
Globally Convergent Newton Methods for Ill-conditioned Generalized Self-concordant Losses

Ulysse Marteau-Ferey
INRIA - École Normale Supérieure
PSL Research University
ulysse.marteau-ferrey@inria.fr

Francis Bach
INRIA - École Normale Supérieure
PSL Research University
francis.bach@inria.fr

Alessandro Rudi
INRIA - École Normale Supérieure
PSL Research University
alessandro.rudi@inria.fr

Abstract

In this paper, we study large-scale convex optimization algorithms based on the Newton method applied to regularized generalized self-concordant losses, which include logistic regression and softmax regression. We first prove that our new simple scheme based on a sequence of problems with decreasing regularization parameters is provably globally convergent, that this convergence is linear with a constant factor which scales only logarithmically with the condition number. In the parametric setting, we obtain an algorithm with the same scaling than regular first-order methods but with an improved behavior, in particular in ill-conditioned problems. Second, in the non-parametric machine learning setting, we provide an explicit algorithm combining the previous scheme with Nyström projection techniques, and prove that it achieves optimal generalization bounds with a time complexity of order $O(ndf_\lambda)$, a memory complexity of order $O(df_\lambda^2)$ and *no dependence on the condition number*, generalizing the results known for least-squares regression. Here n is the number of observations and df_λ is the associated degrees of freedom. In particular, this is the first large-scale algorithm to solve logistic and softmax regressions in the non-parametric setting with large condition numbers and theoretical guarantees.

1 Introduction

Minimization algorithms constitute a crucial algorithmic part of many machine learning methods, with algorithms available for a variety of situations [10]. In this paper, we focus on *finite sum* problems of the form

$$\min_{x \in \mathcal{H}} f_\lambda(x) = f(x) + \frac{\lambda}{2} \|x\|^2, \text{ with } f(x) = \frac{1}{n} \sum_{i=1}^n f_i(x),$$

where \mathcal{H} is a Euclidean or a Hilbert space, and each function is convex and smooth. The running-time of minimization algorithms classically depends on the number of functions n , the explicit (for Euclidean spaces) or implicit (for Hilbert spaces) dimension d of the search space, and the condition number of the problem, which is upper bounded by $\kappa = L/\lambda$, where L characterizes the smoothness of the functions f_i , and λ the regularization parameter.

In the last few years, there has been a strong focus on problems with large n and d , leading to *first-order* (i.e., gradient-based) stochastic algorithms, culminating in a sequence of linearly convergent

algorithms whose running time is favorable in n and d , but scale at best in $\sqrt{\kappa}$ [15, 22, 14, 4]. However, modern problems lead to objective functions with very large condition numbers, i.e., in many learning problems, the regularization parameter that is optimal for test predictive performance may be so small that the scaling above in $\sqrt{\kappa}$ is not practical anymore (see examples in Sect. 5).

These ill-conditioned problems are good candidates for *second-order methods* (i.e., that use the Hessians of the objective functions) such as Newton method. These methods are traditionally discarded within machine learning for several reasons: (1) they are usually adapted to high precision results which are not necessary for generalization to unseen data for machine learning problems [9], (2) computing the Newton step $\Delta_\lambda(x) = \nabla^2 f_\lambda(x)^{-1} \nabla f_\lambda(x)$ requires to form the Hessian and solve the associated linear system, leading to complexity which is at least quadratic in d , and thus prohibitive for large d , and (3) the global convergence properties are not applicable, unless the function is very special, i.e., self-concordant [24] (which includes only few classical learning problems), so they often are only shown to converge in a small area around the optimal x .

In this paper, we argue that the three reasons above for not using Newton method can be circumvented to obtain competitive algorithms: (1) high absolute precisions are indeed not needed for machine learning, but faced with strongly ill-conditioned problems, even a low-precision solution requires second-order schemes; (2) many approximate Newton steps have been designed for approximating the solution of the associated large linear system [1, 27, 25, 8]; (3) we propose a novel second-order method which is globally convergent and which is based on performing approximate Newton methods for a certain class of so-called *generalized self-concordant functions* which includes logistic regression [6]. For these functions, the conditioning of the problem is also characterized by a more *local* quantity: $\kappa_\ell = R^2/\lambda$, where R characterizes the local evolution of Hessians. This leads to second-order algorithms which are competitive with first-order algorithms for well-conditioned problems, while being superior for ill-conditioned problems which are common in practice.

Contributions. We make the following contributions:

- (a) We build a global second-order method for the minimization of f_λ , which relies only on computing approximate Newton steps of the functions $f_\mu, \mu \geq \lambda$. The number of such steps will be of order $O(c \log \kappa_\ell + \log \frac{1}{\epsilon})$ where ϵ is the desired precision, and c is an explicit constant. In the parametric setting ($\mathcal{H} = \mathbb{R}^d$), c can be as bad as $\sqrt{\kappa_\ell}$ in the worst-case but much smaller in theory and practice. Moreover in the non-parametric/kernel machine learning setting (\mathcal{H} infinite dimensional), c does not depend on the local condition number κ_ℓ .
- (b) Together with the appropriate quadratic solver to compute approximate Newton steps, we obtain an algorithm with the same scaling as regular first-order methods but with an improved behavior, in particular in ill-conditioned problems. Indeed, this algorithm matches the performance of the best quadratic solvers but covers any generalized self-concordant function, up to logarithmic terms.
- (c) In the non-parametric/kernel machine learning setting we provide an explicit algorithm combining the previous scheme with Nyström projections techniques. We prove that it achieves optimal generalization bounds with $O(ndf_\lambda)$ in time and $O(df_\lambda^2)$ in memory, where n is the number of observations and df_λ is the associated degrees of freedom. In particular, this is the first large-scale algorithm to solve logistic and softmax regression in the non-parametric setting with large condition numbers and theoretical guarantees.

