

Stankewitz, Bernhard; Mücke, Nicole; Rosasco, Lorenzo

Article — Published Version

From inexact optimization to learning via gradient concentration

Computational Optimization and Applications

Provided in Cooperation with:

Springer Nature

Suggested Citation: Stankewitz, Bernhard; Mücke, Nicole; Rosasco, Lorenzo (2022) : From inexact optimization to learning via gradient concentration, Computational Optimization and Applications, ISSN 1573-2894, Springer US, New York, NY, Vol. 84, Iss. 1, pp. 265-294, <https://doi.org/10.1007/s10589-022-00408-5>

This Version is available at:

<https://hdl.handle.net/10419/313173>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



<https://creativecommons.org/licenses/by/4.0/>



From inexact optimization to learning via gradient concentration

Bernhard Stankewitz¹ · Nicole Mücke² · Lorenzo Rosasco^{3,4}

Received: 20 October 2021 / Accepted: 2 August 2022 / Published online: 25 August 2022
© The Author(s) 2022

Abstract

Optimization in machine learning typically deals with the minimization of empirical objectives defined by training data. The ultimate goal of learning, however, is to minimize the error on future data (test error), for which the training data provides only partial information. In this view, the optimization problems that are practically feasible are based on inexact quantities that are stochastic in nature. In this paper, we show how probabilistic results, specifically gradient concentration, can be combined with results from inexact optimization to derive sharp test error guarantees. By considering unconstrained objectives, we highlight the implicit regularization properties of optimization for learning.

Keywords Implicit regularization · Kernel methods · Statistical learning

1 Introduction

Optimization plays a key role in modern machine learning, and is typically used to define estimators by minimizing empirical objective functions [1]. These objectives are based on a data fit term, suitably penalized, or constrained, to induce an

✉ Bernhard Stankewitz
stankebe@math.hu-berlin.de

Nicole Mücke
nicole.muecke@tu-braunschweig.de

Lorenzo Rosasco
lorenzo.rosasco@unige.it

¹ Department of Mathematics, Humboldt University of Berlin, Unter den Linden 6, 10099 Berlin, Germany

² Institute for Mathematical Stochastics, Technical University Braunschweig, Universitätsplatz 2, 38106 Braunschweig, Lower Saxony, Germany

³ MaLGA, DIBRIS, Università degli Studi di Genova, Via Dodecaneso 35, Genova 16146, Liguria, Italy

⁴ CBMM, MIT & Istituto Italiano di Tecnologia, Genova, Italy

inductive bias in the learning process [2]. The idea is that the empirical objectives should provide an approximation to the error on future data (the test error) which is the quantity that one wishes to minimize in learning. The quality of such an approximation error is typically deferred to a statistical analysis. In this view, optimization and statistical aspects are tackled separately.

Recently, a new perspective has emerged in machine learning showing that optimization itself can in fact directly be used to search for a solution with small test error. Interestingly, no explicit penalties/constraints are needed, since a bias in the search for a solution is implicitly enforced during the optimization process. This phenomenon has been called implicit regularization and it has been shown to possibly play a role in explaining the learning curves observed in deep learning, see for instance [3, 4] and the references therein. Further, implicit regularization has been advocated as a way to improve the efficiency of learning methods by tackling statistical and optimization aspects at once [5–8]. As it turns out, implicit regularization is closely related to the notion of iterative regularization with a long history in inverse problems [9].

The basic example of implicit regularization is gradient descent for linear least squares, which is well known to converge to the minimum norm least squares solution [10, 11]. The learning properties of gradient descent for least squares are now quite well understood [11, 12], including the extensions to non-linear kernelized models [13, 14], stochastic gradients [15–17], accelerated methods [18, 19] and distributed approaches [20–22]. Much less is known when other norms or loss functions are considered. Implicit regularization biased to more general norms has been considered for example in [23, 24]. Implicit regularization for loss functions other than the square loss has been considered in a limited number of works. There is a vast literature on stochastic gradient techniques, see e.g. [17] and the references therein, but these analyses do not apply when (batch) training error gradients are used, which is the focus in this work. The logistic loss function for classification has recently been considered both for linear and non-linear models, see for example [25, 26]. Implicit regularization for general convex Lipschitz loss with linear and kernel models have been first considered in [27] for subgradient methods and in [28] for stochastic gradient methods but only with suboptimal rates. Improved rates have been provided in [6] for strongly convex losses and more recently in [29] with a general but complex analysis. A stability based approach, in the sense of [30], is studied in [31].

In this paper, we further push this line of work considering implicit regularization for linear models with convex, Lipschitz and smooth loss functions based on gradient descent. Indeed, for this setting we derive sharp rates considering both the last and the average iterate. Our approach highlights a proof technique which is less common in learning and is directly based on a combination of optimization and statistical results. The usual approach in learning theory is to derive optimization results for empirical objectives and then use statistical arguments to assess to which extent the empirical objectives approximate the test error that one ideally wished to minimize, see e.g. [2]. Instead, we view the empirical gradient iteration as the inexact version of the gradient iteration for the test error. This allows to apply results from inexact optimization, see e.g. [32, 33], and requires using statistical/

probabilistic arguments to assess the quality of the gradient approximations (rather than that of the objective functions). For this latter purpose, we utilize recent concentration of measure results for vector valued variables to establish gradient concentration [34]. While the idea of combining inexact optimization and concentration results has been considered before [35], here, we illustrate it in a prominent way to highlight its usefulness. Indeed, we show that this approach leads to sharp results for a specific but important setting and we provide some simple numerical results that illustrate and corroborate our findings. By highlighting the key ideas in the proof techniques, we hope to encourage further results combining statistics and optimization, for example considering other forms of gradient approximation or optimization other than the basic gradient descent.

The remainder of the paper is structured as follows: In Sect. 2, we collect some structural assumptions for our setting. In Sect. 3, we formulate the assumptions we put on the loss function and state and discuss the main results of the paper as well as the novel aspects of our approach. Section 4 presents the more technical aspects of the analysis. In particular, we explain in detail how results from inexact optimization and concentration of measure can be combined to come up with a new proof technique for learning rates. Finally, Sect. 5 illustrates the key features of our theoretical results with numerical experiments.

2 Learning with gradient methods and implicit regularization

Let $(\mathcal{H}, \|\cdot\|)$ be a real, separable Hilbert space and \mathcal{Y} a subset of \mathbb{R} . We consider random variables (X, Y) on a probability space $(\Omega, \mathcal{F}, \mathbb{P})$ with values in $\mathcal{H} \times \mathcal{Y}$ and unknown distribution $\mathbb{P}_{(X,Y)}$. The marginal distribution of X is denoted by \mathbb{P}_X . Additionally, we make the standard assumption that X is bounded.

(A1) **(Bound):** We assume $\|X\| \leq \kappa$ almost surely for some $\kappa \in [1, \infty)$.

Based on the observation of n i.i.d. copies $(X_1, Y_1), \dots, (X_n, Y_n)$ of (X, Y) , we want to learn a linear relationship between X and Y expressed as an element $w \in \mathcal{H}$.¹ For an individual observation (X, Y) and the choice $w \in \mathcal{H}$, we suffer the loss $\ell(Y, \langle X, w \rangle)$, where $\ell : \mathcal{Y} \times \mathbb{R} \rightarrow [0, \infty)$ is a product-measurable loss function. Our goal is to find $w \in \mathcal{H}$ such that the *population risk* $\mathcal{L} : \mathcal{H} \rightarrow [0, \infty)$ given by

$$\mathcal{L}(w) := \mathbb{E}_{(X,Y)}[\ell(Y, \langle X, w \rangle)] = \int \ell(y, \langle x, w \rangle) \mathbb{P}_{(X,Y)}(d(x, y)) \quad (1)$$

is small. The observed data represent the training set, while the population risk can be interpreted as an abstraction of the concept of the test error.

In the following, we assume that a minimizer of \mathcal{L} in \mathcal{H} exists.

¹ Note that this includes many settings as special instances. In particular, it includes the standard setting of kernel learning, see Appendix A in [5].

(A2) **(Min)**: We assume there exists some $w_* \in \mathcal{H}$ such that $w_* \in \operatorname{argmin}_{w \in \mathcal{H}} \mathcal{L}(w)$.

Note that the argmin is taken only over \mathcal{H} and not over all measurable functions. Under **(Min)**, minimizing the population risk is equivalent to minimizing the *excess risk* $\mathcal{L}(w) - \mathcal{L}(w_*) \geq 0$.

In this work, we are interested in bounding the excess risk, when our choice of w is based on applying *gradient descent* (GD) to the *empirical risk* $\hat{\mathcal{L}} : \mathcal{H} \rightarrow [0, \infty)$ with

$$\hat{\mathcal{L}}(w) := \frac{1}{n} \sum_{j=1}^n \ell(Y_j, \langle X_j, w \rangle). \tag{2}$$

computed from the training data. We consider a basic gradient iteration, which is well defined when the loss function is differentiable in the second argument with a product-measurable derivative $\ell' : \mathcal{Y} \times \mathbb{R} \rightarrow \mathbb{R}$.

Definition 1 (*Gradient descent algorithm*)

1. Choose $v_0 \in \mathcal{H}$ and a sequence of step sizes $(\gamma_t)_{t \geq 0}$.
2. For $t = 0, 1, 2, \dots$, define the GD-iteration

$$v_{t+1} = v_t - \gamma_t \nabla \hat{\mathcal{L}}(v_t) = v_t - \frac{\gamma_t}{n} \sum_{j=1}^n \ell'(Y_j, \langle X_j, v_t \rangle) X_j. \tag{3}$$

3. For some $T \geq 1$, we consider both the last iterate v_T and the averaged GD-iterate $\bar{v}_T := \frac{1}{T} \sum_{t=1}^T v_t$.

Here, we focus on batch gradient, so that all training points are used in each iteration. Unlike with stochastic gradient methods, the gradients at different iterations are not conditionally independent. Indeed, the analysis of batch gradient is quite different to that of stochastic gradient and could be a first step towards considering minibatching [17, 35, 36]. In our analysis, we always fix a constant step size $\gamma_t = \gamma > 0$ for all $t \geq 0$ and consider both the average and last iterate. Both choices are common in the optimization literature [1] and have also been studied in the context of learning with least squares [14, 16, 17], see also our extended discussion in Sect. 3.2. In the following, we characterize the learning properties of the gradient iteration in Definition 1 in terms of the corresponding excess risk. In particular, we derive learning bounds matching the best known bounds for estimators obtained minimizing the penalized empirical risk. Next, we show that in the considered setting, learning bounds can be derived by studying suitable bias and variance terms controlled by the iteration number and the step size.

3 Main results and discussion

Before stating and discussing our main results, we introduce and comment on the basic assumptions needed in our analysis. We make the following additional assumptions on the loss function.

(A3) **(Conv):** We assume $\ell : \mathcal{Y} \times \mathbb{R} \rightarrow [0, \infty)$ is convex in the second argument.

(A4) **(Lip):** We assume ℓ to be L -Lipschitz, i.e. for some $L > 0$,

$$|\ell(y, a) - \ell(y, b)| \leq L|a - b| \quad \text{for all } y \in \mathcal{Y}, a, b \in \mathbb{R}. \tag{4}$$

(A5) **(Smooth):** We assume ℓ to be M -smooth, i.e. ℓ is differentiable in the second argument with product-measurable derivative $\ell' : \mathcal{Y} \times \mathbb{R} \rightarrow \mathbb{R}$ and for some $M > 0$,

$$|\ell'(y, a) - \ell'(y, b)| \leq M|a - b| \quad \text{for all } y \in \mathcal{Y}, a, b \in \mathbb{R}. \tag{5}$$

Note that Eq. (5) immediately implies that

$$\ell(y, b) \leq \ell(y, a) + \ell'(y, a)(b - a) + \frac{M}{2}|b - a|^2 \quad \text{for all } y \in \mathcal{Y}, a, b \in \mathbb{R}, \tag{6}$$

see e.g. Lemma 3.4 in [37].

