



HAL
open science

Beyond Least-Squares: Fast Rates for Regularized Empirical Risk Minimization through Self-Concordance

Ulysse Marteau-Ferey, Dmitrii M. Ostrovskii, Francis Bach, Alessandro Rudi

► **To cite this version:**

Ulysse Marteau-Ferey, Dmitrii M. Ostrovskii, Francis Bach, Alessandro Rudi. Beyond Least-Squares: Fast Rates for Regularized Empirical Risk Minimization through Self-Concordance. 2019. hal-02011895v1

HAL Id: hal-02011895

<https://inria.hal.science/hal-02011895v1>

Preprint submitted on 8 Feb 2019 (v1), last revised 17 Jun 2019 (v3)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Beyond Least-Squares: Fast Rates for Regularized Empirical Risk Minimization through Self-Concordance

Ulysse Marteau-Ferey Dmitrii Ostrovskii Francis Bach Alessandro Rudi

INRIA - Département d'Informatique de l'École Normale Supérieure
PSL Research University
Paris, France

Abstract

We consider learning methods based on the regularization of a convex empirical risk by a squared Hilbertian norm, a setting that includes linear predictors and non-linear predictors through positive-definite kernels. In order to go beyond the generic analysis leading to convergence rates of the excess risk as $O(1/\sqrt{n})$ from n observations, we assume that the individual losses are self-concordant, that is, their third-order derivatives are bounded by their second-order derivatives. This setting includes least-squares, as well as all generalized linear models such as logistic and softmax regression. For this class of losses, we provide a bias-variance decomposition and show that the assumptions commonly made in least-squares regression, such as the source and capacity conditions, can be adapted to obtain fast non-asymptotic rates of convergence by improving the bias terms, the variance terms or both.

1 Introduction

Regularized empirical risk minimization remains a cornerstone of statistics and supervised learning, from the early days of linear regression (Hoerl and Kennard, 1976) and neural networks (Geman et al., 1992), then to spline smoothing (Wahba, 1990) and more generally kernel-based methods (Shawe-Taylor and Cristianini, 2004). While the regularization by the squared Euclidean norm is applied very widely, the statistical analysis of the resulting learning methods is still not complete.

The main goal of this paper is to provide a sharp non-asymptotic analysis of regularized empirical risk minimization (ERM), or more generally regularized M -estimation, that is estimators obtained as the unique solution of

$$\min_{\theta \in \mathcal{H}} \frac{1}{n} \sum_{i=1}^n \ell_{z_i}(\theta) + \frac{\lambda}{2} \|\theta\|^2, \quad (1)$$

where \mathcal{H} is a Hilbert space (possibly infinite-dimensional) and $\ell_z(\theta)$ is the convex loss associated with an observation z and the estimator $\theta \in \mathcal{H}$. We assume that the observations $z_i, i = 1, \dots, n$ are independent and identically distributed, and that the minimum of the associated unregularized expected risk $L(\theta)$ is attained at a certain $\theta^* \in \mathcal{H}$.

In this paper, we focus on dimension-independent results (thus ultimately extending the analysis in the finite-dimensional setting from Ostrovskii and Bach, 2018). For this class of problems, two main classes of problems have been studied, depending on the regularity assumptions on the loss.

Convex Lipschitz-continuous losses (with respect to the parameter θ), such as for logistic regression or the support vector machine, lead to general *non-asymptotic* bounds for the excess risk of the form (Sridharan et al., 2009):

$$\frac{B^2}{\lambda n} + \lambda \|\theta^*\|^2, \quad (2)$$

where B is a uniform upper bound on the Lipschitz constant for all losses $\theta \mapsto \ell_z(\theta)$. The bound above already has a form that takes into account two separate terms: a *variance term* $B^2/(\lambda n)$ which depends on the sample size n but not on the optimal predictor θ^* , and a *bias term* $\lambda \|\theta^*\|^2$ which depends on the optimal predictor but not on the sample size n . All our bounds will have this form but with smaller quantities (but asking for more assumptions). Without further assumptions, in Eq. (2), λ is taken proportional to $1/\sqrt{n}$, and we get the usual optimal slow rate in excess risk of $O(1/\sqrt{n})$ associated with such a general set-up (see, e.g., Cesa-Bianchi et al., 2015).

For the specific case of quadratic losses of the form $\ell_z(\theta) = \frac{1}{2}(y - \theta \cdot \Phi(x))^2$, where $z = (x, y)$, and $y \in \mathbb{R}$ and $\Phi(x) \in \mathcal{H}$, the situation is much richer. Without further assumptions, the same rate $O(1/\sqrt{n})$ is achieved, but stronger assumptions lead to faster rates (Caponnetto and De Vito, 2007). In particular, the decay of the eigenvalues of the Hessian $\mathbb{E}[\Phi(x) \otimes \Phi(x)]$ (often called the *capacity condition*) leads to an improved variance term, while the finiteness of some bounds on θ^* for norms other than the plain Hilbertian norms $\|\theta^*\|$ (often called the *source condition*) leads to an improved bias term. Both of these assumptions lead to faster rates than $O(1/\sqrt{n})$ for the excess risk, with the proper choice of the regularization parameter λ . For least-squares, these rates are then optimal and provide a better understanding of properties of the problem that influence the generalization capabilities of regularized ERM (see, e.g. Smale and Zhou, 2007; Caponnetto and De Vito, 2007; Steinwart et al., 2009; Fischer and Steinwart, 2017; Blanchard and Mücke, 2018).

Our main goal in this paper is to bridge the gap between Lipschitz-continuous and quadratic losses by improving on slow rates for general classes of losses beyond least-squares. We first note that: (a) there has to be an extra regularity assumption because of lower bounds (Cesa-Bianchi et al., 2015), and (b) asymptotically, we should obtain bounds that approach the local quadratic approximation of $\ell_z(\theta)$ around θ^* with the same optimal behavior as for plain least-squares.

Several frameworks are available for such an extension with extra assumptions on the losses, such as “exp-concavity” (Koren and Levy, 2015; Mehta, 2016), strong convexity (Van de Geer, 2008) or a generalized notion of self-concordance (Bach, 2010; Ostrovskii and Bach, 2018). In this paper, we focus on self-concordance, which links the second and third order derivatives of the loss. This notion is quite general and corresponds to widely used losses in machine learning, and does not suffer from constants which can be exponential in problem parameters (e.g., $\|\theta^*\|$) when applied to generalized linear models like logistic regression. See Sec. 1.1 for a comparison to related work.

With this self-concordance assumption, we will show that our problem behaves like a quadratic problem corresponding to the local approximation around θ^* , in a totally non-asymptotic way, which is the core technical contribution of this paper. As we have already mentioned, this phenomenon is naturally expected in the asymptotic regime, but is hard to capture in the non-asymptotic setting without constants which explode exponentially with the problem parameters.

The paper is organized as follows: in Sec. 2, we present our main assumptions and informal results, as well as our bias-variance decomposition. In order to introduce precise results gradually, we start in Sec. 3 with a result similar to Eq. (2) for our set-up to show that we recover with a simple argument the result from Sridharan et al. (2009), which itself applies more generally. Then, in Sec. 4 we introduce the source condition allowing for a better control of the bias. Finally, in Sec. 5, we detail the capacity condition leading to an improved variance term, which, together with the improved bias leads to fast rates (which are optimal for least-squares).

1.1 Related work

Fast rates for empirical risk minimization. Rates faster than $O(1/\sqrt{n})$ can be obtained with a variety of added assumptions, such as some form of strong convexity (Sridharan et al., 2009; Boucheron and Massart, 2011), noise conditions for classification (Steinwart and Scovel, 2007), or extra conditions on the loss, such as self-concordance (Bach, 2010) or exp-concavity (Koren and Levy, 2015; Mehta, 2016), whose partial goal is to avoid exponential constants. Note that Bach (2010) already considers logistic regression with Hilbert spaces, but only for well-specified models and a fixed design, and without the sharp and simpler results that we obtain in this paper.

Avoiding exponential constants for logistic regression. The problem of exponential constants (i.e., leading factors in the rates scaling as e^{RD} where D is the radius of the optimal predictor, and R the radius of the design) is long known. In fact, Hazan et al. (2014) showed a lower bound, explicitly constructing an adversarial distribution (i.e., an ill-specified model) for which the problem manifests in the finite-sample regime with $n = O(e^{RD})$. Various attempts to address this problem are found in the literature. For example, Ostrovskii and Bach (2018, App. C) prove the optimal d/n rate in the non-regularized d -dimensional setting but, multiplied with the curvature parameter ρ which is at worst exponential but is shown to grow at most as $(RD)^{3/2}$ in the case of Gaussian design. Another approach is due to Foster et al. (2018): they establish “1-mixability” of the logistic loss, then apply Vovk’s aggregating algorithm in the online setting, and then proceed via online-to-batch conversion. While this result allows to obtain the fast $O(d/n)$ rate (and its counterparts in the nonparametric setting) without exponential constants, the resulting algorithm is *improper* (i.e., the canonical parameter $\eta = \Phi(x) \cdot \theta^*$, see below, is estimated by a *non-linear* functional of $\Phi(x)$).

A closely related approach is to use the notion of exp-concavity instead of mixability (Rakhlin and Sridharan, 2015; Koren and Levy, 2015; Mehta, 2016). The two close notions are summarized in the so-called central condition (due to Van Erven et al. (2015)) which fully characterizes when the fast $O(d/n)$ rates (up to log factors and in high probability) are available for improper algorithms. However, when proper learning algorithms are concerned, this analysis requires η -mixability (or η -exp-concavity) of the *overall* loss $\ell_z(\theta)$ for which the η parameter scales with the radius of the set of predictors. This scaling is exponential for the logistic loss, leading to exponential constants.

2 Main Assumptions and Results

Let \mathcal{Z} be a Polish space and Z be a random variable on \mathcal{Z} with distribution ρ . Let \mathcal{H} be a separable (non-necessarily finite-dimensional) Hilbert space, with norm $\|\cdot\|$, and let $\ell : \mathcal{Z} \times \mathcal{H} \rightarrow \mathbb{R}$ be a loss function, we denote by $\ell_z(\cdot)$ the function $\ell(z, \cdot)$. Our goal is to minimize the expected risk with respect to $\theta \in \mathcal{H}$:

$$\inf_{\theta \in \mathcal{H}} L(\theta) = \mathbb{E}[\ell_Z(\theta)].$$

Given $(z_i)_{i=1}^n \in \mathcal{Z}^n$, we will consider the following estimator based on regularized empirical risk minimization given $\lambda > 0$ (note that the minimizer is unique in this case):

$$\hat{\theta}_\lambda^* = \arg \min_{\theta \in \mathcal{H}} \frac{1}{n} \sum_{i=1}^n \ell_{z_i}(\theta) + \frac{\lambda}{2} \|\theta\|^2,$$

where we assume the following.

Assumption 1 (i.i.d. data). *The samples $(z_i)_{1 \leq i \leq n}$ are independently and identically distributed according to ρ .*

The goal of this work is to provide upper bounds in high probability for the so-called *excess risk*

$$L(\hat{\theta}_\lambda^*) - \inf_{\theta \in \mathcal{H}} L(\theta),$$

and thus to provide a general framework to measure the quality of the estimator $\hat{\theta}_\lambda^*$. Algorithms for obtaining such estimators have been extensively studied, in both finite-dimensional regimes, where a direct optimization over θ is performed, typically by gradient descent or stochastic versions thereof (see, e.g., [Bottou and Bousquet, 2008](#); [Shalev-Shwartz et al., 2011](#)) and infinite-dimensional regimes, where kernel-based methods are traditionally used (see, e.g., [Keerthi et al., 2005](#); [Gerfo et al., 2008](#); [Dieuleveut and Bach, 2016](#); [Tu et al., 2016](#); [Rudi et al., 2017](#), and references therein).

Example 1 (Supervised learning). *Although formulated as a general M -estimation problem (see, e.g., [Lehmann and Casella, 2006](#)), our main motivation comes from supervised learning, with $\mathcal{Z} = \mathcal{X} \times \mathcal{Y}$ where \mathcal{X} is the data space and \mathcal{Y} the target space. We will consider, as examples, losses with both real-valued outputs but also the multivariate case. For learning real-valued outputs, consider we have a bounded representation of the input space $\Phi : \mathcal{X} \rightarrow \mathcal{H}$ (potentially implicit when using kernel-based methods, [Aronszajn, 1950](#)). We will provide bounds for the following losses.*

- The square loss $\ell_z(\theta) = \frac{1}{2} (y - \theta \cdot \Phi(x))^2$, which is not Lipschitz-continuous.
- The Huber losses $\ell_z(\theta) = \psi(y - \theta \cdot \Phi(x))$ where $\psi(t) = \sqrt{1 + t^2} - 1$ or $\psi(t) = \log \frac{e^t + e^{-t}}{2}$ ([Hampel et al., 2011](#)), which are Lipschitz-continuous.
- The logistic loss $\ell_z(\theta) = \log(1 + e^{-y\theta \cdot \Phi(x)})$ commonly used in binary classification where $y \in \{-1, 1\}$, which is Lipschitz-continuous.

Our framework goes beyond real-valued outputs, and can be applied to all generalized linear models (GLM) ([McCullagh and Nelder, 1989](#)), including softmax regression: we consider a representation function $\Phi : \mathcal{X} \times \mathcal{Y} \rightarrow \mathcal{H}$ and an a priori measure μ on \mathcal{Y} . The loss we consider in this case is

$$\ell_z(\theta) = -\theta \cdot \Phi(x, y) + \log \int_{\mathcal{Y}} \exp(\theta \cdot \Phi(x, y')) d\mu(y'),$$

which corresponds to the negative conditional log-likelihood when modelling y given x by the distribution $p(y|x, \theta) \sim \frac{\exp(\theta \cdot \Phi(x, y))}{\int_{\mathcal{Y}} \exp(\theta \cdot \Phi(x, y')) d\mu(y')}$. Our framework applies to all of these generalized linear models with almost surely bounded features $\Phi(x, y)$, such as conditional random fields ([Lafferty et al., 2001](#)).

We can now introduce the main technical assumption on the loss ℓ .

Assumption 2 (Generalized self-concordance). *For any $z \in \mathcal{Z}$, the function $\ell_z(\cdot)$ is convex and three times differentiable. Moreover, there exists a set $\varphi(z) \subset \mathcal{H}$ such that it holds:*

$$\forall \theta \in \mathcal{H}, \forall h, k \in \mathcal{H}, |\nabla^3 \ell_z(\theta)[k, h, h]| \leq \sup_{g \in \varphi(z)} |k \cdot g| \nabla^2 \ell_z(\theta)[h, h].$$

This is a generalization of the assumptions introduced by [Bach \(2010\)](#), by allowing a varying term $\sup_{g \in \varphi(z)} |k \cdot g|$ instead of a uniform bound proportional to $\|k\|$. This is crucial for the fast rates we want to show.

Example 2 (Checking assumptions). *For the losses in Example 1, this condition is satisfied with the following corresponding set-function φ .*

- For the square loss $\ell_z(\theta) = \frac{1}{2}(y - \theta \cdot \Phi(x))^2$, $\varphi(z) = \{0\}$.
- For the Huber losses $\ell_z(\theta) = \psi(y - \theta \cdot \Phi(x))$, if $\psi(t) = \sqrt{1+t^2} - 1$, then $\varphi(z) = \{3\Phi(x)\}$ and if $\psi(t) = \log \frac{e^t + e^{-t}}{2}$, then $\varphi(z) = \{2\Phi(x)\}$ (Ostrovskii and Bach, 2018). For the logistic loss $\ell_z(\theta) = \log(1 + e^{-y\theta \cdot \Phi(x)})$, we have $\varphi(z) = \{y\Phi(x)\}$. In these situations, $\varphi(z)$ is reduced to a point.
- For generalized linear models, $\nabla^3 \ell_z(\theta)$ is a third-order cumulant, and thus $|\nabla^3 \ell_z(\theta)[k, h, h]| \leq \mathbb{E}_{p(y|x, \theta)} |k \cdot \Phi(x, y) - k \cdot \mathbb{E}_{p(y'|x, \theta)} \Phi(x, y')| \cdot |h \cdot \Phi(x, y) - h \cdot \mathbb{E}_{p(y'|x, \theta)} \Phi(x, y')|^2 \leq 2 \sup_{y \in \mathcal{Y}} |k \cdot \Phi(x, y)| \nabla^2 \ell_z(\theta)[h, h]$. Therefore $\varphi(z) = \{2\Phi(x, y'), y' \in \mathcal{Y}\}$ (which is not a singleton).

Moreover we require the following two technical assumptions to guarantee that $L(\theta)$ and its first and second derivatives are well defined for any $\theta \in \mathcal{H}$.

Assumption 3 (Boundedness). *There exists $R \geq 0$ such that $\sup_{g \in \varphi(z)} \|g\| \leq R$ almost surely.*

Assumption 4 (Definition in 0). *$|\ell_Z(0)|$, $\|\nabla \ell_Z(0)\|$ and $\text{Tr}(\nabla^2 \ell_Z(0))$ are almost surely bounded.*

The assumptions above are usually easy to check in practice. In particular, if the support of ρ is bounded, the mappings $z \mapsto \ell_z(0), \nabla \ell_z(0), \text{Tr}(\nabla^2 \ell_z(0))$ are continuous, and φ is uniformly bounded on bounded sets, then they hold. The main regularity assumption we make on our statistical problems follows.

Assumption 5 (Existence of a minimizer). *There exists $\theta^* \in \mathcal{H}$ such that $L(\theta^*) = \inf_{\theta \in \mathcal{H}} L(\theta)$.*

While Assumption 3 is standard in the analysis of such models (Caponnetto and De Vito, 2007; Sridharan et al., 2009; Steinwart et al., 2009; Bach, 2014), Assumption 5 imposes that the model is “well-specified”, that is, for supervised learning situations from Example 1, we have chosen a rich enough representation Φ . It is possible to study the non-realizable case in our setting by requiring additional technical assumptions (Steinwart et al., 2009), but this is out of scope of this paper. We can now introduce the main definitions allowing our bias-variance decomposition.

Definition 1 (Hessian, Bias, Degrees of freedom). *Let $L_\lambda(\theta) = L(\theta) + \frac{\lambda}{2}\|\theta\|^2$; define the expected Hessian $\mathbf{H}(\theta)$, the regularized Hessian $\mathbf{H}_\lambda(\theta)$, the bias Bias_λ and the degrees of freedom df_λ as:*

$$\mathbf{H}(\theta) = \mathbb{E} [\nabla^2 \ell_Z(\theta)], \text{ and } \mathbf{H}_\lambda(\theta) = \mathbf{H}(\theta) + \lambda I, \quad (3)$$

$$\text{Bias}_\lambda = \|\mathbf{H}_\lambda(\theta^*)^{-1/2} \nabla L_\lambda(\theta^*)\|, \quad (4)$$

$$\text{df}_\lambda = \mathbb{E} \left[\|\mathbf{H}_\lambda(\theta^*)^{-1/2} \nabla \ell_Z(\theta^*)\|^2 \right]. \quad (5)$$

Note that the bias and degrees of freedom only depend on the optimum $\theta^* \in \mathcal{H}$ and not on the minimizer θ_λ^* of the regularized expected risk. Moreover, the degrees of freedom df_λ correspond to the usual Fisher information term commonly seen in the asymptotic analysis of M -estimation (Van der Vaart, 2000; Lehmann and Casella, 2006), and correspond to the usual quantities introduced in the analysis of least-squares (Caponnetto and De Vito, 2007). Indeed, in the least-squares case, we recover exactly $\text{Bias}_\lambda = \lambda \|\mathbf{C}_\lambda^{-1/2} \theta^*\|$ and $\text{df}_\lambda = \text{Tr}(\mathbf{C} \mathbf{C}_\lambda^{-1})$, where \mathbf{C} is the covariance operator $\mathbf{C} = \mathbb{E} [\Phi(x) \otimes \Phi(x)]$ and $\mathbf{C}_\lambda = \mathbf{C} + \lambda I$.

Our results will rely on the quadratic approximation of the losses around θ^* . Borrowing tools from the analysis of Newton’s method (Nesterov and Nemirovskii, 1994), this will only be possible in the vicinity of θ^* . The proper notion of vicinity is the so-called *radius of the Dikin ellipsoid*, which we define as follows:

$$r_\lambda(\theta) \text{ such that } 1/r_\lambda(\theta) = \sup_{z \in \text{supp}(\rho)} \sup_{g \in \varphi(z)} \|\mathbf{H}_\lambda^{-1/2}(\theta)g\|. \quad (6)$$

Assumptions	Bias	Variance	Optimal λ	Optimal Rate	
None	λ	$\frac{1}{\lambda n}$	$n^{-1/2}$	$n^{-1/2}$	Thm. 2 and Cor. 1
Source	λ^{2r+1}	$\frac{1}{\lambda n}$	$n^{-\frac{1}{2r+2}}$	$n^{-\frac{2r+1}{2r+2}}$	Thm. 3 and Cor. 2
Source + Capacity	λ^{2r+1}	$\frac{1}{\lambda^{1/\alpha} n}$	$n^{-\frac{\alpha}{2r\alpha+\alpha+1}}$	$n^{-\frac{2r\alpha+\alpha}{2r\alpha+\alpha+1}}$	Thm. 4 and Cor. 3

Table 1: Summary of convergence rates, without constants except λ , for source condition (Asm. 6): $\theta^* \in \text{Im}(\mathbf{H}(\theta^*)^r)$, $r \in (0, 1/2]$, capacity condition (Asm. 7): $\text{df}_\lambda = O(\lambda^{-1/\alpha})$, $\alpha \geq 1$.

Our most refined bounds will depend whether the bias term is small enough compared to $r_\lambda(\theta^*)$. The following informal result summarizes all of our results.

Theorem 1 (General bound, informal). *Let $n \in \mathbb{N}$, $\delta \in (0, 1/2]$, $\lambda > 0$. Under Assumptions 1 to 5, whenever*

$$n \geq C_0 \frac{R^2 \text{df}_\lambda \log \frac{2}{\delta}}{\lambda},$$

then with probability at least $1 - 2\delta$, it holds

$$L(\hat{\theta}_\lambda^*) - L(\theta^*) \leq C_{\text{bias}} \text{Bias}_\lambda^2 + C_{\text{var}} \frac{\text{df}_\lambda \log \frac{2}{\delta}}{n},$$

where C_0, C_{bias} and C_{var} are either universal or depend only on $R\|\theta^*\|$.

This mimics a usual bias-variance decomposition, with a bias term Bias_λ^2 and a variance term proportional to df_λ/n . In particular in the rest of the paper we quantify the constants and the rates under various regularity assumptions, and specify the good choices of the regularization parameter λ . In Table 1, we summarize the different assumptions and corresponding rates.

3 Slow convergence rates

Here we bound the quantity of interest without any regularity assumption (e.g., source of capacity condition) beyond some boundedness assumptions on the learning problem. We consider the various bounds on the derivatives of the loss ℓ :

$$\mathbf{B}_1(\theta) = \sup_{z \in \text{supp}(\rho)} \|\nabla \ell_z(\theta)\|, \quad \mathbf{B}_2(\theta) = \sup_{z \in \text{supp}(\rho)} \text{Tr}(\nabla^2 \ell_z(\theta)), \quad \bar{\mathbf{B}}_1 = \sup_{\|\theta\| \leq \|\theta^*\|} \mathbf{B}_1(\theta), \quad \bar{\mathbf{B}}_2 = \sup_{\|\theta\| \leq \|\theta^*\|} \mathbf{B}_2(\theta).$$

Example 3 (Bounded derivatives). *In all the losses considered above, assume the feature representation ($\Phi(x)$ for the Huber losses and the square loss, $y\Phi(x)$ for the logistic loss, and $\Phi(x, y)$ for GLMs) is bounded by \bar{R} . Then the losses considered above apart from the square loss are Lipschitz-continuous and \mathbf{B}_1 is uniformly bounded by \bar{R} . For these losses, \mathbf{B}_2 is also uniformly bounded by \bar{R}^2 . Using Example 2, one can take \bar{R} to be equal to a constant times R (1/2 and 1/3 for the respective Huber losses, 1 for logistic regression and 1/2 for canonical GLMs). For the square loss (where $R = 0$ because the third-order derivative is zero), $\bar{\mathbf{B}}_2 \leq \bar{R}^2$ and $\bar{\mathbf{B}}_1 \leq \bar{R}\|y\|_\infty + \bar{R}^2\|\theta^*\|$, where $\|y\|_\infty$ is an almost sure bound on the output y .*

Theorem 2 (Basic result). *Let $n \in \mathbb{N}$ and $0 < \lambda \leq \bar{\mathbf{B}}_2$. Let $\delta \in (0, 1/2]$. If*

$$n \geq 512 (\|\theta^*\|^2 R^2 \vee 1) \log \frac{2}{\delta}, \quad n \geq 24 \frac{\bar{\mathbf{B}}_2}{\lambda} \log \frac{8\bar{\mathbf{B}}_2}{\lambda\delta}, \quad n \geq 256 \frac{R^2 \bar{\mathbf{B}}_1^2}{\lambda^2} \log \frac{2}{\delta},$$

then with probability at least $1 - 2\delta$,

$$L(\hat{\theta}_\lambda^*) - L(\theta^*) \leq 84 \frac{\bar{\mathbf{B}}_1^2}{\lambda n} \log \frac{2}{\delta} + 2\lambda\|\theta^*\|^2. \quad (7)$$

This result shown in Appendix C.3 as a consequence of Thm. 6 (see also the proof sketch in Sec. 6) matches the one obtained with Lipschitz-continuous losses (Sridharan et al., 2009) and the one for least-squares when only assuming the existence of θ^* (Caponnetto and De Vito, 2007). The following corollary (proved as Thm. 8 in Appendix E) gives the bound optimized in λ , with explicit rates.