1.1 Comparison to related work

We consider two cases for \mathcal{H} and the functions f_i that are common in machine learning: $\mathcal{H} = \mathbb{R}^d$ with linear (in the parameter) models with explicit feature maps, and \mathcal{H} infinite-dimensional, corresponding in machine learning to learning with kernels [32]. Moreover in this section we first consider the quadratic case, for example the squared loss in machine learning (i.e., $f_i(x) = \frac{1}{2}(x^\top z_i - y_i)^2$ for some $z_i \in \mathcal{H}, y_i \in \mathbb{R}$). We first need to introduce the Hessian of the problem, for any $\lambda > 0$, define

$$\mathbf{H}(x) := \nabla^2 f(x), \quad \mathbf{H}_\lambda(x) := \nabla^2 f_\lambda(x) = \mathbf{H}(x) + \lambda \mathbf{I},$$

in particular we denote by \mathbf{H} (and analogously \mathbf{H}_λ) the Hessian at optimum (which in case of squared loss corresponds to the covariance matrix of the inputs).

Quadratic problems and $\mathcal{H} = \mathbb{R}^d$ (ridge regression). The problem then consists in solving a (ill-conditioned) positive semi-definite symmetric linear system of dimension $d \times d$. Methods based on *randomized linear algebra*, *sketching* and suitable *subsampling* [17, 18, 11] are able to find the solution with precision ϵ in time that is $O((nd + \min(n, d)^3) \log(L/\lambda\epsilon))$, so essentially independently of the condition number, because of the logarithmic complexity in λ .

Quadratic problems and \mathcal{H} infinite-dimensional (kernel ridge regression). Here the problem corresponds to solving a (ill-conditioned) infinite-dimensional linear system in a reproducing kernel Hilbert space [32]. Since however the sum defining f is finite, the problem can be projected on a subspace of dimension at most n [5], leading to a linear system of dimension $n \times n$. Solving it with the techniques above would lead to a complexity of the order $O(n^2)$, which is not feasible on massive learning problems (e.g., $n \approx 10^7$). Interestingly these problems are usually approximately low-rank, with the rank represented by the so called *effective-dimension* df_λ [13], counting essentially the eigenvalues of the problem larger than λ ,

$$\text{df}_\lambda = \text{Tr}(\mathbf{H}\mathbf{H}_\lambda^{-1}). \quad (1)$$

Note that df_λ is bounded by $\min\{n, L/\lambda\}$ and in many cases $\text{df}_\lambda \ll \min(n, L/\lambda)$. Using suitable projection techniques, like *Nyström* [34] or *random features* [26] it is possible to further reduce the problem to dimension df_λ , for a total cost to find the solution of $O(n\text{df}_\lambda^2)$. Finally recent methods [29], combining suitable projection methods with refined preconditioning techniques, are able to find the solution with precision compatible with the optimal statistical learning error [13] in time that is $O(n\text{df}_\lambda \log(L/\lambda))$, so being essentially independent of the condition number of the problem.

Convex problems and explicit features (logistic regression). When the loss function is *self-concordant* it is possible to leverage the fast techniques for linear systems in approximate Newton algorithms [25] (see more in Sec. 2), to achieve the solution in essentially $O(nd + \min(n, d)^3)$ time, modulo logarithmic terms. However only few loss functions of interest are self-concordant, in particular the widely used logistic and soft-max losses are not self-concordant, but *generalized-self-concordant* [6]. In such cases we need to use (accelerated/stochastic) first order optimization methods to enter in the quadratic convergence region of Newton methods [2], which leads to a solution in $O(dn + d\sqrt{nL/\lambda} + \min(n, d)^3)$ time, which does not present any improvement on a simple accelerated first-order method. Globally convergent second-order methods have also been proposed to solve such problems [21], but the number of Newton steps needed being bounded only by L/λ , they lead to a solution in $O(L/\lambda (nd + \min(n, d)^3))$. With λ that could be as small as 10^{-12} in modern machine learning problems, this makes both these kind of approaches expensive from a computational viewpoint for ill-conditioned problems. For such problems, with our new global second-order scheme, the algorithm we propose achieves instead a complexity of essentially $O((nd + \min(n, d)^3) \log(R^2/\lambda\epsilon))$ (see Thm. 1).

Convex problems and \mathcal{H} infinite-dimensional (kernel logistic regression). Analogously to the case above, it is not possible to use Newton methods profitably as global optimizers on losses that are not self-concordant as we see in Sec. 3. In such cases by combining projecting techniques developed in Sec. 4 and accelerated first-order optimization methods, it is possible to find a solution in $O(n\text{df}_\lambda + \text{df}_\lambda \sqrt{nL/\lambda})$ time. This can still be prohibitive in the very small regularization scenario, since it strongly depends on the condition number L/λ . In Sec. 4 we suitably combine our optimization algorithm with projection techniques achieving optimal statistical learning error [23] in essentially $O(n\text{df}_\lambda \log(R^2/\lambda))$.

First-order algorithms for finite sums. In dimension d , accelerated algorithms for strongly-convex smooth (not necessarily self-concordant) finite sums, such as K-SVRG [4], have a running time proportional $O((n + \sqrt{nL/\lambda})d)$. This can be improved with preconditioning to $O((n + \sqrt{dL/\lambda})d)$ for large n [2]. Quasi-Newton methods can also be used [20], but typically without the guarantees that we provide in this paper (which are logarithmic in the condition number in natural scenarios).

2 Background: Newton methods and generalized self concordance

In this section we start by recalling the definition of generalized self concordant functions and motivate it with examples. We then recall basic facts about Newton and approximate Newton methods, and

present existing techniques to efficiently compute approximate Newton steps. We start by introducing the definition of generalized self-concordance, that here is an extension of the one in [6].