For notational convenience, we state the assumptions **(Lip)** and **(Smooth)** globally for all $a, b \in \mathbb{R}$. It should be noted, however, that this is not necessary.

Remark 1 (Local formulation of assumptions) In our analysis, we only apply **(Lip)** and **(Smooth)** for arguments of the form $a = \langle v, x \rangle$, where $\|v\| \leq R$ for $R = \max\{1, 3\|w_*\|\}$ and $\|x\| \leq \kappa$ with κ from **(Bound)**. Therefore, all of our results also apply to loss functions which satisfy the above assumptions for all $a, b \in [-\kappa R, \kappa R]$ for constants L and M potentially depending on κ and R .

In light of Remark 1, our analysis is applicable to many widely used loss functions, see e.g. Chapter 2 in [38].

Example 1 (Loss functions satisfying the assumptions)

- (a) **(Squared loss):** If $\mathcal{Y} = [-b, b]$ for some $b > 0$, then checking first and second derivatives yields that the loss $\mathcal{Y} \times [-\kappa R, \kappa R] \ni (y, a) \mapsto (y - a)^2$ is convex, L -Lipschitz with constant $L = 2(b + \kappa R)$ and M -Smooth with constant $M = 2$.
- (b) **(Logistic loss for regression):** If $\mathcal{Y} = \mathbb{R}$, then, analogously, the loss $\mathcal{Y} \times \mathbb{R} \ni (y, a) \mapsto -\log\left(\frac{4e^{y-a}}{(1+e^{y-a})^2}\right)$ is convex, L -Lipschitz with constant $L = 1$ and M -smooth with constant $M = 1$.
- (c) **(Logistic loss for classification):** For classification problems with $\mathcal{Y} = \{-1, 1\}$, analogously, the loss $\mathcal{Y} \times \mathbb{R} \ni (y, a) \mapsto \log(1 + e^{-ya})$ is convex, L -Lipschitz with constant $L = 1$ and M -Smooth with constant $M = 1/4$.

- (d) **(Exponential loss):** For classification problems with $\mathcal{Y} = \{-1, 1\}$, analogously, the loss $\mathcal{Y} \times [-\kappa R, \kappa R] \ni (y, a) \mapsto e^{-ya}$ is convex, L -Lipschitz with constant $L = e^{\kappa R}$ and M -smooth also with $M = e^{\kappa R}$.

Under Assumption **(Smooth)**, the empirical risk $w \mapsto \widehat{\mathcal{L}}(w)$ is differentiable and we have

$$\nabla \widehat{\mathcal{L}}(w) = \frac{1}{n} \sum_{j=1}^n \ell'(Y_j, \langle X_j, w \rangle) X_j. \tag{7}$$

With Assumptions **(Bound)** and **(Lip)**, via dominated convergence, the same is true for the population risk $w \mapsto \mathcal{L}(w)$ and we have

$$\nabla \mathcal{L}(w) = \int \ell'(y, \langle x, w \rangle) x \mathbb{P}_{(X,Y)}(d(x, y)). \tag{8}$$

Further, our assumptions on the loss directly translate into properties of the risks:

- (A3') **(R-Conv):** Under **(Conv)**, both the population and the empirical risk are convex.
- (A4') **(R-Lip):** Under **(Bound)** and **(Lip)**, both the population and the empirical risk are Lipschitz-continuous with constant κL .
- (A5') **(R-Smooth):** Under **(Bound)** and **(Smooth)**, the gradient of both the population and the empirical risk is Lipschitz-continuous with constant $\kappa^2 M$.

The derivation, which is straightforward, is included in Lemma 8 in [Appendix A](#).

3.1 Formulation of main results

A first key result shows that under the above assumptions, we can decompose the excess risk for the averaged GD-iterate \bar{v}_T as well as for the last iterate v_T .

Proposition 1 (Decomposition of the excess risk) *Suppose assumptions **(Bound)**, **(Conv)** and **(Smooth)** are satisfied. Consider the GD-iteration from Definition 1 with $T \in \mathbb{N}$ and constant step size $\gamma \leq 1/(\kappa^2 M)$ and let $w \in \mathcal{H}$ be arbitrary.*

- (i) *The risk of the averaged iterate \bar{v}_T satisfies*

$$\begin{aligned} \mathcal{L}(\bar{v}_T) - \mathcal{L}(w) &\leq \frac{1}{T} \sum_{t=1}^T \mathcal{L}(v_t) - \mathcal{L}(w) \\ &\leq \frac{\|v_0 - w\|^2}{2\gamma T} + \frac{1}{T} \sum_{t=1}^T \langle \nabla \mathcal{L}(v_{t-1}) - \nabla \widehat{\mathcal{L}}(v_{t-1}), v_t - w \rangle. \end{aligned}$$

- (ii) *The risk of the last iterate v_T satisfies*

$$\begin{aligned} \mathcal{L}(v_T) - \mathcal{L}(w) &\leq \frac{1}{T} \sum_{t=1}^T (\mathcal{L}(v_t) - \mathcal{L}(w)) \\ &\quad + \sum_{t=1}^{T-1} \frac{1}{t(t+1)} \sum_{s=T-t+1}^T \langle \nabla \mathcal{L}(v_{s-1}) - \nabla \widehat{\mathcal{L}}(v_{s-1}), v_s - v_{T-t} \rangle. \end{aligned}$$

The proof of Proposition 1 can be found in Appendix A. The above decomposition is derived using ideas from inexact optimization, in particular results studying inexact gradients, see e.g. [32, 33]. Indeed, our descent procedure can be regarded as one in which the population gradients are perturbed by the gradient noise terms

$$e_t := \nabla \widehat{\mathcal{L}}(v_t) - \nabla \mathcal{L}(v_t), \quad t = 1, \dots, T. \tag{9}$$

We further develop this discussion in Sect. 4.1.

Note that the results above apply to any $w \in \mathcal{H}$. Later, we will of course set $w = w_*$ from Assumption (Min). With this choice, Proposition 1 provides decompositions of the excess risk into a deterministic *bias part*

$$\frac{\|v_0 - w_*\|^2}{2\gamma T}, \tag{10}$$

which can be seen as an optimization error, and a stochastic *variance part*, which is an average of the terms

$$\begin{aligned} \langle -e_{t-1}, v_t - w_* \rangle \quad \text{and} \quad \langle -e_{s-1}, v_s - v_{T-t} \rangle, \\ t = 1, \dots, T, s = T - t + 1, \dots, T. \end{aligned} \tag{11}$$

Note that Proposition 1 (i) can be applied to the first sum on the right-hand side in (ii). In order to control the bias part, it is sufficient to choose γT large enough. Controlling the variance part is more subtle and requires some care. By Cauchy-Schwarz inequality,

$$\langle -e_{t-1}, v_t - w_* \rangle \leq \|e_{t-1}\| \|v_t - w_*\| \quad \text{for all } t = 1, \dots, T. \tag{12}$$

A similar estimate holds for the terms $\langle -e_{s-1}, v_s - v_{T-t} \rangle$, $s = T - t + 1, \dots, T$. This shows that in order to upper bound the excess risk of the gradient iteration, it is sufficient to solve two problems:

1. Bound the *gradient noise* terms $e_{t-1} = \nabla \widehat{\mathcal{L}}(v_{t-1}) - \nabla \mathcal{L}(v_{t-1})$ in norm;
2. Bound the *gradient path* $(v_t)_{t \geq 0}$ in a ball around w_* .

Starting from this observation, in Proposition 5, we state a general gradient concentration result which, for fixed $R > 0$, allows to derive

$$\sup_{\|v\| \leq R} \|\nabla \mathcal{L}(v) - \nabla \widehat{\mathcal{L}}(v)\| \leq 20\kappa^2 R(L + M) \sqrt{\frac{\log(4/\delta)}{n}} \tag{13}$$

with high probability in δ when n is sufficiently large. If we could prove that the gradient path $(v_t)_{t \geq 0}$ stays bounded, this would allow to control the gradient noise terms. Interestingly, the result in Eq. (13) itself is enough to directly derive a bound for the gradient path. In Proposition 7, we show how gradient concentration can be used to inductively prove that with high probability, $\|v_t - w_*\|$ stays bounded by $R = \max\{1, 3\|w_*\|\}$ for $t \leq T$ sufficiently large. Importantly, gradient concentration thereby allows to control the generalization error of the excess risk and the deviation of the gradient path at the same time. This makes this proof technique particularly appealing comparative to other approaches in the literature, see the discussion in Sects. 3.2 and 4. Taken together, the arguments above are sufficient to prove sharp rates for the excess risk.

Theorem 2 (Excess Risk) *Suppose Assumptions (Bound), (Conv), (Lip), (Smooth) and (Min) are satisfied. Let $v_0 = 0$, $T \geq 3$ and choose a constant step size $\gamma \leq \min\{1/(\kappa^2 M), 1\}$ in the GD-iteration from Definition 1. Then, for any $\delta \in (0, 1]$, such that*

$$\sqrt{n} \geq \max\{1, 90\gamma T \kappa^2(1 + \kappa L)(M + L)\} \sqrt{\log(4/\delta)}, \tag{14}$$

the average iterate \bar{v}_T and the last iterate v_T satisfy with probability at least $1 - \delta$ that

$$\begin{aligned} \mathcal{L}(\bar{v}_T) - \mathcal{L}(w_*) &\leq \frac{\|w_*\|^2}{2\gamma T} + 180 \max\{1, \|w_*\|^2\} \kappa^2(M + L) \sqrt{\frac{\log(4/\delta)}{n}}, \\ \mathcal{L}(v_T) - \mathcal{L}(w_*) &\leq \frac{\|w_*\|^2}{2\gamma T} + 425 \max\{1, \|w_*\|^2\} \kappa^2(M + L) \log(T) \sqrt{\frac{\log(4/\delta)}{n}}. \end{aligned}$$

In particular, setting $\gamma T = \sqrt{n}/(90\kappa^2(1 + \kappa L)(M + L)\sqrt{\log(4/\delta)})$ yields

$$\begin{aligned} \mathcal{L}(\bar{v}_T) - \mathcal{L}(w_*) &\leq 225 \max\{1, \|w_*\|^2\} \kappa^2(1 + \kappa L)(M + L) \sqrt{\frac{\log(4/\delta)}{n}}, \\ \mathcal{L}(v_T) - \mathcal{L}(w_*) &\leq 470 \max\{1, \|w_*\|^2\} \kappa^2(1 + \kappa L)(M + L) \log(T) \sqrt{\frac{\log(4/\delta)}{n}}. \end{aligned}$$

The proof of Theorem 2 is in Appendix A. To the best of our knowledge, it is not known if the above given rate of convergence is minimax optimal as there are no lower bounds so far in the literature for our set of assumptions on the class of loss functions. We emphasize, however, that the above bound for averaged GD with constant stepsize matches the minimax optimal rate for the least squares loss, see [14].

The gradient concentration inequality allows to derive an explicit estimate for the variance part. As expected, the latter improves as the number of samples increases, but interestingly it stays bounded, provided that γT is not too large, see Eq. (14). Optimizing the choice of γT leads to the final excess risk bound. The estimate is sharp in the sense that it matches the best available bounds for other estimation schemes based on empirical risk minimization with ℓ_2 -penalties, see e.g. [2, 38] and

the references therein. We note that the average and last iterates have essentially the same performance, up to constants and logarithmic terms.

A number of different choices for the stopping time T and the step size γ are possible, as long as their product stays constant. Assuming that κ from **(Bound)** is known, the user may choose the step size γ a priori when M from **(Smooth)** is known, see Example 1 (a), (b), (c). When M depends on the bound $R = \max\{1, 3\|w_*\|\}$, see Proposition 7, the choice of γ must be adapted to the norm of the minimizer w_* , see e.g. Example 1 (d) and the discussion in Remark 1. In this sense, it is indeed the product γT that plays the role of a regularization parameter, see also the simulations in Sect. 5.