Corollary 1 (Basic Rates). *Let $\delta \in (0, 1/2]$. Under Assumptions 1 to 5, when $n \geq N$, $\lambda = C_0 \sqrt{\log(2/\delta)}/n$, then with probability at least $1 - 2\delta$,*

$$L(\widehat{\theta}_\lambda^*) - L(\theta^*) \leq C_1 n^{-1/2} \log^{1/2} \frac{2}{\delta}.$$

with $C_0 = 16\overline{B}_1 \max(1, R)$, $C_1 = 48\overline{B}_1 \max(1, R) \max(1, \|\theta^*\|^2)$ and with N defined in Eq. (41) and satisfying $N = O(\text{poly}(\overline{B}_1, \overline{B}_2, R\|\theta^*\|))$ where poly denotes a polynomial function of the inputs.

Both bias and variance terms are of order $O(1/\sqrt{n})$ and we recover up to constants terms the result of Sridharan et al. (2009). In the next section, we will improve both bias and variance terms to obtain faster rates.

4 Faster Rates with Source Conditions

Here we provide a more refined bound, where we introduce a *source condition* on θ^* allowing to improve the bias term and to achieve learning rates as fast as $O(n^{-2/3})$. We first define the localized versions of $\overline{B}_1, \overline{B}_2$:

$$B_1^* = B_1(\theta^*), \quad B_2^* = B_2(\theta^*),$$

and recall the definition of the bias

$$\text{Bias}_\lambda = \|\mathbf{H}_\lambda(\theta^*)^{-1/2} \nabla L_\lambda(\theta^*)\|. \quad (8)$$

Note that since θ^* is the minimizer of L , we have $\nabla L(\theta^*) = 0$, so that $\nabla L_\lambda(\theta^*) = \nabla L(\theta^*) + \lambda\theta^* = \lambda\theta^*$, and $\text{Bias}_\lambda = \lambda\|\mathbf{H}_\lambda(\theta^*)^{-1/2}\theta^*\|$. This characterization is always bounded by $\lambda\|\theta^*\|^2$, but allows a finer control of the regularity of θ^* , leading to improved rates compared to Sec. 3, as shown below.

Note that in the least-squares case, we recover exactly the bias of ridge regression $\text{Bias}_\lambda = \lambda\|\mathbf{C}_\lambda^{-1/2}\theta^*\|$, where \mathbf{C} is the covariance operator $\mathbf{C} = \mathbb{E}[\Phi(x) \otimes \Phi(x)]$.

Using self-concordance, we will relate quantities at θ^* to quantities at θ_λ^* using:

$$t_\lambda = \sup_{z \in \text{supp}(\rho)} \sup_{g \in \varphi(z)} |(\theta_\lambda^* - \theta^*) \cdot g|.$$

The following theorem, proved in Appendix D.4, relates Bias_λ to the excess risk.

Theorem 3 (Decomposition with refined bias). *Let $n \in \mathbb{N}$, $\delta \in (0, 1/2]$, $0 < \lambda \leq B_2^*$. Whenever*

$$n \geq \Delta_1 \frac{B_2^*}{\lambda} \log \frac{8\Box_1^2 B_2^*}{\lambda\delta}, \quad n \geq \Delta_2 \frac{(B_1^* R)^2}{\lambda^2} \log \frac{2}{\delta},$$

then with probability at least $1 - 2\delta$, it holds

$$L(\widehat{\theta}_\lambda^*) - L(\theta^*) \leq C_{\text{bias}} \text{Bias}_\lambda^2 + C_{\text{var}} \frac{(B_1^*)^2}{\lambda n} \log \frac{2}{\delta}, \quad (9)$$

where $\Box_1 \leq e^{t_\lambda/2}$, $\Delta_1 \leq 2304e^{4t_\lambda} (1/2 \vee R\|\theta^*\|)$, $\Delta_2 \leq 256e^{2t_\lambda}$, $C_{\text{bias}} \leq 6e^{2t_\lambda}$, $C_{\text{var}} \leq 256e^{3t_\lambda}$.

It turns out that the *radius of the Dikin ellipsoid* $r_\lambda(\theta^*)$ defined in Eq. (6) provides the sufficient control over the constants above: when the bias is of the same order of the radius of the Dikin ellipsoid, the quantities $C_{\text{bias}}, C_{\text{var}}, \Delta_1, \Delta_2$ become universal constants instead of depending exponentially on $R\|\theta^*\|$, as shown by the lemma below, proved in Lemma 4 in Appendix D.

Lemma 1. *When $\text{Bias}_\lambda \leq \frac{r_\lambda(\theta^*)}{2}$ then $t_\lambda \leq \log 2$ else $t_\lambda \leq 2R\|\theta^*\|$.*

Interestingly, regularity of θ^* , like the source condition below, can induce this effect, allowing a better dependence on λ for the bias term.

Assumption 6 (Source condition). *There exists $r \in (0, 1/2]$ and $v \in \mathcal{H}$ such that $\theta^* = \mathbf{H}(\theta^*)^r v$.*

In particular we denote by $L := \|v\|$. Assumption 6 is commonly made in least-squares regression (Caponnetto and De Vito, 2007; Steinwart et al., 2009; Blanchard and Mücke, 2018) and is equivalent to requiring that, when expressing θ^* with respect to the eigenbasis of $\mathbf{H}(\theta^*)$, i.e., $\theta^* = \sum_{j \in \mathbb{N}} \alpha_j u_j$, where λ_j, u_j is the eigendecomposition of $\mathbf{H}(\theta^*)$, and $\alpha_j = \theta^* \cdot u_j$, then α_j decays as λ_j^r . In particular, with this assumption, defining $\beta_j = v \cdot u_j$,

$$\text{Bias}_\lambda^2 = \lambda^2 \sum_j \frac{\alpha_j^2}{\lambda_j + \lambda} = \lambda^2 \sum_j \frac{\lambda_j^{2r} \beta_j^2}{\lambda_j + \lambda} \leq \lambda^2 \left(\sup_j \frac{\lambda_j^{2r}}{\lambda_j + \lambda} \right) \sum_j \beta_j^2 \leq \lambda^{1+2r} \|v\|^2.$$

Note moreover that $\mathbf{H}(\theta^*) \preceq B_2^* C$, meaning that the usual sufficient conditions leading to the source conditions for least-squares also apply here. For example, for logistic regression, if the log-odds ratio is smooth enough, then it is in \mathcal{H} . So, when \mathcal{H} corresponds to a Sobolev space of smoothness m and the marginal of ρ on the input space is a density bounded away from 0 and infinity with bounded support, then the source condition corresponds essentially to requiring θ^* to be $(1 + 2r)m$ -times differentiable (see discussion after Thm. 9 of Steinwart et al., 2009, for more details).

In conclusion, the effect of additional regularity for θ^* as Assumption 6, has two beneficial effects: (a) on one side it allows to obtain faster rates as shown in the next corollary, (b) as mentioned before, somewhat surprisingly, it reduces the constants to universal, since it allows the bias to go to zero faster than the Dikin radius (indeed, the squared radius $r_\lambda^2(\theta^*)$ is always larger than λ/R^2 , which is strictly larger than $\lambda^{1+2r}\|v\|^2$ if $r > 0$ and λ small enough). This is why we do not get the exponential constants imposed by Hazan et al. (2014).

Corollary 2 (Rates with source condition). *Let $\delta \in (0, 1/2]$. Under Assumptions 1 to 5 and Assumption 6, whenever $n \geq N$ and $\lambda = (C_0/n)^{1/(2+2r)}$, then with probability at least $1 - 2\delta$, it holds*

$$L(\widehat{\theta}_\lambda^*) - L(\theta^*) \leq C_1 n^{-\frac{1+2r}{2+2r}} \log \frac{2}{\delta},$$

with $C_0 = 256 (B_1^*/L)^2$, $C_1 = 8 (256)^\gamma ((B_1^*)^\gamma L^{1-\gamma})^2$, $\gamma = \frac{1+2r}{2+2r}$ and with N defined in Eq. (48) and satisfying $N = O(\text{poly}(B_1^*, B_2^*, L, R, \log(1/\delta)))$.

The corollary above, derived in Appendix F, is obtained by minimizing in λ the r.h.s. side of Eq. (9) in Thm. 3, and considering that when θ^* satisfies the source condition, then $\text{Bias}_\lambda \leq \lambda^{1+2r} L$, while the variance is still of the form $1/(\lambda n)$. When r is close to 0, the rate $1/\sqrt{n}$ is recovered. When instead the target function is more regular, implying $r = 1/2$, a rate of $n^{-2/3}$ is achieved. Two considerations are in order: (a) the obtained rate is the same as least-squares and minimax optimal (Caponnetto and De Vito, 2007; Steinwart et al., 2009; Blanchard and Mücke, 2018), (b) the fact that regularized ERM is adaptive to the regularity of the function up to $r = 1/2$ is a byproduct of Tikhonov regularization as already shown for the least-squares case by Gerfo et al. (2008). Using different regularization techniques may remove the limit $r = 1/2$.

5 Fast Rates with both Source and Capacity Conditions

In this section, we consider improved results with a finer control of the effective dimension df_λ (often called degrees of freedom), which, together with the source condition allows to achieve rates as fast as $1/n$:

$$\text{df}_\lambda = \mathbb{E} \left[\left\| \mathbf{H}_\lambda(\theta^*)^{-1/2} \nabla \ell_Z(\theta^*) \right\|^2 \right],$$

As mentioned earlier this definition of df_λ corresponds to the usual asymptotic term in M -estimation. Moreover, in the case of least-squares, it corresponds to the standard notion of effective dimension $\text{df}_\lambda = \text{Tr}(\mathbf{C}\mathbf{C}_\lambda^{-1})$ (Caponnetto and De Vito, 2007; Blanchard and Mücke, 2018). Note that by definition, we always have $\text{df}_\lambda \leq \mathbf{B}_1^{*2}/\lambda$, but we can have in general a much finer control. For example, for least-squares, $\text{df}_\lambda = O(\lambda^{-1/\alpha})$ if the eigenvalues of the covariance operator \mathbf{C} decay as $\lambda_j(\mathbf{C}) = O(j^{-\alpha})$, for $\alpha \geq 1$. Moreover note that since \mathbf{C} is trace-class, by Asm. 3, the eigenvalues form a summable sequence and so \mathbf{C} satisfies $\lambda_j(\mathbf{C}) = O(j^{-\alpha})$ with α always larger than 1.

Example 4 (Generalized linear models). *For generalized linear models, an extra assumption makes the degrees of freedom particularly simple: if the probabilistic model is well-specified, that is, there exists θ^* such that almost surely, $p(y|x) = p(y|x, \theta^*) = \frac{\exp(\theta^* \cdot \Phi(x, y))}{\int_{\mathcal{Y}} \exp(\theta^* \cdot \Phi(x, y')) d\mu(y')}$, then from the usual Bartlett identities (Bartlett, 1953) relating the expected squared derivatives and Hessians, we have $\mathbb{E} [\nabla \ell_Z(\theta^*) \otimes \nabla \ell_Z(\theta^*)] = \mathbf{H}(\theta^*)$, leading to $\text{df}_\lambda = \text{Tr}(\mathbf{H}_\lambda(\theta^*)^{-1} \mathbf{H}(\theta^*))$.*

As we have seen in the previous example there are interesting problems for which $\text{df}_\lambda = \text{Tr}(\mathbf{H}(\theta^*) + \lambda I)^{-1} \mathbf{H}(\theta^*)$. Since we have $\mathbf{H}(\theta^*) \preceq \mathbf{B}_2^* \mathbf{C}$, df_λ still enjoys a polynomial decay depending on the eigenvalue decay of \mathbf{C} as observed for least-squares. Now we are ready to state our result in the most general form, proved in Appendix D.4.

Theorem 4 (General bound). *Let $n \in \mathbb{N}$, $\delta \in (0, 1/2]$, $0 < \lambda \leq \mathbf{B}_2^*$. Whenever*

$$n \geq \Delta_1 \frac{\mathbf{B}_2^*}{\lambda} \log \frac{8\Box_1^2 \mathbf{B}_2^*}{\lambda \delta}, \quad n \geq \Delta_2 \frac{\text{df}_\lambda \vee (\mathbf{Q}^*)^2}{r_\lambda(\theta^*)^2} \log \frac{2}{\delta},$$

with $(\mathbf{Q}^*)^2 = \mathbf{B}_1^{*2}/\mathbf{B}_2^*$, then with probability at least $1 - 2\delta$, it holds

$$L(\widehat{\theta}_\lambda^*) - L(\theta^*) \leq C_{\text{bias}} \text{Bias}_\lambda^2 + C_{\text{var}} \frac{\text{df}_\lambda \vee (\mathbf{Q}^*)^2}{n} \log \frac{2}{\delta}, \quad (10)$$

where, $C_{\text{bias}}, C_{\text{var}}, \Box_1 \leq 414$, $\Delta_1, \Delta_2 \leq 5184$ when $\text{Bias}_\lambda \leq r_\lambda(\theta^*)/2$; otherwise $C_{\text{bias}}, C_{\text{var}}, \Box_1 \leq 256e^{6R\|\theta^*\|}$, $\Delta_1, \Delta_2 \leq 2304(1 + R\|\theta^*\|)^2 e^{8R\|\theta^*\|}$.

As shown in the theorem above, the variance term depends on df_λ/n , implying that, when df_λ has a better dependence in λ than $1/\lambda$, it is possible to achieve faster rates. We quantify this with the following assumption.

Assumption 7 (Capacity condition). *There exists $\alpha > 0$ and $\mathbf{Q} \geq 0$ such that $\text{df}_\lambda \leq \mathbf{Q}\lambda^{-1/\alpha}$.*

Assumption 7 is standard in the context of least-squares, (Caponnetto and De Vito, 2007) and in many interesting settings is implied by the eigenvalue decay order of $\mathbf{H}(\theta^*)$, or \mathbf{C} as discussed above. In the following corollary we quantify the effect of df_λ in the learning rates.

Corollary 3. *Let $\delta \in (0, 1/2]$. Under Assumptions 1 to 5, Assumption 6 and Assumption 7, when $n \geq N$ and $\lambda = (C_0/n)^{\alpha/(1+\alpha(1+2r))}$, then with probability at least $1 - 2\delta$,*

$$L(\widehat{\theta}_\lambda^*) - L(\theta^*) \leq C_1 n^{-\frac{\alpha(1+2r)}{1+\alpha(1+2r)}} \log \frac{2}{\delta},$$

with $C_0 = 256(Q/L)^2$, $C_1 = 8(256)^\gamma (Q^\gamma L^{1-\gamma})^2$, $\gamma = \frac{\alpha(1+2r)}{1+\alpha(1+2r)}$ and N defined in Eq. (48) and satisfying $N = O(\text{poly}(B_1^*, B_2^*, L, Q, R, \log(1/\delta)))$.

The result above is derived in Cor. 4 in Appendix F and is obtained by bounding Bias_λ with $\lambda^{1+2r}L$ due to the source condition, and df_λ with $\lambda^{-1/\alpha}$ due to the capacity condition and then optimizing the r.h.s. of Eq. (10) in λ . Note that (a) the learning rate under the considered assumptions is the same as least-squares and minimax optimal (Caponnetto and De Vito, 2007; Lin et al., 2018; Blanchard and Mücke, 2018), and (b) when $\alpha = 1$ the same rate of Cor. 2 is achieved, which can be as fast as $n^{-2/3}$, otherwise, when $\alpha \gg 1$, we achieve a learning rate in the order of $1/n$, for $\lambda = n^{-1/(1+2r)}$.

6 Sketch of the proof

In this section we will use the notation $\|v\|_{\mathbf{A}} := \|\mathbf{A}^{1/2}v\|$, with $v \in \mathcal{H}$ and \mathbf{A} a bounded positive semi-definite operator on \mathcal{H} . Here we prove that the excess risk decomposes using the bias term Bias_λ defined in Eq. (8) and a variance term V_λ , where V_λ is defined as

$$V_\lambda := \|\nabla \widehat{L}_\lambda(\theta_\lambda^*)\|_{\mathbf{H}_\lambda^{-1}(\theta_\lambda^*)}, \quad \text{with} \quad \widehat{L}_\lambda(\cdot) = \frac{1}{n} \sum_{i=1}^n \ell_{z_i}(\cdot) + \frac{\lambda}{2} \|\cdot\|^2,$$

which in turn is a random variable that concentrate in high probability to $\sqrt{\text{df}_\lambda/n}$.

Required tools. To proceed with the proof we need two main tools. The first is a result on the equivalence of norms of the empirical Hessian $\widehat{\mathbf{H}}_\lambda(\theta) = \nabla^2 \widehat{L}_\lambda(\theta)$ w.r.t. the true Hessian $\mathbf{H}_\lambda(\theta) = \nabla^2 L_\lambda(\theta)$ for $\lambda > 0$ and $\theta \in \mathcal{H}$. The result is proven in Lemma 6 of Appendix D.3, using Bernstein inequalities for Hermitian operators (Tropp, 2012), and essentially states that for $\delta \in (0, 1]$, whenever $n \geq \frac{24B_2(\theta)}{\lambda} \log \frac{8B_2(\theta)}{\lambda\delta}$, then with probability $1 - \delta$, it holds

$$\|\cdot\|_{\mathbf{H}_\lambda(\theta)} \leq 2\|\cdot\|_{\widehat{\mathbf{H}}_\lambda(\theta)}, \quad \|\cdot\|_{\widehat{\mathbf{H}}_\lambda^{-1}(\theta)} \leq 2\|\cdot\|_{\mathbf{H}_\lambda^{-1}(\theta)}. \quad (11)$$

The second result is about localization properties induced by generalized self-concordance on the risk. We express the result with respect to a generic probability μ (we will use it with $\mu = \rho$ and $\mu = \frac{1}{n} \sum_{i=1}^n \delta_{z_i}$). Let μ be a probability distribution with support contained in the support of ρ . Denote by $L_\mu(\theta)$ the risk $L_\mu(\theta) = \mathbb{E}_{z \sim \mu}[\ell_z(\theta)]$ and by $L_{\mu,\lambda}(\theta) = L_\mu(\theta) + \frac{\lambda}{2} \|\theta\|^2$ (then $L_{\mu,\lambda} = L_\lambda$ when $\mu = \rho$, or \widehat{L}_λ when $\mu = \frac{1}{n} \sum_{i=1}^n \delta_{z_i}$).

Proposition 1. *Under Assumptions 2 to 4, the following holds: (a) $L_{\mu,\lambda}(\theta), \nabla L_{\mu,\lambda}(\theta), \mathbf{H}_{\mu,\lambda}(\theta)$ are defined for all $\theta \in \mathcal{H}, \lambda \geq 0$, (b) for all $\lambda > 0$, there exists a unique $\theta_{\mu,\lambda}^* \in \mathcal{H}$ minimizing $L_{\mu,\lambda}$ over \mathcal{H} , and (c) for all $\lambda > 0$ and $\theta \in \mathcal{H}$,*

$$\mathbf{H}_{\mu,\lambda}(\theta) \preceq e^{t_0} \mathbf{H}_{\mu,\lambda}(\theta_{\mu,\lambda}^*), \quad (12)$$

$$L_{\mu,\lambda}(\theta) - L_{\mu,\lambda}(\theta_{\mu,\lambda}^*) \leq \psi(t_0) \|\theta - \theta_{\mu,\lambda}^*\|_{\mathbf{H}_{\mu,\lambda}(\theta_{\mu,\lambda}^*)}^2, \quad (13)$$

$$\underline{\phi}(t_0) \|\theta - \theta_{\mu,\lambda}^*\|_{\mathbf{H}_{\mu,\lambda}(\theta)} \leq \|\nabla L_{\mu,\lambda}(\theta)\|_{\mathbf{H}_{\mu,\lambda}^{-1}(\theta)}, \quad (14)$$

(d) Eqs. (12) and (13) hold also for $\lambda = 0$, provided that $\theta_{\mu,0}^*$ exists. Here $t_0 := t(\theta - \theta_{\mu,\lambda}^*)$ and $\underline{\phi}(t) = (1 - e^{-t})/t$, $\psi(t) = (e^t - t - 1)/t^2$.

The result above is proved in Appendix B.1 and is essentially an extension of results by Bach (2010) applied to $L_{\mu,\lambda}$ under Assumptions 2 to 4.

Sketch of the proof. Now we are ready to decompose the excess risk using our bias and variance terms. In particular we will sketch the decomposition without studying the terms that lead to constants terms. For the complete proof of the decomposition see Thm. 7 in Appendix D.1. Since θ^* exists by Assumption 5, using Eq. (13), applied with $\mu = \underline{\rho}$ and $\lambda = 0$, we have $L(\theta) - L(\theta^*) \leq \psi(\mathfrak{t}(\theta - \theta^*)) \|\theta - \theta^*\|_{\mathbf{H}(\theta^*)}^2$ for any $\theta \in \mathcal{H}$. By setting $\theta = \hat{\theta}_\lambda^*$, we obtain

$$L(\hat{\theta}_\lambda^*) - L(\theta^*) \leq \psi(\mathfrak{t}(\hat{\theta}_\lambda^* - \theta^*)) \|\hat{\theta}_\lambda^* - \theta^*\|_{\mathbf{H}(\theta^*)}^2.$$

The term $\psi(\mathfrak{t}(\hat{\theta}_\lambda^* - \theta^*))$ will become a constant. For the sake of simplicity, in this sketch of proof we will not deal with it nor with other terms of the form $\mathfrak{t}(\cdot)$ leading to constants. On the other hand, the term $\|\hat{\theta}_\lambda^* - \theta^*\|_{\mathbf{H}(\theta^*)}^2$ will yield our bias and variance terms. Using the fact that $\mathbf{H}(\theta^*) \preceq \mathbf{H}(\theta^*) + \lambda I =: \mathbf{H}_\lambda(\theta^*)$, by adding and subtracting θ_λ^* , we have

$$\|\theta_\lambda^* - \theta^*\|_{\mathbf{H}(\theta^*)} \leq \|\theta_\lambda^* - \theta^*\|_{\mathbf{H}_\lambda(\theta^*)} \leq \|\theta_\lambda^* - \theta^*\|_{\mathbf{H}_\lambda(\theta^*)} + \|\hat{\theta}_\lambda^* - \theta_\lambda^*\|_{\mathbf{H}_\lambda(\theta^*)},$$

so

$$L(\hat{\theta}_\lambda^*) - L(\theta^*) \leq \text{const.} \times (\|\theta_\lambda^* - \theta^*\|_{\mathbf{H}_\lambda(\theta^*)}^2 + \|\hat{\theta}_\lambda^* - \theta_\lambda^*\|_{\mathbf{H}_\lambda(\theta^*)}^2).$$

By applying Eq. (12) with $\mu = \rho$ and $\theta = \theta^*$, we have $\mathbf{H}_\lambda(\theta^*) \preceq e^{\mathfrak{t}_\lambda} \mathbf{H}_\lambda(\theta_\lambda^*)$ and so we further bound $\|\hat{\theta}_\lambda^* - \theta_\lambda^*\|_{\mathbf{H}_\lambda(\theta^*)}$ with $e^{\mathfrak{t}_\lambda/2} \|\hat{\theta}_\lambda^* - \theta_\lambda^*\|_{\mathbf{H}_\lambda(\theta_\lambda^*)}$ obtaining

$$L(\hat{\theta}_\lambda^*) - L(\theta^*) \leq \text{const.} \times (\|\theta_\lambda^* - \theta^*\|_{\mathbf{H}_\lambda(\theta^*)} + e^{\mathfrak{t}_\lambda/2} \|\hat{\theta}_\lambda^* - \theta_\lambda^*\|_{\mathbf{H}_\lambda(\theta_\lambda^*)})^2.$$

The term $\|\theta_\lambda^* - \theta^*\|_{\mathbf{H}_\lambda(\theta^*)}$ will lead to the *bias terms*, while the term $\|\hat{\theta}_\lambda^* - \theta_\lambda^*\|_{\mathbf{H}_\lambda(\theta_\lambda^*)}$ will lead to the *variance term*.