Definition 1 (generalized self-concordant (GSC) function). *Let \mathcal{H} be a Hilbert space. We say that f is a generalized self-concordant function on $\mathcal{G} \subset \mathcal{H}$, when \mathcal{G} is a bounded subset of \mathcal{H} and f is a convex and three times differentiable mapping on \mathcal{H} such that*

$$\forall x \in \mathcal{H}, \forall h, k \in \mathcal{H}, \nabla^{(3)} f(x)[h, k, k] \leq \sup_{g \in \mathcal{G}} |g \cdot h| \nabla^2 f(x)[k, k].$$

We will usually denote by R the quantity $\sup_{g \in \mathcal{G}} \|g\| < \infty$ and often omit \mathcal{G} when it is clear from the context (for simplicity think of \mathcal{G} as the ball in \mathcal{H} centered in zero and with radius $R > 0$, then $\sup_{g \in \mathcal{G}} |g \cdot h| = R\|h\|$). The globally convergent second-order scheme we present in Sec. 3 is specific to losses which satisfy this generalized self-concordance property. The following loss functions, which are widely used in machine learning, are generalized-self-concordant, and motivate this work.

Example 1 (Application to finite-sum minimization). *The following loss functions are generalized self-concordant functions, but not self-concordant:*

- (a) *Logistic regression: $f_i(x) = \log(1 + \exp(-y_i w_i^\top x))$, where $x, w_i \in \mathbb{R}^d$ and $y_i \in \{-1, 1\}$.*
- (b) *Softmax regression: $f_i(x) = \log(\sum_{j=1}^k \exp(x_j^\top w_i)) - x_{y_i}^\top w_i$, where now $x \in \mathbb{R}^{d \times k}$ and $y_i \in \{1, \dots, k\}$ and x_j denotes the j -th column of x .*
- (c) *Generalized linear models with bounded features (see details in [7, Sec. 2.1]), which include conditional random fields [33].*
- (d) *Robust regression: $f_i(x) = \varphi(y_i - w_i^\top x)$ with $\varphi(u) = \log(e^u + e^{-u})$.*

Note that these losses are not *self-concordant* in the sense of [25]. Moreover, even if the losses f_i are self-concordant, the objective function f is not necessarily self-concordant, making any attempt to prove the self-concordance of the objective function f almost impossible.

Newton method (NM). Given $x_0 \in \mathcal{H}$, the Newton method consists in doing the following update:

$$x_{t+1} = x_t - \Delta_\lambda(x_t), \quad \Delta_\lambda(x_t) := \mathbf{H}_\lambda^{-1}(x_t) \nabla f_\lambda(x_t). \quad (2)$$

The quantity $\Delta_\lambda(x) := \mathbf{H}_\lambda^{-1}(x) \nabla f_\lambda(x)$ is called the Newton step at point x , and $x - \Delta_\lambda(x)$ is the minimizer of the second order approximation of f_λ around x . Newton methods enjoy the following key property: if x_0 is close enough to the optimum, the convergence to the optimum is quadratic and the number of iterations required to a given precision is independent of the condition number of the problem [12].

However Newton methods have two main limitations: (a) the region of quadratic convergence can be quite small and reaching the region can be computationally expensive, since it is usually done via first order methods [2] that converge linearly depending on the condition number of the problem, (b) the cost of computing the Hessian can be really expensive when n, d are large, and also (c) the cost of computing $\Delta_\lambda(x_t)$ can be really prohibitive. In the rest of the section we recall some ways to deal with (b) and (c). Our main result of Sec. 3 is to provide globalization scheme for the Newton method to tackle problem (a), which is easily integrable with approximate techniques to deal with (b) and (c), to make second-order technique competitive.

Approximate Newton methods (ANM) and approximate solutions to linear systems. Computing exactly the Newton increment $\Delta_\lambda(x_t)$, which corresponds essentially to the solution of a linear system, can be too expensive when n, d are large. A natural idea is to approximate the Newton iteration, leading to *approximate Newton methods*,

$$x_{t+1} = x_t - \tilde{\Delta}_\lambda(x_t), \quad \tilde{\Delta}_\lambda \approx \Delta_\lambda(x_t). \quad (3)$$

In this paper, more generally we consider any technique to compute $\tilde{\Delta}_\lambda(x_t)$ that provides a *relative approximation* [16] of $\Delta_\lambda(x_t)$ defined as follows.

Definition 2 (relative approximation). *Let $\rho < 1$, let \mathbf{A} be an invertible positive definite Hermitian operator on \mathcal{H} and b in \mathcal{H} . We denote by $\text{LinApprox}(\mathbf{A}, b, \rho)$ the set of all ρ -relative approximations of $z^* = \mathbf{A}^{-1}b$, i.e., $\text{LinApprox}(\mathbf{A}, b, \rho) = \{z \in \mathcal{H} \mid \|z - z^*\|_{\mathbf{A}} \leq \rho \|z^*\|_{\mathbf{A}}\}$.*

Sketching and subsampling for approximate Newton methods. Many techniques for approximating linear systems have been used to compute $\tilde{\Delta}_\lambda$, in particular *sketching* of the Hessian matrix via fast transforms and *subsampling* (see [25, 8, 2] and references therein). Assuming for simplicity that $f_i = \ell_i(w_i^\top x)$, with $\ell_i : \mathbb{R} \rightarrow \mathbb{R}$ and $w_i \in \mathcal{H}$, it holds:

$$\mathbf{H}(x) = \frac{1}{n} \sum_{i=1}^n \ell_i^{(2)}(w_i^\top x) w_i w_i^\top = V_x^\top V_x, \quad (4)$$

with $V_x \in \mathbb{R}^{n \times d} = D_x W$, where $D_x \in \mathbb{R}^{n \times n}$ is a diagonal matrix defined as $(D_x)_{ii} = (\ell_i^{(2)}(w_i^\top x))^{1/2}$ and $W \in \mathbb{R}^{n \times d}$ defined as $W = (w_1, \dots, w_n)^\top$.