The excess risk bound in the Theorem 2 matches the best bound for least squares, obtained with an ad hoc analysis [11, 12]. The obtained bound improves the results obtained in [27] and matches the rates for SGD in Theorem 4 of [29]. These latter results are more general and allow to derive fast rates. The generality is payed for, however, in terms of a considerably more complex analysis. In particular, our analysis allows to get explicit constants and keep the step size constant. More importantly, the proof we consider follows a different path, highlighting the connection to inexact optimization. We further develop this point of view next.

3.2 Discussion of related work

Comparison to the classical approach. In order to better locate our work in the machine learning and statistical literature, we compare it with the most important related line of research.

We contrast our approach with the one typically used to study learning with gradient descent and general loss functions. We briefly review this latter and more classical approach. The following decomposition is often considered to analyze the excess risk at v_t :

$$\mathcal{L}(v_t) - \mathcal{L}(w_*) = \underbrace{\mathcal{L}(v_t) - \widehat{\mathcal{L}}(v_t)}_{=(I)} + \underbrace{\widehat{\mathcal{L}}(v_t) - \widehat{\mathcal{L}}(w_*)}_{=(II)} + \underbrace{\widehat{\mathcal{L}}(w_*) - \mathcal{L}(w_*)}_{=(III)}, \tag{15}$$

see e.g. [2, 39]. The second term in the decomposition can be seen as an *optimization error* and treated by deterministic results from “exact” optimization. The first and last terms are stochastic and are bounded using probabilistic tools. In particular, the first term, often called *generalization error*, needs some care. The two more common approaches are based on stability, see e.g. [30, 31], or empirical process theory [38, 40]. The latter is considered in [27, 29]. In this case, the key quantity is the empirical process defined as

$$\sup_{\|v\| \leq R} |\widehat{\mathcal{L}}(v) - \mathcal{L}(v)|. \tag{16}$$

Here, a main complication is that the iterates norm/path needs to be bounded, which is a delicate point, as discussed in detail in Sect. 4.2. In our approach, gradient concentration allows to find a sharp bound on the gradient path and *at the same time* to

directly derive an excess risk bound, avoiding the decomposition in (15) and further empirical process bounds.

Inexact optimization and gradient concentration. We are not the first to employ tools from inexact optimization to treat learning problems, see [41] and [6]. A similar decomposition as in Proposition 1 together with a peeling argument instead of gradient concentration is used in [6]. There, the authors derive a bound for a “conditional excess risk”. More specifically, the risk is the conditional expectation, conditioned on the covariates, and is thus still a random quantity. The minimizer considered is the minimizer with respect to this random risk and therefore is a random quantity too. Additionally, their analysis requires strong convexity of the conditional risk with respect to the empirical norm. Our approach allows to overcome these two restrictions.

Also gradient concentration has been considered before, see e.g. [42, 43]. In [42], an analysis is developed under the assumption that minimization of the risk is constrained over a closed, convex and bounded set $\mathcal{W} \subset \mathbb{R}^d$, which effectively acts as an explicit regularization. During their gradient iteration, a projection step is considered to enforce this constraint. As a consequence, the dimension d and the diameter of \mathcal{W} appear as key quantities that determine the error behavior of their algorithm. The same is essentially true for [43]. In comparison, our analysis is dimension free. More importantly, we do not constrain the minimization problem. Hence, we consider implicit rather than explicit regularization. Also from a technical point of view, this is a key difference. As we discuss in Sect. 4.2, bounding the gradient path is required in the absence of explicit constraints. The main contribution of our paper, as we see it, is to show that the combination of optimization and concentration of measure techniques presented allow to seamlessly control the excess risk and the length of the gradient path at the same time, whereas in other analyses, e.g. [29], these two tasks have to be separated and are much more involved.

Finally, we discuss the results in [35], of which we had not been aware until after having finished this work. This paper also combines inexact optimization and gradient concentration, albeit in a different way. In Theorem G.1., the authors consider stochastic gradient descent for a convex and smooth objective function on \mathbb{R}^d , notably also on an unbounded domain. For their analysis, they introduce clipped versions of the stochastic gradients. They also borrow a decomposition of the excess risk from inexact optimization, although a different one. In particular, it is not straightforward that their decomposition would also yield results for the last gradient iteration. In a second step, they then use the conditional independence of gradient batches and a Bernstein-type inequality for Martingale differences to derive concentration for several terms involving the gradient noise. In comparison, instead of concentration based on individual batches, we use the full empirical gradients together with a uniform concentration result based on Rademacher complexities of Hilbert space valued function classes, see Sect. 4.2. On the one hand, our setting is more general, since we consider a Hilbert space instead of \mathbb{R}^d . On the other hand, [35] are notably able to forgo property **(R-Lip)**, i.e. their gradients can be unbounded. This is the main aspect of their analysis. As a consequence, their result is tailored to this setting and does not contain ours as a special case. With property **(R-Lip)**, even on

\mathbb{R}^d , our result is much sharper. We avoid an additional log-factor and, more importantly, we are able to freely choose a large, fixed step size $\gamma > 0$. In Theorem G.1. of [35], the step size has to depend both on the number of iterations and the high probability guarantee of the result. Further, our results in Theorem 2 are particularly sharp with explicit constants and one clear regularization parameter γT that can, in principle, be chosen via sample splitting and early stopping. Conversely, in order to control the unbounded gradients, [35] have to introduce two additional hyperparameters: the gradient clipping threshold λ and the batch size m . In their analysis, both of these have to be chosen in dependence of the true minimizer. Notably, the clipping threshold λ de facto regularizes the problem based on a priori knowledge of the true solution, the same way a bounded domain would. Developing these observations further would be an interesting venue for future research.

Last iterate vs. averaged iterates convergence. We compare our results to other high probability bounds for gradient descent. High probability bounds for both last iterate and (tail-)averaged gradient descent with constant stepsize for least squares regression in Hilbert spaces are well established. Indeed, the former follows from [14, 44] as gradient descent belongs to the broader class of spectral regularization methods. This is well known in the context of inverse problems, see e.g. [10]. As observed in [17], also average gradient descent can be cast and analyzed in the spectral filtering framework. Average and last iterates can be seen to share essentially the same excess risk bound. The proof, however, is heavily tailored to least squares. Compared to these results, for smooth losses, we establish a high probability bound of order $\mathcal{O}(1/\gamma T)$ for uniform averaging and $\mathcal{O}(\log(T)/\gamma T)$ for last iterate GD, for any n sufficiently large, with *constant* stepsize, worsened only by a factor $\log(T)$. We note that it was shown in [45] that the $\log(T)$ factor is in fact necessary for Lipschitz functions for last iterate SGD and GD with *decaying* stepsizes. The authors derive a sharp high probability bound of order $\mathcal{O}(\log(T)/\sqrt{T})$ for last iterate (S) GD, while uniform averaging achieves a faster rate of $\mathcal{O}(1/\sqrt{T})$. Notably, this work even shows the stronger statement: Any convex combination of the last k iterates must incur a $\log(T/k)$ factor. Finally, we note that [27] derive finite sample bounds

Table 1 Comparison of high probability last iterate bounds for gradient descent

Loss	Step size	Excess risk	Optimality
Least squares [14, 44]	$\gamma = \text{const.}$	$\mathcal{O}\left(\frac{1}{\gamma T}\right)$	Yes
Lipschitz & strongly convex [45]	$\gamma_t \simeq \frac{1}{t}$	$\mathcal{O}\left(\frac{\log(T)}{T}\right)$	Yes
Lipschitz [45]	$\gamma_t \simeq \frac{1}{t}$	$\mathcal{O}\left(\frac{\log(T)}{\sqrt{T}}\right)$	Yes
Convex [27]	$\gamma_t \simeq \frac{1}{\sqrt{t}}$	$\mathcal{O}\left(\frac{\log(T^{3/2})}{\sqrt{T}}\right)$	No
Convex & smooth [27]	$\gamma = \text{const.}$	$\mathcal{O}\left(\frac{1}{\sqrt{T}}\right)$	No
Convex & Lipschitz & smooth (this work)	$\gamma = \text{const.}$	$\mathcal{O}\left(\frac{\log(T)}{\gamma T}\right)$	Best known

Table 2 Comparison of high probability bounds for averaged gradient descent

Loss	Step size	Excess risk	Optimality
Least squares [17]	$\gamma = \text{const.}$	$\mathcal{O}\left(\frac{1}{\gamma T}\right)$	Yes
Lipschitz & strongly convex [45]	$\gamma_t \simeq \frac{1}{t}$	$\mathcal{O}\left(\frac{1}{T}\right)$	Yes
Lipschitz [45]	$\gamma_t \simeq \frac{1}{t}$	$\mathcal{O}\left(\frac{1}{\sqrt{T}}\right)$	Yes
Convex [27]	$\gamma_t \simeq \frac{1}{\sqrt{t}}$	$\mathcal{O}\left(\frac{\log(T)}{T^{1/4}}\right)$	No
Convex & Lipschitz & smooth (this work)	$\gamma = \text{const.}$	$\mathcal{O}\left(\frac{1}{\gamma T}\right)$	Best known

for subgradient descent for convex losses considering the last iterate. In this work, early stopping gives a suboptimal rate with decaying stepsize and also an additional logarithmic factor. This vanishes under additional differentiability and smoothness for constant stepsize. We condense an overview about the rates of convergence for different variants of GD under specific assumptions in Tables 1 and 2.

4 From inexact optimization to learning

In this section, we further discuss the important elements of the proof. The alternative error decomposition we presented in Proposition 1 follows from taking the point of view of optimization with inexact gradients [32]. The idea is to consider an ideal GD-iteration subject to noise, i.e.

$$v_{t+1} = v_t - \gamma(\nabla\mathcal{L}(v_t) + e_t) \quad t = 0, 1, 2, \dots, \tag{17}$$

where, the $(e_t)_{t \geq 0}$ are gradient noise terms. In Eq. (17), very general choices for e_t may be considered. Clearly, in our setting, we have

$$e_t = \nabla\widehat{\mathcal{L}}(v_t) - \nabla\mathcal{L}(v_t) \quad t = 0, 1, 2, \dots \tag{18}$$

From this perspective, the empirical GD-iteration can be seen as performing gradient descent directly on the *population* risk, where the gradient is corrupted with noise and convergence has to be balanced out with a control of the stability of the iterates. Next, we see how these ideas can be applied to the learning problem.

4.1 Inexact gradient descent

From the point of view discussed above, it becomes essential to relate both the risk and the norm of a fixed GD-iteration to the gradient noise. In the following, we provide two technical Lemmas which do exactly that. Both results could also be formulated for general gradient noise terms $(e_t)_{t \geq 0}$. For the sake of simplicity, however, we

opt for the more explicit formulation in terms of the gradients. The proofs are based on entirely deterministic arguments and can be found in [Appendix B](#).

Lemma 3 (Inexact gradient descent: Risk) *Suppose assumptions **(Bound)**, **(Conv)**, and **(Smooth)** are satisfied. Consider the GD-iteration from Definition 1 with constant step size $\gamma \leq 1/(\kappa^2 M)$ and let $w \in \mathcal{H}$. Then, for any $t \geq 1$, the risk of the iterate v_t satisfies*

$$\mathcal{L}(v_t) - \mathcal{L}(w) \leq \frac{1}{2\gamma} (\|v_{t-1} - w\|^2 - \|v_t - w\|^2) + \langle \nabla \mathcal{L}(v_{t-1}) - \widehat{\nabla} \mathcal{L}(v_{t-1}), v_t - w \rangle.$$

Lemma 3 is the key component to obtain the decomposition of the excess risk in Proposition 1 for the averaged GD-iteration. This *online to batch conversion* easily follows by exploiting the convexity of the population risk **(R-Conv)**.

The next Lemma is crucial in providing a high probability guarantee for the boundedness of the gradient path in Proposition 7, which is necessary to apply gradient concentration to the decomposition of the excess risk in Proposition 1.