Bounding the bias terms. Recall the definition of bias $\text{Bias}_\lambda = \|\nabla L_\lambda(\theta^*)\|_{\mathbf{H}_\lambda^{-1}(\theta^*)}$ and of the constant $\mathfrak{t}_\lambda := \mathfrak{t}(\theta^* - \theta_\lambda^*)$. We bound $\|\theta^* - \theta_\lambda^*\|_{\mathbf{H}_\lambda(\theta^*)}$ by applying Eq. (14) with $\mu = \rho$ and $\theta = \theta^*$

$$\|\theta^* - \theta_\lambda^*\|_{\mathbf{H}_\lambda(\theta^*)} \leq 1/\underline{\phi}(\mathfrak{t}_\lambda) \|\nabla L_\lambda(\theta^*)\|_{\mathbf{H}_\lambda^{-1}(\theta^*)} = 1/\underline{\phi}(\mathfrak{t}_\lambda) \text{Bias}_\lambda.$$

Bounding the variance terms. To bound the term $\|\hat{\theta}_\lambda^* - \theta_\lambda^*\|_{\mathbf{H}_\lambda(\theta_\lambda^*)}$, we assume n large enough to apply Eq. (11) in high probability. Thus, we obtain

$$\|\hat{\theta}_\lambda^* - \theta_\lambda^*\|_{\mathbf{H}_\lambda(\theta_\lambda^*)} \leq 2\|\hat{\theta}_\lambda^* - \theta_\lambda^*\|_{\hat{\mathbf{H}}_\lambda(\theta_\lambda^*)}.$$

Applying Eq. (14) with $\mu = \frac{1}{n} \sum_{i=1}^n \delta_{z_i}$ and $\theta = \hat{\theta}_\lambda^*$, since $L_{\mu,\lambda} = \hat{L}_\lambda$ for the given choice of μ ,

$$\|\theta_\lambda^* - \hat{\theta}_\lambda^*\|_{\hat{\mathbf{H}}_\lambda(\theta_\lambda^*)} \leq \|\nabla \hat{L}_\lambda(\theta_\lambda^*)\|_{\hat{\mathbf{H}}_\lambda^{-1}(\theta_\lambda^*)} / \underline{\phi}(\mathfrak{t}(\theta_\lambda^* - \hat{\theta}_\lambda^*)),$$

and applying Eq. (11) in high probability again, we obtain

$$\|\nabla \hat{L}_\lambda(\theta_\lambda^*)\|_{\hat{\mathbf{H}}_\lambda^{-1}(\theta_\lambda^*)} \leq 2\|\nabla \hat{L}_\lambda(\theta_\lambda^*)\|_{\mathbf{H}_\lambda^{-1}(\theta_\lambda^*)}.$$

Bias-variance decomposition. A technical part of the proof relates $\|\nabla \hat{L}_\lambda(\theta_\lambda^*)\|_{\mathbf{H}_\lambda^{-1}(\theta_\lambda^*)}$ with $\|\nabla \hat{L}_\lambda(\theta^*)\|_{\mathbf{H}_\lambda^{-1}(\theta^*)} =: V_\lambda$, by many applications of Prop. 1. Here we assume it is done, obtaining

$$L(\hat{\theta}_\lambda^*) - L(\theta^*) \leq \text{const.} \times (\text{Bias}_\lambda^2 + V_\lambda^2).$$

From V_λ to $\sqrt{\text{df}_\lambda/n}$. By construction, $\nabla \widehat{L}_\lambda(\theta_\lambda^*) = \frac{1}{n} \sum_{i=1}^n \zeta_i$, with $\zeta_i := \nabla \ell_{z_i}(\theta_\lambda^*) + \lambda \theta_\lambda^*$. Moreover since the z_i 's are i.i.d. samples from ρ , $\mathbb{E}[\zeta_i] = \nabla L_\lambda(\theta_\lambda^*)$. Finally since θ_λ^* is the minimizer of L_λ , $\nabla L_\lambda(\theta_\lambda^*) = 0$. Thus $\nabla \widehat{L}_\lambda(\theta_\lambda^*)$ is the average of n i.i.d. zero-mean random vectors, and so the variance of V_λ is exactly

$$\mathbb{E}[V_\lambda^2] = \frac{1}{n} \mathbb{E} \left[\|\mathbf{H}_\lambda^{-1/2}(\theta^*) \nabla \ell_Z(\theta^*)\|^2 \right] = \frac{\text{df}_\lambda}{n}.$$

Finally, by using Bernstein inequality for random vectors (e.g. Thm. 3.3.4 [Yurinsky, 1995](#)), we bound V_λ roughly with $\sqrt{\text{df}_\lambda \log(2/\delta)/n}$ in high probability.

7 Conclusion

In this paper we have presented non-asymptotic bounds with faster rates than $O(1/\sqrt{n})$, for regularized empirical risk minimization with self-concordant losses such as the logistic loss. It would be interesting to extend our work to algorithms used to minimize the empirical risk, in particular stochastic gradient descent or Newton's method.

Acknowledgments

This work was supported by the European Research Council (project Sequoia 724063).

References

- Nachman Aronszajn. Theory of reproducing kernels. *Transactions of the American mathematical society*, 68(3):337–404, 1950.
- Francis Bach. Self-concordant analysis for logistic regression. *Electronic Journal of Statistics*, 4: 384–414, 2010.
- Francis Bach. Adaptivity of averaged stochastic gradient descent to local strong convexity for logistic regression. *Journal of Machine Learning Research*, 15(1):595–627, 2014.
- M. S. Bartlett. Approximate confidence intervals. *Biometrika*, 40(1/2):12–19, 1953.
- Gilles Blanchard and Nicole Mücke. Optimal rates for regularization of statistical inverse learning problems. *Foundations of Computational Mathematics*, 18(4):971–1013, 2018.
- Léon Bottou and Olivier Bousquet. The trade-offs of large scale learning. In *Advances in Neural Information Processing systems*, pages 161–168, 2008.
- Stéphane Boucheron and Pascal Massart. A high-dimensional Wilks phenomenon. *Probability Theory and Related Fields*, 150(3-4):405–433, 2011.
- A. Caponnetto and E. De Vito. Optimal rates for the regularized least-squares algorithm. *Found. Comput. Math.*, 7(3):331–368, July 2007.
- Nicolo Cesa-Bianchi, Yishay Mansour, and Ohad Shamir. On the complexity of learning with kernels. In *Conference on Learning Theory*, pages 297–325, 2015.
- Aymeric Dieuleveut and Francis Bach. Nonparametric stochastic approximation with large step-sizes. *The Annals of Statistics*, 44(4):1363–1399, 2016.

- Simon Fischer and Ingo Steinwart. Sobolev norm learning rates for regularized least-squares algorithm. *arXiv preprint arXiv:1702.07254*, 2017.
- Dylan J. Foster, Satyen Kale, Haipeng Luo, Mehryar Mohri, and Karthik Sridharan. Logistic regression: The importance of being improper. *Proceedings of COLT*, 2018.
- Stuart Geman, Elie Bienenstock, and René Doursat. Neural networks and the bias/variance dilemma. *Neural Computation*, 4(1):1–58, 1992.
- L. Lo Gerfo, Lorenzo Rosasco, Francesca Odone, E. De Vito, and Alessandro Verri. Spectral algorithms for supervised learning. *Neural Computation*, 20(7):1873–1897, 2008.
- Frank R. Hampel, Elvezio M. Ronchetti, Peter J. Rousseeuw, and Werner A. Stahel. *Robust statistics: the approach based on influence functions*, volume 196. John Wiley & Sons, 2011.
- Elad Hazan, Tomer Koren, and Kfir Y. Levy. Logistic regression: tight bounds for stochastic and online optimization. In *Proceedings of The 27th Conference on Learning Theory*, volume 35, pages 197–209, 2014.
- Arthur E. Hoerl and Robert W. Kennard. Ridge regression iterative estimation of the biasing parameter. *Communications in Statistics-Theory and Methods*, 5(1):77–88, 1976.
- S. Sathiya Keerthi, K. B. Duan, Shirish Krishnaji Shevade, and Aun Neow Poo. A fast dual algorithm for kernel logistic regression. *Machine learning*, 61(1-3):151–165, 2005.
- Tomer Koren and Kfir Levy. Fast rates for exp-concave empirical risk minimization. In *Advances in Neural Information Processing Systems*, pages 1477–1485, 2015.
- John Lafferty, Andrew McCallum, and Fernando Pereira. Conditional random fields: Probabilistic models for segmenting and labeling sequence data. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2001.
- Erich L. Lehmann and George Casella. *Theory of Point Estimation*. Springer Science & Business Media, 2006.
- Junhong Lin, Alessandro Rudi, Lorenzo Rosasco, and Volkan Cevher. Optimal rates for spectral algorithms with least-squares regression over hilbert spaces. *Applied and Computational Harmonic Analysis*, 2018.
- P. McCullagh and J.A. Nelder. *Generalized Linear Models*. Chapman and Hall, 1989.
- Nishant A. Mehta. Fast rates with high probability in exp-concave statistical learning. Technical Report 1605.01288, ArXiv, 2016.
- Yurii Nesterov and Arkadii Nemirovskii. *Interior-point polynomial algorithms in convex programming*, volume 13. SIAM, 1994.
- Dmitrii Ostrovskii and Francis Bach. Finite-sample analysis of M-estimators using self-concordance. Technical Report 1810.06838, arXiv, 2018.
- Alexander Rakhlin and Karthik Sridharan. Online nonparametric regression with general loss functions. *arXiv preprint arXiv:1501.06598*, 2015.
- Alessandro Rudi and Lorenzo Rosasco. Generalization properties of learning with random features. In *Advances in Neural Information Processing Systems 30*, pages 3215–3225. Curran Associates, Inc., 2017.

- Alessandro Rudi, Luigi Carratino, and Lorenzo Rosasco. Falkon: An optimal large scale kernel method. In *Advances in Neural Information Processing Systems*, pages 3888–3898, 2017.
- Shai Shalev-Shwartz, Yoram Singer, Nathan Srebro, and Andrew Cotter. Pegasos: Primal estimated sub-gradient solver for SVM. *Mathematical programming*, 127(1):3–30, 2011.
- John Shawe-Taylor and Nello Cristianini. *Kernel Methods for Pattern Analysis*. Cambridge University Press, 2004.
- Steve Smale and Ding-Xuan Zhou. Learning theory estimates via integral operators and their approximations. *Constructive Approximation*, 26(2):153–172, 2007.
- Karthik Sridharan, Shai Shalev-Shwartz, and Nathan Srebro. Fast rates for regularized objectives. In *Advances in Neural Information Processing Systems 21*, pages 1545–1552. 2009.
- Ingo Steinwart and Clint Scovel. Fast rates for support vector machines using gaussian kernels. *The Annals of Statistics*, 35(2):575–607, 2007.
- Ingo Steinwart, Don R. Hush, and Clint Scovel. Optimal rates for regularized least squares regression. In *Proc. COLT*, 2009.
- Joel A. Tropp. User-friendly tail bounds for sums of random matrices. *Foundations of computational mathematics*, 12(4):389–434, 2012.
- Stephen Tu, Rebecca Roelofs, Shivaram Venkataraman, and Benjamin Recht. Large scale kernel learning using block coordinate descent. *arXiv preprint arXiv:1602.05310*, 2016.
- Sara A. Van de Geer. High-dimensional generalized linear models and the Lasso. *The Annals of Statistics*, 36(2):614–645, 2008.
- Aad W. Van der Vaart. *Asymptotic Statistics*, volume 3. Cambridge University Press, 2000.
- Tim Van Erven, Peter D. Grünwald, Nishant A. Mehta, Mark D. Reid, and Robert C. Williamson. Fast rates in statistical and online learning. *Journal of Machine Learning Research*, 16:1793–1861, 2015.
- Grace Wahba. *Spline Models for Observational Data*, volume 59. SIAM, 1990.
- Vadim Yurinsky. *Sums and Gaussian vectors, volume 1617 of Lecture Notes in Mathematics*. Springer-Verlag, Berlin, 1995.

Organization of the Appendix

A *Setting, definitions, assumptions*

B *Preliminary results on self concordant losses*

B.1 *Basic results on self-concordance (proof of Proposition 1)*

B.2 *Localization properties for t_λ (proof of Lemma 1)*

C *Main result, simplified*

C.1 *Analytic decomposition of the risk*

C.2 *Concentration lemmas*

C.3 *Final result (proof of Thm. 2)*

D *Main result, refined analysis*

D.1 *Analytic decomposition of the risk*

D.2 *Analytic decomposition of terms related to the variance*

D.3 *Concentration lemmas*

D.4 *Final result (proof of Thms. 3 and 4)*

E *Explicit bounds for the simplified case (proof of Cor. 1)*

F *Explicit bounds for the refined case (proof of Cors. 2 and 3)*

G *Additional lemmas*

G.1 *Self-concordance and sufficient conditions to define L*

G.2 *Bernstein inequalities for operators*

A **Setting, definitions, assumptions**

Let \mathcal{Z} be a Polish space and Z a random variable on \mathcal{Z} with law ρ . Let \mathcal{H} be a separable (non-necessarily finite) Hilbert space and let $\ell : \mathcal{Z} \times \mathcal{H} \rightarrow \mathbb{R}$ be a loss function; we denote by $\ell_z(\cdot)$ the function $\ell(z, \cdot)$. Our goal is to solve

$$\inf_{\theta \in \mathcal{H}} L(\theta), \quad \text{with} \quad L(\theta) = \mathbb{E}[\ell_Z(\theta)].$$

Given $(z_i)_{i=1}^n$ we will consider the following estimator

$$\hat{\theta}_\lambda^* = \arg \min_{\theta \in \mathcal{H}} \hat{L}_\lambda(\theta), \quad \text{with} \quad \hat{L}_\lambda(\theta) := \frac{1}{n} \sum_{i=1}^n \ell_{z_i}(\theta) + \frac{\lambda}{2} \|\theta\|^2.$$

The goal of this work is to give upper bounds in high probability to the so called *excess risk*

$$L(\hat{\theta}_\lambda^*) - \inf_{\theta \in \mathcal{H}} L(\theta).$$

In the rest of this introduction we will introduce the basic assumptions required to make $\hat{\theta}_\lambda^*$ and the *excess risk* well defined, and we will introduce basic objects that are needed for the proofs.

First we introduce some notation we will use in the rest of the appendix: let $\lambda \geq 0$, $\theta \in \mathcal{H}$ and \mathbf{A} be a bounded positive semidefinite Hermitian operator on \mathcal{H} , we denote by \mathbf{I} , the identity operator and

$$\|f\|_{\mathbf{A}} := \|\mathbf{A}^{1/2}f\|, \quad (15)$$

$$\mathbf{A}_\lambda := \mathbf{A} + \lambda\mathbf{I}, \quad (16)$$

$$\ell_z^\lambda(\theta) := \ell_z(\theta) + \frac{\lambda}{2}\|\theta\|^2, \quad (17)$$

$$L_\lambda(\theta) := L(\theta) + \frac{\lambda}{2}\|\theta\|^2. \quad (18)$$

Now we recall the assumptions we require on the loss function $\ell, \rho, (z_i)_{1 \leq i \leq n}$.

Assumption 1 (i.i.d. data). *The samples $(z_i)_{1 \leq i \leq n}$ are independently and identically distributed according to ρ .*

Assumption 8 (Generalized self-concordance). *The mapping $z \mapsto \ell_z(\theta)$ is measurable for all $\theta \in \mathcal{H}$ and for any $z \in \mathcal{Z}$, the function ℓ_z is convex and three times differentiable. Moreover, there exists a set $\varphi(z) \subset \mathcal{H}$ such that it holds:*

$$\forall \theta \in \mathcal{H}, \forall h, k \in \mathcal{H}, \quad |\nabla^3 \ell_z(\theta)[k, h, h]| \leq \sup_{g \in \varphi(z)} |k \cdot g| \nabla^2 \ell_z(\theta)[h, h].$$

Assumption 3 (Boundedness). *There exists $R \geq 0$ such that $\sup_{g \in \varphi(z)} \|g\| \leq R$ almost surely.*

Assumption 4 (Definition in 0). *$|\ell_Z(0)|, \|\nabla \ell_Z(0)\|$ and $\text{Tr}(\nabla^2 \ell_Z(0))$ are almost surely bounded.*

Introduce the following definitions.

Definition 2. *Let $\lambda > 0$, $\theta \in \mathcal{H}$. We introduce*

$$\mathbf{B}_1(\theta) = \sup_{z \in \text{supp}(\rho)} \|\nabla \ell_z(\theta)\|, \quad \mathbf{B}_2(\theta) = \sup_{z \in \text{supp}(\rho)} \text{Tr}(\nabla^2 \ell_z(\theta)). \quad (19)$$

$$\mathbf{H}(\theta) = \mathbb{E}[\nabla^2 \ell_Z(\theta)], \quad \mathbf{H}_\lambda(\theta) = \mathbf{H}(\theta) + \lambda\mathbf{I}. \quad (20)$$

$$\theta_\lambda^* = \arg \min_{\theta \in \mathcal{H}} L_\lambda(\theta). \quad (21)$$

Proposition 2. *Under Assumptions 3, 4 and 8, $\mathbf{B}_1(\theta), \mathbf{B}_2(\theta), L(\theta), \nabla L(\theta), \mathbf{H}(\theta), \theta_\lambda^*$ exist for any $\theta \in \mathcal{H}, \lambda > 0$. Moreover $\nabla L = \mathbb{E}[\nabla \ell_Z(\theta)], \mathbf{H}(\theta) = \nabla^2 L(\theta)$ and $\mathbf{H}(\theta)$ is trace class.*

Proof. We start by proving, using the assumptions, that $\mathbf{B}_2, \mathbf{B}_1$ and $\theta \mapsto \sup_{z \in \text{supp}(\rho)} |\ell_z(\theta)|$ are all locally bounded (see Lemmas 11 to 13). This allows us to show that $\ell_z(\theta), \nabla \ell_z(\theta)$ and $\text{Tr}(\nabla^2 \ell_z(\theta))$ are uniformly integrable on any ball of finite radius. The fact that θ_λ^* exists is due to the strong convexity of the function L_λ .

Proposition 3. *Under Assumptions 1, 4 and 8, when $\lambda > 0$, $\widehat{\theta}_\lambda^*$ exists and is unique.*

Proof. By Assumption 1 we know that z_1, \dots, z_n are in the support of ρ . Thus, by Assumption 4, $\frac{1}{n} \sum_{i=1}^n \ell_{z_i}$ is finite valued in 0. Since $\frac{1}{n} \sum_{i=1}^n \ell_{z_i}$ is convex three times differentiable as a sum of such functions, it is real-valued on \mathcal{H} and hence \widehat{L}_λ is real-valued on \mathcal{H} ; by strong convexity, $\widehat{\theta}_\lambda^*$ exists and is unique.

Recall that we also make the following regularity assumption.

Assumption 5 (Existence of a minimizer). *There exists $\theta^* \in \mathcal{H}$ such that $L(\theta^*) = \inf_{\theta \in \mathcal{H}} L(\theta)$.*

Finally we conclude with the following definitions that will be used later.

Definition 3. For $\theta \in \mathcal{H}$, denote by $\mathfrak{t}(\theta)$ the function

$$\mathfrak{t}(\theta) = \sup_{z \in \text{supp}(\rho)} \left(\sup_{g \in \varphi(z)} |\theta \cdot g| \right),$$

and define

$$\text{Bias}_\lambda = \|\nabla L_\lambda(\theta^*)\|_{\mathbf{H}_\lambda^{-1}(\theta^*)}, \quad (22)$$

$$\widehat{\text{Var}}_\lambda = \|\mathbf{H}_\lambda^{1/2}(\theta_\lambda^*) \widehat{\mathbf{H}}_\lambda^{-1/2}(\theta_\lambda^*)\|^2 \|\nabla \widehat{L}_\lambda(\theta_\lambda^*)\|_{\mathbf{H}_\lambda^{-1}(\theta_\lambda^*)}, \quad (23)$$

$$\text{df}_\lambda = \mathbb{E} \left[\|\nabla \ell_Z(\theta^*)\|_{\mathbf{H}_\lambda^{-1}(\theta^*)}^2 \right], \quad (24)$$

$$\mathfrak{t}_\lambda = \mathfrak{t}(\theta^* - \theta_\lambda^*), \quad (25)$$

$$r_\lambda(\theta) \quad \text{such that} \quad 1/r_\lambda(\theta) = \sup_{z \in \text{supp}(\rho)} \left(\sup_{g \in \varphi(z)} \|g\|_{\mathbf{H}_\lambda^{-1}(\theta)} \right). \quad (26)$$

B Preliminary results on self concordant losses

In this section, we show how our definition/assumption of self concordance (see Assumption 8) enables a fine control on the excess risk. In particular, we clearly relate the difference in function values to the quadratic approximation at the optimum as well as the renormalized gradient. We start by presenting a general bounds in Appendix B.1 before applying them to the problem of localizing the optimum Appendix B.2.

B.1 Basic results on self-concordance

In this section, as in the rest of the appendix, we are under the conditions of Assumption 8. **In this section only**, we give ourselves a probability measure μ on \mathcal{Z} . We will apply the results of this section to $\mu = \rho, \widehat{\rho}, \delta_z$, where $\widehat{\rho} = \frac{1}{n} \sum_{i=1}^n \delta_{z_i}$ and z is sampled from ρ .

First of all, let us introduce the following notation. For any probability measure μ on \mathcal{Z} and any $\theta \in \mathcal{H}$, define

- $R^\mu = \sup_{z \in \text{supp}(\mu)} \left(\sup_{g \in \varphi(z)} \|g\| \right)$,
- $\mathfrak{t}^\mu(\theta) = \sup_{z \in \text{supp}(\mu)} \left(\sup_{g \in \varphi(z)} |\theta \cdot g| \right)$.

In order to be able to define $L_\mu(\theta) = \mathbb{E}_\mu[\ell_z(\theta)]$ and to derive under the expectation, we assume that Assumptions 3 and 4 are satisfied for μ (replace ρ by μ in the assumption).

Since μ and ℓ satisfy Assumptions 3, 4 and 8, Proposition 8 ensures that we can define $L_\mu(\theta) = \mathbb{E}_\mu[\ell_z(\theta)]$ and $L_{\mu,\lambda}(\theta) = L_\mu(\theta) + \frac{\lambda}{2} \|\theta\|^2$, as well as their respective Hessians $\mathbf{H}_\mu(\theta)$ and $\mathbf{H}_{\mu,\lambda}(\theta)$.

The following result is greatly inspired from results in (Bach, 2010) on generalized self concordant losses, and their refinement in (Ostrovskii and Bach, 2018). However, while Eqs. (27), (29) and (30) appear more or less explicitly, Eq. (28) provides an easier way to deal with certain bounds afterwards and was not used in this form before.

Proposition 4 (using the self-concordance of ℓ). *Let $\theta_0, \theta_1 \in \mathcal{H}$ and $\lambda \geq 0$. Assume that $(\ell_z)_z$ and μ satisfy Assumptions 3, 4 and 8. We have the following inequalities:*

- *Bounds on Hessians*

$$\mathbf{H}_{\mu,\lambda}(\theta_1) \preceq \exp(\mathbf{t}^\mu(\theta_1 - \theta_0)) \mathbf{H}_{\mu,\lambda}(\theta_0). \quad (27)$$

- *Bounds on gradients (if $\lambda > 0$)*

$$\underline{\phi}(\mathbf{t}^\mu(\theta_1 - \theta_0)) \|\theta_1 - \theta_0\|_{\mathbf{H}_{\mu,\lambda}(\theta_0)} \leq \|\nabla L_{\mu,\lambda}(\theta_1) - \nabla L_{\mu,\lambda}(\theta_0)\|_{\mathbf{H}_{\mu,\lambda}^{-1}(\theta_0)}, \quad (28)$$

$$\|\nabla L_{\mu,\lambda}(\theta_1) - \nabla L_{\mu,\lambda}(\theta_0)\|_{\mathbf{H}_{\mu,\lambda}^{-1}(\theta_0)} \leq \overline{\phi}(\mathbf{t}^\mu(\theta_1 - \theta_0)) \|\theta_1 - \theta_0\|_{\mathbf{H}_{\mu,\lambda}(\theta_0)}, \quad (29)$$

where $\overline{\phi}(t) = (e^t - 1)/t$ and $\underline{\phi}(t) = (1 - e^{-t})/t$.

- *Bounds on function values*

$$L_{\mu,\lambda}(\theta_1) - L_{\mu,\lambda}(\theta_0) - \nabla L_{\mu,\lambda}(\theta_0)(\theta_1 - \theta_0) \leq \psi(\mathbf{t}^\mu(\theta_1 - \theta_0)) \|\theta_1 - \theta_0\|_{\mathbf{H}_{\mu,\lambda}(\theta_0)}^2, \quad (30)$$

where $\psi(t) = (e^t - t - 1)/t^2$.