Both sketching and subsampling methods approximate $z^* = \mathbf{H}_\lambda(x)^{-1} \nabla f_\lambda(x)$ with $\tilde{z} = \tilde{\mathbf{H}}_\lambda(x)^{-1} \nabla f_\lambda(x)$, in particular, in the case of subsampling $\tilde{\mathbf{H}}_\lambda(x) = \sum_{j=1}^Q p_j w_{i_j} w_{i_j}^\top$ where $Q \ll \min(n, d)$, $(p_j)_{j=1}^Q$ are suitable weights and $(i_j)_{j=1}^Q$ are indices selected at random from $\{1, \dots, n\}$ with suitable probabilities. Sketching methods instead use $\tilde{\mathbf{H}}_\lambda(x) = \tilde{V}_x^\top \tilde{V}_x$, with $\tilde{V}_x = \Omega V_x$ with $\Omega \in \mathbb{R}^{Q \times n}$ a structured matrix such that computing \tilde{V}_x has a cost in the order of $O(nd \log n)$; to this end usually Ω is based on fast Fourier or Hadamard transforms [25]. Note that essentially all the techniques used in approximate Newton methods guarantee relative approximation. In particular the following results can be found in the literature (see Lemmas 28 and 29 in Appendix I and [25], Lemma 2 for more details).

Lemma 1. *Let $x, b \in \mathcal{H}$ and assume that $\ell_i^{(2)} \leq a$ for $a > 0$. With probability $1 - \delta$ the following methods output an element in $\text{LinApprox}(\mathbf{H}_\lambda(x), b, \rho)$, in $O(Q^2 d + Q^3 + c)$ time, $O(Q^2 + d)$ space:*

- (a) *Subsampling with uniform sampling (see [27, 28]), where $Q = O(\rho^{-2} a / \lambda \log \frac{1}{\lambda \delta})$ and $c = O(1)$.*
- (b) *Subsampling with approximate leverage scores [27, 3, 28]), where $Q = O(\rho^{-2} \text{df}_\lambda \log 1 / \lambda \delta)$, $c = O(\min(n, a / \lambda) \text{df}_\lambda^2)$ and $\text{df}_\lambda = \text{Tr}(W^\top W (W^\top W + \lambda / a I)^{-1})$ [30]. Note that $\text{df}_\lambda \leq \min(n, d)$.*
- (c) *Sketching with fast Hadamard transform [25], where $Q = O(\rho^{-2} \text{df}_\lambda \log a / \lambda \delta)$, $c = O(nd \log n)$.*

3 Globally convergent scheme for ANM algorithms on GSC functions

The algorithm is based on the observation that when f_λ is generalized self concordant, there exists a region where t steps of ANM converge as fast as 2^{-t} . Our idea is to start from a very large regularization parameter λ_0 , such that we are sure that x_0 is in the convergence region and perform some steps of ANM such that the solution enters in the convergence region of f_{λ_1} , with $\lambda_1 = q \lambda_0$ with $q < 1$, and to iterate this procedure until we enter the convergence region of f_λ . First we define the region of interest and characterize the behavior of NM and ANM in the region, then we analyze the globalization scheme.

Preliminary results: the Dikin ellipsoid. We consider the following region that we prove to be contained in the region of quadratic convergence for the Newton method and that will be useful to build the globalization scheme. Let $c, R > 0$ and f_λ be generalized self-concordant with coefficient R , we call *Dikin ellipsoid* and denote by $D_\lambda(c)$ the region

$$D_\lambda(c) := \{x \mid \nu_\lambda(x) \leq c\sqrt{\lambda}/R\}, \quad \text{with} \quad \nu_\lambda(x) := \|\nabla f_\lambda(x)\|_{\mathbf{H}_\lambda^{-1}(x)},$$

where $\nu_\lambda(x)$ is usually called the *Newton decrement* and $\|x\|_{\mathbf{A}}$ stands for $\|\mathbf{A}^{1/2}x\|$.

Lemma 2. *Let $\lambda > 0, c \leq 1/7$, let f_λ be generalized self-concordant and $x \in D_\lambda(c)$. Then it holds: $\frac{1}{4}\nu_\lambda(x)^2 \leq f_\lambda(x) - f_\lambda(x_\lambda^*) \leq \nu_\lambda(x)^2$. Moreover Newton method starting from x_0 has quadratic convergence, i.e., let x_t be obtained via $t \in \mathbb{N}$ steps of Newton method in Eq. (2), then $\nu_\lambda(x_t) \leq 2^{-(2^t - 1)}\nu_\lambda(x_0)$. Finally, approximate Newton methods starting from x_0 have a linear convergence rate, i.e., let x_t given by Eq. (3), with $\tilde{\Delta}_t \in \text{LinApprox}(\mathbf{H}_\lambda(x_t), \nabla f_\lambda(x_t), \rho)$ and $\rho \leq 1/7$, then $\nu_\lambda(x_t) \leq 2^{-t}\nu_\lambda(x_0)$.*

This result is proved in Lemma 11 in Appendix B.3. The crucial aspect of the result above is that when $x_0 \in D_\lambda(c)$, the convergence of the approximate Newton method is linear and does not depend on the condition number of the problem. However $D_\lambda(c)$ itself can be very small depending on $\sqrt{\lambda}/R$. In the next subsection we see how to enter in $D_\lambda(c)$ in an efficient way.