Lemma 4 (Inexact gradient descent: Gradient path) *Suppose assumptions **(Bound)**, **(Conv)**, **(Lip)**, **(Smooth)** and **(Min)** are satisfied and choose a constant step size $\gamma \leq \min\{1/(\kappa^2 M), 1\}$ in Definition 1. Then, for any $t \geq 0$, the norm of the GD-iterate v_{t+1} is recursively bounded by*

$$\begin{aligned} \|v_{t+1} - w_*\|^2 &\leq \|v_0 - w_*\|^2 \\ &\quad + 2\gamma \sum_{s=0}^t \left(\langle \nabla \mathcal{L}(v_s) - \widehat{\nabla} \mathcal{L}(v_s), v_s - w_* \rangle + \kappa L \|\nabla \mathcal{L}(v_s) - \widehat{\nabla} \mathcal{L}(v_s)\| \right). \end{aligned}$$

Assuming that for some fixed $R > 0$, $\|v_s - w_*\| \leq R$ for all $s \leq t$, Lemma 4 guarantees that

$$\|v_{t+1} - w_*\|^2 \leq \|v_0 - w_*\|^2 + 2\gamma(R + \kappa L) \sum_{s=1}^t \|\nabla \mathcal{L}(v_s) - \widehat{\nabla} \mathcal{L}(v_s)\|, \tag{19}$$

which, in combination with gradient concentration, allows for an inductive bound on $\|v_{t+1}\|$. Summarizing, Lemmas 3 and 4 can be regarded as tools to study our learning problem using gradient concentration directly.

4.2 Gradient concentration

In this section, we discuss how the gradient concentration inequality in Eq. (13) is derived using results from [34]. We use a gradient concentration result which is expressed in terms of the Rademacher complexity of a function class defined by the gradients $w \mapsto \nabla \mathcal{L}(w)$ with

$$\nabla \mathcal{L}(w) = \int \ell'(y, \langle x, w \rangle) x \mathbb{P}_{(X,Y)}(d(x, y)), \quad w \in \mathcal{H}. \tag{20}$$

Since the gradients above are elements of the Hilbert space \mathcal{H} , the notion of Rademacher complexities has to be stated for Hilbert space-valued function classes, see [46].

Definition 2 (Rademacher complexities) Let $(\mathcal{H}, \|\cdot\|)$ be a real, separable Hilbert space. Further, let \mathcal{G} be a class of maps $g : \mathcal{Z} \rightarrow \mathcal{H}$ and $Z = (Z_1, \dots, Z_n) \in \mathcal{Z}^n$ be a vector of i.i.d. random variables. We define the *empirical* and *population* Rademacher complexities of \mathcal{G} by

$$\widehat{\mathcal{R}}_n(\mathcal{G}) := \mathbb{E}_\varepsilon \left[\sup_{g \in \mathcal{G}} \left\| \frac{1}{n} \sum_{j=1}^n \varepsilon_j g(Z_j) \right\| \right] \quad \text{and} \quad \mathcal{R}_n(\mathcal{G}) := \mathbb{E}_Z [\widehat{\mathcal{R}}_n(\mathcal{G})] \tag{21}$$

respectively, where $\varepsilon = (\varepsilon_1, \dots, \varepsilon_n) \in \{-1, +1\}^n$ is a vector of i.i.d. Rademacher random variables independent of Z .

In our setting, $(Z_1, \dots, Z_n) = ((X_1, Y_1), \dots, (X_n, Y_n))$. Fix some $R > 0$ and consider the scalar function class

$$\mathcal{F}_R := \{f_v = \langle \cdot, v \rangle : \|v\| \leq R\} \subset L^2(\mathbb{P}_X) \tag{22}$$

and more importantly, the \mathcal{H} -valued, composite function class

$$\mathcal{G}_R := \nabla \ell \circ \mathcal{F}_R := \{\mathcal{Y} \times \mathcal{H} \ni (x, y) \mapsto \ell'(y, f(x))x : f \in \mathcal{F}_R\}. \tag{23}$$

Under **(Bound)** and **(Lip)**, we have

$$G_R := \sup_{g \in \mathcal{G}_R} \|g\|_\infty = \sup_{f \in \mathcal{F}_R} \|\ell'(Y, f(X))X\|_\infty \leq \kappa L, \tag{24}$$

where $\|\cdot\|_\infty$ denotes the ∞ -norm on the underlying probability space $(\Omega, \mathcal{F}, \mathbb{P})$. The gradient concentration result can now be formulated in terms of the empirical Rademacher complexity of \mathcal{G}_R .

Proposition 5 (Gradient concentration) Suppose assumption **(Bound)** **(Lip)** and **(Smooth)** are satisfied and let $R > 0$. Then, for any $\delta > 0$,

$$\sup_{\|v\| \leq R} \|\nabla \mathcal{L}(v) - \nabla \widehat{\mathcal{L}}(v)\| \leq 4\widehat{\mathcal{R}}_n(\mathcal{G}_R) + G_R \sqrt{\frac{2 \log(4/\delta)}{n}} + G_R \frac{4 \log(4/\delta)}{n}$$

with probability at least $1 - \delta$, where G_R is defined in Eq. (24).

The proof of Proposition 5 is stated in Appendix B. To apply Proposition 5, we need to bound $\widehat{\mathcal{R}}_n(\mathcal{G}_R)$. This can be done via relating the empirical Rademacher complexity of the composite function class \mathcal{G}_R to the complexity of the scalar function class \mathcal{F}_R .

Lemma 6 (Bounds on the empirical Rademacher complexities) *Fix $R > 0$.*

- (i) Under **(Bound)**, we have $\widehat{\mathcal{R}}(\mathcal{F}_R) \leq \frac{\kappa R}{\sqrt{n}}$.
- (ii) Under **(Bound)**, **(Lip)** and **(Smooth)**, we have

$$\widehat{\mathcal{R}}(\mathcal{G}_R) \leq 2\sqrt{2} \left(\frac{\kappa L}{\sqrt{n}} + \kappa M \widehat{\mathcal{R}}(\mathcal{F}_R) \right) \leq \frac{2\sqrt{2}(\kappa L + \kappa^2 MR)}{\sqrt{n}}.$$

Note that since the bounds in Lemma 6 do not depend on the sample $(X_1, Y_1), \dots, (X_n, Y_n)$, they also hold for the population Rademacher complexities. Lemma 6 (i) is a classic result, which we restate for completeness. Lemma 6 (ii) is more involved and requires combining a vector-contraction inequality from [46] with additional more classical contraction arguments to disentangle the concatenation in the function class \mathcal{G}_R . The Proof of Lemma 6 is stated in Appendix B. Note that the arguments for both Proposition 5 and Lemma 6 are essentially contained in [34]. Here, we provide a self-contained derivation for our setting.

Together with Lemma 4, the gradient concentration result provides an immediate high probability guarantee for the gradient path not to diverge too far from the minimizer w_* .

Proposition 7 (Bounded gradient path) *Suppose assumptions **(Bound)**, **(Conv)**, **(Lip)**, **(Smooth)** and **(Min)** are satisfied, set $v_0 = 0$ and choose a constant step size $\gamma \leq \min\{1/(\kappa^2 M), 1\}$ in Definition 1. Fix $\delta \in (0, 1]$ such that*

$$\sqrt{n} \geq \max\{1, 90\gamma T \kappa^2 (1 + \kappa L)(M + L)\} \sqrt{\log(4/\delta)} \tag{25}$$

and $R = \max\{1, 3\|w_*\|\}$. Then, on the gradient concentration event from Proposition 5 with probability at least $1 - \delta$ for the above choice of R , we have

$$\|v_t\| \leq R \quad \text{and} \quad \|v_t - w_*\| \leq \frac{2R}{3} \quad \text{for all } t = 1, \dots, T.$$

The proof of Proposition 7 is stated in Appendix B. In a learning setting, bounding the gradient path is essential to the analysis of gradient descent based estimation procedures. Either one has to guarantee its boundedness a priori, e.g. by projecting back onto a ball of known radius $R > 0$ or making highly restrictive additional assumptions, see [47], or one has to make usually involved arguments to guarantee its boundedness up to a sufficiently large iteration number, see e.g. [6, 29]. Our numerical illustrations in Sect. 5 show that from a practical perspective, such a boundedness result is indeed necessary to control the variance. Additionally, if the boundedness of the gradient path was already controlled by the optimization procedure for arbitrarily large iterations T , then the decomposition in Proposition 1 together with our gradient concentration result in Proposition 5 would guarantee that

for $T \rightarrow \infty$, the deterministic bias part $\|w_*\|^2/(2\gamma T)$ vanishes completely, while the stochastic variance part

$$\frac{1}{T} \sum_{t=1}^T \langle \nabla \mathcal{L}(v_{t-1}) - \nabla \widehat{\mathcal{L}}(v_{t-1}), v_t - w_* \rangle \tag{26}$$

would remain of order $\sqrt{\log(4/\delta)/n}$ independently of T . This would suggest that for large T , there is no tradeoff between reducing the bias of the estimation method and its variance anymore, which, in that form, should be surprising for learning, see the discussion in [48]. From this perspective, to analyze gradient descent for learning, it seems necessary to establish a result like Proposition 7.

As stated above, several other results rely on bounding the gradient path to obtain bounds on the excess risk. We compare our result in Proposition 7 with the techniques used in [29], that are the most recent in the literature. Under the *self-boundedness* assumption

$$|\ell'(y, a)|^2 \lesssim \ell(y, a) + 1 \quad \text{for all } y, a \in \mathbb{R}, \tag{27}$$

the authors relate the stochastic gradient descent iteration v_t to the Tikhonov regularizer w_λ , whose norm can be controlled, and obtain a uniform bound over $t = 1, \dots, T$ of the form

$$\|v_{t+1}\|^2 \lesssim \sum_{s=1}^t \gamma_s \max\{0, \widehat{\mathcal{L}}(w_\lambda) - \widehat{\mathcal{L}}(v_s)\} + \log(2T/\delta)(\|w_\lambda\|^2 + 1) \tag{28}$$

with high probability in δ . Later, the risk quantities in Eq. (28) are related to the approximation error of a kernel space. Inductively, this guarantees that the stochastic gradient path stays sufficiently bounded. For the bound in Eq. (28), the authors in [29] have to choose a decaying sequence of step sizes γ_t with $\sum_{t=1}^T \gamma_t^2 < \infty$. In comparison, the result in Proposition 7 allows for a fixed step size $\gamma > 0$. Since sharp rates essentially require that $\sum_{t=0}^T \gamma_t$ is of order \sqrt{n} , we may therefore stop the algorithm earlier. In this regard, our result is slightly sharper. At the same time, the result in [29] is more general. Under a capacity condition, the authors adapt the bound in Eq. (28) to allow for fast rates. However, both the proof of Eq. (28) and its adaptation to the capacity dependent setting are complex and quite technical. In comparison, Proposition 7 is an immediate corollary of Proposition 5. In particular, if under additional assumptions, a sharper concentration result for the gradients is possible, our proof technique would immediately translate this to the bound on the gradient path that is needed to guarantee this sharper rate for the excess risk. Indeed, we think these ideas can be fruitfully developed to get new improved results.