Proof. First of all, note that for any μ and λ , given $\theta \in \mathcal{H}$ and $k, h \in \mathcal{H}$,

$$\begin{aligned} |\nabla^3 L_{\mu,\lambda}(\theta)[h, k, k]| &= \left| \mathbb{E}_{z \sim \mu} \left[\nabla^3 \ell_z^\lambda(\theta)[h, k, k] \right] \right| \\ &\leq \mathbb{E}_{z \sim \mu} \left[\left| \nabla^3 \ell_z(\theta)[h, k, k] \right| \right] \\ &\leq \mathbb{E}_{z \sim \mu} \left[\sup_{g \in \varphi(z)} |h \cdot g| \nabla^2 \ell_z(\theta)[k, k] \right] \\ &\leq \mathbf{t}^\mu(h) \mathbb{E}_{z \sim \mu} \left[\nabla^2 \ell_z(\theta)[k, k] \right] = \mathbf{t}^\mu(h) \nabla^2 L_{\mu,\lambda}(\theta)[k, k]. \end{aligned}$$

This yields the following fundamental inequality :

$$|\nabla^3 L_{\mu,\lambda}(\theta)[h, k, k]| \leq \mathbf{t}^\mu(h) \nabla^2 L_{\mu,\lambda}(\theta)[k, k]. \quad (31)$$

We now define, for any $t \in \mathbb{R}$, $\theta_t := \theta_0 + t(\theta_1 - \theta_0)$.

Point 1. For the first inequality, let $h \in \mathcal{H}$ be a fixed vector, and consider the function $\varphi : t \in \mathbb{R} \mapsto \nabla^2 L_{\mu,\lambda}(\theta_t)[h, h]$. Since $\varphi'(t) = \nabla^3 L_{\mu,\lambda}(\theta_t)[\theta_1 - \theta_0, h, h]$, using Eq. (31), we get that $\varphi'(t) \leq \mathbf{t}^\mu(\theta_1 - \theta_0) \varphi(t)$. Using Lemma 10, we directly find that $\varphi(1) \leq \exp(\mathbf{t}^\mu(\theta_1 - \theta_0))\varphi(0)$, which, rewriting the definition of φ , yields

$$\nabla^2 L_{\mu,\lambda}(\theta_1)[h, h] \leq \exp(\mathbf{t}^\mu(\theta_1 - \theta_0)) \nabla^2 L_{\mu,\lambda}(\theta_0)[h, h].$$

This being true for any direction h , we have (27).

Point 2. To prove Eq. (28), let us look at the quantity $(\theta_1 - \theta_0) \cdot (\nabla L_{\mu,\lambda}(\theta_1) - \nabla L_{\mu,\lambda}(\theta_0))$. Since $\nabla L_{\mu,\lambda}(\theta_1) - \nabla L_{\mu,\lambda}(\theta_0) = \int_0^1 \nabla^2 L_{\mu,\lambda}(\theta_t)(\theta_1 - \theta_0) dt$, we have

$$(\theta_1 - \theta_0) \cdot (\nabla L_{\mu,\lambda}(\theta_1) - \nabla L_{\mu,\lambda}(\theta_0)) = \int_0^1 \nabla^2 L_{\mu,\lambda}(\theta_t)[\theta_1 - \theta_0, \theta_1 - \theta_0] dt.$$

Applying Eq. (27) to θ_0 and θ_t and the reverse, we find that

$$\forall t \in [0, 1], e^{-t\mathbf{t}^\mu(\theta_1 - \theta_0)} \nabla^2 L_{\mu,\lambda}(\theta_0) \preceq \nabla^2 L_{\mu,\lambda}(\theta_t).$$

Hence, integrating the previous equation, we have

$$(\theta_1 - \theta_0) \cdot (\nabla L_{\mu,\lambda}(\theta_1) - \nabla L_{\mu,\lambda}(\theta_0)) \geq \underline{\phi}(\mathbf{t}^\mu(\theta_1 - \theta_0)) \|\theta_1 - \theta_0\|_{\mathbf{H}_{\mu,\lambda}(\theta_0)}^2.$$

Finally, bounding $(\theta_1 - \theta_0) \cdot (\nabla L_{\mu,\lambda}(\theta_1) - \nabla L_{\mu,\lambda}(\theta_0))$ by $\|\theta_1 - \theta_0\|_{\mathbf{H}_{\mu,\lambda}(\theta_0)} \|\nabla L_{\mu,\lambda}(\theta_1) - \nabla L_{\mu,\lambda}(\theta_0)\|_{\mathbf{H}_{\mu,\lambda}^{-1}(\theta_0)}$, and simplifying by $\|\theta_1 - \theta_0\|_{\mathbf{H}_{\mu,\lambda}(\theta_0)}$, we obtain Eq. (28).

Point 3. To prove Eq. (29), first write

$$\begin{aligned} \|\nabla L_{\mu,\lambda}(\theta_1) - \nabla L_{\mu,\lambda}(\theta_0)\|_{\mathbf{H}_{\mu,\lambda}^{-1}(\theta_0)} &= \left\| \int_0^1 \mathbf{H}_{\mu,\lambda}^{-1/2}(\theta_0) \mathbf{H}_{\mu,\lambda}(\theta_t) (\theta_1 - \theta_0) dt \right\| \\ &= \left\| \int_0^1 \mathbf{H}_{\mu,\lambda}^{-1/2}(\theta_0) \mathbf{H}_{\mu,\lambda}(\theta_t) \mathbf{H}_{\mu,\lambda}^{-1/2}(\theta_0) \mathbf{H}_{\mu,\lambda}^{1/2}(\theta_0) (\theta_1 - \theta_0) dt \right\| \\ &\leq \left(\int_0^1 \|\mathbf{H}_{\mu,\lambda}^{-1/2}(\theta_0) \mathbf{H}_{\mu,\lambda}(\theta_t) \mathbf{H}_{\mu,\lambda}^{-1/2}(\theta_0)\| dt \right) \|\theta_1 - \theta_0\|_{\mathbf{H}_{\mu,\lambda}(\theta_0)}. \end{aligned}$$

Then apply Eq. (27) to have

$$\forall t \in [0, 1], \mathbf{H}_{\mu,\lambda}(\theta_t) \preceq e^{t\mathfrak{t}^\mu(\theta_1 - \theta_0)} \mathbf{H}_{\mu,\lambda}(\theta_0).$$

This implies

$$\forall t \in [0, 1], \mathbf{H}_{\mu,\lambda}^{-1/2}(\theta_0) \mathbf{H}_{\mu,\lambda}(\theta_t) \mathbf{H}_{\mu,\lambda}^{-1/2}(\theta_0) \preceq e^{t\mathfrak{t}^\mu(\theta_1 - \theta_0)} I.$$

And hence in particular

$$\forall t \in [0, 1], \|\mathbf{H}_{\mu,\lambda}^{-1/2}(\theta_0) \mathbf{H}_{\mu,\lambda}(\theta_t) \mathbf{H}_{\mu,\lambda}^{-1/2}(\theta_0)\| \leq e^{t\mathfrak{t}^\mu(\theta_1 - \theta_0)}.$$

Finally, integrating this, we get

$$\int_0^1 \|\mathbf{H}_{\mu,\lambda}^{-1/2}(\theta_0) \mathbf{H}_{\mu,\lambda}(\theta_t) \mathbf{H}_{\mu,\lambda}^{-1/2}(\theta_0)\| dt \leq \bar{\phi}(\mathfrak{t}^\mu(\theta_1 - \theta_0)).$$

Thus Eq. (29) is proved.

Point 4. To prove Eq. (30), define $\forall t \in \mathbb{R}$, $\varphi(t) = L_{\mu,\lambda}(\theta_t) - L_{\mu,\lambda}(\theta_0) - t \nabla L_{\mu,\lambda}(\theta_0)(\theta_1 - \theta_0)$. We have $\varphi''(t) = \|\theta_1 - \theta_0\|_{\mathbf{H}_{\mu,\lambda}(\theta_t)}^2 \leq e^{t\mathfrak{t}^\mu(\theta_1 - \theta_0)} \varphi''(0)$. Then using the fact that $\varphi(0), \varphi'(0) = 0$ and integrating this inequality two times, we get the result.

Proof of Proposition 1. First note that since the support of μ is included in the support of ρ , Assumption 3 and Assumption 4 also hold for μ . Hence, since Assumptions 2 to 4 are satisfied, by Proposition 8, $L_{\mu,\lambda}$, $\nabla L_{\mu,\lambda}$ and $\nabla^2 L_{\mu,\lambda}$ are well-defined.

Assuming the existence of a minimizer $\theta_{\mu,\lambda}^*$ of $L_{\mu,\lambda}$, the reported equations are the same than those of Proposition 4 when taking $\theta_1 = \theta$ and $\theta_0 = \theta_{\mu,\lambda}^*$, with the fact that $\mathfrak{t}^\mu(v) \leq \mathfrak{t}(v)$ for any $v \in \mathcal{H}$ since the support of μ is a subset of the support of ρ , and $\nabla L_{\mu,\lambda}(\theta_{\mu,\lambda}^*) = 0$. Note that since $L_{\mu,\lambda}$ is defined on \mathcal{H} , if $\lambda > 0$, then $\theta_{\mu,\lambda}^*$ always exists and is unique by strong convexity.

B.2 Localization properties for \mathfrak{t}_λ

The aim of this section is to localize the optima θ_λ^* and $\hat{\theta}_\lambda^*$ using the re-normalized gradient. This type of result is inspired by Proposition 2 of (Bach, 2010) or Proposition 3.5 of (Ostrovskii and Bach, 2018). However, their proof is based on a slightly different result, namely Eq. (28), and its formulation is slightly different. Indeed, while the two propositions mentioned above concentrate on performing a quadratic approximation directly, we bound the term that could have been too large in that quadratic approximation.

Proposition 5 (localisation). *Let $\theta \in \mathcal{H}$, then the following holds*

$$\|\nabla L_\lambda(\theta)\|_{\mathbf{H}_\lambda^{-1}(\theta)} \leq \frac{r_\lambda(\theta)}{2} \implies \mathfrak{t}(\theta - \theta_\lambda^*) = \mathfrak{t}_\lambda \leq \log 2, \quad (32)$$

$$\|\widehat{\nabla L}_\lambda(\theta)\|_{\mathbf{H}_\lambda^{-1}(\theta)} \|\widehat{\mathbf{H}}_\lambda^{-1/2}(\theta) \mathbf{H}_\lambda^{1/2}(\theta)\|^2 \leq \frac{r_\lambda(\theta)}{2} \implies \mathfrak{t}(\theta - \hat{\theta}_\lambda^*) \leq \log 2. \quad (33)$$

Proof. To prove Eq. (32), we first write

$$\mathfrak{t}(\theta - \theta_\lambda^*) = \sup_{z \in \text{supp}(\rho)} \sup_{g \in \varphi(z)} |(\theta - \theta_\lambda^*) \cdot g| \leq \|\theta - \theta_\lambda^*\|_{\mathbf{H}_\lambda(\theta)} \sup_{z \in \text{supp}(\rho)} \sup_{g \in \varphi(z)} \|g\|_{\mathbf{H}_\lambda^{-1}(\theta)}.$$

Now we use Eq. (14) to bound $\|\theta - \theta_\lambda^*\|_{\mathbf{H}_\lambda(\theta)}$, and putting things together, we get

$$\mathfrak{t}(\theta - \theta_\lambda^*) \underline{\phi}(\mathfrak{t}(\theta - \theta_\lambda^*)) \leq \frac{\|\nabla L_\lambda(\theta)\|_{\mathbf{H}_\lambda^{-1}(\theta)}}{r_\lambda(\theta)}.$$

Using the fact that $\underline{\phi}(t) = 1 - e^{-t}$ is an increasing function, we see that if $\underline{\phi}(t) \leq 1/2$, then $t \leq \log 2$ hence the result.

To prove Eq. (33), we use the same reasoning. First, we bound

$$\mathfrak{t}(\theta - \widehat{\theta}_\lambda^*) = \sup_{z \in \text{supp}(\rho)} \sup_{g \in \varphi(z)} |(\theta - \widehat{\theta}_\lambda^*) \cdot g| \leq \|\theta - \widehat{\theta}_\lambda^*\|_{\widehat{\mathbf{H}}_\lambda(\theta)} \|\widehat{\mathbf{H}}_\lambda^{-1/2}(\theta) \mathbf{H}_\lambda^{1/2}(\theta)\| \sup_{z \in \text{supp}(\rho)} \sup_{g \in \varphi(z)} \|g\|_{\mathbf{H}_\lambda^{-1}(\theta)}.$$

Now using Eq. (14) to the function \widehat{L}_λ , we get

$$\mathfrak{t}(\theta - \widehat{\theta}_\lambda^*) \underline{\phi}(\mathfrak{t}^\rho(\theta - \widehat{\theta}_\lambda^*)) \leq \|\nabla \widehat{L}_\lambda(\theta)\|_{\widehat{\mathbf{H}}_\lambda^{-1}(\theta)} \|\widehat{\mathbf{H}}_\lambda^{-1/2}(\theta) \mathbf{H}_\lambda^{1/2}(\theta)\| \frac{1}{r_\lambda(\theta)}.$$

Now using the fact that $\mathfrak{t}^\rho(\theta - \widehat{\theta}_\lambda^*) \leq \mathfrak{t}(\theta - \widehat{\theta}_\lambda^*)$ and that $\underline{\phi}$ is a decreasing function, and that $\|\nabla \widehat{L}_\lambda(\theta)\|_{\widehat{\mathbf{H}}_\lambda^{-1}(\theta)} \leq \|\widehat{\mathbf{H}}_\lambda^{-1/2}(\theta) \mathbf{H}_\lambda^{1/2}(\theta)\| \|\nabla \widehat{L}_\lambda(\theta)\|_{\mathbf{H}_\lambda^{-1}(\theta)}$, this yields

$$\mathfrak{t}(\theta - \widehat{\theta}_\lambda^*) \underline{\phi}(\mathfrak{t}(\theta - \widehat{\theta}_\lambda^*)) \leq \|\nabla \widehat{L}_\lambda(\theta)\|_{\mathbf{H}_\lambda^{-1}(\theta)} \|\widehat{\mathbf{H}}_\lambda^{-1/2}(\theta) \mathbf{H}_\lambda^{1/2}(\theta)\|^2 \frac{1}{r_\lambda(\theta)}.$$

We conclude using the same argument as before.

C Main result, simplified

In this section, we perform a simplified analysis in the case where we assume nothing on Bias_λ more than just the fact that θ^* exists. In this section we assume that ℓ_z and ρ satisfy Assumptions 3 to 5 and 8.

Definition 4 (Definition of $\overline{\mathbf{B}}_1$, $\overline{\mathbf{B}}_2$ and $\overline{\text{df}}_\lambda$). *Under assumptions Assumptions 3 to 5 and 8, the following quantities are well-defined and real-valued.*

$$\overline{\mathbf{B}}_1 = \sup_{\|\theta\| \leq \|\theta^*\|} \mathbf{B}_1(\theta) \quad \overline{\mathbf{B}}_2 = \sup_{\|\theta\| \leq \|\theta^*\|} \mathbf{B}_2(\theta), \quad \overline{\text{df}}_\lambda = \mathbb{E} \left[\|\nabla \ell_z(\theta_\lambda^*)\|_{\mathbf{H}_\lambda^{-1}(\theta_\lambda^*)}^2 \right].$$

Proposition 6. *The quantities in Definition 4 are finite and moreover*

$$\overline{\text{df}}_\lambda \leq \frac{\overline{\mathbf{B}}_1^2}{\lambda}.$$

Proof. These are well defined thanks to Lemmas 11 and 12.

Definition 5 (Constants). *In this section, we will use the following constants.*

$$K_{\text{var}} = \frac{1 + \psi(\log 2)}{\underline{\phi}(\log 2)^2} \leq 4, \quad \Delta = 2\sqrt{2} \left(1 + \frac{1}{2\sqrt{3}} \right) \leq 4,$$

$$C_{\text{bias}} = 1 + \frac{K_{\text{var}}}{8} \leq 2, \quad C_{\text{var}} = 2K_{\text{var}}\Delta^2 \leq 84.$$

C.1 Analytic results

Theorem 5 (Analytic decomposition). *For any $\lambda > 0$ and $n \in \mathbb{N}$, if $\frac{R}{\sqrt{\lambda}} \widehat{\text{Var}}_\lambda \leq \frac{1}{2}$,*

$$L(\widehat{\theta}_\lambda^*) - L(\theta^*) \leq K_{\text{var}} \widehat{\text{Var}}_\lambda^2 + \lambda \|\theta^*\|^2, \quad (34)$$

where K_{var} is defined in Definition 5.

Proof.

First decompose the excess risk of $\widehat{\theta}_\lambda^*$ in the following way:

$$L(\widehat{\theta}_\lambda^*) - L(\theta^*) = \underbrace{L_\lambda(\widehat{\theta}_\lambda^*) - L_\lambda(\theta_\lambda^*)}_{\text{variance}} + \underbrace{L(\theta_\lambda^*) - L(\theta^*)}_{\text{bias}} + \underbrace{\frac{\lambda}{2} \left(\|\theta_\lambda^*\|^2 - \|\widehat{\theta}_\lambda^*\|^2 \right)}_{\text{mixed}}.$$

1) Variance term: For the variance term, use Eq. (13)

$$L_\lambda(\widehat{\theta}_\lambda^*) - L_\lambda(\theta_\lambda^*) \leq \psi \left(\mathfrak{t}(\theta_\lambda^* - \widehat{\theta}_\lambda^*) \right) \|\widehat{\theta}_\lambda^* - \theta_\lambda^*\|_{\mathbf{H}_\lambda(\theta_\lambda^*)}^2.$$

2) Bias term: For the bias term, note that since $\|\theta_\lambda^*\| \leq \|\theta^*\|$,

$$L(\theta_\lambda^*) - L(\theta^*) = L_\lambda(\theta_\lambda^*) - L_\lambda(\theta^*) + \frac{\lambda}{2} \|\theta^*\|^2 - \frac{\lambda}{2} \|\theta_\lambda^*\|^2 \leq \frac{\lambda}{2} \|\theta^*\|^2.$$

3) Mixed term: For the mixed term, since $\|\theta_\lambda^*\|_{\mathbf{H}_\lambda(\theta_\lambda^*)}^{-1} \leq \|\mathbf{H}_\lambda(\theta_\lambda^*)\|^{-1/2} \|\theta_\lambda^*\| \leq \lambda^{-1/2} \|\theta_\lambda^*\| \leq \lambda^{-1/2} \|\theta^*\|$, we have

$$\begin{aligned} \frac{\lambda}{2} \left(\|\theta_\lambda^*\|^2 - \|\widehat{\theta}_\lambda^*\|^2 \right) &= \frac{\lambda}{2} \left(\theta_\lambda^* - \widehat{\theta}_\lambda^* \right) \cdot \left(\theta_\lambda^* + \widehat{\theta}_\lambda^* \right) \\ &\leq \frac{\lambda}{2} \|\theta_\lambda^* - \widehat{\theta}_\lambda^*\|_{\mathbf{H}_\lambda(\theta_\lambda^*)} \left(\|\widehat{\theta}_\lambda^* - \theta_\lambda^*\|_{\mathbf{H}_\lambda(\theta_\lambda^*)}^{-1} + 2\|\theta_\lambda^*\|_{\mathbf{H}_\lambda(\theta_\lambda^*)}^{-1} \right) \\ &\leq \frac{1}{2} \|\theta_\lambda^* - \widehat{\theta}_\lambda^*\|_{\mathbf{H}_\lambda(\theta_\lambda^*)}^2 + \sqrt{\lambda} \|\theta^*\| \|\theta_\lambda^* - \widehat{\theta}_\lambda^*\|_{\mathbf{H}_\lambda(\theta_\lambda^*)} \\ &\leq \|\theta_\lambda^* - \widehat{\theta}_\lambda^*\|_{\mathbf{H}_\lambda(\theta_\lambda^*)}^2 + \frac{\lambda}{2} \|\theta^*\|^2. \end{aligned}$$

where we get the last inequality by using $ab \leq \frac{a^2}{2} + \frac{b^2}{2}$.

4) Putting things together

$$L(\widehat{\theta}_\lambda^*) - L(\theta^*) \leq \left(1 + \psi \left(\mathfrak{t}(\theta_\lambda^* - \widehat{\theta}_\lambda^*) \right) \right) \|\theta_\lambda^* - \widehat{\theta}_\lambda^*\|_{\mathbf{H}_\lambda(\theta_\lambda^*)}^2 + \lambda \|\theta^*\|^2.$$

By using Eq. (14) we have

$$\begin{aligned} \|\theta_\lambda^* - \widehat{\theta}_\lambda^*\|_{\mathbf{H}_\lambda(\theta_\lambda^*)} &\leq \|\mathbf{H}_\lambda^{1/2}(\theta_\lambda^*) \widehat{\mathbf{H}}_\lambda^{-1/2}(\theta_\lambda^*)\| \|\theta_\lambda^* - \widehat{\theta}_\lambda^*\|_{\widehat{\mathbf{H}}_\lambda(\theta_\lambda^*)} \\ &\leq \frac{1}{\underline{\phi} \left(\mathfrak{t}^\rho(\theta_\lambda^* - \widehat{\theta}_\lambda^*) \right)} \|\mathbf{H}_\lambda^{1/2}(\theta_\lambda^*) \widehat{\mathbf{H}}_\lambda^{-1/2}(\theta_\lambda^*)\| \|\nabla \widehat{L}_\lambda(\theta_\lambda^*)\|_{\widehat{\mathbf{H}}_\lambda^{-1}(\theta_\lambda^*)}. \end{aligned}$$

Note that by multiplying and dividing for $\mathbf{H}_\lambda^{1/2}(\theta_\lambda^*)$,

$$\begin{aligned} \|\nabla \widehat{L}_\lambda(\theta_\lambda^*)\|_{\widehat{\mathbf{H}}_\lambda^{-1}(\theta_\lambda^*)} &= \|\widehat{\mathbf{H}}_\lambda^{-1/2}(\theta_\lambda^*) \nabla \widehat{L}_\lambda(\theta_\lambda^*)\| = \|\widehat{\mathbf{H}}_\lambda^{-1/2}(\theta_\lambda^*) \mathbf{H}_\lambda^{-1/2}(\theta_\lambda^*) \mathbf{H}_\lambda^{1/2}(\theta_\lambda^*) \nabla \widehat{L}_\lambda(\theta_\lambda^*)\| \\ &\leq \|\widehat{\mathbf{H}}_\lambda^{-1/2}(\theta_\lambda^*) \mathbf{H}_\lambda^{-1/2}(\theta_\lambda^*)\| \|\mathbf{H}_\lambda^{1/2}(\theta_\lambda^*) \nabla \widehat{L}_\lambda(\theta_\lambda^*)\| \\ &= \|\widehat{\mathbf{H}}_\lambda^{-1/2}(\theta_\lambda^*) \mathbf{H}_\lambda^{-1/2}(\theta_\lambda^*)\| \|\nabla \widehat{L}_\lambda(\theta_\lambda^*)\|_{\mathbf{H}_\lambda^{-1}(\theta_\lambda^*)}. \end{aligned}$$

Then,

$$\begin{aligned} \|\theta_\lambda^* - \widehat{\theta}_\lambda^*\|_{\mathbf{H}_\lambda(\theta_\lambda^*)} &\leq \frac{1}{\underline{\phi}\left(\mathfrak{t}(\theta_\lambda^* - \widehat{\theta}_\lambda^*)\right)} \|\mathbf{H}_\lambda^{1/2}(\theta_\lambda^*) \widehat{\mathbf{H}}_\lambda^{-1/2}(\theta_\lambda^*)\|^2 \|\nabla \widehat{L}_\lambda(\theta_\lambda^*)\|_{\mathbf{H}_\lambda^{-1}(\theta_\lambda^*)} \\ &= \frac{1}{\underline{\phi}\left(\mathfrak{t}(\theta_\lambda^* - \widehat{\theta}_\lambda^*)\right)} \widehat{\text{Var}}_\lambda. \end{aligned}$$

Now we know that using Eq. (33), if $\widehat{\text{Var}}_\lambda \leq \frac{r_\lambda(\theta_\lambda^*)}{2}$, then $\mathfrak{t}(\theta_\lambda^* - \widehat{\theta}_\lambda^*) \leq \log 2$, which yields the following bound:

$$L(\widehat{\theta}_\lambda^*) - L(\theta^*) \leq \frac{(1 + \psi(\log 2))}{\underline{\phi}(\log 2)^2} \widehat{\text{Var}}_\lambda + \lambda \|\theta^*\|^2.$$

Finally, we can bound $\frac{1}{r_\lambda(\theta_\lambda^*)} \leq \frac{R}{\lambda^{1/2}}$ to have the final form of the proposition.

C.2 Probabilistic results

Lemma 2 (bounding $\|\nabla \widehat{L}_\lambda(\theta_\lambda^*)\|_{\mathbf{H}_\lambda(\theta_\lambda^*)}$). *Let $n \in \mathbb{N}$, $\lambda > 0$ and $\delta \in (0, 1]$. For $k \geq 1$, if*

$$n \geq 24 \frac{\overline{\mathbf{B}}_2}{\lambda} \log \frac{2}{\delta}, n \geq k^2 2 \log \frac{2}{\delta},$$

then with probability at least $1 - \delta$,

$$\|\mathbf{H}_\lambda(\theta_\lambda^*)^{-1/2} \nabla \widehat{L}_\lambda(\theta_\lambda^*)\| \leq \Delta/2 \sqrt{\frac{\overline{\text{df}}_\lambda \vee (\overline{\mathbf{B}}_1^2 / \overline{\mathbf{B}}_2) \log \frac{2}{\delta}}{n}} + \frac{2}{k} \sqrt{\lambda} \|\theta^*\|$$

where Δ is defined in Definition 5.