Entering the Dikin ellipsoid using a second-order scheme. The lemma above shows that $D_\lambda(c)$ is a good region where to use the approximate Newton algorithm on GSC functions. However the region itself is quite small, since it depends on $\sqrt{\lambda}/R$. Some other globalization schemes arrive to regions of interest by first-order methods or back-tracking schemes [2, 1]. However such approaches require a number of steps that is usually proportional to $\sqrt{L/\lambda}$ making them non-beneficial in machine learning contexts. Here instead we consider the following simple scheme where $\text{ANM}_\rho(f_\lambda, x, t)$ is the result of a ρ -relative approximate Newton method performing t steps of optimization starting from x .

The main ingredient to guarantee the scheme to work is the following lemma (see Lemma 13 in Appendix C.1 for a proof).

Lemma 3. *Let $\mu > 0$, $c < 1$ and $x \in \mathcal{H}$. Let $s = 1 + R\|x\|/c$, then for $q \in [1 - 2/(3s), 1)$*

$$D_\mu(c/3) \subseteq D_{q\mu}(c).$$

Now we are ready to show that we can guarantee the loop invariant $x_k \in D_{\mu_k}(c)$. Indeed assume that $x_{k-1} \in D_{\mu_{k-1}}(c)$. Then $\nu_{\mu_{k-1}}(x_{k-1}) \leq c\sqrt{\mu_{k-1}}/R$. By taking $t = 2$, $\rho = 1/7$, and performing $x_k = \text{ANM}_\rho(f_{\mu_{k-1}}, x_{k-1}, t)$, by Lemma 2, $\nu_{\mu_{k-1}}(x_k) \leq 1/4\nu_{\mu_{k-1}}(x_{k-1}) \leq c/4\sqrt{\mu_{k-1}}/R$, i.e., $x_k \in D_{\mu_{k-1}}(c/4)$. If q_k is large enough, this implies that $x_k \in D_{q_k\mu_{k-1}}(c) = D_{\mu_k}(c)$, by Lemma 3. Now we are ready to state our main theorem of this section.

Proposed Globalization Scheme

Phase I: Getting in the Dikin ellipsoid of f_λ

Start with $x_0 \in \mathcal{H}$, $\mu_0 > 0$, $t, T \in \mathbb{N}$ and $(q_k)_{k \in \mathbb{N}} \in (0, 1]$.

For $k \in \mathbb{N}$

$$x_{k+1} \leftarrow \text{ANM}_\rho(f_{\mu_k}, x_k, t)$$

$$\mu_{k+1} \leftarrow q_{k+1}\mu_k$$

Stop when $\mu_{k+1} < \lambda$ and set $x_{last} \leftarrow x_k$.

Phase II: reach a certain precision starting from inside the Dikin ellipsoid

Return $\hat{x} \leftarrow \text{ANM}_\rho(f_\lambda, x_{last}, T)$

Fully adaptive method. The scheme presented above converges with the following parameters.

Theorem 1. *Let $\epsilon > 0$. Set $\mu_0 = 7R\|\nabla f(0)\|$, $x_0 = 0$, and perform the globalization scheme above for $\rho \leq 1/7$, $t = 2$, and $q_k = \frac{1/3+7R\|x_k\|}{1+7R\|x_k\|}$, $T = \lceil \log_2 \sqrt{1 \vee (\lambda\epsilon^{-1}/R^2)} \rceil$. Then denoting by K the number of steps performed in the Phase I, it holds:*

$$f_\lambda(\hat{x}) - f_\lambda(x_\lambda^*) \leq \epsilon, \quad K \leq \lfloor (3 + 11R\|x_\lambda^*\|) \log(7R\|\nabla f(0)\|/\lambda) \rfloor.$$

Note that the theorem above (proven in Appendix C.3) guarantees a solution with error ϵ with K steps of ANM each performing 2 iterations of approximate linear system solving, plus a final step of ANM which performs T iterations of approximate linear system solving. In case of $f_i(x) = \ell_i(w_i^\top x)$, with $\ell_i: \mathbb{R} \rightarrow \mathbb{R}$, $w_i \in \mathcal{H}$ with $\ell_i^{(2)} \leq a$, for $a > 0$, the final runtime cost of the proposed scheme to achieve precision ϵ , when combined with of the methods for approximate linear system solving from Lemma 1 (i.e. sketching), is $O(Q^2 + d)$ in memory and

$$O\left((nd \log n + dQ^2 + Q^3) \left(R\|x_\lambda^*\| \log \frac{R}{\lambda} + \log \frac{\lambda}{R\epsilon} \right)\right) \text{ in time,} \quad Q = O\left(d\bar{f}_\lambda \log \frac{1}{\lambda\delta}\right),$$

where $d\bar{f}_\lambda$, defined in Lemma 1, measures the *effective dimension* of the correlation matrix $W^\top W$ with $W = (w_1, \dots, w_n)^\top \in \mathbb{R}^{n \times d}$, corresponding essentially to the number of eigenvalues of $W^\top W$ larger than λ/a . In particular note that $d\bar{f}_\lambda \leq \min(n, d, \text{rank}(W), ab^2/\lambda)$, with $b := \max_i \|w_i\|$, and usually way smaller than such quantities.

Remark 1. *The proposed method does not depend on the condition number of the problem L/λ , but on the term $R\|x_\lambda^*\|$ which can be in the order of $R/\sqrt{\lambda}$ in the worst case, but usually way smaller. For example, it is possible to prove that this term is bounded by an absolute constant not depending on λ , if at least one minimum for f exists. In the regularity (see Proposition 7), we show a variant of this adaptive method which can leverage the regularity of the solution with respect to the Hessian, i.e., depending on the smaller quantity $R\sqrt{\lambda}\|x_\lambda^*\|_{\mathbf{H}_\lambda^{-1}(x_\lambda^*)}$ instead of $R\|x_\lambda^*\|$.*

Finally note that it is possible to use $q_k = q$ fixed for all the iterations and way smaller than the one in Thm. 1, depending on some regularity properties of \mathbf{H} (see Proposition 8 in Appendix C.2).