5 Numerics

In this section, we provide empirical illustrations of the effects described in Sects. 3 and 4. In particular, we consider the logistic loss for regression from Example 1 (b) and the exponential loss for classification Example 1 (d). We concentrate on two aspects: The (un)bounded gradient path for the averaged iterates and the interplay between step size and stopping time. Our experiments are conducted on synthetic data with $d = 100$ dimensions, generated as follows: We set the covariance matrix $\Sigma \in \mathbb{R}^{d \times d}$ as a diagonal matrix with entries $\Sigma_{jj} = j^{-2}$, $j = 1, \dots, d$ and choose $w_* = \Sigma e$, with $e = (1, \dots, 1)^\top \in \mathbb{R}^d$. We generate $n_{\text{train}} = 1000$ training data, where the covariates X_j are drawn from a Gaussian distribution with zero mean and covariance Σ . For the logistic loss, the labels follow the model

$$Y_j = \langle X_j, w_* \rangle + \epsilon_j, \quad j = 1, \dots, n \tag{29}$$

with $\epsilon_j \sim N(0, 5)$ i.i.d. For the exponential loss, it is a well known fact that the risk is minimized by half the conditional log-odds. Therefore, we choose the labels as independent observations of

$$Y_j \sim 2(\text{Ber}(p_j) - 0.5), \quad \text{with} \quad \log\left(\frac{p_j}{1 - p_j}\right) = 2\langle X w_* \rangle_j, \quad j = 1, \dots, n \tag{30}$$

such that Assumption **(Min)** is satisfied. Each experiment is repeated 1000 times and we report the average. The results are presented in Figs. 1 and 2.

Our first experiment illustrates the behavior of the path $t \mapsto \|v_t - w_*\|$ for a fixed step size. We report the average path length together with the minimum and maximum path lengths. As Proposition 7 suggests, the path becomes unbounded when the number of iterations grows large.

In a second experiment, we choose a grid of step sizes γ and stopping times T and report the average excess test risk with $n_{\text{test}} = \lceil n_{\text{train}}/3 \rceil$ test data. Note that the grid

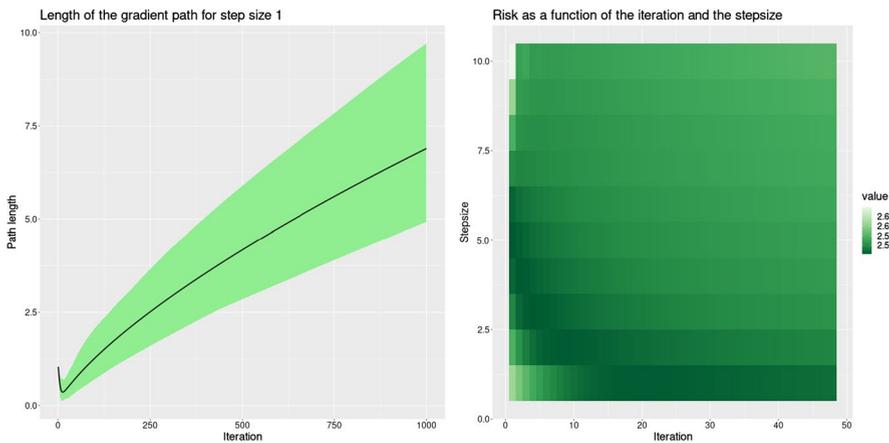


Fig. 1 Simulation results for the logistic loss

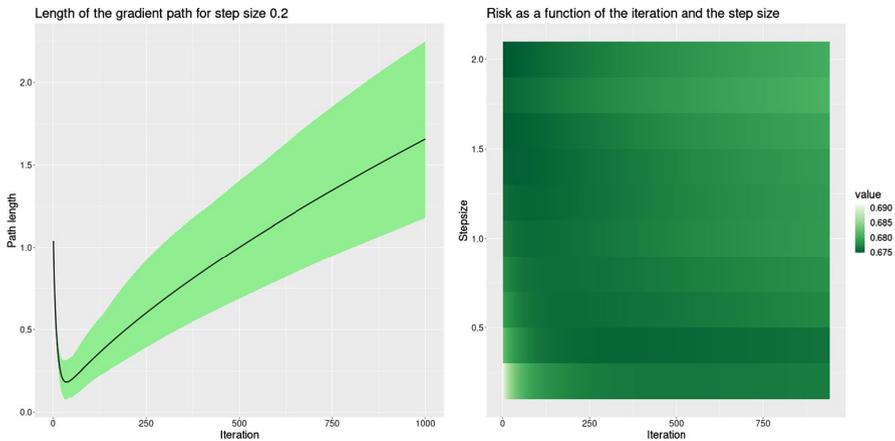


Fig. 2 Simulation results for the exponential loss

of step sizes is chosen differently for the individual loss functions, since larger values of the Lipschitz constant M of the gradient require smaller step sizes. As Theorem 2 predicts, for fixed n_{train} , the performance of averaged GD remains roughly constant as $\gamma \cdot T$ remains constant.

6 Conclusion

In this paper, we studied implicit/iterative regularization for possibly infinite dimensional, linear models, where the error cost is a convex, differentiable loss function. Our main contribution is a sharp high probability bound on the excess risk of the averaged and last iterate of batch gradient descent. We derive these results combining ideas and results from optimization and statistics. Indeed, we show how it is possible to leverage results from inexact optimization together with concentration inequalities for vector valued functions. The theoretical results are illustrated to see how the step size and the iteration number control the bias and the stability of the solution.

A number of research directions can further be developed. In our study, we favored a simple analysis to illustrate the main ideas, and as a consequence our results are limited to a basic setting. It would be interesting to develop the analysis we presented to get faster learning rates under further assumptions, for example considering capacity conditions or even finite dimensional models. Another possible research direction is to consider less regular loss functions, dropping the differentiability assumption. Along similar lines it would be interesting to consider other forms of implicit bias or non linear models. Finally, other forms of optimization, including stochastic and accelerated methods, could be considered.

Appendix A: Proofs for Sect. 3

Lemma 8 (*Properties of the risks*)

(i) Under **(Conv)**, the population risk is convex, i.e., for all $v, w \in \mathcal{H}$, we have

$$\mathcal{L}(v) \leq \mathcal{L}(w) - \langle \nabla \mathcal{L}(v), w - v \rangle. \tag{A1}$$

(ii) Under **(Bound)** and **(Lip)**, the population risk is Lipschitz-continuous with constant κL , i.e., for all $v, w \in \mathcal{H}$, we have

$$|\mathcal{L}(v) - \mathcal{L}(w)| \leq \kappa L \|v - w\| \tag{A2}$$

(iii) Under **(Bound)** and **(Smooth)**, the gradient of the population risk is Lipschitz-continuous with constant $\kappa^2 M$, i.e., for all $v, w \in \mathcal{H}$, we have

$$\|\nabla \mathcal{L}(v) - \nabla \mathcal{L}(w)\| \leq \kappa^2 M \|v - w\|. \tag{A3}$$

Note that this implies that

$$\mathcal{L}(w) \leq \mathcal{L}(v) + \langle \nabla \mathcal{L}(v), w - v \rangle + \frac{\kappa^2 M}{2} \|w - v\|^2. \tag{A4}$$

Moreover, (i), (ii) and (iii) also hold for the empirical risk $\widehat{\mathcal{L}}$ with the same constants.

Proof

(i) This follows directly from **(Conv)** and the linearity of the expectation.

(ii) For $v, w \in \mathcal{H}$, we have

$$\begin{aligned} |\mathcal{L}(w) - \mathcal{L}(v)| &= |\mathbb{E}_{(X,Y)}[\mathcal{L}(Y, \langle X, w \rangle)] - \mathbb{E}_{(X,Y)}[\mathcal{L}(Y, \langle X, v \rangle)]| \\ &\leq L \mathbb{E}_{(X,Y)}[|\langle X, w - v \rangle|] \leq \kappa L \|w - v\|, \end{aligned} \tag{A5}$$

where the first inequality follows from **(Lip)** and Cauchy-Schwarz inequality and the second inequality follows from **(Bound)**.

(iii) For $v, w \in \mathcal{H}$, we have

$$\begin{aligned} \|\nabla \mathcal{L}(w) - \nabla \mathcal{L}(v)\| &= \|\mathbb{E}_{(X,Y)}[\mathcal{L}'(Y, \langle X, w \rangle)X] - \mathbb{E}_{(X,Y)}[\mathcal{L}'(Y, \langle X, v \rangle)X]\| \\ &\leq \kappa |\mathbb{E}_{(X,Y)}[\mathcal{L}'(Y, \langle X, w \rangle)] - \mathbb{E}_{(X,Y)}[\mathcal{L}'(Y, \langle X, v \rangle)]| \\ &\leq \kappa M |\mathbb{E}_{(X,Y)}[\langle X, w - v \rangle]| \leq \kappa^2 M \|w - v\|, \end{aligned} \tag{A6}$$

where the first and the third inequality follow from **(Bound)** and Cauchy-Schwarz inequality and the second one follows from **(Smooth)**. □

For the proof of the second part of Proposition 1, we need the following simple Lemma. A different version of this was put forward in a blog post by Francesco Orabona with a reference to the convergence proof of the last iterate of SGD in [27]. Since our version is different and for the sake of completeness, we give a full proof.

Lemma 9 Let $(q_t)_{t=1,\dots,T}$ be a sequence of real numbers. Then,

$$q_T = \frac{1}{T} \sum_{t=1}^T q_t + \sum_{t=1}^{T-1} \frac{1}{t(t+1)} \sum_{s=T-t+1}^T (q_s - q_{T-t}).$$

Proof Define

$$S_t := \frac{1}{t} \sum_{s=T-t+1}^T q_s, \quad t = 1, \dots, T. \tag{A7}$$

Then, any $t \leq T - 1$ satisfies

$$\begin{aligned} tS_t &= (t+1)S_{t+1} - q_{T-t} = tS_{t+1} + S_{t+1} - q_{T-t} \\ &= tS_{t+1} + \frac{1}{t+1} \sum_{s=T-t}^T (q_s - q_{T-t}), \end{aligned} \tag{A8}$$

which implies

$$S_t = S_{t+1} + \frac{1}{t(t+1)} \sum_{s=T-t}^T (q_s - q_{T-t}). \tag{A9}$$

Inductively applying (A9), we obtain

$$\begin{aligned} q_T = S_1 = S_T + \sum_{t=1}^{T-1} \frac{1}{t(t+1)} \sum_{s=T-t}^T (q_s - q_{T-t}). \\ = \frac{1}{T} \sum_{t=1}^T q_t + \sum_{t=1}^{T-1} \frac{1}{t(t+1)} \sum_{s=T-t+1}^T (q_s - q_{T-t}). \end{aligned} \tag{A10}$$

□

Proof of Proposition 1 (Decomposition of the excess risk)

(i) From **(R-Conv)** and Lemma 3, we obtain

$$\begin{aligned} &\frac{1}{T} \sum_{t=1}^T \mathcal{L}(v_t) - \mathcal{L}(w) \\ &\leq \frac{1}{T} \sum_{t=1}^T \frac{\|v_{t-1} - w\|^2 - \|v_t - w\|^2}{2\gamma} + \frac{1}{T} \sum_{t=1}^T \langle \nabla \mathcal{L}(v_{t-1}) - \nabla \widehat{\mathcal{L}}(v_{t-1}), v_t - w \rangle \tag{A11} \\ &\leq \frac{\|v_0 - w\|^2}{2\gamma T} + \frac{1}{T} \sum_{t=1}^T \langle \nabla \mathcal{L}(v_{t-1}) - \nabla \widehat{\mathcal{L}}(v_{t-1}), v_t - w \rangle, \end{aligned}$$

where we have resolved the telescopic sum, to obtain the last inequality.

