results of optimizing over P in both the noiseless setting ($\sigma^2 = 0$) and the case where $\sigma^2 = 1$. We observe that minimizing the exact bound on the excess risk is highly effective — especially when using a full-matrix preconditioner. In this latter setting, PGD finds an interpolating solution with near optimal risk. In contrast, the preconditioners obtained by minimizing the upper-bounds given by the Frobenius and 2-operator norms yield interpolating solutions that generalize only as well as the min-norm solution. This suggests that these bounds are too loose to be useful for learning better preconditioners. Such a conclusion is corroborated by Table 1, which shows that the Operator and Frobenius preconditioners do not improve the excess-risk bound over the identity matrix. Lastly, it is highly interesting to note that while natural gradient descent converges very quickly, it obtains the worst generalization performance out of all methods considered. We speculate that this is because of the discrepancy in rank between the empirical and true covariance matrices: $\text{rank}(\hat{\Sigma}) \leq n \ll \text{rank}(\Sigma) = d$.

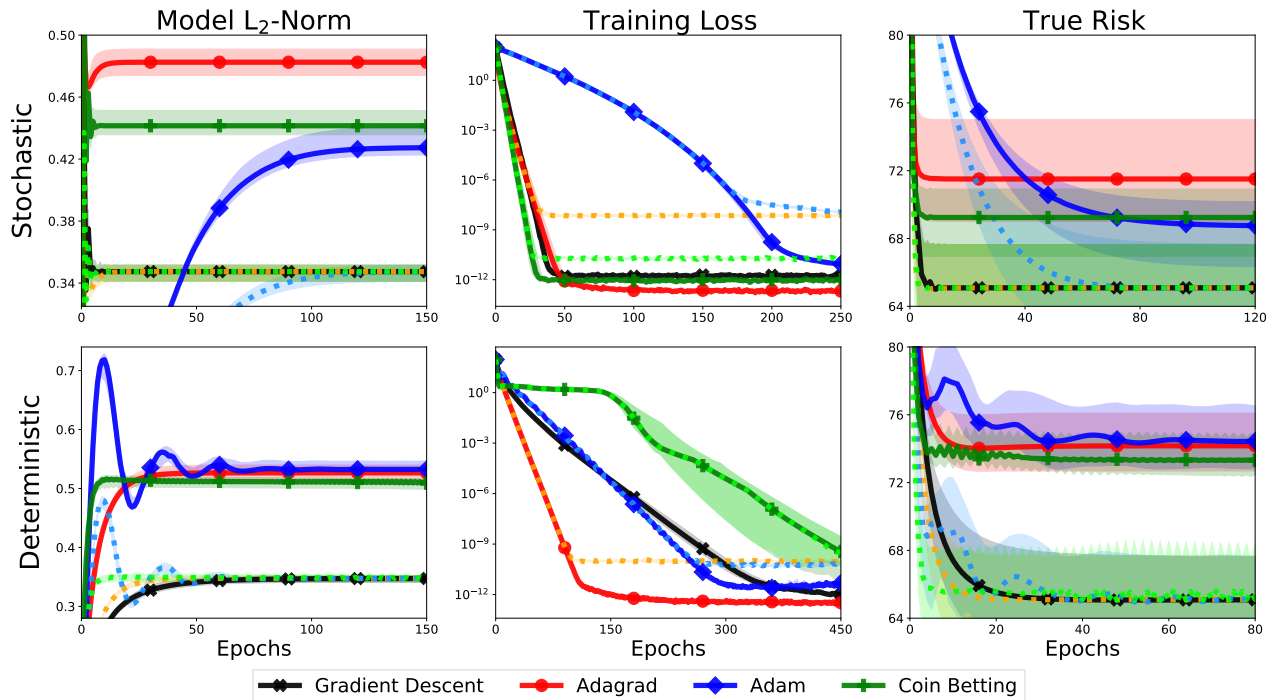


Figure 14: Effects of projecting onto $\text{span}(\mathbf{X}^T)$ after every iteration when using the Adagrad, Adam, and Coin Betting optimizers. Solid lines with markers denote standard optimizers, while dotted lines shows the same algorithms with projections onto the span of the data at each iteration. We note that projecting onto the data span introduces a precision floor for the training loss, but otherwise does not affect optimization dynamics of the original optimizers. Projected methods converge to the min-norm solution and display similar generalization to gradient descent, as predicted by Proposition 1.