4 Application to the non-parametric setting: Kernel methods

In supervised learning the goal is to predict well on future data, given the observed training dataset. Let \mathcal{X} be the input space and $\mathcal{Y} \subseteq \mathbb{R}^p$ be the output space. We consider a probability distribution P over $\mathcal{X} \times \mathcal{Y}$ generating the data and the goal is to estimate $g^* : \mathcal{X} \rightarrow \mathcal{Y}$ solving the problem

$$g^* = \arg \min_{g: \mathcal{X} \rightarrow \mathcal{Y}} \mathcal{L}(g), \quad \mathcal{L}(g) = \mathbb{E}[\ell(g(x), y)], \quad (5)$$

for a given loss function $\ell : \mathcal{Y} \times \mathcal{Y} \rightarrow \mathbb{R}$. Note that P is not known, and accessible only via the dataset $(x_i, y_i)_{i=1}^n$, with $n \in \mathbb{N}$, independently sampled from P . A prototypical estimator for g^* is the regularized minimizer of the empirical risk $\widehat{\mathcal{L}}(g) = \frac{1}{n} \sum_{i=1}^n \ell(g(x_i), y_i)$ over a suitable space of functions \mathcal{G} . Given $\phi : \mathcal{X} \rightarrow \mathcal{H}$ a common choice is to select \mathcal{G} as the set of linear functions of $\phi(x)$, that is, $\mathcal{G} = \{w^\top \phi(\cdot) \mid w \in \mathcal{H}\}$. Then the regularized minimizer of $\widehat{\mathcal{L}}$, denoted by \widehat{g}_λ , corresponds to

$$\widehat{g}_\lambda(x) = \widehat{w}_\lambda^\top \phi(x), \quad \widehat{w}_\lambda = \arg \min_{w \in \mathcal{H}} \frac{1}{n} \sum_{i=1}^n f_i(w) + \lambda \|w\|^2, \quad f_i(w) = \ell(w^\top \phi(x_i), y_i). \quad (6)$$

Learning theory guarantees how fast \widehat{g}_λ converges to the best possible estimator g^* with respect to the number of observed examples, in terms of the so called *excess risk* $\mathcal{L}(\widehat{g}_\lambda) - \mathcal{L}(g^*)$. The following theorem recovers the minimax optimal learning rates for squared loss and extend them to any generalized self-concordant loss function.

Note on df_λ . In this section, we always denote with df_λ the effective dimension of the problem in Eq. (5). When the loss belongs to the family of generalized linear models (see Example 1) and if the model is well-specified, then df_λ is defined exactly as in Eq. (1) otherwise we need a more refined definition (see [23] or Eq. (30) in Appendix D).

Theorem 2 (from [23], Thm. 4). *Let $\lambda > 0, \delta \in (0, 1]$. Let ℓ be generalized self-concordant with parameter $R > 0$ and $\sup_{x \in \mathcal{X}} \|\phi(x)\| \leq C < \infty$. Assume that there exists g^* minimizing \mathcal{L} . Then there exists c_0 not depending on $n, \lambda, \delta, \text{df}_\lambda, C, g^*$, such that if $\sqrt{\text{df}_\lambda/n}, \text{b}_\lambda \leq \lambda^{1/2}/R$, and $n \geq C/\lambda \log(\delta^{-1}C/\lambda)$ the following holds with probability $1 - \delta$:*

$$\mathcal{L}(\widehat{g}_\lambda) - \mathcal{L}(g^*) \leq c_0 \left(\frac{\text{df}_\lambda}{n} + \text{b}_\lambda^2 \right) \log(1/\delta), \quad \text{b}_\lambda := \lambda \|g^*\|_{\mathbf{H}_\lambda^{-1}(g^*)}. \quad (7)$$

Under standard regularity assumptions of the learning problems [23], i.e., (a) the *capacity condition* $\sigma_j(\mathbf{H}(g^*)) \leq Cj^{-\alpha}$, for $\alpha \geq 1, C > 0$ (i.e., a decay of eigenvalues $\sigma_j(\mathbf{H}(g^*))$ of the Hessian at the optimum), and (b) the *source condition* $g^* = \mathbf{H}(g^*)^r v$, with $v \in \mathcal{H}$ and $r > 0$ (i.e., the control of the optimal g^* for a specific Hessian-dependent norm), $\text{df}_\lambda \leq C'\lambda^{-1/\alpha}$ and $\text{b}_\lambda^2 \leq C''\lambda^{1+2r}$, leading to the following optimal learning rate,

$$\mathcal{L}(\widehat{g}_\lambda) - \mathcal{L}(g^*) \leq c_1 n^{-\frac{1+2r\alpha}{1+\alpha+2r\alpha}} \log(1/\delta), \quad \text{when } \lambda = n^{-\frac{\alpha}{1+\alpha+2r\alpha}}. \quad (8)$$

Now we propose an algorithmic scheme to compute efficiently an approximation of \widehat{g}_λ that achieves the same optimal learning rates. First we need to introduce the technique we are going to use.