(ii) Applying Lemma 9 with $q_t = \mathcal{L}(v_t) - \mathcal{L}(w)$, we find

$$\begin{aligned} \mathcal{L}(v_T) - \mathcal{L}(w) &= \frac{1}{T} \sum_{t=1}^T (\mathcal{L}(v_t) - \mathcal{L}(w)) \\ &+ \sum_{t=1}^{T-1} \frac{1}{t(t+1)} \sum_{s=T-t+1}^T (\mathcal{L}(v_s) - \mathcal{L}(v_{T-t})). \end{aligned} \tag{A12}$$

We aim at bounding the last sum in the above equality. Summing the bound in Lemma 3 from $T - t + 1$ to T yields that for all $v \in \mathcal{H}$,

$$\begin{aligned} \sum_{s=T-t+1}^T (\mathcal{L}(v_s) - \mathcal{L}(v)) &\leq \frac{1}{2\gamma} \|v_{T-t} - v\|^2 \\ &+ \sum_{s=T-t+1}^T \langle \nabla \mathcal{L}(v_{s-1}) - \nabla \hat{\mathcal{L}}(v_{s-1}), v_s - v \rangle. \end{aligned} \tag{A13}$$

Hence, setting $v = v_{T-t}$ yields

$$\sum_{s=T-t+1}^T (\mathcal{L}(v_s) - \mathcal{L}(v_{T-t})) \leq \sum_{s=T-t+1}^T \langle \nabla \mathcal{L}(v_{s-1}) - \nabla \hat{\mathcal{L}}(v_{s-1}), v_s - v \rangle. \tag{A14}$$

The result follows by plugging the last inequality into (A12). □

Proof of Theorem 2 (Excess risk) We initially consider the case of the averaged GD-iterate. By convexity, Proposition 1 and an application of Cauchy-Schwarz inequality, we have

$$\begin{aligned} \mathcal{L}(\bar{v}_T) - \mathcal{L}(w_*) &\leq \frac{1}{T} \sum_{t=1}^T \mathcal{L}(v_t) - \mathcal{L}(w) \\ &\leq \frac{\|w_*\|^2}{2\gamma T} + \frac{1}{T} \sum_{t=1}^T \langle \nabla \mathcal{L}(v_t) - \nabla \hat{\mathcal{L}}(v_t), v_t - w_* \rangle \\ &\leq \frac{\|w_*\|^2}{2\gamma T} + \frac{1}{T} \sum_{t=1}^T \|\nabla \mathcal{L}(v_t) - \nabla \hat{\mathcal{L}}(v_t)\| \|v_t - w_*\|. \end{aligned} \tag{A15}$$

The assumptions of Theorem 2 are chosen exactly as in Proposition 7. Therefore, on the gradient concentration event with probability at least $1 - \delta$ from Proposition 5 and the choice R as above, we have

$$\begin{aligned} \|v_t - w_*\| &\leq \frac{2R}{3}, \quad \|v_t\| \leq R \quad \text{and} \\ \|\nabla \mathcal{L}(v_t) - \nabla \hat{\mathcal{L}}(v_t)\| &\leq 20\kappa^2 R(L + M) \sqrt{\frac{\log(4/\delta)}{n}} \\ &\text{for all } t = 0, 1, \dots, T, \end{aligned} \tag{A16}$$

where the last inequality is derived in exactly the same way as in the proof of Proposition 7. Plugging this into the inequality in Eq. (A15), we obtain

$$\begin{aligned} \mathcal{L}(\bar{v}_T) - \mathcal{L}(w_*) &\leq \frac{\|w_*\|^2}{2\gamma T} + 20\kappa^2 R^2(M + L)\sqrt{\frac{\log(4/\delta)}{n}}. \\ &\leq \frac{\|w_*\|^2}{2\gamma T} + 180 \max\{1, \|w_*\|^2\}\kappa^2(M + L)\sqrt{\frac{\log(4/\delta)}{n}}. \end{aligned} \tag{A17}$$

For the last iterate, we set $e_t := \nabla \hat{\mathcal{L}}(v_t) - \hat{\mathcal{L}}(v_t)$, $t = 1, \dots, T$ to reduce the notation. Proposition 1 with an application of Cauchy-Schwarz yields

$$\begin{aligned} \mathcal{L}(v_T) - \mathcal{L}(w_*) &\leq \frac{1}{T} \sum_{t=1}^T (\mathcal{L}(v_t) - \mathcal{L}(w_*)) + \sum_{t=1}^{T-1} \frac{1}{t(t+1)} \sum_{s=T-t+1}^T \langle -e_{s-1}, v_s - v_{T-t} \rangle \\ &\leq \frac{\|w_*\|^2}{2\gamma T} + \frac{1}{T} \sum_{t=1}^T \langle -e_{s-1}, v_s - v_{T-t} \rangle \\ &\quad + \sum_{t=1}^{T-1} \frac{1}{t(t+1)} \sum_{s=T-t+1}^T \langle -e_{s-1}, v_s - v_{T-t} \rangle \\ &\leq \frac{\|w_*\|^2}{2\gamma T} + \frac{1}{T} \sum_{t=1}^T \|e_s\| \|v_s - v_*\| \\ &\quad + \sum_{t=1}^{T-1} \frac{1}{t(t+1)} \sum_{s=T-t+1}^T \|e_s\| \|v_s - v_{T-t}\|. \end{aligned} \tag{A18}$$

Now, by Propositions 5 and 7, if

$$\sqrt{n} \geq 90\gamma T \kappa^2 (1 + \kappa L)(M + L)\sqrt{\log(4/\delta)}, \tag{A19}$$

we find with probability at least $1 - \delta$ for all $t = 0, \dots, T$, that

$$\|w_*\| \leq \frac{2R}{3} \leq R, \quad \|v_t\| \leq R, \tag{A20}$$

$$\|e_t\| \leq \sup_{v \in \mathcal{F}_R} \|\nabla \mathcal{L}(v) - \nabla \hat{\mathcal{L}}(v)\| \leq 20\kappa^2 R(L + M)\sqrt{\frac{\log(4/\delta)}{n}}. \tag{A21}$$

In particular,

$$\|v_s - v_{T-t}\| \leq \frac{4R}{3} \tag{A22}$$

for any $s = T - t + 1, \dots, T, t = 1, \dots, T$. Hence, with probability at least $1 - \delta$,

$$\begin{aligned} \mathcal{L}(v_T) - \mathcal{L}(w_*) &\leq \frac{\|w_*\|^2}{2\gamma T} + 20\kappa^2 R^2(L + M)\sqrt{\frac{\log(4/\delta)}{n}} \\ &\quad + \frac{80}{3}\kappa^2 R^2(L + M)\sqrt{\frac{\log(4/\delta)}{n}} \left(\sum_{t=1}^{T-1} \frac{1}{t}\right). \end{aligned}$$

Finally, since

$$\sum_{t=1}^{T-1} \frac{1}{t} \leq \log(T - 1) \leq \log T, \tag{A23}$$

we arrive at

$$\mathcal{L}(v_T) - \mathcal{L}(w_*) \leq \frac{\|w_*\|^2}{2\gamma T} + \frac{140}{3}\kappa^2 R^2(L + M)\sqrt{\frac{\log(4/\delta)}{n}} \log T. \tag{A24}$$

Plugging in for R^2 and simplifying the constant completes the proof. □

Appendix B: Proofs for Sect. 4

Proof of Lemma 3 (Inexact gradient descent: Risk) By **(R-Conv)** (Eq. (A1)) and **(R-Smooth)** (Eq. (A4)), the population risk is convex and $\kappa^2 M$ -smooth. We have

$$\begin{aligned} \mathcal{L}(v_t) &\leq \mathcal{L}(v_{t-1}) + \langle \nabla \mathcal{L}(v_{t-1}), v_t - v_{t-1} \rangle + \frac{\kappa^2 M}{2} \|v_t - v_{t-1}\|^2 \\ &\leq \mathcal{L}(w) + \langle \nabla \mathcal{L}(v_{t-1}), v_{t-1} - w \rangle + \langle \nabla \mathcal{L}(v_{t-1}), v_t - v_{t-1} \rangle + \frac{\kappa^2 M}{2} \|v_t - v_{t-1}\|^2 \\ &\leq \mathcal{L}(w) + \langle \nabla \mathcal{L}(v_{t-1}), v_t - w \rangle + \frac{1}{2\gamma} \|v_t - v_{t-1}\|^2, \end{aligned} \tag{B25}$$

where the last inequality uses the fact that $\gamma \leq 1/(\kappa^2 M)$. The statement now follows from

$$\|v_t - v_{t-1}\|^2 = \|v_{t-1} - w\|^2 - \|v_t - w\|^2 - 2\gamma \langle \widehat{\nabla \mathcal{L}}(v_{t-1}), v_t - w \rangle. \tag{B26}$$

□

Proof of Lemma 4 (Inexact gradient descent: Gradient path) For $v, w \in \mathcal{H}$, the smoothness of the risk **(R-Smooth)** (Eq. (A3)) implies that

$$\|\nabla \mathcal{L}(v) - \nabla \mathcal{L}(w)\|^2 \leq \kappa^2 M \langle v - w, \nabla \mathcal{L}(v) - \nabla \mathcal{L}(w) \rangle, \tag{B27}$$

see e.g. Equation (3.6) in [37]. In particular, since $\nabla \mathcal{L}(w_*) = 0$, we have

$$\|\nabla \mathcal{L}(v)\|^2 \leq \kappa^2 M \langle v - w_*, \nabla \mathcal{L}(v) \rangle \quad \text{for all } v \in \mathcal{H}. \tag{B28}$$

Setting $e_s := \nabla \widehat{\mathcal{L}}(v_s) - \nabla \mathcal{L}(v_s)$, we obtain that for any $s \geq 0$,

$$\begin{aligned} \|v_{s+1} - w_*\|^2 &= \|v_s - w_*\|^2 - 2\gamma \langle \nabla \widehat{\mathcal{L}}(v_s), v_s - w_* \rangle + \gamma^2 \|\nabla \widehat{\mathcal{L}}(v_s)\|^2 \\ &= \|v_s - w_*\|^2 - 2\gamma \langle e_s, v_s - w_* \rangle \\ &\quad \underbrace{- 2\gamma \langle \nabla \mathcal{L}(v_s), v_s - w_* \rangle + \gamma^2 \|\nabla \mathcal{L}(v_s)\|^2}_{=(I)} \\ &\quad \underbrace{+ \gamma^2 \|\nabla \widehat{\mathcal{L}}(v_s)\|^2 - \gamma^2 \|\nabla \mathcal{L}(v_s)\|^2}_{=(II)}. \end{aligned} \tag{B29}$$

We treat the terms (I) and (II) separately: By Eq. (B27) and our choice of $\gamma \leq 1/(\kappa^2 M)$, we have

$$\begin{aligned} (I) &= -2\gamma \langle \nabla \mathcal{L}(v_s), v_s - w_* \rangle + \gamma^2 \|\nabla \mathcal{L}(v_s)\|^2 \\ &\leq \left(\frac{-2\gamma}{\kappa^2 M} + \gamma^2 \right) \|\nabla \mathcal{L}(v_s)\|^2 \leq 0. \end{aligned} \tag{B30}$$

Further, by (R-Lip) (Eq. (A2)), Cauchy-Schwarz inequality and the fact that $\gamma \leq 1$,

$$\begin{aligned} (II) &= \gamma^2 \|\nabla \widehat{\mathcal{L}}(v_s)\|^2 - \gamma^2 \|\nabla \mathcal{L}(v_s)\|^2 = \gamma^2 \langle \nabla \widehat{\mathcal{L}}(v_s) + \nabla \mathcal{L}(v_s), e_s \rangle \\ &\leq \gamma^2 \|\nabla \widehat{\mathcal{L}}(v_s) + \nabla \mathcal{L}(v_s)\| \|e_s\| \leq 2\gamma \kappa L \|e_s\|. \end{aligned} \tag{B31}$$

Together, Eqs. (B30) and (B31) yield

$$\|v_{s+1} - w_*\|^2 - \|v_s - w_*\|^2 \leq -2\gamma \langle v_s - w_*, e_s \rangle + 2\gamma \kappa L \|e_s\|. \tag{B32}$$

Summing over s then yields the result. □

Proof of Lemma 6 (Bounds on the empirical Rademacher complexities) The first statement of Lemma 6 is a classical result, see e.g. [49].