Noise Level	Identity	Exact	Exact (FM)	Frobenius	Frobenius (FM)	Operator	Sigma
0	64.45 ± 3.39	14.13 ± 1.07	0.02 ± 0.02	64.47 ± 3.37	64.7 ± 3.49	71.34 ± 2.64	81.31 ± 0.85
1	64.82 ± 3.39	15.05 ± 1.66	0.57 ± 0.02	64.85 ± 3.39	65.28 ± 3.4	71.66 ± 4.77	81.45 ± 0.85

Table 1: Mean and standard deviations for evaluations of the excess risk bound in Lemma 11 at preconditioners obtained by minimizing the excess risk bound (“Exact”), upper-bounds on this bound using the 2-operator and Frobenius norms (“Operator” and “Frobenius”) and the inverse covariance of the data-generating distribution (“Sigma”). FM denotes that we optimize over a full-matrix preconditioner (default is diagonal matrix). Directly minimizing the bound from Lemma 11 is highly effective and leads to near-optimal risk for the solution found by PGD, while the operator and Frobenius learning rules do not improve upon the identity matrix.

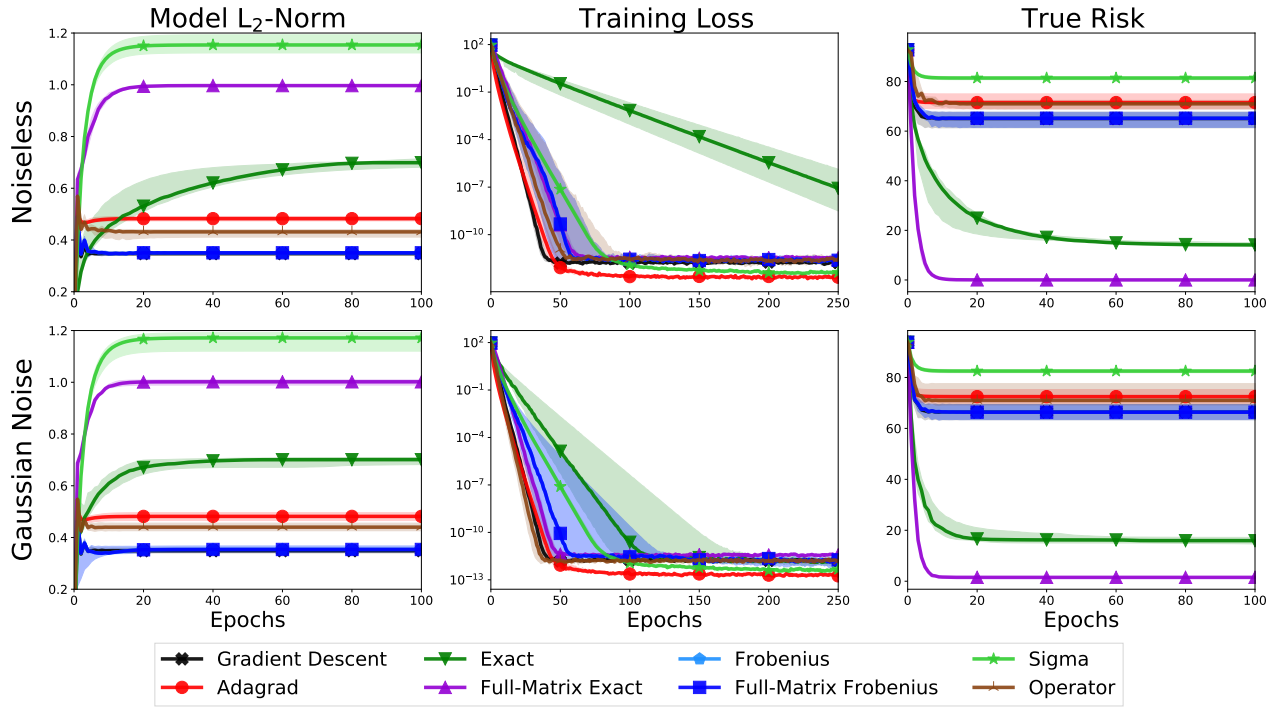


Figure 15: Generalization performance of stochastic PGD where the preconditioner is selected by optimizing bounds on the excess risk. The preconditioner for "Exact" minimizes the excess risk bound given in Lemma 11, while "Operator" and "Frobenius" minimize upper-bounds Operator and Frobenius , respectively. "Sigma" uses the true precision matrix of the data, Σ^{-1} , as the preconditioner. The full-matrix and diagonal preconditioners learned by minimizing the exact excess-risk bound greatly improve generalization performance over the min-norm solution, while "Frobenius" and "Operator" fail to outperform SGD and only minorly improve upon Adagrad.