Nyström projection. It consists in suitably selecting $\{\bar{x}_1, \dots, \bar{x}_M\} \subset \{x_1, \dots, x_n\}$, with $M \ll n$ and computing $\bar{g}_{M,\lambda}$, i.e., the solution of Eq. (6) over $\mathcal{H}_M = \text{span}\{\phi(\bar{x}_1), \dots, \phi(\bar{x}_M)\}$ instead of \mathcal{H} . In this case the problem can be reformulated as a problem in \mathbb{R}^M as

$$\bar{g}_{M,\lambda} = \bar{\alpha}_{M,\lambda}^\top \mathbf{T}^{-1} v(x), \quad \bar{\alpha}_{M,\lambda} = \arg \min_{\alpha \in \mathbb{R}^M} \bar{f}_\lambda(\alpha), \quad \bar{f}_\lambda(\alpha) = \frac{1}{n} \sum_{i=1}^n \bar{f}_i(\alpha) + \lambda \|\alpha\|^2, \quad (9)$$

where $\bar{f}_i(\alpha) = \ell(v(x_i)^\top \mathbf{T}^{-1} \alpha, y_i)$ and $v(x) \in \mathbb{R}^M$, $v(x) = (k(x, \bar{x}_1), \dots, k(x, \bar{x}_M))$ with $k(x, x') = \phi(x)^\top \phi(x')$ the associated positive-definite kernel [32], while \mathbf{T} is the upper triangular matrix such that $\mathbf{K} = \mathbf{T}^\top \mathbf{T}$, with $\mathbf{K} \in \mathbb{R}^{M \times M}$ with $\mathbf{K}_{ij} = k(\bar{x}_i, \bar{x}_j)$. In the next theorem we characterize the sufficient M to achieve minimax optimal rates, for two standard techniques of choosing the Nyström points $\{\bar{x}_1, \dots, \bar{x}_M\}$.

Theorem 3 (Optimal rates for learning with Nyström). *Let $\lambda > 0, \delta \in (0, 1]$. Assume the conditions of Thm. 2. Then the excess risk of $\bar{g}_{M,\lambda}$ is bounded with prob. $1 - 2\delta$ as in Eq. (7) (with $c'_1 \propto c_1$), when*

- (1) *Uniform Nyström method [28, 29] is used and $M \geq C_1/\lambda \log(C_2/\lambda\delta)$.*
- (2) *Approximate leverage score method [3, 28, 29] is used and $M \geq C_3 \text{df}_\lambda \log(C_4/\lambda\delta)$.*

Here C, C_1, C_2, C_4 do not depend on $\lambda, n, M, \text{df}_\lambda, \delta$.

Thm. 3 generalizes results for learning with Nyström and squared loss [28], to GSC losses. It is proved in Thm. 6, in Appendix D.4. As in [28], Thm. 3 shows that Nyström is a valid technique for dimensionality reduction. Indeed it is essentially possible to project the learning problem on a subspace \mathcal{H}_M of dimension $M = O(c/\lambda)$ or even as small as $M = O(\text{df}_\lambda)$ and still achieve the optimal rates of Thm. 2. Now we are ready to introduce our algorithm.

Proposed algorithm. The algorithm conceptually consists in (a) performing a projection step with Nyström, and (b) solving the resulting optimization problem with the globalization scheme proposed in Sec. 3 based on ANM in Eq. (3). In particular, we want to avoid to apply explicitly \mathbf{T}^{-1} to each $v(x_i)$ in Eq. (9), which would require $O(nM^2)$ time. Then we will use the following approximation technique based only on matrix vector products, so we can just apply \mathbf{T}^{-1} to α at each iteration, with a total cost proportional only to $O(nM + M^2)$ per iteration. Given $\alpha, \nabla \bar{f}_\lambda(\alpha)$, we approximate $z^* = \bar{\mathbf{H}}_\lambda(\alpha)^{-1} \nabla \bar{f}_\lambda(\alpha)$, where $\bar{\mathbf{H}}_\lambda$ is the Hessian of $\bar{f}_\lambda(\alpha)$, with \tilde{z} defined as

$$\tilde{z} = \text{prec-conj-grad}_t(\bar{\mathbf{H}}_\lambda(\alpha), \nabla \bar{f}_\lambda(\alpha)),$$

where prec-conj-grad_t corresponds to performing t steps of preconditioned conjugate gradient [19] with preconditioner computed using a subsampling approach for the Hessian among the ones presented in Sec. 2, in the paragraph starting with Eq. (4). The pseudocode for the whole procedure is presented in Alg. 1, Appendix E. This technique of approximate linear system solving has been studied in [29] in the context of empirical risk minimization for squared loss.

Lemma 4 ([29]). *Let $\lambda > 0, \alpha, b \in \mathbb{R}^M$. The previous method, applied with $t = O(\log 1/\rho)$, outputs an element of $\text{LinApprox}(\bar{\mathbf{H}}_\lambda(\alpha), b, \rho)$, with probability $1 - \delta$ with complexity $O((nM + M^2Q + M^3 + c)t)$ in time and $O(M^2 + n)$ in space, with $Q = O(C_1/\lambda \log(C_1/\lambda\delta))$, $c = O(1)$ if uniform sub-sampling is used or $Q = O(C_2 \text{df}_\lambda \log(C_1/\lambda\delta))$, $c = O(\text{df}_\lambda^2 \min(n, \frac{1}{\lambda}))$ if sub-sampling with leverage scores is used [30].*

A more complete version of this lemma is shown in Proposition 12 in Appendix D.5.1. We conclude this section with a result proving the learning properties of the proposed algorithm.