For the second statement, recall that for any $w \in \mathcal{H}$, we have $\|w\| = \sup_{\|v\|=1} \langle v, w \rangle$, since \mathcal{H} is assumed to be real. Thus, we may write

$$\begin{aligned} \widehat{\mathcal{R}}_\varepsilon(\mathcal{G}_R) &= \mathbb{E}_\varepsilon \left[\sup_{\nabla \ell \circ f \in \mathcal{G}_R} \left\| \frac{1}{n} \sum_{j=1}^n \varepsilon_j \ell'(Y_j, f(X_j)) X_j \right\| \right] \\ &= \mathbb{E}_\varepsilon \left[\sup_{f \in \mathcal{F}_R} \sup_{\|v\|=1} \frac{1}{n} \sum_{j=1}^n \varepsilon_j \ell'(Y_j, f(X_j)) \langle X_j, v \rangle \right]. \end{aligned} \tag{B33}$$

In order to bound the right-hand side in Eq. (B33), we apply Theorem 2 from [46], which states that for functions $\psi_i : \mathcal{S} \rightarrow \mathbb{R}$, $\phi_i : \mathcal{S} \rightarrow \ell^2$, $i = 1, \dots, n$, from a countable set \mathcal{S} ,

$$\mathbb{E}_\varepsilon \left[\sup_{s \in \mathcal{S}} \sum_{j=1}^n \varepsilon_j \psi_j(s) \right] \leq \sqrt{2} \mathbb{E}_\varepsilon \left[\sup_{s \in \mathcal{S}} \sum_{j=1}^n \sum_{k=1}^\infty \varepsilon_{jk} \phi_j^{(k)}(s) \right] \tag{B34}$$

whenever it is satisfied that

$$|\psi_i(s') - \psi_i(s)| \leq \|\phi_i(s') - \phi_i(s)\|_{\ell^2} \quad \text{for all } s, s' \in \mathcal{S}. \tag{B35}$$

Here, the $(\epsilon_{j,k})$ are i.i.d. copies of the (ϵ_j) .

Adopting the notation from this result, we may restrict the supremum in Eq. (B33) to a countable dense subset \mathcal{S} of $\mathcal{F}_R \times \{v \in \mathcal{H} : \|v\| \leq 1\}$. Note that this is possible, since by **(Smooth)**, we have that ℓ' is continuous in the second argument. Further, set

$$\begin{aligned} \psi_j &: \mathcal{S} \rightarrow \mathbb{R}, & \psi_j(f, v) &:= \ell'(Y_j, f(X_j))\langle X_j, v \rangle; \\ \phi_j^{(1)} &: \mathcal{S} \rightarrow \mathbb{R}, & \phi_j^{(1)}(f, v) &:= L\langle X_j, v \rangle; \\ \phi_j^{(2)} &: \mathcal{S} \rightarrow \mathbb{R}, & \phi_j^{(2)}(f, v) &:= \kappa \ell'(Y_j, f(X_j)); \\ \phi_j &: \mathcal{S} \rightarrow \mathbb{R}^2, & \phi_j(f, v) &= (\phi_j^{(1)}(f, v), \phi_j^{(2)}(f, v)). \end{aligned} \tag{B36}$$

Then, for any $j = 1, \dots, n$, and $(f, v), (g, w) \in \mathcal{S}$, we use that $\|\ell'\|_\infty \leq L$ by **(Lip)**, $\|X_j\| \leq \kappa$ by **(Bound)** and $\|w\| \leq 1$ to obtain

$$\begin{aligned} |\psi_j(f, v) - \psi_j(g, w)| &= |\ell'(Y_j, f(X_j))\langle X_j, v \rangle - \ell'(Y_j, g(X_j))\langle X_j, w \rangle| \\ &\leq |\ell'(Y_j, f(X_j))\langle X_j, v - w \rangle| \\ &\quad + |(\ell'(Y_j, f(X_j)) - \ell'(Y_j, g(X_j)))\langle X_j, w \rangle| \\ &\leq |L\langle X_j, v \rangle - L\langle X_j, w \rangle| + |\kappa \ell'(Y_j, f(X_j)) - \kappa \ell'(Y_j, g(X_j))| \\ &= \|\phi_j(f, v) - \phi_j(g, w)\|_{1, \mathbb{R}^2} \\ &\leq 2\|\phi_j(f, v) - \phi_j(g, w)\|_{2, \mathbb{R}^2}, \end{aligned} \tag{B37}$$

where $\|\cdot\|_{p, \mathbb{R}^2}$ denotes the p -norm on \mathbb{R}^2 . Equation (B37) shows that Theorem 2 from [46] is in fact applicable, which yields

$$\begin{aligned} \widehat{\mathcal{R}}(\mathcal{G}_R) &\leq \underbrace{2\sqrt{2}\mathbb{E}_\epsilon \sup_{f \in \mathcal{F}_R} \sup_{\|v\|=1} \frac{1}{n} \sum_{j=1}^n \epsilon_j \phi_j^{(1)}(f, v)}_{=: (I)} \\ &\quad + \underbrace{2\sqrt{2}\mathbb{E}_\epsilon \sup_{f \in \mathcal{F}_R} \sup_{\|v\|=1} \frac{1}{n} \sum_{j=1}^n \epsilon_j \phi_j^{(2)}(f, v)}_{=: (II)}. \end{aligned} \tag{B38}$$

We proceed by bounding each term individually. By Jensen’s inequality,

$$\begin{aligned}
 (I) &= 2\sqrt{2}L\mathbb{E}_\varepsilon \left[\sup_{f \in \mathcal{F}_R} \sup_{\|v\|=1} \frac{1}{n} \sum_{j=1}^n \varepsilon_j \langle X_j, v \rangle \right] = 2\sqrt{2}L\mathbb{E}_\varepsilon \left\| \frac{1}{n} \sum_{j=1}^n \varepsilon_j X_j \right\| \\
 &\leq \frac{2\sqrt{2}L}{n} \sqrt{\sum_{j=1}^n \|X_j\|^2} \leq \frac{2\sqrt{2}\kappa L}{\sqrt{n}},
 \end{aligned}
 \tag{B39}$$

where we have used again that $\|X_j\| \leq \kappa$.

For the second summand, Talagrand’s contraction principle, see e.g. Exercise 6.7.7 in [50], together with the fact that by (Smooth), ℓ' is M -Lipschitz yields the bound

$$(II) \leq 2\sqrt{2}\kappa M \widehat{\mathcal{R}}(\mathcal{F}) \leq \frac{2\sqrt{2}\kappa^2 MR}{\sqrt{n}},
 \tag{B40}$$

due to the first part of this Lemma. Together, Eqs. (B39) and (B40) yield the result. \square

Proof of Proposition 5 (Gradient concentration) For $(x, y) \in \mathcal{H} \times \mathcal{Y}$, $f \in \mathcal{F}_R$ denote

$$g_f(x, y) = (\nabla \ell \circ f)(x, y) = \ell'(y, f(x))x.
 \tag{B41}$$

Then $g_f \in \mathcal{G}_R$. For any $x_j, x'_j \in \mathcal{H}$ and $y_j, y'_j \in \mathbb{R}$, $j = 1, \dots, n$, we have

$$\begin{aligned}
 &\left| \sup_{g \in \mathcal{G}_R} \left\| \mathbb{E}[g(X, Y)] - \frac{1}{n} \sum_{j=1}^n g(x'_j, y'_j) \right\| - \sup_{g \in \mathcal{G}_R} \left\| \mathbb{E}[g(X, Y)] - \frac{1}{n} \sum_{j=1}^n g(x_j, y_j) \right\| \right| \\
 &\leq \sup_{g \in \mathcal{G}_R} \left\| \frac{1}{n} \sum_{j=1}^n g(x'_j, y'_j) - \frac{1}{n} \sum_{j=1}^n g(x_j, y_j) \right\| \leq 2G_R.
 \end{aligned}
 \tag{B42}$$

Therefore, McDiarmid’s bounded difference inequality, see e.g. Corollary 2.21 in [51], yields that on an event with probability at least $1 - \delta$,

$$\begin{aligned}
 &\sup_{f \in \mathcal{F}_R} \|\nabla \mathcal{L}(f) - \nabla \widehat{\mathcal{L}}(f)\| = \sup_{g \in \mathcal{G}_R} \left\| \mathbb{E}[g_f(X, Y)] - \frac{1}{n} \sum_{j=1}^n g_f(X_j, Y_j) \right\| \\
 &\leq \mathbb{E} \left[\sup_{g \in \mathcal{G}_R} \left\| \mathbb{E}[g_f(X, Y)] - \frac{1}{n} \sum_{j=1}^n g_f(X_j, Y_j) \right\| \right] + G_R \sqrt{\frac{2 \log(2/\delta)}{n}}.
 \end{aligned}
 \tag{B43}$$

The expectation above can now be bounded by the result in Lemma 4 from [34], which states that

$$\mathbb{E} \left[\sup_{g \in \mathcal{G}_R} \left\| \mathbb{E}[g_f(X, Y)] - \frac{1}{n} \sum_{j=1}^n g_f(X_j, Y_j) \right\| \right] \leq 4\widehat{\mathcal{R}}_n(\mathcal{G}_R) + 4G_R \frac{\log(2/\delta)}{n}
 \tag{B44}$$

on an event with probability at least $1 - \delta$. A union bound finally yields

$$\sup_{f \in \mathcal{F}_R} \|\nabla \mathcal{L}(f) - \nabla \widehat{\mathcal{L}}(f)\| \leq 4\widehat{\mathcal{R}}_n(\mathcal{G}_R) + G_R \sqrt{\frac{2 \log(4/\delta)}{n}} + G_R \frac{4 \log(4/\delta)}{n} \tag{B45}$$

on an event of at least probability $1 - \delta$. □

Proof of Proposition 7 (Bounded gradient path) Firstly, note that

$$\|v_t\| \leq \|v_t - w_*\| + \|w_*\| \leq \|v_t - w_*\| + \frac{R}{3}. \tag{B46}$$

Therefore, it suffices to prove that $\|v_t - w_*\| \leq 2R/3$ on the gradient concentration event. We proceed via induction over $t \leq T$. For $t = 0$, this is trivially satisfied, since $\|v_1 - w_*\| = \|w_*\| \leq R/3$. Now, assume that the result is true for $s = 0, \dots, t < T$. From Lemma 4, we have

$$\begin{aligned} & \|v_{t+1} - w_*\|^2 \\ & \leq \|w_*\|^2 + 2\gamma \sum_{s=0}^t \left(\langle \nabla \mathcal{L}(v_s) - \nabla \widehat{\mathcal{L}}(v_s), v_s - w_* \rangle + \kappa L \|\nabla \mathcal{L}(v_s) - \nabla \widehat{\mathcal{L}}(v_s)\| \right) \\ & \leq \frac{R^2}{9} + 2\gamma \sum_{s=0}^t (\|v_s - w_*\| + \kappa L) \|\nabla \mathcal{L}(v_s) - \nabla \widehat{\mathcal{L}}(v_s)\| \\ & \leq \frac{R^2}{9} + 2\gamma T \left(\frac{2R}{3} + \kappa L \right) \sup_{v \in \mathcal{F}_R} \|\nabla \mathcal{L}(v) - \nabla \widehat{\mathcal{L}}(v)\|, \end{aligned} \tag{B47}$$

where \mathcal{F}_R is defined in Eq. (22).