C.2. Experimental Details

In this section we give additional details for the regression experiments presented in the main paper and the additional results shown in Appendix C.1.

C.2.1. DATASETS

Synthetic Regression Datasets: We generate synthetic regression problems by first sampling a normalized ground-truth weight vector $\mathbf{w}^* = \mathbf{v} / \|\mathbf{v}\|$, where $\mathbf{v} \sim \mathcal{N}(0, 1)$. We then sample features from a diagonal, mean-zero Gaussian distribution and computing targets as the inner product with the ground-truth vector:

$$\mathbf{x}_i \sim \mathcal{N}(0, \Sigma), \quad y_i = \mathbf{x}_i^\top \mathbf{w}^* + \epsilon,$$

where $\epsilon \sim \mathcal{N}(0, \sigma^2)$ is the target noise. To control the hardness of the optimization problem, we generate ill-conditioned, positive-definite covariances by perturbing the identity matrix with squared Gaussian noise. In particular, we compute

$$\Sigma = \mathbf{I} + \text{Diag}(\delta^2), \quad \delta \sim \mathcal{N}(0, \zeta^2 \mathbf{I}).$$

The setting $\zeta^2 = 10$ is used in all experiments but Figure 12, where $\zeta^2 = 1$ is chosen. We use a training set of 100 observations in all synthetic regression experiments. For a model \mathbf{w} , we report the true risk

$$\mathbb{E}_{\mathbf{x}, y} [\mathbf{x}^\top \mathbf{w} - y]^2 = (\mathbf{w} - \mathbf{w}^*)^\top \Sigma (\mathbf{w} - \mathbf{w}^*) + \sigma^2,$$

instead of using a test set when using the original data features. When using features from a neural tangent kernel, we instead sample a test of 400 examples and use this to evaluate the model performance. We repeat all experiments ten times with the *same* ground-truth weights w^* and data covariance Σ to control for randomness in the generation of the training and test sets. Figures show the median and inter-quartile range. All stochastic experiments on synthetic regression use mini-batches of 5 examples.

UCI Datasets: We use the `wine` and `mushroom` datasets from the UCI dataset repository [13]. We use the training and validation splits⁴ created by Fernández-Delgado et al. [15] and used by Arora et al. [4]. In all regression experiments, we randomly subset fifty examples from the training set to fit our models and evaluate on the full validation split. All experiments are repeated ten times to control for the effects of sub-setting the training set. Figures show the median and inter-quartile range. Stochastic experiments use mini-batches of two examples.

Computing Neural Tangent Kernels: We standardize the data before computing the neural tangent kernel. Neural networks are initialized with standard normal weights and use the so-called NTK parameterization [21, 33] as well as sigmoid activations. We use the BackPACK library [12] to compute the Jacobian of the network output.

C.2.2. REGRESSION WITH NEURAL TANGENT KERNELS

Here we provide specific details for the experiments shown in Figures 1, 9, and 10. For Figures 1 and 10, we generate a training set of 100 examples with dimension $d = 20$ as described above. We use minibatches of 5 examples in the stochastic case. Figure 9 considers the `mushroom` and `wine` datasets, where we use training sets of size 50 and mini-batches of 2 examples. We experiment with neural tangent kernels for single-layer feed-forward neural networks of width 50 and 100. For all three datasets, we plot the step-size for SGD which maximized the convergence rate on the training loss while still converging on all ten repeats. The grid-search uses the following grid:

$$\eta \in \{20, 10, 5, 2.5, 0.1, 0.5, 0.25, 0.1\}.$$

The final step-size chosen for all three datasets was $\eta = 2.5$. We use the same step-size grid for Adagrad and plot all step-sizes which do not diverge or under-fit.

4. Publicly available at <http://persoal.citius.usc.es/manuel.fernandez.delgado/papers/jmlr/>.

C.2.3. VERIFICATION OF LEMMA 6

Here we provide specific details for the experiment shown in Figure 11. We generate a synthetic regression dataset as described above and sample three random, diagonal positive-definite preconditioners as $\mathbf{P} = \text{diag}(\mathbf{v}^2)$, where $\mathbf{v} \sim \mathcal{N}(0, \mathbf{I})$. We select step-sizes individually for each restart by grid-search over the set

$$\eta \in \{0.1, 0.01, 0.005, 0.001, 5 \times 10^{-4}, 1 \times 10^{-4}, 5 \times 10^{-5}, 1 \times 10^{-5}, 5 \times 10^{-6}, 1 \times 10^{-6}\}. \quad (73)$$

Each step-size is picked to minimize the average of the training loss halfway through and at the end of training, subject to the constraint that the optimizer does not diverge initially (and then recovers) or at the end of optimization.