Proof.1 First use Bernstein inequality for random vectors (e.g. Thm. 3.3.4 of [Yurinsky, 1995](#)): for any $n \in \mathbb{N}$ and $\delta \in (0, 1]$, with probability at least $1 - \delta$, we have

$$\|\mathbf{H}_\lambda(\theta_\lambda^*)^{-1/2} \nabla \widehat{L}_\lambda(\theta_\lambda^*)\| \leq \frac{2M \log \frac{2}{\delta}}{n} + \sigma \sqrt{\frac{2 \log \frac{2}{\delta}}{n}},$$

where $M = \sup_{z \in \text{supp}(\rho)} \|\nabla \ell_z^\lambda(\theta_\lambda^*)\|_{\mathbf{H}_\lambda^{-1}(\theta_\lambda^*)}$ and $\sigma = \mathbb{E} \left[\|\nabla \ell_z^\lambda(\theta_\lambda^*)\|_{\mathbf{H}_\lambda^{-1}(\theta_\lambda^*)}^2 \right]^{1/2}$.

2) Using the fact that $\nabla \ell_z^\lambda(\theta_\lambda^*) = \nabla \ell_z(\theta_\lambda^*) + \lambda \theta_\lambda^*$, we bound M as follows:

$$M = \sup_{z \in \text{supp}(\rho)} \|\nabla \ell_z^\lambda(\theta_\lambda^*)\|_{\mathbf{H}_\lambda(\theta_\lambda^*)} \leq \sup_{z \in \text{supp}(\rho)} \|\nabla \ell_z(\theta_\lambda^*)\|_{\mathbf{H}_\lambda(\theta_\lambda^*)} + \lambda \|\theta_\lambda^*\|_{\mathbf{H}_\lambda^{-1}(\theta_\lambda^*)} \leq \frac{\overline{\mathbf{B}}_1}{\sqrt{\lambda}} + \sqrt{\lambda} \|\theta^*\|,$$

where in the last inequality, we use the fact that $\|\theta_\lambda^*\|_{\mathbf{H}_\lambda^{-1}(\theta_\lambda^*)} \leq \frac{1}{\sqrt{\lambda}} \|\theta^*\| \leq \frac{1}{\sqrt{\lambda}} \|\theta^*\|$. Similarly, we bound σ

$$\sigma \leq \mathbb{E} \left[\|\nabla \ell_z(\theta_\lambda^*)\|_{\mathbf{H}_\lambda^{-1}(\theta_\lambda^*)}^2 \right]^{1/2} + \lambda \|\theta_\lambda^*\|_{\mathbf{H}_\lambda^{-1}(\theta_\lambda^*)} \leq \overline{\text{df}}_\lambda^{1/2} + \sqrt{\lambda} \|\theta^*\|.$$

3) Injecting these bounds in the concentration inequality,

$$\begin{aligned} \|\mathbf{H}_\lambda(\theta_\lambda^*)^{-1/2} \nabla \widehat{L}_\lambda(\theta_\lambda^*)\| &\leq \sqrt{\frac{2\overline{\mathbf{B}}_2 \log \frac{2}{\delta}}{\lambda n}} \sqrt{\frac{2(\overline{\mathbf{B}}_1^2 / \overline{\mathbf{B}}_2) \log \frac{2}{\delta}}{n}} + \sqrt{\frac{2\overline{\text{df}}_\lambda \log \frac{2}{\delta}}{n}} \\ &\quad + \sqrt{\lambda} \|\theta^*\| \left(\frac{2 \log \frac{2}{\delta}}{n} + \sqrt{\frac{2 \log \frac{2}{\delta}}{n}} \right), \end{aligned}$$

where we have decomposed $\frac{2\bar{\mathbf{B}}_1^2 \log \frac{2}{\delta}}{\sqrt{\lambda n}} = \sqrt{\frac{2\bar{\mathbf{B}}_2 \log \frac{2}{\delta}}{\lambda n}} \sqrt{\frac{2(\bar{\mathbf{B}}_1^2/\bar{\mathbf{B}}_2) \log \frac{2}{\delta}}{n}}$ for the first term. Reordering the terms, this yields

$$\begin{aligned} \|\mathbf{H}_\lambda(\theta_\lambda^*)^{-1/2} \nabla \widehat{L}_\lambda(\theta_\lambda^*)\| &\leq \left(1 + \sqrt{\frac{2\bar{\mathbf{B}}_2 \log \frac{2}{\delta}}{\lambda n}}\right) \sqrt{\frac{2\overline{\text{df}}_\lambda \vee (\bar{\mathbf{B}}_1^2/\bar{\mathbf{B}}_2) \log \frac{2}{\delta}}{n}} \\ &\quad + \sqrt{\lambda} \|\theta^*\| \left(\frac{2 \log \frac{2}{\delta}}{n} + \sqrt{\frac{2 \log \frac{2}{\delta}}{n}}\right). \end{aligned}$$

4) Now assuming that

$$n \geq 24 \frac{\bar{\mathbf{B}}_2}{\lambda} \log \frac{2}{\delta}, n \geq k^2 2 \log \frac{2}{\delta},$$

this yields

$$\|\mathbf{H}_\lambda(\theta_\lambda^*)^{-1/2} \nabla \widehat{L}_\lambda(\theta_\lambda^*)\| \leq \left(1 + \frac{1}{2\sqrt{3}}\right) \sqrt{\frac{2\overline{\text{df}}_\lambda \vee (\bar{\mathbf{B}}_1^2/\bar{\mathbf{B}}_2) \log \frac{2}{\delta}}{n}} + \frac{2}{k} \sqrt{\lambda} \|\theta^*\|.$$

Combining the two previous lemmas, we get:

Lemma 3 (Bounding $\widehat{\text{Var}}_\lambda$). *Let $n \in \mathbb{N}$ and $0 < \lambda \leq \bar{\mathbf{B}}_2$. Let $\delta \in (0, 1]$. If for $k \geq 1$*

$$n \geq 24 \frac{\bar{\mathbf{B}}_2}{\lambda} \log \frac{8\bar{\mathbf{B}}_2}{\lambda\delta}, \quad n \geq 2k^2 \log \frac{2}{\delta},$$

then with probability at least $1 - 2\delta$,

$$\widehat{\text{Var}}_\lambda \leq \Delta \sqrt{\frac{\overline{\text{df}}_\lambda \vee (\bar{\mathbf{B}}_1^2/\bar{\mathbf{B}}_2) \log \frac{2}{\delta}}{n}} + \frac{4}{k} \sqrt{\lambda} \|\theta^*\|,$$

where Δ is a constant defined in Definition 5.

Proof. Recall that $\widehat{\text{Var}}_\lambda = \|\mathbf{H}_\lambda^{1/2}(\theta_\lambda^*) \widehat{\mathbf{H}}_\lambda^{-1/2}(\theta_\lambda^*)\|^2 \|\nabla \widehat{L}_\lambda(\theta_\lambda^*)\|_{\widehat{\mathbf{H}}_\lambda^{-1}(\theta_\lambda^*)}$. Using Lemma 6, under the conditions of this lemma, we have $\|\mathbf{H}_\lambda^{1/2}(\theta_\lambda^*) \widehat{\mathbf{H}}_\lambda^{-1/2}(\theta_\lambda^*)\|^2 \leq 2$. Combining this with the bound for $\|\nabla \widehat{L}_\lambda(\theta_\lambda^*)\|_{\widehat{\mathbf{H}}_\lambda^{-1}(\theta_\lambda^*)}$ obtained in Lemma 2, we get the result (the probability $1 - 2\delta$ comes from the fact that we perform a union bound).

C.3 Final result

Theorem 6 (General bound, simplified setting). *Let $n \in \mathbb{N}$ and $0 < \lambda \leq \bar{\mathbf{B}}_2$. Let $\delta \in (0, 1]$. If*

$$n \geq 512 (\|\theta^*\|^2 R^2 \vee 1) \log \frac{2}{\delta}, \quad n \geq 24 \frac{\bar{\mathbf{B}}_2}{\lambda} \log \frac{8\bar{\mathbf{B}}_2}{\lambda\delta}, \quad n \geq 16\Delta^2 R^2 \frac{\overline{\text{df}}_\lambda \vee (\bar{\mathbf{B}}_1^2/\bar{\mathbf{B}}_2)}{\lambda} \log \frac{2}{\delta},$$

then with probability at least $1 - 2\delta$,

$$L(\widehat{\theta}_\lambda^*) - L(\theta^*) \leq C_{\text{var}} \frac{\overline{\text{df}}_\lambda \vee (\bar{\mathbf{B}}_1^2/\bar{\mathbf{B}}_2)}{n} \log \frac{2}{\delta} + C_{\text{bias}} \lambda \|\theta^*\|^2,$$

where $\Delta, C_{\text{bias}}, C_{\text{var}}$ are defined in Definition 5.

Proof.1) Recall the analytical decomposition in Thm. 5. For any $\lambda > 0$ and $n \in \mathbb{N}$, if $\frac{R}{\sqrt{\lambda}} \widehat{\text{Var}}_\lambda \leq \frac{1}{2}$,

$$L(\widehat{\theta}_\lambda^*) - L(\theta^*) \leq K_{\text{var}} \widehat{\text{Var}}_\lambda^2 + \lambda \|\theta^*\|^2,$$

where K_{var} is defined in Definition 5.

2) Now apply Lemma 3 for a given $k \geq 1$. If

$$n \geq 24 \frac{\overline{\text{B}}_2}{\lambda} \log \frac{8\overline{\text{B}}_2}{\lambda\delta}, \quad n \geq 2k^2 \log \frac{2}{\delta},$$

then with probability at least $1 - 2\delta$,

$$\widehat{\text{Var}}_\lambda \leq \Delta \sqrt{\frac{\overline{\text{df}}_\lambda \vee (\overline{\text{B}}_1^2 / \overline{\text{B}}_2) \log \frac{2}{\delta}}{n}} + \frac{4}{k} \sqrt{\lambda} \|\theta^*\|,$$

where Δ is a constant defined in Definition 5.

In order to satisfy the condition to have the analytical decomposition, namely $\frac{R}{\sqrt{\lambda}} \widehat{\text{Var}}_\lambda \leq \frac{1}{2}$, it is therefore sufficient to have

$$\Delta R \sqrt{\frac{\overline{\text{df}}_\lambda \vee (\overline{\text{B}}_1^2 / \overline{\text{B}}_2) \log \frac{2}{\delta}}{\lambda n}} \leq \frac{1}{4}, \quad \frac{4}{k} R \|\theta^*\| \leq \frac{1}{4}.$$

3) Thus, if we choose $k = 16(R\|\theta^*\| \vee 1)$, we have both $k \geq 1$ and the second condition in the previous equation. Moreover, the condition $n \geq 2k^2 \log \frac{2}{\delta}$ becomes $n \geq 512(R^2\|\theta^*\|^2 \vee 1) \log \frac{2}{\delta}$. Hence, under the conditions of this theorem, we can apply the analytical decomposition :

$$L(\widehat{\theta}_\lambda^*) - L(\theta^*) \leq K_{\text{var}} \widehat{\text{Var}}_\lambda^2 + \lambda \|\theta^*\|^2 \leq 2K_{\text{var}} \Delta^2 \frac{\overline{\text{df}}_\lambda \vee (\overline{\text{B}}_1^2 / \overline{\text{B}}_2) \log \frac{2}{\delta}}{n} + \left(1 + K_{\text{var}} \frac{32}{k^2}\right) \lambda \|\theta^*\|^2.$$

In the last inequality, we have used $(a + b)^2 \leq 2a^2 + 2b^2$ to separate the terms coming from $\widehat{\text{Var}}_\lambda^2$.

Finally, using the fact that $k \geq 16$ and hence that $\frac{32}{k^2} \leq \frac{1}{8}$, we get the constants in the theorem.

Proof.of Thm. 2 Since $\forall \lambda > 0$, $\overline{\text{df}}_\lambda \leq \frac{\overline{\text{B}}_1^2}{\lambda}$, and since $\lambda \leq \overline{\text{B}}_2$, $\overline{\text{df}}_\lambda \vee \overline{\text{B}}_1^2 / \overline{\text{B}}_2 \leq \frac{\overline{\text{B}}_1^2}{\lambda}$. From Definition 5, we get that $\Delta \leq 4$, $C_{\text{bias}} \leq 2$, $C_{\text{var}} \leq 84$. Thus, we can use these bounds in Thm. 6 to obtain the result.

D Main result, refined analysis

In subsection Appendix D.1 we split the excess risk in terms of bias and variance, that will be controlled in Appendix D.3, the final result is Thm. 4 in Appendix D.4, while in Appendix F a version with explicit dependence in λ, n is reported.

Constants First, we introduce three constants that will be crucial for the final bound.

Definition 6.

$$\text{B}_1^* = \text{B}_1(\theta^*), \quad \text{B}_2^* = \text{B}_2(\theta^*), \quad \text{Q}^* = \text{B}_1^* / \sqrt{\text{B}_2^*}.$$

In the following sections, we also will use the following functions of t_λ and $\widetilde{\text{t}}_\lambda$ which we will treat as constants (see Proposition 7).

Definition 7.

$$\begin{aligned}
K_{\text{bias}}(\mathbf{t}_\lambda) &= 2 \frac{\psi(\mathbf{t}_\lambda + \log 2)}{\underline{\phi}(\mathbf{t}_\lambda)^2} \leq 2e^{3\mathbf{t}_\lambda}, & K_{\text{var}}(\mathbf{t}_\lambda) &= 2 \frac{\psi(\mathbf{t}_\lambda + \log 2)e^{\mathbf{t}_\lambda}}{\underline{\phi}(\log 2)^2} \leq 8e^{2\mathbf{t}_\lambda}, \\
\Box_1(\mathbf{t}_\lambda) &= e^{\mathbf{t}_\lambda/2}, & \Box_2(\mathbf{t}_\lambda) &= e^{\mathbf{t}_\lambda/2} (1 + e^{\mathbf{t}_\lambda}) \leq 2e^{3\mathbf{t}_\lambda/2} \\
C_{\text{bias}} &= \psi(\mathbf{t}_\lambda + \log 2) \left(\frac{2}{\underline{\phi}(\mathbf{t}_\lambda)} + \frac{e^{\mathbf{t}_\lambda}}{\underline{\phi}(\log 2)^2} \right) \leq 6e^{2\mathbf{t}_\lambda}, & C_{\text{var}} &= \frac{64\psi(\mathbf{t}_\lambda + \log 2)e^{2\mathbf{t}_\lambda}}{\underline{\phi}(\log 2)^2} \leq 256e^{3\mathbf{t}_\lambda} \\
\Delta_1 &= 576\Box_1^2\Box_2^2(1/2 \vee \tilde{\mathbf{t}}_\lambda)^2 \leq 2304e^{4\mathbf{t}_\lambda}(\tilde{\mathbf{t}}_\lambda \vee 1/2)^2, & \Delta_2 &= 256\Box_1^4 \leq 256e^{2\mathbf{t}_\lambda}.
\end{aligned}$$

Note that these functions are all increasing in \mathbf{t}_λ and $\tilde{\mathbf{t}}_\lambda$, and are lower bounded by strictly positive constants.

For the second bounds, we use the fact that $\psi(t) \leq \frac{e^t}{2}$ and $1/\underline{\phi}(t) \leq e^t$ to bound all the quantities using only exponentials of \mathbf{t}_λ .

A priori, these constants will depend on λ . However, we can always bound \mathbf{t}_λ and $\tilde{\mathbf{t}}_\lambda$ in the following way.

Lemma 4. *Recall the definitions of $\mathbf{t}_\lambda := \mathbf{t}(\theta_\lambda^* - \theta^*)$ and $\tilde{\mathbf{t}}_\lambda := \frac{\text{Bias}_\lambda}{r_\lambda(\theta^*)}$. We have the following cases.*

- If $\tilde{\mathbf{t}}_\lambda \leq \frac{1}{2}$, then $\mathbf{t}_\lambda \leq \log 2$,
- else, $\tilde{\mathbf{t}}_\lambda \leq R\|\theta^*\|$ and $\mathbf{t}_\lambda \leq 2R\|\theta^*\|$.

Proof. The first point is a direct application of Eq. (32). One can obtain the second by noting that $\mathbf{t}(\theta_\lambda^* - \theta^*) \leq R\|\theta_\lambda^* - \theta^*\|$. Since $\|\theta_\lambda^*\| \leq \|\theta^*\|$, we have the bound on \mathbf{t}_λ . For the bound on $\tilde{\mathbf{t}}_\lambda$, since $\text{Bias}_\lambda \leq \sqrt{\lambda}\|\theta^*\|$ and $\frac{1}{r_\lambda(\theta^*)} \leq \frac{R}{\sqrt{\lambda}}$, we have the wanted bound.

Hence, we can always bound the constants in Definition 7 by constants independent of λ .

Proposition 7. *If $\tilde{\mathbf{t}}_\lambda \leq 1/2$, then $\mathbf{t}_\lambda \leq \log 2$ and*

$$\begin{aligned}
K_{\text{bias}}(\mathbf{t}_\lambda) &\leq 4, & K_{\text{var}}(\mathbf{t}_\lambda) &\leq 7, & \Box_1(\mathbf{t}_\lambda) &\leq 2, & \Box_2(\mathbf{t}_\lambda) &\leq 5 \\
\Delta_1(\mathbf{t}_\lambda, \tilde{\mathbf{t}}_\lambda) &\leq 5184, & \Delta_2(\mathbf{t}_\lambda) &\leq 1024, & C_{\text{bias}} &\leq 6, & C_{\text{var}} &\leq 414.
\end{aligned}$$

Else,

$$\begin{aligned}
K_{\text{bias}}(\mathbf{t}_\lambda) &\leq 2e^{6R\|\theta^*\|}, & K_{\text{var}}(\mathbf{t}_\lambda) &\leq 8e^{4R\|\theta^*\|}, & \Box_1(\mathbf{t}_\lambda) &\leq e^{R\|\theta^*\|}, \\
\Box_2(\mathbf{t}_\lambda) &\leq 2e^{3R\|\theta^*\|}, & \Delta_1(\mathbf{t}_\lambda, \tilde{\mathbf{t}}_\lambda) &\leq 2304(R\|\theta^*\|)^2 e^{8R\|\theta^*\|}, & \Delta_2(\mathbf{t}_\lambda) &\leq 256e^{4R\|\theta^*\|}, \\
C_{\text{bias}} &\leq 6e^{4R\|\theta^*\|}, & C_{\text{var}} &\leq 256e^{6R\|\theta^*\|}.
\end{aligned}$$

Proof. For the first bound, we use the fact that $\mathbf{t}_\lambda \leq \log 2$ and plug that in the expressions above as these functions are increasing in \mathbf{t}_λ . We compute them numerically from the definition.

For the second set of bounds, we simply inject the bounds for \mathbf{t}_λ and $\tilde{\mathbf{t}}_\lambda$ in the second bounds of Definition 7.

D.1 Analytic decomposition of the risk

In this section, we make use of self-concordance to control certain quantities required to control the variance, with respect to our main quantities Bias_λ , r_λ and df_λ . The excess risk has been already decomposed in Sec. 6.

Theorem 7 (Analytic decomposition). *Let $\lambda > 0$ and K_{bias} and K_{var} be the increasing functions of t_λ described in Eq. (37). When $\widehat{\text{Var}}_\lambda \leq r_\lambda(\theta_\lambda^*)/2$, then*

$$L(\widehat{\theta}_\lambda^*) - L(\theta^*) \leq K_{\text{bias}}(t_\lambda) \text{Bias}_\lambda^2 + K_{\text{var}}(t_\lambda) \widehat{\text{Var}}_\lambda^2. \quad (35)$$

Moreover $K_{\text{bias}}(t_\lambda), K_{\text{var}}(t_\lambda) \leq 7$ if $\text{Bias}_\lambda \leq \frac{1}{2}r_\lambda(\theta^*)$, otherwise $K_{\text{bias}}(t_\lambda), K_{\text{var}}(t_\lambda) \leq 8e^{6\|\theta^*\| R}$ (see Proposition 7 in Appendix D for more precise bounds).

Proof. Since θ^* exists by Assumption 5, using Eq. (13), applied with $\mu = \rho$ and $\lambda = 0$, we have $L(\theta) - L(\theta^*) \leq \psi(t(\theta - \theta^*))\|\theta - \theta^*\|_{\mathbf{H}(\theta^*)}^2$, for any $\theta \in \mathcal{H}$. By setting $\theta = \widehat{\theta}_\lambda^*$, we obtain

$$L(\widehat{\theta}_\lambda^*) - L(\theta^*) \leq \psi(t(\widehat{\theta}_\lambda^* - \theta^*))\|\widehat{\theta}_\lambda^* - \theta^*\|_{\mathbf{H}(\theta^*)}^2.$$

Using the fact that $\mathbf{H}(\theta^*) \preceq \mathbf{H}(\theta^*) + \lambda I =: \mathbf{H}_\lambda(\theta^*)$, by adding and subtracting θ_λ^* , we have

$$\|\theta_\lambda^* - \theta^*\|_{\mathbf{H}(\theta^*)} \leq \|\theta_\lambda^* - \theta^*\|_{\mathbf{H}_\lambda(\theta^*)} \leq \|\theta_\lambda^* - \theta^*\|_{\mathbf{H}_\lambda(\theta^*)} + \|\widehat{\theta}_\lambda^* - \theta_\lambda^*\|_{\mathbf{H}_\lambda(\theta^*)},$$

and analogously since $t(\cdot)$ is a (semi)norm, $t(\widehat{\theta}_\lambda^* - \theta^*) \leq t_\lambda + t(\widehat{\theta}_\lambda^* - \theta_\lambda^*)$, so

$$L(\widehat{\theta}_\lambda^*) - L(\theta^*) \leq \psi(t_\lambda + t(\widehat{\theta}_\lambda^* - \theta_\lambda^*)) (\|\theta_\lambda^* - \theta^*\|_{\mathbf{H}_\lambda(\theta^*)} + \|\widehat{\theta}_\lambda^* - \theta_\lambda^*\|_{\mathbf{H}_\lambda(\theta^*)})^2.$$

By applying Eq. (12) with $\mu = \rho$ and $\theta = \theta^*$, we have $\mathbf{H}_\lambda(\theta^*) \preceq e^{t_\lambda} \mathbf{H}_\lambda(\theta_\lambda^*)$ and so

$$L(\widehat{\theta}_\lambda^*) - L(\theta^*) \leq \psi(t_\lambda + t(\widehat{\theta}_\lambda^* - \theta_\lambda^*)) (\|\theta_\lambda^* - \theta^*\|_{\mathbf{H}_\lambda(\theta^*)} + e^{t_\lambda/2} \|\widehat{\theta}_\lambda^* - \theta_\lambda^*\|_{\mathbf{H}_\lambda(\theta_\lambda^*)})^2. \quad (36)$$

The terms t_λ and $\|\theta_\lambda^* - \theta^*\|_{\mathbf{H}_\lambda(\theta^*)}$ are related to the *bias terms*, while the terms $t(\widehat{\theta}_\lambda^* - \theta_\lambda^*)$ and $\|\widehat{\theta}_\lambda^* - \theta_\lambda^*\|_{\mathbf{H}_\lambda(\theta_\lambda^*)}$ are related to the *variance term*.