On the gradient concentration event from Proposition 5, by Lemma 6 (ii), we have the bound

$$\sup_{v \in \mathcal{F}_R} \|\nabla \mathcal{L}(v) - \nabla \widehat{\mathcal{L}}(v)\| \leq \frac{8\sqrt{2}(\kappa L + \kappa^2 MR)}{\sqrt{n}} + G_R \sqrt{\frac{2 \log(4/\delta)}{n}} + G_R \frac{4 \log(4/\delta)}{n}, \tag{B48}$$

where $G_R \leq \kappa L$ by Eq. (24). Equation (25) guarantees that

$$\frac{\log(4/\delta)}{n} \leq \sqrt{\frac{\log(4/\delta)}{n}} \tag{B49}$$

and hence with $1 \leq \sqrt{\log(4/\delta)}$, on the gradient concentration event, we obtain that

$$\sup_{v \in \mathcal{F}_R} \|\nabla \mathcal{L}(v) - \nabla \widehat{\mathcal{L}}(v)\| \leq 20\kappa^2 R(L + M) \sqrt{\frac{\log(4/\delta)}{n}}. \tag{B50}$$

Plugging the last bound into Eq. (B47) yields

$$\|v_t - w_*\|^2 \leq \frac{R^2}{9} + 40\gamma T \kappa^2 \left(\frac{2R^2}{3} + \kappa RL \right) (M + L) \sqrt{\frac{\log(4/\delta)}{n}}. \quad (\text{B51})$$

Hence, we obtain our result when the second term above is smaller than $4R^2/9$, which is satisfied when

$$\sqrt{n} \geq 90\gamma T \kappa^2 (1 + \kappa L)(M + L) \sqrt{\log(4/\delta)}, \quad (\text{B52})$$

where we have used the fact that $R \geq 1$. This completes the proof. \square

Acknowledgements The authors would like to thank Silvia Villa, Francesco Orabona, Markus Reiß and Martin Wahl for useful discussions. The authors are also grateful for the comments of the two referees that helped to improve the manuscript.

Funding Open Access funding enabled and organized by Projekt DEAL. The research of B.S. has been partially funded by the Deutsche Forschungsgemeinschaft (DFG)- Project-ID 318763901 - SFB1294. N.M. acknowledges funding by the Deutsche Forschungsgemeinschaft (DFG) under Excellence Strategy *The Berlin Mathematics Research Center MATH+* (EXC-2046/1, project ID:390685689). L.R. acknowledges support from the Center for Brains, Minds and Machines (CBMM), funded by NSF STC award CCF-1231216. L.R. also acknowledges the financial support of the European Research Council (grant SLING 819789), the AFOSR projects FA9550-18-1-7009, FA9550-17-1-0390 and BAA-AFRL-AFOSR-2016-0007 (European Office of Aerospace Research and Development), and the EU H2020-MSCA-RISE project NoMADS - DLV-777826.

Data availability statement The source code of the simulation study is available from the authors upon request.

Declarations

Ethics approval (Not applicable).

Consent to participate (Not applicable).

Consent for publication (Not applicable).

Conflict of interest The authors have no relevant financial or non-financial interests to disclose.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

1. Sra, S., Nowozin, S., Wright, S.J.: Optimization for Machine Learning. Neural Information Processing. The MIT Press, Cambridge (2011)
2. Shalev-Shwartz, S., Ben-David, S.: Understanding Machine Learning: From Theory to Algorithms. Cambridge University Press, New York (2014)
3. Gunasekar, S., Lee, J., Soudry, D., Srebro, N.: Characterizing Implicit Bias in Terms of Optimization Geometry. In: International Conference on Machine Learning, pp. 1832–1841. PMLR (2018)
4. Neyshabur, B.: Implicit Regularization in Deep Learning. [arXiv:1709.01953](https://arxiv.org/abs/1709.01953) [stat.ML] (2017)
5. Rosasco, L., Villa, S.: Learning with incremental iterative regularization. In: Advances in Neural Information Processing Systems, vol. 28. Curran Associates, Inc., Redhook, NY (2015)
6. Yang, F., Wei, Y., Wainwright, M.J.: Early stopping for kernel boosting algorithms: a general analysis with localized complexities. *IEEE Trans. Inf. Theory* **65**(10), 6685–6703 (2019)
7. Blanchard, G., Hoffmann, M., Reiß, M.: Early stopping for statistical inverse problems via truncated SVD estimation. *Electron. J. Stat.* **12**(2), 3204–3231 (2018)
8. Celisse, A., Wahl, M.: Analyzing the discrepancy principle for kernelized spectral filter learning algorithms. *J. Mach. Learn. Res.* **22**(76), 1–59 (2021)
9. Landweber, L.: An iteration formula for fredholm integral equations of the first kind. *Am J Math* **73**, 615–624 (1951)
10. Engl, H., Hanke, M., Neubauer, A.: Regularization of Inverse Problems. Mathematics and its Applications, vol. 375. Kluwer Academic Publishers, Dordrecht (1996)
11. Yao, Y., Caponetto, A., Rosasco, L.: On early stopping in gradient descent learning. *Constr. Approx.* **26**, 289–315 (2007)
12. Raskutti, G., Yu, B., Wainwright, M.J.: Early stopping and non-parametric regression: an optimal data-dependent stopping rule. *J. Mach. Learn. Res.* **15**, 335–366 (2014)
13. Bauer, F., Pereverzev, S., Rosasco, L.: On regularization algorithms in learning theory. *J. Complex.* **23**(1), 52–72 (2007)
14. Blanchard, G., Mücke, N.: Optimal rates for regularization of statistical inverse learning problems. *Found. Comput. Math.* **18**, 971–1013 (2018)
15. Dieuleveut, A., Flammarion, N., Bach, F.: Harder, better, faster, stronger convergence rates for least-squares regression. *J. Mach. Learn. Res.* **18**(1), 3520–3570 (2017)
16. Dieuleveut, A., Bach, F.: Nonparametric stochastic approximation with large step-sizes. *Ann. Stat.* **44**(4), 1363–1399 (2016)
17. Mücke, N., Neu, G., Rosasco, L.: Beating SGD Saturation with Tail-averaging and Minibatching. In: Advances in Neural Information Processing Systems, vol. 32. Curran Associates, Inc., Redhook, NY (2019)
18. Blanchard, G., Krämer, N.: Convergence rates of kernel conjugate gradient for random design regression. *Anal. Appl.* **14**(06), 763–794 (2016)
19. Pagliana, N., Rosasco, L.: Implicit Regularization of Accelerated Methods in Hilbert Spaces. In: Advances in Neural Information Processing Systems, vol. 32. Curran Associates, Inc., Redhook, NY (2019)
20. Zhang, Y., Duchi, J., Wainwright, M.J.: Divide and conquer kernel ridge regression: a distributed algorithm with minimax optimal rates. *J. Mach. Learn. Res.* **16**(102), 3299–3340 (2015)
21. Mücke, N., Blanchard, G.: Parallelizing Spectrally Regularized Kernel Algorithms. *J. Mach. Learn. Res.* **19**(1), 1069–1097 (2018)
22. Richards, D., Rebeschini, P.: Graph-dependent implicit regularisation for distributed stochastic sub-gradient descent. *J. Mach. Learn. Res.* **21**(2020), 34–13444 (2020)
23. Vaškevičius, T., Kanade, V., Rebeschini, P.: The Statistical Complexity of Early Stopped Mirror Descent. [arXiv:2002.00189](https://arxiv.org/abs/2002.00189) [stat.ML] (2020)
24. Villa, S., Matet, S., Vu, B.C., Rosasco, L.: Don't relax: early stopping for convex regularization. [arXiv:1707.05422](https://arxiv.org/abs/1707.05422) [stat.ML] (2017)
25. Soudry, D., Hoffer, E., Nacson, M.S., Gunasekar, S., Srebro, N.: The implicit bias of gradient descent on separable data. *J. Mach. Learn. Res.* **19**(1), 2822–2878 (2018)
26. Ji, Z., Telgarsky, M.: The implicit bias of gradient descent on nonseparable data. In: Conference on learning theory, vol. 99, pp. 1772–1798. PMLR, (2019)
27. Lin, J., Rosasco, L., Zhou, D.: Iterative regularization for learning with convex loss functions. *J. Mach. Learn. Res.* **17**(1), 2718–2755 (2016)

28. Lin, J., Camoriano, R., Rosasco, L.: Generalization properties and implicit regularization for multiple passes SGM. In: International Conference on Machine Learning, pp. 2340–2348. PMLR, (2016)
29. Lei, Y., Hu, T., Tang, K.: Generalization performance of multi-pass stochastic gradient descent with convex loss functions. *J. Mach. Learn. Res.* **22**(25), 1–41 (2021)
30. Bousquet, O., Elisseeff, A.: Stability and generalization. *J. Mach. Learn. Res.* **2**, 499–526 (2002)
31. Chen, Y., Jin, C., Yu, B.: Stability and convergence trade-off of iterative optimization algorithms. [arXiv:1804.01619](https://arxiv.org/abs/1804.01619) [stat.ML] (2018)
32. Bertsekas, D.P., Tsitsiklis, J.N.: Gradient convergence in gradient methods with errors. *SIAM J. Optim.* **10**(3), 627–642 (2000)
33. Schmidt, M., Roux, N., Bach, F.: Convergence rates of inexact proximal-gradient methods for convex optimization. In: Advances in Neural Information Processing Systems, vol. 24. Curran Associates, Inc., Redhook, NY (2011)
34. Foster, D.J., Sekhari, A., Sridharan, K.: Uniform convergence of gradients for non-convex learning and optimization. In: Advances in Neural Information Processing Systems, vol. 31. Curran Associates, Inc., Redhook, NY (2018)
35. Gorbunov, E., Danilova, M., Gasnikov, A.: Stochastic optimization with heavy-tailed noise via accelerated gradient clipping. In: Advances in Neural Information Processing Systems. Curran Associates, Inc., Redhook, NY (2020)
36. Lin, J., Rosasco, L.: Optimal rates for multi-pass stochastic gradient methods. *J. Mach. Learn. Res.* **18**(1), 3375–3421 (2017)
37. Bubeck, S.: Convex optimization: algorithms and complexity. *Found. Trends Mach. Learn.* **8**(3–4), 231–357 (2015)
38. Steinwart, I., Christmann, A.: Support Vector Machines. Information Science and Statistics. Springer, New York (2008)
39. Bottou, L., Bousquet, O.: The tradeoffs of large scale learning. In: Optimization for Machine Learning, pp. 351–368. MIT Press (2011)
40. Boucheron, S., Lugosi, G., Massart, P.: Concentration Inequalities: A Non-asymptotic Theory of Independence. Oxford University Press, Oxford (2013)
41. Balakrishnan, S., Wainwright, M.J., Yu, B.: Statistical guarantees for the EM algorithm: from population to sample-based analysis. *Ann. Stat.* **45**(1), 77–120 (2017)
42. Holland, M.J., Ikeda, K.: Efficient learning with robust gradient descent. [arXiv:1706.00182](https://arxiv.org/abs/1706.00182) [stat.ML] (2018)
43. Prasad, A., Suggala, A.S., Balakrishnan, S., Ravikumar, P.: Robust estimation via robust gradient estimation. [arXiv:1802.06485](https://arxiv.org/abs/1802.06485) [stat.ML] (2018)
44. Lin, J., Rudi, A., Rosasco, L., Cevher, V.: Optimal rates for spectral algorithms with least-squares regression over Hilbert spaces. *Appl. Comput. Harmonic Anal.* **48**(3), 868–890 (2020)
45. Harvey, N., Liaw, C., Plan, Y., Randhawa, S.: Tight analyses for non-smooth stochastic gradient descent. In: Conference on Learning Theory, pp. 1579–1613. PMLR (2019)
46. Maurer, A.: A Vector-contraction Inequality for Rademacher Complexities. In: Algorithmic Learning Theory, vol. 9925, pp. 3–17. Springer, (2016)
47. Lei, Y., Tang, K.: Stochastic composite mirror descent: optimal bounds with high probability. In: Advances in Neural Information Processing Systems. Curran Associates, Inc., Redhook, NY (2018)
48. Derumigny, A., Schmidt-Hieber, J.: On lower bounds for the bias-variance trade-off. [arXiv:2006.00278](https://arxiv.org/abs/2006.00278) [stat.ML] (2020)
49. Bartlett, P.L., Mendelson, S.: Rademacher and gaussian complexities: risk bounds and structural results. *J. Mach. Learn. Res.* **3**, 463–482 (2002)
50. Vershynin, R.: High-dimensional Probability. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press, Cambridge (2018)
51. Wainwright, M.J.: High-dimensional Statistics: A Non-asymptotic Viewpoint. Cambridge University Press, Cambridge (2019)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.