C.2.4. VERIFICATION OF LEMMAS 8 AND 9

Here we provide specific details for Figure 12, which verifies that Newton’s method and full-matrix Adagrad converge to the min-norm solution for over-parameterized linear regression problems. Unlike the other synthetic regression experiments, here we generate a well-conditioned problem ($\zeta^2 = 1$; see above) to avoid complications when computing the Newton step. As in other experiments, we choose step-sizes independently for each repeat (i.e. randomly sampled training set) by grid-search of the following set

$$\eta \in \{1, 0.1, 0.01, 0.005, 0.001, 5 \times 10^{-4}, 1 \times 10^{-4}, 5 \times 10^{-5}, 1 \times 10^{-5}, 5 \times 10^{-6}, 1 \times 10^{-6}\}.$$

We use the same rule to select step-sizes as described in Appendix C.2.3.

C.2.5. VERIFICATION OF LEMMA 4

Here we provide specific details for Figure 13, which verifies that each interpolating solution has a corresponding preconditioner P such that PGD with this preconditioner converges to the same solution. We generate a synthetic regression dataset as described above and consider the interpolating solutions found by three optimizers: Adagrad, Adam, and Coin Betting. We repeat the step-size selection procedure given in Appendix C.2.3 to select step-sizes for these optimizers and then run them until they have converged to interpolating solutions \mathbf{w}_{opt} . We then compute the corresponding preconditioners as described in Lemma 4. However, rather than sampling a random vector in the data span, we use the setting

$$\nu = \frac{\|\mathbf{w}_{\text{opt}}\|}{\|\mathbf{X}^\top \mathbf{y}\|} \mathbf{X}^\top \mathbf{y},$$

which can be interpreted as a normalized, one-step approximation to the min-norm solution. In practice, this leads to well-conditioned preconditioners, unlike naive random sampling. We repeat the step-size grid-search to select step-sizes for the PGD optimizers using these preconditioners.

C.2.6. VERIFICATION OF PROPOSITION 1

Figure 14 considers improving the generalization of popular optimizers like Adam, Adagrad, and Coin Betting, which do not converge to the min-norm solution, by projecting their iterations onto $\text{span}(\mathbf{X}^\top)$ after each iteration. Once again, we generate a synthetic regression dataset and use a grid-search to select step-sizes for each algorithm independently on each sampled training set. We use the step-size grid in Equation 73. The same step-size is used for the projected and unprojected variants of all optimizers. Projections onto the data span do not change the optimization dynamics for linear models, since any model components lost in the projection are orthogonal to the gradient.

C.2.7. IMPROVED GENERALIZATION VIA BETTER PRECONDITIONERS

We optimize the excess-risk bound and both upper-bounds (Operator and Frobenius) on the excess-risk using gradient descent with a fixed step-size. We select the step-size independently for each randomly generated train/test split by grid search. The grid considered is

$$\eta \in \{5 \times 10^{-1}, 10^{-1}, 10^{-2}, 5 \times 10^{-3}, 10^{-3}, 10^{-4}\},$$

and final step-sizes are chosen to minimize the value of the optimized risk bound. We run the optimization procedure for 7500 iterations or until the gradient norm is smaller than 1×10^{-7} . Note that we optimize the risk bound using the

combined train/test data, as this does not use knowledge of the test labels. This can be viewed as a form of unsupervised learning, where information from the unlabeled test examples, data covariance Σ , and target variance σ^2 are leveraged to obtain a better preconditioner. We select step-sizes for PGD with the optimized preconditioners using a search over the grid

$$\eta \in \{10^{-1}, 10^{-2}, 5 \times 10^{-3}, 10^{-3}, 5 \times 10^{-4}, 10^{-4}, 5 \times 10^{-5}, 10^{-5}, 5 \times 10^{-6}, 10^{-6}, 5 \times 10^{-7}, 10^{-7}\}.$$

We choose final step-sizes as described above in Appendix C.2.3. In practice, we found that the preconditioners obtained by optimizing the exact bound on the excess risk were highly ill-conditioned or indefinite and so required very small step-sizes. Results are provided only for the stochastic setting, as we found the deterministic setting to be virtually identical. We use mini-batches of five examples for all optimizers.