Bounding the bias terms. Recall the definition of the bias $\text{Bias}_\lambda = \|\nabla L_\lambda(\theta^*)\|_{\mathbf{H}_\lambda^{-1}(\theta^*)}$. We bound $t_\lambda = t(\theta_\lambda^* - \theta^*)$, by Lemma 1 and the term $\|\theta^* - \theta_\lambda^*\|_{\mathbf{H}_\lambda(\theta^*)}$ by applying Eq. (14) with $\mu = \rho$ and $\theta = \theta^*$

$$\|\theta^* - \theta_\lambda^*\|_{\mathbf{H}_\lambda(\theta^*)} \leq 1/\underline{\phi}(t_\lambda) \|\nabla L_\lambda(\theta^*)\|_{\mathbf{H}_\lambda^{-1}(\theta^*)} = 1/\underline{\phi}(t_\lambda) \text{Bias}_\lambda.$$

Bounding the variance terms. First we bound the term $\|\widehat{\theta}_\lambda^* - \theta_\lambda^*\|_{\mathbf{H}_\lambda(\theta_\lambda^*)} := \|\mathbf{H}_\lambda(\theta_\lambda^*)^{1/2}(\widehat{\theta}_\lambda^* - \theta_\lambda^*)\|$, by multiplying and dividing for $\widehat{\mathbf{H}}_\lambda(\theta_\lambda^*)^{-1/2}$, we have

$$\begin{aligned} \|\widehat{\theta}_\lambda^* - \theta_\lambda^*\|_{\mathbf{H}_\lambda(\theta_\lambda^*)} &= \|\mathbf{H}_\lambda(\theta_\lambda^*)^{1/2} \widehat{\mathbf{H}}_\lambda(\theta_\lambda^*)^{-1/2} \widehat{\mathbf{H}}_\lambda(\theta_\lambda^*)^{1/2} (\widehat{\theta}_\lambda^* - \theta_\lambda^*)\| \\ &\leq \|\mathbf{H}_\lambda(\theta_\lambda^*)^{1/2} \widehat{\mathbf{H}}_\lambda(\theta_\lambda^*)^{-1/2}\| \|\widehat{\theta}_\lambda^* - \theta_\lambda^*\|_{\widehat{\mathbf{H}}_\lambda(\theta_\lambda^*)}. \end{aligned}$$

Applying Eq. (14) with $\mu = \frac{1}{n} \sum_{i=1}^n \delta_{z_i}$ and $\theta = \widehat{\theta}_\lambda^*$, since $L_{\mu,\lambda} = \widehat{L}_\lambda$ for the given choice of μ , we have

$$\|\theta_\lambda^* - \widehat{\theta}_\lambda^*\|_{\widehat{\mathbf{H}}_\lambda(\theta_\lambda^*)} \leq \|\nabla \widehat{L}_\lambda(\theta_\lambda^*)\|_{\widehat{\mathbf{H}}_\lambda^{-1}(\theta_\lambda^*)} / \underline{\phi}(t(\theta_\lambda^* - \widehat{\theta}_\lambda^*))$$

and since $\|\nabla \widehat{L}_\lambda(\theta_\lambda^*)\|_{\widehat{\mathbf{H}}_\lambda^{-1}(\theta_\lambda^*)} := \|\widehat{\mathbf{H}}_\lambda^{-1/2}(\theta_\lambda^*) \nabla \widehat{L}_\lambda(\theta_\lambda^*)\|$, by multiplying and dividing by $\mathbf{H}_\lambda(\theta_\lambda^*)$, we have:

$$\begin{aligned} \|\widehat{\mathbf{H}}_\lambda^{-1/2}(\theta_\lambda^*) \nabla \widehat{L}_\lambda(\theta_\lambda^*)\| &= \|\widehat{\mathbf{H}}_\lambda^{-1/2}(\theta_\lambda^*) \mathbf{H}_\lambda(\theta_\lambda^*)^{1/2} \mathbf{H}_\lambda(\theta_\lambda^*)^{-1/2} \nabla \widehat{L}_\lambda(\theta_\lambda^*)\| \\ &\leq \|\widehat{\mathbf{H}}_\lambda^{-1/2}(\theta_\lambda^*) \mathbf{H}_\lambda(\theta_\lambda^*)^{1/2}\| \|\nabla \widehat{L}_\lambda(\theta_\lambda^*)\|_{\mathbf{H}_\lambda^{-1}(\theta_\lambda^*)}. \end{aligned}$$

Then

$$\|\theta_\lambda^* - \widehat{\theta}_\lambda^*\|_{\mathbf{H}_\lambda(\theta_\lambda^*)} \leq \frac{1}{\underline{\phi}(t(\theta_\lambda^* - \widehat{\theta}_\lambda^*))} \|\mathbf{H}_\lambda^{1/2}(\theta_\lambda^*) \mathbf{H}_\lambda^{-1/2}(\theta_\lambda^*)\|^2 \|\nabla \widehat{L}_\lambda(\theta_\lambda^*)\|_{\widehat{\mathbf{H}}_\lambda^{-1}(\theta_\lambda^*)} = \frac{\widehat{\text{Var}}_\lambda}{\underline{\phi}(t(\theta_\lambda^* - \widehat{\theta}_\lambda^*))}.$$

To conclude this part of the proof we need to bound $t(\widehat{\theta}_\lambda^* - \theta_\lambda^*)$. Since we require $\widehat{\text{Var}}_\lambda/r_\lambda(\theta_\lambda^*) \leq 1/2$, by Proposition 5 we have $t(\widehat{\theta}_\lambda^* - \theta_\lambda^*) \leq \log 2$.

Gathering the terms. By gathering the results of the previous paragraphs

$$L(\widehat{\theta}_\lambda^*) - L(\theta^*) \leq \psi(t_\lambda + \log 2) \left(1/\underline{\phi}(t_\lambda) \text{Bias}_\lambda + e^{t_\lambda/2}/\underline{\phi}(\log 2) \widehat{\text{Var}}_\lambda \right)^2$$

Using the fact that $(a + b)^2 \leq 2a^2 + 2b^2$, we have the desired result, with

$$K_{\text{bias}}(t_\lambda) = 2\psi(t_\lambda + \log 2)/\underline{\phi}(t_\lambda)^2, \quad K_{\text{var}}(t_\lambda) = 2\psi(t_\lambda + \log 2)e^{t_\lambda}/\underline{\phi}(\log 2)^2. \quad (37)$$

which are bounded in Definition 7 and Proposition 7 of Appendix D.

D.2 Analytic bounds for terms related to the variance

In this lemma, we aim to control the essential supremum and the variance of the random vector $\mathbf{H}_\lambda^{-1/2}(\theta_\lambda^*) \nabla \ell_z^\lambda(\theta_\lambda^*)$ relating it to quantities at θ^* . The results will be used to control the variance via Bernstein concentration inequalities, so we are going to control its essential supremum and its variance.

Lemma 5 (Control of $\mathbf{H}_\lambda(\theta_\lambda^*)^{-1/2} \nabla \ell_z^\lambda(\theta_\lambda^*)$). *For any $0 < \lambda \leq \mathbf{B}_2^*$, we have*

1. *A bound on the essential supremum:*

$$\sup_{z \in \text{supp}(\rho)} \|\nabla \ell_z(\theta_\lambda^*)\|_{\mathbf{H}_\lambda^{-1}(\theta_\lambda^*)} \leq \square_1 \frac{\mathbf{B}_1^*}{\sqrt{\lambda}} + 2\square_2 \frac{\mathbf{B}_2^*}{\lambda} \text{Bias}_\lambda.$$

2. *A bound on the variance*

$$\mathbb{E} \left[\|\nabla \ell_z^\lambda(\theta_\lambda^*)\|_{\mathbf{H}_\lambda^{-1}(\theta_\lambda^*)}^2 \right]^{1/2} \leq \square_1 \sqrt{\text{df}_\lambda} + \sqrt{2}\square_2 \sqrt{\frac{\mathbf{B}_2^*}{\lambda}} \text{Bias}_\lambda,$$

where \square_1, \square_2 are increasing functions of t_λ : $\square_1(t_\lambda) = e^{t_\lambda/2}$ $\square_2(t_\lambda) = e^{t_\lambda/2} (1 + e^{t_\lambda})$.

Proof. Start by noting that if $\lambda \leq \mathbf{B}_2^*$, then $\sup_{z \in \text{supp}(\rho)} \|\mathbf{H}_\lambda^{-1/2}(\theta^*) \nabla^2 \ell_z^\lambda(\theta^*)^{1/2}\|^2 \leq 1 + \frac{\mathbf{B}_2^*}{\lambda} \leq 2\frac{\mathbf{B}_2^*}{\lambda}$. Moreover, note that for any vector $h \in \mathcal{H}$, multiplying and dividing by $\nabla^2 \ell_z(\theta^*)^{1/2}$,

$$\begin{aligned} \|h\|_{\mathbf{H}_\lambda^{-1}(\theta^*)} &:= \|\mathbf{H}_\lambda^{-1/2}(\theta^*) h\| = \|\mathbf{H}_\lambda^{-1/2}(\theta^*) \nabla^2 \ell_z(\theta^*)^{1/2} \nabla^2 \ell_z(\theta^*)^{-1/2} h\| \\ &\leq \|\mathbf{H}_\lambda^{-1/2}(\theta^*) \nabla^2 \ell_z(\theta^*)^{1/2}\| \|\nabla^2 \ell_z(\theta^*)^{-1/2} h\| \\ &\leq \sqrt{\frac{2\mathbf{B}_2^*}{\lambda}} \|\nabla^2 \ell_z(\theta^*)^{-1/2} h\| \\ &= \sqrt{\frac{2\mathbf{B}_2^*}{\lambda}} \|h\|_{\nabla^2 \ell_z(\theta^*)^{-1}}, \end{aligned}$$

where the last bound is mentioned at the beginning of the proof. Similarly, we can show

$$\|h\|_{\nabla^2 \ell_z(\theta^*)} \leq \sqrt{\frac{2\mathbf{B}_2^*}{\lambda}} \|h\|_{\mathbf{H}_\lambda(\theta^*)}, \quad \|h\|_{\mathbf{H}_\lambda^{-1}(\theta^*)} \leq \sqrt{\frac{2\mathbf{B}_2^*}{\lambda}} \|h\|_{\nabla^2 \ell_z(\theta^*)^{-1}}. \quad (38)$$

Essential supremum. Let $z \in \text{supp}(\rho)$. First note that using Eq. (27), we have

$$\|\nabla \ell_z^\lambda(\theta_\lambda^*)\|_{\mathbf{H}_\lambda^{-1}(\theta_\lambda^*)} \leq e^{t_\lambda/2} \|\nabla \ell_z^\lambda(\theta_\lambda^*)\|_{\mathbf{H}_\lambda^{-1}(\theta^*)}.$$

Now bound

$$\|\nabla \ell_z^\lambda(\theta_\lambda^*)\|_{\mathbf{H}_\lambda^{-1}(\theta^*)} \leq \|\nabla \ell_z^\lambda(\theta_\lambda^*) - \nabla \ell_z^\lambda(\theta^*)\|_{\mathbf{H}_\lambda^{-1}(\theta^*)} + \|\nabla \ell_z^\lambda(\theta^*)\|_{\mathbf{H}_\lambda^{-1}(\theta^*)}.$$

Since $\nabla \ell_z^\lambda(\theta^*) = \nabla \ell_z(\theta^*) + \lambda \theta^*$, the last term is bounded by

$$\text{Bias}_\lambda + \sup_{z \in \text{supp}(\rho)} \|\nabla \ell_z(\theta^*)\|_{\mathbf{H}_\lambda^{-1}(\theta^*)} \leq \text{Bias}_\lambda + \frac{\mathbf{B}_1^*}{\sqrt{\lambda}}.$$

For the first term, start by using Eq. (38).

$$\|\nabla \ell_z^\lambda(\theta_\lambda^*) - \nabla \ell_z^\lambda(\theta^*)\|_{\mathbf{H}_\lambda^{-1}(\theta^*)} \leq \sqrt{\frac{2\mathbf{B}_2^*}{\lambda}} \|\nabla \ell_z^\lambda(\theta_\lambda^*) - \nabla \ell_z^\lambda(\theta^*)\|_{\nabla^2 \ell_z^\lambda(\theta^*)^{-1}}.$$

Using Eq. (29) on ℓ_z^λ , we find

$$\|\nabla \ell_z^\lambda(\theta_\lambda^*) - \nabla \ell_z^\lambda(\theta^*)\|_{\nabla^2 \ell_z^\lambda(\theta^*)^{-1}} \leq \bar{\phi}(t_\lambda) \|\theta_\lambda^* - \theta^*\|_{\nabla^2 \ell_z^\lambda(\theta^*)}.$$

Applying once again Eq. (38), we bound

$$\|\theta_\lambda^* - \theta^*\|_{\nabla^2 \ell_z^\lambda(\theta^*)} \leq \sqrt{\frac{2\mathbf{B}_2^*}{\lambda}} \|\theta_\lambda^* - \theta^*\|_{\mathbf{H}_\lambda(\theta^*)}.$$

Finally, using Eq. (28) on L_λ , we get

$$\|\theta_\lambda^* - \theta^*\|_{\mathbf{H}_\lambda(\theta^*)} \leq \frac{1}{\underline{\phi}(t_\lambda)} \text{Bias}_\lambda.$$

Hence, putting things together, we get

$$\|\nabla \ell_z^\lambda(\theta_\lambda^*) - \nabla \ell_z^\lambda(\theta^*)\|_{\mathbf{H}_\lambda^{-1}(\theta^*)} \leq \frac{2\mathbf{B}_2^* \bar{\phi}(t_\lambda)}{\lambda \underline{\phi}(t_\lambda)} \text{Bias}_\lambda = \frac{2\mathbf{B}_2^*}{\lambda} e^{t_\lambda} \text{Bias}_\lambda.$$

We combine all our different computation to get the bound.

Variance. We start by using Eq. (27) to show that

$$\mathbb{E} \left[\|\nabla \ell_z^\lambda(\theta_\lambda^*)\|_{\mathbf{H}_\lambda^{-1}(\theta_\lambda^*)}^2 \right]^{1/2} \leq e^{t_\lambda/2} \mathbb{E} \left[\|\nabla \ell_z^\lambda(\theta_\lambda^*)\|_{\mathbf{H}_\lambda^{-1}(\theta^*)}^2 \right]^{1/2}.$$

Then we use the triangle inequality

$$\mathbb{E} \left[\|\nabla \ell_z^\lambda(\theta_\lambda^*)\|_{\mathbf{H}_\lambda^{-1}(\theta^*)}^2 \right]^{1/2} \leq \mathbb{E} \left[\|\nabla \ell_z^\lambda(\theta_\lambda^*) - \nabla \ell_z^\lambda(\theta^*)\|_{\mathbf{H}_\lambda^{-1}(\theta^*)}^2 \right]^{1/2} + \mathbb{E} \left[\|\nabla \ell_z^\lambda(\theta^*)\|_{\mathbf{H}_\lambda^{-1}(\theta^*)}^2 \right]^{1/2}.$$

We can easily bound the last term on the right hand side by $\text{Bias}_\lambda + \text{df}_\lambda$. For the first term, we proceed as in the previous case to obtain

$$\forall z \in \text{supp}(\rho), \|\nabla \ell_z^\lambda(\theta_\lambda^*) - \nabla \ell_z^\lambda(\theta^*)\|_{\mathbf{H}_\lambda^{-1}(\theta^*)} \leq \sqrt{\frac{2\mathbf{B}_2^*}{\lambda}} \bar{\phi}(t_\lambda) \|\theta_\lambda^* - \theta^*\|_{\nabla^2 \ell_z^\lambda(\theta^*)}.$$

Now taking the expectancy of this inequality squared,

$$\begin{aligned}\mathbb{E} \left[\|\nabla \ell_z^\lambda(\theta_\lambda^*) - \nabla \ell_z^\lambda(\theta^*)\|_{\mathbf{H}_\lambda^{-1}(\theta^*)}^2 \right]^{1/2} &\leq \sqrt{\frac{2\mathbf{B}_2^*}{\lambda} \bar{\phi}(t_\lambda)} \mathbb{E} \left[\|\theta_\lambda^* - \theta^*\|_{\nabla^2 \ell_z^\lambda(\theta^*)}^2 \right]^{1/2} \\ &= \sqrt{\frac{2\mathbf{B}_2^*}{\lambda} \bar{\phi}(t_\lambda)} \|\theta_\lambda^* - \theta^*\|_{\mathbf{H}_\lambda(\theta^*)},\end{aligned}$$

where the last equality comes from $\mathbb{E} [\nabla^2 \ell_z^\lambda(\theta^*)] = \mathbf{H}_\lambda(\theta^*)$. Now applying Eq. (28) to L_λ , we obtain

$$\|\theta_\lambda^* - \theta^*\|_{\mathbf{H}_\lambda(\theta^*)} \leq \frac{1}{\underline{\phi}(t_\lambda)} \text{Bias}_\lambda.$$

Regrouping all these bounds, we obtain

$$\mathbb{E} \left[\|\nabla \ell_z^\lambda(\theta_\lambda^*) - \nabla \ell_z^\lambda(\theta^*)\|_{\mathbf{H}_\lambda^{-1}(\theta^*)}^2 \right]^{1/2} \leq e^{t_\lambda} \sqrt{\frac{2\mathbf{B}_2^*}{\lambda}} \text{Bias}_\lambda.$$

Hence the final bound is proved, regrouping all our computations.

D.3 Concentration lemmas

Here we concentrate in high probability the quantities obtained in the analytical decomposition. Details on the proof technique are given in Sec. 6 of the paper.

Lemma 6 (Equivalence of empirical and expected Hessian). *Let $\theta \in \mathcal{H}$ and $n \in \mathbb{N}$. For any $\delta \in (0, 1]$, $\lambda > 0$, if*

$$n \geq 24 \frac{\mathbf{B}_2(\theta)}{\lambda} \log \frac{8\mathbf{B}_2(\theta)}{\lambda\delta}, \quad (39)$$

then with probability at least $1 - \delta$: $\mathbf{H}_\lambda(\theta) \preceq 2\widehat{\mathbf{H}}_\lambda(\theta)$, or equivalently

$$\|\mathbf{H}_\lambda^{1/2}(\theta) \widehat{\mathbf{H}}_\lambda^{-1/2}(\theta)\|^2 \leq 2.$$

Proof. By Remark 4 and the definition of $\mathbf{B}_2(\theta)$, the condition we require on n is sufficient to apply Proposition 10, in particular Eq. (51), to $\mathbf{H}_\lambda(\theta)$, $\widehat{\mathbf{H}}_\lambda(\theta)$, for $t = 1/2$, which provides the desired result.

Lemma 7 (Concentration of the empirical gradient). *Let $n \in \mathbb{N}$, $\delta \in (0, 1]$, $0 < \lambda \leq \mathbf{B}_2^*$. For any $k \geq 4$, if $n \geq k^2 \square_2^2 \frac{\mathbf{B}_2^*}{\lambda} \log \frac{2}{\delta}$, then with probability at least $1 - \delta$, we have*

$$\|\nabla \widehat{L}_\lambda(\theta_\lambda^*)\|_{\mathbf{H}_\lambda^{-1}(\theta_\lambda^*)} \leq \frac{2\sqrt{3}}{k} \text{Bias}_\lambda + 2\square_1 \sqrt{\frac{\text{df}_\lambda \vee (\mathbf{Q}^*)^2 \log \frac{2}{\delta}}{n}}. \quad (40)$$

Here, \square_1, \square_2 are defined in Lemma 5 in Appendix D and $(\mathbf{Q}^*)^2 = (\mathbf{B}_1^*)^2 / \mathbf{B}_2^*$.

Proof.1) First let us concentrate $\mathbf{H}_\lambda(\theta_\lambda^*)^{-1/2} \nabla \widehat{L}_\lambda(\theta_\lambda^*)$ using a Bernstein-type inequality.

We can see $\mathbf{H}_\lambda(\theta_\lambda^*)^{-1/2} \nabla \widehat{L}_\lambda(\theta_\lambda^*)$ as the mean of n i.i.d. random variables distributed from the law of the vector $\mathbf{H}_\lambda(\theta_\lambda^*)^{-1/2} \nabla \ell_z(\theta_\lambda^*)$.

As we have shown in Lemma 5, the essential supremum and variance of this vector is bounded, then we can use Bernstein inequality for random vectors (e.g. Thm. 3.3.4 of [Yurinsky, 1995](#)): for any $\lambda > 0$, any $n \in \mathbb{N}$ and $\delta \in (0, 1]$, with probability at least $1 - \delta$, we have

$$\|\mathbf{H}_\lambda(\theta_\lambda^*)^{-1/2} \nabla \widehat{L}_\lambda(\theta_\lambda^*)\| \leq \frac{2M \log \frac{2}{\delta}}{n} + \sigma \sqrt{\frac{2 \log \frac{2}{\delta}}{n}},$$

where $M = \sup_{z \in \text{supp}(\rho)} \|\nabla \ell_z(\theta_\lambda^*)\|_{\mathbf{H}_\lambda^{-1}(\theta_\lambda^*)}$ and $\sigma = \mathbb{E} \left[\|\nabla \ell_z(\theta_\lambda^*)\|_{\mathbf{H}_\lambda^{-1}(\theta_\lambda^*)}^2 \right]^{1/2}$.

2) Using the bounds obtained in Lemma 5,

$$M \leq \square_1 \frac{B_1^*}{\sqrt{\lambda}} + 2\square_2 \frac{B_2^*}{\lambda} \text{Bias}_\lambda, \quad \sigma \leq \square_1 \sqrt{\text{df}_\lambda} + \sqrt{2}\square_2 \frac{\sqrt{B_2^*}}{\sqrt{\lambda}} \text{Bias}_\lambda.$$

3) Injecting these in the Bernstein inequality,

$$\begin{aligned} \|\mathbf{H}_\lambda(\theta_\lambda^*)^{-1/2} \nabla \widehat{L}_\lambda(\theta_\lambda^*)\| &\leq \frac{2 \left(\square_1 B_1^*/\sqrt{\lambda} + 2\square_2 (B_2^*/\lambda) \text{Bias}_\lambda \right) \log \frac{2}{\delta}}{n} \\ &+ \left(\square_1 \text{df}_\lambda^{1/2} + \sqrt{2}\square_2 \sqrt{B_2^*/\lambda} \text{Bias}_\lambda \right) \sqrt{\frac{2 \log \frac{2}{\delta}}{n}} \\ &= \left[\frac{4\square_2 B_2^* \log \frac{2}{\delta}}{\lambda n} + \sqrt{\frac{4\square_2^2 B_2^* \log \frac{2}{\delta}}{\lambda n}} \right] \text{Bias}_\lambda \\ &+ \sqrt{\frac{2\square_1^2 \text{df}_\lambda \log \frac{2}{\delta}}{n}} + \sqrt{\frac{2 B_2^* \log \frac{2}{\delta}}{\lambda n}} \sqrt{\frac{2\square_1^2 (B_1^*)^2/B_2^* \log \frac{2}{\delta}}{n}}. \end{aligned}$$

In the last inequality, we have regrouped the terms with a factor Bias_λ and we have separated the first term of the decomposition in the following way :

$$\frac{2\square_1 B_1^* \log \frac{2}{\delta}}{\sqrt{\lambda n}} = \sqrt{\frac{2 B_2^* \log \frac{2}{\delta}}{\lambda n}} \sqrt{\frac{2\square_1^2 (Q^*)^2 \log \frac{2}{\delta}}{n}}.$$

Hence, we can bound the second line of the last inequality:

$$\sqrt{\frac{2\square_1^2 \text{df}_\lambda \log \frac{2}{\delta}}{n}} + \sqrt{\frac{2B_2^* \log \frac{2}{\delta}}{\lambda n}} \sqrt{\frac{2\square_1^2 (Q^*)^2 \log \frac{2}{\delta}}{n}} \leq \left(1 + \sqrt{\frac{2B_2^* \log \frac{2}{\delta}}{\lambda n}} \right) \sqrt{\frac{2\square_1^2 \text{df}_\lambda \vee (Q^*)^2 \log \frac{2}{\delta}}{n}}.$$

Thus, if we assume that $n \geq k^2 \square_2^2 \frac{B_2^*}{\lambda} \log \frac{2}{\delta}$,

$$\|\mathbf{H}_\lambda(\theta_\lambda^*)^{-1/2} \nabla \widehat{L}_\lambda(\theta_\lambda^*)\| \leq \left(\frac{4}{k^2} + \frac{2}{k} \right) \text{Bias}_\lambda + \left(1 + \frac{\sqrt{2}}{k} \right) \sqrt{\frac{2\square_1^2 \text{df}_\lambda \vee (B_1^*)^2/B_2^* \log \frac{2}{\delta}}{n}}.$$

In particular, for $k \geq 4$,

$$\|\mathbf{H}_\lambda(\theta_\lambda^*)^{-1/2} \nabla \widehat{L}_\lambda(\theta_\lambda^*)\| \leq \frac{3}{k} \text{Bias}_\lambda + 2\square_1 \sqrt{\frac{\text{df}_\lambda \vee (Q^*)^2 \log \frac{2}{\delta}}{n}}.$$

Lemma 8 (control of $\widehat{\text{Var}}_\lambda$). *Let $n \in \mathbb{N}$, $\delta \in (0, 1]$ and $0 < \lambda \leq B_2^*$. Assume that for a certain $k \geq 5$,*

$$n \geq k^2 \square_2^2 \frac{B_2^*}{\lambda} \log \frac{\square_1^2 B_2^*}{\lambda \delta}.$$

Then with probability at least $1 - 2\delta$, we have

$$\widehat{\text{Var}}_\lambda \leq \frac{6}{k} \text{Bias}_\lambda + 4\square_1 \sqrt{\frac{\text{df}_\lambda \vee (Q^*)^2 \log \frac{2}{\delta}}{n}}.$$

Here, \square_1, \square_2 are defined in Lemma 5

Proof.

- First we apply Lemma 6 to $\theta = \theta_\lambda^*$. Since $B_2(\theta_\lambda^*) \leq e^{t_\lambda} B_2^* = \square_1^2 B_2^*$, we see that the condition

$$n \geq 24 \frac{B_2(\theta_\lambda^*)}{\lambda} \log \frac{8B_2(\theta_\lambda^*)}{\lambda\delta}$$

is satisfied if

$$n \geq 24 \square_1^2 \frac{B_2^*}{\lambda} \log \frac{8 \square_1^2 B_2^*}{\lambda\delta}.$$

Because $k \geq 5$ and $\square_2 \geq \square_1$, and we see that the assumption of this lemma imply the conditions above and hence Lemma 6 is satisfied. In particular, $\|\mathbf{H}_\lambda(\theta_\lambda^*)^{1/2} \widehat{\mathbf{H}}_\lambda(\theta_\lambda^*)^{-1/2}\|^2 \leq 2$.

- Note that the condition of this proposition also imply the conditions of Lemma 7, because $\lambda \leq B_2^*$ and $\square_1 \geq 1$ imply $\frac{\square_1^2 B_2^*}{\lambda\delta} \geq \frac{1}{\delta}$.

D.4 Final results

First, we find conditions on n such that the hypothesis $\widehat{\text{Var}}_\lambda \leq \frac{r_\lambda(\theta_\lambda^*)}{2}$ is satisfied.

Lemma 9. Let $n \in \mathbb{N}$, $\delta \in (0, 1]$, $0 < \lambda \leq B_2^*$ and

$$n \geq \Delta_1 \frac{B_2^*}{\lambda} \log \frac{8 \square_1^2 B_2^*}{\lambda\delta}, \quad n \geq \Delta_2 \frac{\text{df}_\lambda \vee (\mathbf{Q}^*)^2}{r_\lambda(\theta^*)^2} \log \frac{2}{\delta},$$

then with probability at least $1 - 2\delta$

$$\frac{\widehat{\text{Var}}_\lambda}{r_\lambda(\theta_\lambda^*)} \leq \square_1 \frac{\widehat{\text{Var}}_\lambda}{r_\lambda(\theta^*)} \leq \frac{1}{2},$$

where $\square_1, \Delta_1, \Delta_2$ are constants defined in Definition 7.

Proof. Recall that $\tilde{t}_\lambda = \frac{\text{Bias}_\lambda}{r_\lambda(\theta^*)}$.

Using Lemma 8, we see that under the conditions of this lemma, we have

$$\square_1 \frac{\widehat{\text{Var}}_\lambda}{r_\lambda(\theta^*)} \leq \frac{6 \square_1 \text{Bias}_\lambda}{k r_\lambda(\theta^*)} + 4 \square_1^2 \sqrt{\frac{\text{df}_\lambda \vee (\mathbf{Q}^*)^2 \log \frac{2}{\delta}}{n r_\lambda(\theta^*)^2}}.$$

Thus, taking $k = 24 \square_1 (1/2 \vee \tilde{t}_\lambda)$ and $n \geq 256 \square_1^4 \frac{\text{df}_\lambda \vee \mathbf{Q}^*}{r_\lambda(\theta^*)^2} \log \frac{2}{\delta}$, both terms in the sum are bounded by $1/4$ hence the result.

Note that here, we have defined

$$\Delta_1 = 576 \square_1^2 \square_2^2 (1/2 \vee \tilde{t}_\lambda)^2, \quad \Delta_2 = 256 \square_1^4,$$

hence the constants in the definition above.

Proof of Thm. 4 First we recall that $\Delta_1, \Delta_2, \square_1, C_{\text{bias}}$ and C_{var} are defined in Definition 7, and bounded in Proposition 7.

First note that, given the requirements on n , by Lemma 9, we have $\widehat{\text{Var}}_\lambda \leq \frac{r_\lambda(\theta_\lambda^*)}{2}$ with probability at least $1 - 2\delta$. Thus, we are in a position to apply Thm. 7 :

$$L(\widehat{\theta}_\lambda^*) - L(\theta^*) \leq K_{\text{bias}} \text{Bias}_\lambda^2 + K_{\text{var}} \widehat{\text{Var}}_\lambda^2,$$

with $K_{\text{bias}}, K_{\text{var}}$ defined in the proof of the theorem. Note that in the proof of Lemma 9, we have taken $k = 24\Box_1(1/2 \vee \tilde{t}_\lambda) \geq 12$. Hence, using Lemma 8, we find

$$\widehat{\text{Var}}_\lambda \leq \frac{1}{2}\text{Bias}_\lambda + 4\Box_1 \sqrt{\frac{\text{df}_\lambda \vee (\mathbf{Q}^*)^2 \log \frac{2}{\delta}}{n}}.$$

Hence,

$$\widehat{\text{Var}}_\lambda^2 \leq \frac{1}{2}\text{Bias}_\lambda^2 + 32\Box_1^2 \frac{\text{df}_\lambda \vee (\mathbf{Q}^*)^2 \log \frac{2}{\delta}}{n},$$

which yields the wanted result with $C_{\text{bias}} = K_{\text{bias}} + \frac{1}{2}K_{\text{var}}$ and $C_{\text{var}} = 32\Box_1^2 K_{\text{var}}$.

Proof of Thm. 3

We get this theorem as a corollary of Thm. 4. Indeed, $\forall \lambda \leq B_2^*$, $\text{df}_\lambda \vee (\mathbf{Q}^*)^* \leq \frac{(B_1^*)^2}{\lambda}$, hence the result.

E Explicit bounds for the simplified case

In this section, assume that Assumptions 1, 3 to 5 and 8 hold.

Define the following constant N :

$$N = 36A^2 \log^2 \left(6A^2 \frac{1}{\delta} \right) \vee 256 \frac{1}{A^2} \log \frac{2}{\delta} \vee 512 (\|\theta^*\|^2 R^2 \vee 1) \log \frac{2}{\delta}, \quad (41)$$

where $A = \frac{\bar{B}_2}{\bar{B}_1}$.

We have the following slow rates theorem.

Theorem 8 (Quantitative slow rates result). *Let $n \in \mathbb{N}$. Let $\delta \in (0, 1]$. Setting*

$$\lambda = 16((R \vee 1)\bar{B}_1) \frac{1}{\sqrt{n}} \log^{1/2} \frac{2}{\delta},$$

if $n \geq N$, with probability at least $1 - 2\delta$,

$$L(\widehat{\theta}_\lambda^*) - L(\theta^*) \leq 48 \max(R, 1) \max(\|\theta^*\|^2, 1) \bar{B}_1 \frac{1}{\sqrt{n}} \log^{1/2} \frac{2}{\delta}, \quad (42)$$

and $N = O(\text{poly}(\bar{B}_1, \bar{B}_2, R\|\theta^*\|))$ is given explicitly in Eq. (41). Here, poly denotes a certain rational function of the inputs.

Proof. Note that $\bar{\text{df}}_\lambda \leq \frac{\bar{B}_1^2}{\lambda}$. Hence, if $\lambda \leq \bar{B}_2$, then $\bar{\text{df}}_\lambda \vee (\bar{B}_1^2/\bar{B}_2) \leq \frac{\bar{B}_1^2}{\lambda}$.

1) Let us reformulate Thm. 6. Let $n \in \mathbb{N}$ and $0 < \lambda \leq \bar{B}_2$. Let $\delta \in (0, 1]$. If

$$n \geq 512 (\|\theta^*\|^2 R^2 \vee 1) \log \frac{2}{\delta}, \quad n \geq 24 \frac{\bar{B}_2}{\lambda} \log \frac{8\bar{B}_2}{\lambda\delta}, \quad n \geq 16\Delta^2 \frac{R^2 \bar{B}_1^2}{\lambda^2} \log \frac{2}{\delta},$$

then with probability at least $1 - 2\delta$,

$$L(\widehat{\theta}_\lambda^*) - L(\theta^*) \leq C_{\text{var}} \frac{\bar{B}_1^2}{\lambda n} \log \frac{2}{\delta} + C_{\text{bias}} \lambda \|\theta^*\|^2,$$

where $\Delta, C_{\text{bias}}, C_{\text{var}}$ are defined in Definition 5.

2) Now setting $\lambda = 16R\bar{B}_1 \log^{1/2} \frac{2}{\delta} \frac{1}{n^{1/2}}$, we see that the inequality

$$n \geq 16\Delta^2 \frac{R^2\bar{B}_1^2}{\lambda^2} \log \frac{2}{\delta}$$

is automatically satisfied since $\Delta \leq 4$. Hence, if

$$n \geq 512 (\|\theta^*\|^2 R^2 \vee 1) \log \frac{2}{\delta}, \quad n \geq 24 \frac{\bar{B}_2}{\lambda} \log \frac{8\bar{B}_2}{\lambda\delta}, \quad 0 < \lambda \leq \bar{B}_2,$$

then

$$L(\hat{\theta}_\lambda^*) - L(\theta^*) \leq \frac{C_{\text{var}}}{256} \frac{1}{R^2} \lambda + C_{\text{bias}} \lambda \|\theta^*\|^2 \leq \left(\frac{C_{\text{var}}}{256} + C_{\text{bias}} \right) \max\left(\frac{1}{R^2}, \|\theta^*\|^2\right) \lambda.$$

Since by Definition 5, $C_{\text{var}} \leq 84$ and $C_{\text{bias}} \leq 2$, we get

$$L(\hat{\theta}_\lambda^*) - L(\theta^*) \leq 3 \max\left(\frac{1}{R^2}, \|\theta^*\|^2\right) \lambda.$$

3) Having our fixed $\lambda = 16 \frac{\bar{B}_1 R \log^{1/2} \frac{2}{\delta}}{n^{1/2}}$, let us look for conditions for which

$$n \geq 512 (\|\theta^*\|^2 R^2 \vee 1) \log \frac{2}{\delta}, \quad n \geq 24 \frac{\bar{B}_2}{\lambda} \log \frac{8\bar{B}_2}{\lambda\delta}, \quad 0 < \lambda \leq \bar{B}_2,$$

are satisfied.

To deal with $n \geq 24 \frac{\bar{B}_2}{\lambda} \log \frac{8\bar{B}_2}{\lambda\delta}$, bound

$$\frac{\bar{B}_2}{\lambda} \leq \frac{1}{16} \frac{\bar{B}_2}{R\bar{B}_1 \log^{1/2} \frac{2}{\delta}} n^{1/2} \leq \frac{1}{8} \frac{\bar{B}_2}{R\bar{B}_1} n^{1/2},$$

where we have used the fact that $\log^{1/2} \frac{2}{\delta} \geq \frac{1}{2}$. apply Lemma 14 with $a_1 = 3, a_2 = 1, A = \frac{\bar{B}_2}{R\bar{B}_1}$ to get the following condition:

$$n \geq 4a_1^2 A^2 \log^2 \left(\frac{2a_1 a_2 A^2}{\delta} \right),$$

which we express as

$$n \geq 36A^2 \log^2 \left(6A^2 \frac{1}{\delta} \right).$$

To deal with the bound $\lambda < \bar{B}_2$, we need only apply the definition to obtain

$$n \geq 256 \frac{R^2 \bar{B}_1^2}{\bar{B}_2^2} \log \frac{2}{\delta}.$$

Thus, we can concentrate all these bounds as $n \geq N$ where

$$N = 36A^2 \log^2 \left(6A^2 \frac{1}{\delta} \right) \vee 256 \frac{1}{A^2} \log \frac{2}{\delta} \vee 512 (\|\theta^*\|^2 R^2 \vee 1) \log \frac{2}{\delta},$$

where $A = \frac{\bar{B}_2}{R\bar{B}_1}$.

4) Since R is only an upper bound, we can replace R by $R \vee 1$. In this case, we see that $A \leq \frac{\bar{B}_2}{\bar{B}_1}$ and $\max(\frac{1}{R\bar{B}_1}, (R \vee 1)\|\theta^*\|^2) \leq (R \vee 1)(\|\theta^*\| \vee 1)^2$ hence the final bounds.

F Explicit bounds for the refined case

In this part, we continue to assume Assumptions 1, 3 to 5 and 8. We present a classification of distributions ρ and show that we can achieve better rates than the classical slow rates.

Definition 8 (class of distributions). *Let $\alpha \in [1, +\infty]$ and $r \in [0, 1/2]$.*

We denote with $\mathcal{P}_{\alpha,r}$ the set of probability distributions ρ such that there exists $L, Q \geq 0$,

- $\text{Bias}_\lambda \leq L \lambda^{\frac{1+2r}{2}}$
- $\text{df}_\lambda \leq Q^2 \lambda^{-1/\alpha}$,

where this holds for any $0 < \lambda \leq 1$. For simplicity, if $\alpha = +\infty$, we assume that $Q \geq Q^$.*

Note that given our assumptions, we always have

$$\rho \in \mathcal{P}_{1,0}, \quad L = \|\theta^*\|, \quad Q = B_1^*. \quad (43)$$

We also define

$$\lambda_1 = \left(\frac{Q}{Q^*} \right)^{2\alpha} \wedge 1, \quad (44)$$

such that

$$\forall \lambda \leq \lambda_1, \quad \text{df}_\lambda \vee (Q^*)^2 \leq \frac{Q^2}{\lambda^{1/\alpha}}.$$

Interpretation of the classes

- The bias term Bias_λ characterizes the regularity of the objective θ^* . In a sense, if r is big, then this means θ^* is very regular and will be easier to estimate. The following results reformulates this intuition.

Remark 1 (source condition). *Assume there exists $0 \leq r \leq 1/2$ and $v \in \mathcal{H}$ such that*

$$P_{\mathbf{H}(\theta^*)} \theta^* = \mathbf{H}(\theta^*)^r v.$$

Then we have

$$\forall \lambda > 0, \quad \text{Bias}_\lambda \leq L \lambda^{\frac{1+2r}{2}}, \quad L = \|\mathbf{H}(\theta^*)^{-r} \theta^*\|.$$

- The effective dimension df_λ characterizes the size of the space \mathcal{H} with respect to the problem. The higher α , the smaller the space. If \mathcal{H} is finite dimensional for instance, $\alpha = +\infty$.

We will give explicit bounds for the performance of $\hat{\theta}_\lambda^*$ depending on which class ρ belongs to, i.e., as a function of α, r .

Well-behaved problems $r_\lambda(\theta^*)$ has a limiting role. However, as soon as we have some sort of regularity, this role is no longer limiting, i.e. this quantity does not appear in the final rates and the constants in these rates have no dependence on the problem. This motivates the following definition.

We say that a problem is well behaved if the following equation holds.

$$\forall \delta \in (0, \frac{1}{2}], \quad \exists \lambda_0(\delta) \in (0, 1], \quad \forall 0 < \lambda \leq \lambda_0(\delta), \quad \frac{L \lambda^{1/2+r}}{r_\lambda(\theta^*)} \log \frac{2}{\delta} \leq \frac{1}{2}. \quad (45)$$

Remark 2 (well-behaved problems). *Note that Eq. (45) is satisfied if one of the following holds.*

- *If $R = 0$, then the condition holds for $\lambda_0 = 1$.*
- *If $r > 0$, then the condition holds for $\lambda_0 = (2LR \log \frac{2}{\delta})^{-1/r} \wedge 1$.*
- *If there exists $\mu \in [0, 1)$ and $F \geq 0$ such that $r_\lambda(\theta^*) \geq \frac{1}{F} \lambda^{\mu/2}$, then this holds for $\lambda_0 = (2RF \log \frac{2}{\delta})^{-2/(1-\mu+2r)} \wedge 1$.*

Moreover, if Eq. (45) is satisfied, then for any $\lambda \leq \lambda_0$, $t_\lambda \leq \log 2$.

Note that the first possible condition corresponds to the case where the loss functions are quadratic in θ (if the loss is the square loss for instance). The second condition corresponds to having a strict source condition, i.e. something strictly better than just $\theta^* \in \mathcal{H}$. Finally, the third condition corresponds to the fact that the radius r_λ decreases slower than the original bound of $r_\lambda \geq \frac{\lambda^{1/2}}{R}$, and hence it is not limiting.

Note that a priori, using only the assumptions, our problems do not satisfy Eq. (45) (see Eq. (43), and the fact that $r_\lambda \geq \frac{\sqrt{\lambda}}{R}$).

F.1 Quantitative bounds

In this section, for any given pair (α, r) characterizing the regularity and size of the problem, we associate

$$\beta = \frac{1}{1 + 2r + 1/\alpha}, \quad \gamma = \frac{\alpha(1 + 2r)}{\alpha(1 + 2r) + 1}.$$

In what follows, we define

$$N = \frac{256Q^2}{L^2} (B_2^* \wedge \lambda_0 \wedge \lambda_1)^{-1/\beta} \vee \left(1296 \frac{1}{1-\beta} A \log \left(5184 \frac{1}{1-\beta} A^2 \frac{1}{\delta} \right) \right)^{1/(1-\beta)}, \quad (46)$$

where $A = \frac{B_2^* L^{2\beta}}{Q^{2\beta}}$, λ_0 is given by Eq. (45) and λ_1 is given by Eq. (44) : $\lambda_1 = \frac{Q^{2\alpha}}{(Q^*)^{2\alpha}}$.

Theorem 9 (Quantitative results when Eq. (45) is satisfied and $\alpha < \infty$ or $r > 0$). *Let $\rho \in \mathcal{P}_{\alpha,r}$ and that we have either $\alpha < \infty$ or $r > 0$. Let $\delta \in (0, \frac{1}{2}]$.*

If Eq. (45) is satisfied, and

$$n \geq N, \quad \lambda = \left(256 \left(\frac{Q}{L} \right)^2 \frac{1}{n} \right)^\beta,$$

then with probability at least $1 - 2\delta$,

$$L(\widehat{\theta}_\lambda^*) - L(\theta^*) \leq 8 (256)^\gamma (Q^\gamma L^{1-\gamma})^2 \frac{1}{n^\gamma} \log \frac{2}{\delta},$$

where N is defined in Eq. (46).

Proof.

Using the definition of λ_1 , as soon as $\lambda \leq \lambda_1$ we have $\text{df}_\lambda \vee (Q^*)^2 \leq Q^2 \lambda^{-1/\alpha}$.

Let us formulate Thm. 4 using the fact that $\rho \in \mathcal{P}_{\alpha,r}$.

Let $\delta \in (0, 1]$, $0 < \lambda \leq B_2^* \wedge \lambda_1$ and $n \in \mathbb{N}$ such that

$$n \geq \Delta_1 \frac{B_2^*}{\lambda} \log \frac{8\Box_1^2 B_2^*}{\lambda\delta}, \quad n \geq \Delta_2 \frac{Q^2}{\lambda^{1/\alpha} r_\lambda(\theta^*)^2} \log \frac{2}{\delta},$$

then with probability at least $1 - 2\delta$

$$L(\widehat{\theta}_\lambda^*) - L(\theta^*) \leq C_{\text{bias}} L^2 \lambda^{1+2r} + C_{\text{var}} \frac{Q^2}{\lambda^{1/\alpha} n} \log \frac{2}{\delta},$$

where $C_{\text{bias}}, C_{\text{var}}$ are defined in Definition 7. Now let us distinguish the two cases of our theorem.

Assume that ρ satisfies Eq. (45). In this case the proof proceeds as follows. Note that as soon as $\lambda \leq \lambda_0$, we have $\frac{C_{\text{bias}}}{r_\lambda(\theta^*)} \leq \frac{1}{2}$ and hence the bounds in Proposition 7 apply.

1) First, we find a simple condition to guarantee

$$r_\lambda(\theta^*)^2 \lambda^{1/\alpha} \geq \Delta_2 Q^2 \frac{1}{n} \log \frac{2}{\delta}.$$

Using the fact that Eq. (45) is satisfied, we see that if $\lambda \leq \lambda_0$, then $r_\lambda \geq 2L\lambda^{1/2+r} \log \frac{2}{\delta}$. Hence, this condition is satisfied if

$$\lambda \leq \lambda_0, \quad 4L^2 \lambda^{1+2r+1/\alpha} \geq \Delta_2 Q^2 \frac{1}{n}.$$

2) Now fix $C_\lambda = 256 \geq \Delta_2/4$ (see Proposition 7) and fix

$$\lambda^{1+2r+1/\alpha} = C_\lambda \frac{Q^2}{L^2} \frac{1}{n} \iff \lambda = \left(C_\lambda \frac{Q^2}{L^2} \frac{1}{n} \right)^\beta.$$

where $\beta = 1/(1 + 2r + 1/\lambda) \in [1/2, 1)$.

Using our restatement of Thm. 4, we have that with probability at least $1 - 2\delta$,

$$L(\widehat{\theta}_\lambda^*) - L(\theta^*) \leq \left(C_{\text{bias}} + \frac{1}{C_\lambda} C_{\text{var}} \log \frac{2}{\delta} \right) L^2 \lambda^{1+2r} \leq K \log \frac{2}{\delta} L^2 \lambda^{1+2r},$$

where we have set $K = (C_{\text{bias}} + \frac{1}{256} C_{\text{var}}) \leq 8$ (see Proposition 7).

This result holds provided

$$0 < \lambda \leq B_2^* \wedge \lambda_0 \wedge \lambda_1, \quad n \geq \Delta_1 \frac{B_2^*}{\lambda} \log \frac{8\Box_1^2 B_2^*}{\lambda\delta}. \quad (47)$$

Indeed, we have shown in the previous point that since $C_\lambda \geq \frac{\Delta_2}{4}$, $r_\lambda(\theta^*)^2 \lambda^{1/\alpha} \geq \Delta_2 Q^2 \frac{1}{n} \log \frac{2}{\delta}$.

3) Let us now work to guarantee the conditions in Eq. (47).

First, to guarantee $n \geq \Delta_1 \frac{B_2^*}{\lambda} \log \frac{8\Box_1^2 B_2^*}{\lambda\delta}$, bound

$$\frac{B_2^*}{\lambda} = \frac{B_2^* L^{2\beta} n^\beta}{C_\lambda^\beta Q^{2\beta} \log^\beta \frac{2}{\delta}} \leq \frac{2}{C_\lambda^\beta} \frac{B_2^* L^{2\beta}}{Q^{2\beta}} n^\beta.$$

Then apply Lemma 15 with $a_1 = \frac{2\Delta_1}{C_\lambda^\beta}$, $a_2 = \frac{16\Box_1^2}{C_\lambda^\beta}$, $A = \frac{B_2^* L^{2\beta}}{Q^{2\beta}}$. Since $\beta \geq 1/2$, using the bounds in Proposition 7, we find $a_1 \leq 648$ and $a_2 \leq 4$, hence the following sufficient condition:

$$n \geq \left(1296 \frac{1}{1-\beta} A \log \left(5184 \frac{1}{1-\beta} A^2 \frac{1}{\delta} \right) \right)^{1/(1-\beta)}.$$

Then, to guarantee the condition

$$\lambda \leq B_2^* \wedge \lambda_0 \wedge \lambda_1,$$

we simply need

$$n \geq \frac{256Q^2}{L^2} (B_2^* \wedge \lambda_0 \wedge \lambda_1)^{-1/\beta}.$$

Hence, defining

$$N = \frac{256Q^2}{L^2} (B_2^* \wedge \lambda_0 \wedge \lambda_1)^{-1/\beta} \vee \left(1296 \frac{1}{1-\beta} A \log \left(5184 \frac{1}{1-\beta} A^2 \frac{1}{\delta} \right) \right)^{1/(1-\beta)},$$

where $A = \frac{B_2^* L^{2\beta}}{Q^{2\beta}}$, we see that as soon as $n \geq N$, Eq. (47) holds.

We now state the following corollary, for $r > 0$. We define N in the following way:

$$N = \frac{256Q^2}{L^2} (B_2^* \wedge \lambda_0 \wedge \lambda_1)^{-1/\beta} \vee \left(1296 \frac{1}{1-\beta} A \log \left(5184 \frac{1}{1-\beta} A^2 \frac{1}{\delta} \right) \right)^{1/(1-\beta)} \quad (48)$$

where $A = \frac{B_2^* L^{2\beta}}{Q^{2\beta}}$, $\lambda_0 = (2LR \log \frac{2}{\delta})^{-1/r} \wedge 1$ and $\lambda_1 = \frac{Q^{2\alpha}}{(Q^*)^{2\alpha}}$.

Corollary 4. Assume $\rho \in \mathcal{P}_{\alpha,r}$ with $r > 0$. Let $\delta \in (0, 0.5]$ and $n \geq N$, where N is defined in Eq. (48). For

$$\lambda = \left(256 \left(\frac{Q}{L} \right)^2 \frac{1}{n} \right)^\beta,$$

with probability at least $1 - 2\delta$,

$$L(\widehat{\theta}_\lambda^*) - L(\theta^*) \leq 8 (256)^\gamma (Q^\gamma L^{1-\gamma})^2 \frac{1}{n^\gamma} \log \frac{2}{\delta},$$

Moreover, $N = O(\text{poly}(B_1^*, B_2^*, L, Q, R, \log \frac{1}{\delta}))$, which means that N is bounded by a rational function of the arguments of poly.

Proof of Cor. 2 We simply apply Cor. 4 for $\alpha = 1$ and $Q = B_1^*$.

G Additional lemmas

G.1 Self-concordance, sufficient conditions to define L and related quantities

In this section, we will consider an arbitrary probability measure μ on \mathcal{Z} . We assume that ℓ_z satisfies Assumption 8 with a certain given function φ . Recall that $R^\mu = \sup_{z \in \text{supp}(\mu)} \sup_{g \in \varphi(z)} \|g\|$. In this section, we will also assume that $R^\mu < \infty$.

Lemma 10 (Gronwall lemma). Let $\varphi : \mathbb{R} \rightarrow \mathbb{R}$ be a differentiable function such that

$$\forall t \in \mathbb{R}, \varphi'(t) \leq C\varphi(t).$$

Then

$$\forall (t_0, t_1) \in \mathbb{R}^2, \varphi(t_1) \leq e^{C|t_1-t_0|} \varphi(t_0).$$

Lemma 11. Assume that there exists θ_0 such that $\sup_{z \in \text{supp}(\mu)} \text{Tr}(\nabla^2 \ell_z(\theta_0)) < \infty$

- $\sup_{z \in \text{supp}(\mu)} \text{Tr}(\nabla^2 \ell_z(\theta)) < \infty$ for any $\theta \in \mathcal{H}$;
- For any given radius $T > 0$, and any $\|\theta_0\| \leq T$, we have

$$\forall \|\theta\| \leq T, \forall z \in \mathcal{Z}, \text{Tr}(\nabla^2 \ell_z(\theta)) \leq \exp(2R^\mu T) \text{Tr}(\nabla^2 \ell_z(\theta_0)) < \infty.$$

Proof.

Let $z \in \text{supp}(\mu)$ be fixed. Using the same reasoning as in the proof of Eq. (27), we can show

$$\forall \theta_0, \theta_1 \in \mathcal{H}, \nabla^2 \ell_z(\theta_1) \preceq \exp\left(\sup_{g \in \varphi(z)} |g \cdot (\theta_1 - \theta_0)|\right) \nabla^2 \ell_z(\theta_0) \preceq \exp(R^\mu \|\theta_1 - \theta_0\|) \nabla^2 \ell_z(\theta_0)$$

Where we have used the fact that $R^\mu = \sup_{z \in \text{supp}(\mu)} \sup_{g \in \varphi(z)} \|g\| < \infty$. Thus, in particular

$$\forall z \in \text{supp}(\mu), \forall \theta_0, \theta_1 \in \mathcal{H}, \text{Tr}(\nabla^2 \ell_z(\theta_1)) \leq \exp(R^\mu \|\theta_1 - \theta_0\|) \text{Tr}(\nabla^2 \ell_z(\theta_0)),$$

which leads to the desired bounds.

Lemma 12. Assume that there exists θ_0 such that

$$\sup_{z \in \text{supp}(\mu)} \text{Tr}(\nabla^2 \ell_z(\theta_0)) < \infty, \quad \sup_{z \in \text{supp}(\mu)} \|\nabla \ell_z(\theta_0)\| < \infty.$$

Then

- $\sup_{z \in \text{supp}(\mu)} \|\nabla \ell_z(\theta)\| < \infty$ for any $\theta \in \mathcal{H}$
- For any $T > 0$ and any $\|\theta_0\|, \|\theta\| \leq T, z \in \text{supp}(\mu)$,

$$\begin{aligned} \|\nabla \ell_z(\theta)\| &\leq \|\nabla \ell_z(\theta_0)\| + 2T \text{Tr}(\nabla^2 \ell_z(\theta_0)) \\ &\quad + 4R^\mu \psi(2R^\mu T) \text{Tr}(\nabla^2 \ell_z(\theta_0)) R^2. \end{aligned}$$

Proof.

Fix $z \in \mathcal{Z}, \theta_0, \theta_1 \in \mathcal{H}$ and $h \in \mathcal{H}$. Let us look at the function

$$f : t \in [0, 1] \mapsto (\nabla \ell_z(\theta_t) - \nabla \ell_z(\theta_0) - t \nabla^2 \ell_z(\theta_0)(\theta_1 - \theta_0)) \cdot h.$$

We have $f''(t) = \nabla^3 \ell_z(\theta_t)[\theta_1 - \theta_0, \theta_1 - \theta_0, h]$. By the self-concordant assumption, we have

$$\begin{aligned} |f''(t)| &\leq \sup_{g \in \varphi(z)} |g \cdot h| \nabla^2 \ell_z(\theta_t)[\theta_1 - \theta_0, \theta_1 - \theta_0] \\ &\leq \sup_{g \in \varphi(z)} |g \cdot h| \exp\left(t \sup_{g \in \varphi(z)} |g \cdot \theta_1 - \theta_0|\right) \|\theta_1 - \theta_0\|_{\nabla^2 \ell_z(\theta_0)}^2. \end{aligned}$$

Integrating this knowing $f'(0) = f(0) = 0$ yields

$$|f(1)| \leq \sup_{g \in \varphi(z)} |g \cdot h| \psi\left(\sup_{g \in \varphi(z)} |g \cdot (\theta_1 - \theta_0)|\right) \|\theta_1 - \theta_0\|_{\nabla^2 \ell_z(\theta_0)}^2.$$

Hence :

$$\|\nabla \ell_z(\theta_1) - \nabla \ell_z(\theta_0)\| \leq \|\nabla^2 \ell_z(\theta_0)\| \|\theta_1 - \theta_0\| + \|\varphi(z)\| \psi\left(\sup_{g \in \varphi(z)} |g \cdot (\theta_1 - \theta_0)|\right) \|\nabla^2 \ell_z(\theta_0)\| \|\theta_1 - \theta_0\|^2$$

where $\psi(t) = (e^t - t - 1)/t^2$. Then, noting that $\|\nabla^2 \ell_z(\theta)\| \leq \text{Tr}(\nabla^2 \ell_z(\theta))$, we have proved our lemma.

Lemma 13. Assume that there exists θ_0 such that

$$\sup_{z \in \text{supp}(\mu)} \text{Tr}(\nabla^2 \ell_z(\theta_0)) < \infty, \quad \sup_{z \in \text{supp}(\mu)} \|\nabla \ell_z(\theta_0)\| < \infty, \quad \sup_{z \in \text{supp}(\mu)} |\ell_z(\theta_0)| < \infty.$$

Then

- For any $\theta \in \mathcal{H}$, $\sup_{z \in \text{supp}(\mu)} |\ell_z(\theta)| < \infty$
- For any $\theta_0 \in \mathcal{H}$, $T \geq \|\theta_0\|$, $\|\theta\| \leq T$, $z \in \text{supp}(\mu)$, we have:

$$|\ell_z(\theta)| \leq |\ell_z(\theta_0)| + 2\|\nabla \ell_z(\theta_0)\|T + \psi(2R^\mu T) \text{Tr}(\nabla^2 \ell_z(\theta_0)) T^2.$$

Proof. Proceeding as in the proof of Eq. (30), we get

$$\forall z \in \mathcal{Z}, \forall \theta_0, \theta_1 \in \mathcal{H}, 0 \leq \ell_z(\theta_1) - \ell_z(\theta_0) - \nabla \ell_z(\theta_0)(\theta_1 - \theta_0) \leq \psi\left(\sup_{g \in \varphi(z)} |g \cdot (\theta_1 - \theta_0)|\right) \|\theta_1 - \theta_0\|_{\nabla^2 \ell_z(\theta_0)}^2$$

where $\psi(t) = (e^t - t - 1)/t^2$.

To conclude, we give the following result.

Proposition 8. Let $\lambda \geq 0$. If a probability measure μ and ℓ satisfy Assumptions 3, 4 and 8, the function $L_{\mu, \lambda}(\theta) := \mathbb{E}_\mu[\ell_z(\theta)] + \lambda \|\theta\|^2$ and $\nabla L_{\mu, \lambda}(\theta)$, $\nabla^2 L_{\mu, \lambda}(\theta)$ are well-defined for any $\theta \in \mathcal{H}$, and we can differentiate under the expectation. Moreover,

$$\forall \theta \in \mathcal{H}, \quad \sup_{z \in \text{supp}(\rho)} |\ell_z(\theta)|, \quad \sup_{z \in \text{supp}(\rho)} \|\nabla \ell_z(\theta)\|, \quad \sup_{z \in \text{supp}(\rho)} \text{Tr}(\nabla^2 \ell_z(\theta)) < \infty.$$

Proof. We combine the results given in Lemmas 11 to 13.

G.2 Bernstein inequalities for operators

We start by proposing a slight modification of Proposition 6 in (Rudi and Rosasco, 2017). First we need to introduce the following quantity and some notation for Hermitian operators. We denote by \preceq is the partial order between positive semidefinite Hermitian operators. Let A, B be bounded Hermitian operators on \mathcal{H} ,

$$A \preceq B \iff v \cdot (Av) \leq v \cdot (Bv), \quad \forall v \in \mathcal{H} \iff B - A \text{ is positive semidefinite.}$$

Let q be a random positive semi-definite operator and let $\mathbf{Q} := \mathbb{E}[q]$, denote by $\mathcal{F}(\lambda)$ the function of λ defined as

$$\mathcal{F}(\lambda) := \text{ess sup Tr} \left(\mathbf{Q}_\lambda^{-1/2} q \mathbf{Q}_\lambda^{-1/2} \right),$$

where ess sup is the *essential support* of q .

Remark 3. Note that if $\text{Tr}(q) \leq c_0$, for a $c_0 > 0$ almost surely, then $\mathcal{F}(\lambda) \leq c_0/\lambda$. Vice versa, if $\mathcal{F}(\lambda_0) < \infty$ for a given $\lambda_0 > 0$, then $\text{Tr}(q) \leq (\|\mathbf{Q}\| + \lambda_0)\mathcal{F}(\lambda_0)$ almost surely, moreover $\mathcal{F}(\lambda) < \frac{\|\mathbf{Q}\| + \lambda_0}{\|\mathbf{Q}\| + \lambda} \mathcal{F}(\lambda_0)$ for any $\lambda > 0$.

Proposition 9 (Prop. 6 of (Rudi and Rosasco, 2017)). Let q_1, \dots, q_n be identically distributed random positive semi-definite operators on a separable Hilbert space \mathcal{H} such that the q are trace class and $\mathbf{Q} = \mathbb{E}[q]$. Let $\mathbf{Q}_n = \frac{1}{n} \sum_{i=1}^n q_i$ and take $0 < \lambda \leq \|\mathbf{Q}\|$ and assume $\mathcal{F}(\lambda) < \infty$. For any $\delta > 0$, the following holds with probability at least $1 - \delta$:

$$\|\mathbf{Q}_\lambda^{-1/2} (\mathbf{Q} - \mathbf{Q}_n) \mathbf{Q}_\lambda^{-1/2}\| \leq \frac{2\beta(1 + \mathcal{F}(\lambda))}{3n} + \sqrt{\frac{2\beta\mathcal{F}(\lambda)}{n}}, \quad \beta = \log \frac{8\mathcal{F}(\lambda)}{\delta}$$

Proof. Use Proposition 3 of (Rudi and Rosasco, 2017) and proceed as in the proof of Proposition 6 of (Rudi and Rosasco, 2017) except that we bound $\text{Tr}(\mathbf{Q}_\lambda^{-1}\mathbf{Q}) \leq \mathcal{F}_\infty(\lambda)$ instead of bounding $\text{Tr}(\mathbf{Q}_\lambda^{-1}\mathbf{Q}) \leq \frac{\text{Tr}(\mathbf{Q})}{\lambda}$, we find this result.

Here we slightly extend the results of Prop. 8 and Prop. 6 of (Rudi and Rosasco, 2017), to extend the range of λ for which the result on the partial order between operators holds, from $0 < \lambda < \|\mathbf{Q}\|$ to $\lambda > 0$.

Proposition 10 (Prop. 8 together with Prop. 6 of (Rudi and Rosasco, 2017)). *Let q_1, \dots, q_n be identically distributed random positive semi-definite operators on a separable Hilbert space \mathcal{H} such that the q are trace class and $\mathbf{Q} = \mathbb{E}[q]$. Let $\mathbf{Q}_n = \frac{1}{n} \sum_{i=1}^n q_i$. Let any $\delta \in (0, 1]$, $t > 0$, $0 < \lambda \leq \|\mathbf{Q}\|$ and assume $\mathcal{F}(\lambda) < \infty$, when*

$$n \geq 8\mathcal{F}_\infty(\lambda) \log \frac{8\mathcal{F}_\infty(\lambda)}{\delta} \left(\frac{1}{4t^2} + \frac{1}{t} \right) \quad (49)$$

then the following holds with probability at least $1 - \delta$:

$$\|\mathbf{Q}_\lambda^{-1/2} (\mathbf{Q} - \mathbf{Q}_n) \mathbf{Q}_\lambda^{-1/2}\| \leq t. \quad (50)$$

Moreover let $\lambda > 0, \delta \in (0, 1]$ and Eq. (49) is satisfied for $t \leq 1/2$, then the following holds with probability at least $1 - \delta$,

$$\mathbf{Q}_\lambda \preceq 2\mathbf{Q}_{n,\lambda}, \quad \iff \quad \|\mathbf{Q}_{n,\lambda}^{-1/2} \mathbf{Q}_\lambda^{1/2}\|^2 \leq 2. \quad (51)$$

Finally, let $\lambda > 0, \delta \in (0, 1]$, Eq. (49) is satisfied for $t \leq 1/2$ and

$$n \geq 16 \frac{c_0^2}{\|\mathbf{Q}\|^2} \log \frac{2}{\delta},$$

with $c_0 = \text{ess sup Tr}(q)$, then the following holds with probability at least $1 - \delta$,

$$\mathbf{Q}_{n,\lambda} \preceq \frac{3}{2} \mathbf{Q}_\lambda, \quad \iff \quad \|\mathbf{Q}_{n,\lambda}^{1/2} \mathbf{Q}_\lambda^{-1/2}\|^2 \leq 3/2. \quad (52)$$

Proof.

Point 1) Let $\delta \in (0, 1]$ and $0 < \lambda \leq \|\mathbf{Q}\|$. Using Proposition 9, we have that with probability at least $1 - \delta$,

$$\|\mathbf{Q}_\lambda^{-1/2} (\mathbf{Q} - \mathbf{Q}_n) \mathbf{Q}_\lambda^{-1/2}\| \leq \frac{2\beta(1 + \mathcal{F}_\infty(\lambda))}{3n} + \sqrt{\frac{2\beta\mathcal{F}_\infty(\lambda)}{n}}, \quad \beta = \log \frac{8\mathcal{F}_\infty(\lambda)}{\delta}.$$

Now note that if $\lambda \leq \|\mathbf{Q}\|$, we have

$$\frac{1}{2} \leq \frac{\|\mathbf{Q}\|}{\|\mathbf{Q}\| + \lambda} = \|\mathbf{Q}_\lambda^{-1}\mathbf{Q}\| \leq \text{Tr}(\mathbf{Q}_\lambda^{-1}\mathbf{Q}) \leq \mathcal{F}_\infty(\lambda).$$

Thus we can bound $1 + \mathcal{F}_\infty(\lambda) \leq 3\mathcal{F}_\infty(\lambda)$, and we rewrite the previous bound

$$\|\mathbf{Q}_\lambda^{-1/2} (\mathbf{Q} - \mathbf{Q}_n) \mathbf{Q}_\lambda^{-1/2}\| \leq \frac{2\beta\mathcal{F}_\infty(\lambda)}{n} + \sqrt{\frac{2\beta\mathcal{F}_\infty(\lambda)}{n}}, \quad \beta = \log \frac{8\mathcal{F}_\infty(\lambda)}{\delta}.$$

Point 2) Now let $t > 0, \delta \in (0, 1]$ and $0 < \lambda \leq \|\mathbf{Q}\|$. If

$$n \geq 8\mathcal{F}_\infty(\lambda)\beta \left(\frac{1}{4t^2} + \frac{1}{t} \right),$$

then

$$\|\mathbf{Q}_\lambda^{-1/2} (\mathbf{Q} - \mathbf{Q}_n) \mathbf{Q}_\lambda^{-1/2}\| \leq t.$$

Indeed, assume we want to find $n_0 > 0$ for which for all $n \geq n_0$, $\frac{A}{n} + \sqrt{\frac{B}{n}} \leq \frac{1}{2}$ where $A, B \geq 0$. setting $x = \sqrt{n}$, this is equivalent to finding x_0 such that $\forall x \geq x_0$, $\frac{x^2}{2} - \sqrt{B}x - A \geq 0$. A sufficient condition for this is that $x \geq \sqrt{B} + \sqrt{B + 2A}$. Thus, since $A, B \geq 0$, the condition $x \geq 2\sqrt{B + 2A}$ is sufficient, hence the condition $n \geq 4(B + 2A)$. Then we apply this to the following A and B to obtain the condition.

$$A = \frac{\beta \mathcal{F}_\infty(\lambda)}{t}, \quad B = \frac{\beta \mathcal{F}_\infty(\lambda)}{2t^2}.$$

Point 3) When $\lambda > \|\mathbf{Q}\|$, the result is obtained noting that

$$\|\mathbf{Q}_\lambda^{1/2} \mathbf{Q}_{n,\lambda}^{-1/2}\|^2 \leq \frac{\|\mathbf{Q}\| + \lambda}{\lambda} = 1 + \frac{\|\mathbf{Q}\|}{\lambda} \leq 2.$$

When, on the other hand $0 < \lambda \leq \|\mathbf{Q}\|$, the final result is obtained by applying Prop. 6 and Prop. 8 of (Rudi and Rosasco, 2017), or equivalently applying Eq. (50), with $t = 1/2$, for which the following holds with probability $1 - \delta$: $\|\mathbf{Q}_\lambda^{-1/2} (\mathbf{Q} - \mathbf{Q}_n) \mathbf{Q}_\lambda^{-1/2}\| \leq t$ and noting that,

$$\|\mathbf{Q}_\lambda^{1/2} \mathbf{Q}_{n,\lambda}^{-1/2}\|^2 \leq \frac{1}{1 - \|\mathbf{Q}_\lambda^{-1/2} (\mathbf{Q} - \mathbf{Q}_n) \mathbf{Q}_\lambda^{-1/2}\|} \leq 2.$$

To conclude this point, we recall that, given two Hermitian operators A, B and $t > 0$, the inequality $A \preceq tB$ is equivalent to $B^{-1/2} A B^{-1/2} \preceq tI$, when B is invertible. Since $B^{-1/2} A B^{-1/2}$ and tI are commutative, then $B^{-1/2} A B^{-1/2} \preceq tI$ is equivalent to $v \cdot (B^{-1/2} A B^{-1/2} v) \leq t\|v\|^2$ for any $v \in \mathcal{H}$, which in turn is equivalent to $\|B^{-1/2} A B^{-1/2}\| \leq t$. So

$$\|A^{1/2} B^{-1/2}\|^2 \leq t \iff A \preceq tB.$$

Point 4) First note that

$$\|\mathbf{Q}_\lambda^{-1/2} \mathbf{Q}_{n,\lambda}^{1/2}\|^2 \leq 1 + \|\mathbf{Q}_\lambda^{-1/2} (\mathbf{Q} - \mathbf{Q}_n) \mathbf{Q}_\lambda^{-1/2}\|. \quad (53)$$

When $0 < \lambda \leq \|\mathbf{Q}\|$, by applying Eq. (50) with $t = 1/2$, we have with probability $1 - \delta$: $\|\mathbf{Q}_\lambda^{-1/2} (\mathbf{Q} - \mathbf{Q}_n) \mathbf{Q}_\lambda^{-1/2}\| \leq t$, moreover by Eq. (53) we have

$$\|\mathbf{Q}_\lambda^{-1/2} \mathbf{Q}_{n,\lambda}^{1/2}\|^2 \leq 1 + t \leq 3/2.$$

When instead $\lambda > \|\mathbf{Q}\|$, we consider the following decomposition

$$\|\mathbf{Q}_\lambda^{-1/2} (\mathbf{Q} - \mathbf{Q}_n) \mathbf{Q}_\lambda^{-1/2}\| \leq \frac{1}{\lambda} \|\mathbf{Q} - \mathbf{Q}_n\| \leq \frac{1}{\lambda} \|\mathbf{Q} - \mathbf{Q}_n\|_{HS},$$

where we denote by $\|\cdot\|_{HS}$, the Hilbert-Schmidt norm (i.e. $\|A\|_{HS}^2 = \text{Tr}(A^* A)$) and $\|\mathbf{Q} - \mathbf{Q}_n\|_{HS}$ is well defined since both \mathbf{Q}, \mathbf{Q}_n are trace class. Now since the space of Hilbert-Schmidt operators on a separable Hilbert space is itself a separable Hilbert space and q are bounded almost surely by $c_0 := \text{ess sup Tr}(q)$, we can concentrate $\|\mathbf{Q} - \mathbf{Q}_n\|_{HS}$ via Bernstein inequality for random vectors (e.g. Thm. 3.3.4 of Yurinsky, 1995), obtaining with probability at least $1 - \delta$

$$\|\mathbf{Q} - \mathbf{Q}_n\|_{HS} \leq \frac{2c_0 \log \frac{2}{\delta}}{n} + \sqrt{\frac{2c_0^2 \log \frac{2}{\delta}}{n}} \leq \|\mathbf{Q}\|/2,$$

where the last step is due to the fact that we require $n \geq 16c_0^2(\log \frac{2}{\delta})/\|\mathbf{Q}\|^2$, and the fact that by construction $\|\mathbf{Q}\| \leq B$. Then,

$$\|\mathbf{Q}_\lambda^{-1/2} \mathbf{Q}_{n,\lambda}^{1/2}\|^2 \leq 1 + \frac{\|\mathbf{Q}\|}{2\lambda} \leq 3/2.$$

The final result on \preceq is obtained as for Point 5.

Remark 4. Let $\text{Tr}(q) \leq c_0$ almost surely, for a $c_0 > 0$. Then $\mathcal{F}(\lambda) \leq c_0/\lambda$. So Eq. (49) is satisfied when

$$n \geq \frac{8c_0}{\lambda} \log \frac{8c_0}{\lambda \delta} \left(\frac{1}{4t^2} + \frac{1}{t} \right),$$

since $\mathcal{F}(\lambda) \leq c_0/\lambda$ as observed in Remark 3. In particular, when $t = 1/2$, Eq. (49) is satisfied when

$$n \geq \frac{24c_0}{\lambda} \log \frac{8c_0}{\lambda \delta}.$$

G.3 Sufficient conditions to bound n in order to guarantee $n \geq C_1 n^p \log \frac{C_2 n^p}{\delta}$

Lemma 14. Let $a_1, a_2, A \geq 0$ and $\delta > 0$. If

$$n \geq 4a_1^2 A^2 \log^2 \left(\frac{2a_1 a_2 A^2}{\delta} \right),$$

then $n \geq a_1 A n^{1/2} \log \frac{a_2 A n^{1/2}}{\delta}$.

Proof. Indeed, note that

$$n \geq a_1 A n^{1/2} \log \frac{a_2 A n^{1/2}}{\delta} \iff \frac{a_1 A}{n^{1/2}} \log \frac{a_2 A n^{1/2}}{\delta} \leq 1.$$

Now use the fact that for $A, B \geq 0$, $k \geq 2A \log(2AB)$ implies $\frac{A}{k} \log(Bk) \leq 1$. Indeed, $\log(Bk) = \log(2AB) + \log \frac{Bk}{2AB} = \log(2AB) + \log \frac{k}{2A} \leq \log(2AB) + \frac{k}{2A}$. Hence, multiplying by $\frac{A}{k}$, we get the result.

We apply this to $A = a_1 A$, $B = \frac{a_2 A}{\delta}$ and $k = n^{1/2}$ to get the bound.

Lemma 15. Let $a_1, a_2, A \geq 0$ and $\delta > 0$. Let $p \in [\frac{1}{2}, 1)$. If

$$n^{1-p} \geq 2 \frac{1}{1-p} a_1 A \log \left(2a_1 (a_2 \vee 1) \frac{1}{1-p} A^2 \frac{1}{\delta} \right),$$

then

$$n \geq a_1 A n^p \log \frac{a_2 A n^p}{\delta}.$$

Proof.1) Let $C_1, C_2 \geq 0$, and $p \in [0, 1)$. Then

$$n \geq C_1 n^p \log(C_2 n^p) \iff \frac{C_1 \frac{p}{1-p}}{n^{1-p}} \log \left(C_2^{(1-p)/p} n^{1-p} \right) \leq 1.$$

Now use the fact that for $A, B \geq 0$, $k \geq 2A \log(2AB)$ implies $\frac{A}{k} \log(Bk) \leq 1$ (see proof of Lemma 14).

Thus, $n^{1-p} \geq 2C_1 \frac{p}{1-p} \log \left(2C_1 \frac{p}{1-p} C_2^{(1-p)/p} \right)$ is a sufficient condition.

2) Now taking $C_1 = a_1 A$ and $C_2 = \frac{a_2 A}{\delta}$, we find that

$$n^{1-p} \geq 2 \frac{p}{1-p} a_1 A \log \left(2a_1 a_2^{(1-p)/p} \frac{p}{1-p} A^{1/p} \left(\frac{1}{\delta} \right)^{(1-p)/p} \right).$$

Since $0.5 \leq p \leq 1$, we see that $\frac{1-p}{p} \leq 1$ and $\frac{1}{p} \leq 2$ and thus we get our final sufficient condition.

$$n^{1-p} \geq 2 \frac{1}{1-p} a_1 A \log \left(2a_1 (a_2 \vee 1) \frac{1}{1-p} A^2 \frac{1}{\delta} \right).$$