Theorem 4 (Optimal rates for the proposed algorithms). *Let $\lambda > 0$ and $\epsilon < \lambda/R^2$. Under the hypotheses of Thm. 3, if we set M as in Thm. 3, Q as in Lemma 4 and setting the globalization scheme as in Thm. 1, then the proposed algorithm (Alg. 1, Appendix E) finishes in a finite number of newton steps $N_{ns} = O(R\|g^*\| \log(C/\lambda) + \log(C/\epsilon))$ and returns a predictor $g_{Q,M,\lambda}$ of the form $g_{Q,M,\lambda} = \alpha^\top \mathbf{T}^{-1}v(x)$. With probability at least $1 - \delta$, this predictor satisfies:*

$$\mathcal{L}(g_{Q,M,\lambda}) - \mathcal{L}(g^*) \leq c_0 \left(\frac{\text{df}_\lambda}{n} + b_\lambda^2 + \epsilon \right) \log(1/\delta), \quad b_\lambda := \lambda \|g^*\|_{\mathbf{H}_\lambda^{-1}(g^*)}. \quad (10)$$

The theorem above (see Proposition 14, Appendix D.6 for exact quantifications) shows that the proposed algorithm is able to achieve the same learning rates of plain empirical risk minimization as in Thm. 2. The total complexity of the procedure, including the cost of computing the preconditioner, the selection of the Nyström points via approximate leverage scores and also the computation of the leverage scores [30] is then

$$O(R\|g^*\| \log(R^2/\lambda) (n \text{df}_\lambda \log(C\lambda^{-1}\delta^{-1}) c_X + \text{df}_\lambda^3 \log^3(C\lambda^{-1}\delta^{-1}) + \min(n, C/\lambda) \text{df}_\lambda^2))$$

in time and $O(\text{df}_\lambda^2 \log^2(C\lambda^{-1}\delta^{-1}))$ in space, where c_X is the cost of computing the inner product $k(x, x')$ (in the kernel setting assumed when the input space X is $X = \mathbb{R}^p$ it is $c = O(p)$). As noted in [30], under the standard regularity assumptions on the learning problem seen above, $\text{df}_\lambda^2 \leq \text{df}_\lambda/\lambda \leq n$ when the optimal λ is chosen. So the total computational complexity is

$$O(R \log(R^2/\lambda) \log^3(C\lambda^{-1}\delta^{-1}) \|g^*\| \cdot n \cdot \text{df}_\lambda \cdot c_X) \text{ in time, } O(\text{df}_\lambda^2 \cdot \log^2(C\lambda^{-1}\delta^{-1})) \text{ in space.}$$

First note, the fact that due to the statistical properties of the problem the complexity does not depend even implicitly on $\sqrt{C/\lambda}$, but only on $\log(C/\lambda)$, so the algorithm runs in essentially $O(n\text{df}_\lambda)$, compared to $O(\text{df}_\lambda \sqrt{nC/\lambda})$ of the accelerated first-order methods we develop in Appendix F and the $O(n\text{df}_\lambda \sqrt{C/\lambda})$ of other Newton schemes (see Sec. 1.1). To our knowledge, this is the first algorithm to achieve optimal statistical learning rates for generalized self-concordant losses and with complexity only $\tilde{O}(n\text{df}_\lambda)$. This generalizes similar results for squared loss [29, 30].




Figure 1: Training loss and test error as a function of the number of passes on the data for our algorithm vs. K-SVRG. on the **(left)** Susy and **(right)** Higgs data sets.

5 Experiments

The code necessary to reproduce the following experiments is available on GitHub at <https://github.com/umarteau/Newton-Method-for-GSC-losses->.

We compared the performances of our algorithm for kernel logistic regression on two large scale classification datasets ($n \approx 10^7$), Higgs and Susy, pre-processed as in [29]. We implemented the algorithm in pytorch and performed the computations on 1 Tesla P100-PCIE-16GB GPU. For Susy ($n = 5 \times 10^6, p = 18$): we used Gaussian kernel with $k(x, x') = e^{-\|x-x'\|^2/(2\sigma^2)}$, with $\sigma = 5$, which we obtained through a grid search (in [29], $\sigma = 4$ is taken for the ridge regression); $M = 10^4$ Nyström centers and a subsampling $Q = M$ for the preconditioner, both obtained with uniform sampling. Analogously for Higgs ($n = 1.1 \times 10^7, p = 28$): we used a Gaussian kernel with $\sigma = 5$ and $M = 2.5 \times 10^4$ and $Q = M$, using again uniform sampling. To find reasonable λ for supervised learning applications, we cross-validated λ finding the minimum test error at $\lambda = 10^{-10}$ for Susy and $\lambda = 10^{-9}$ for Higgs (see Figs. 2 and 3 in Appendix F) for such values our algorithm and the competitor achieve an error of 19.5% on the test set for Susy, comparable to the state of the art (19.6% [29]) and analogously for Higgs (see Appendix F). We then used such λ 's as regularization parameters and compared our algorithm with a well known accelerated stochastic gradient technique *Katyusha SVRG* (K-SVRG) [4], tailored to our problem using mini batches. In Fig. 1 we show the convergence of the training loss and classification error with respect to the number of passes on the data, of our algorithm compared to K-SVRG. It is possible to note our algorithm is order of magnitude faster in achieving convergence, validating empirically the fact that the proposed algorithm scales as $O(ndf_\lambda)$ in learning settings, while accelerated first order methods go as $O((n + \sqrt{nL/\lambda})df_\lambda)$. Moreover, as mentioned in the introduction, this highlights the fact that precise optimization is necessary to achieve a good performance in terms of test error. Finally, note that since a pass on the data is much more expensive for K-SVRG than for our second order method (see Appendix F for details), the difference in computing time between the second order scheme and K-SVRG is even more in favour of our second order method (see Figs. 4 and 5 in Appendix F).

Acknowledgments

We acknowledge support from the European Research Council (grant SEQUOIA 724063).

References

- [1] Murat A. Erdogdu and Andrea Montanari. Convergence rates of sub-sampled Newton methods. Technical Report 1508.02810, ArXiv, 2015.
- [2] Naman Agarwal, Brian Bullins, and Elad Hazan. Second-order stochastic optimization for machine learning in linear time. *J. Mach. Learn. Res.*, 18(1):4148–4187, January 2017.
- [3] Ahmed Alaoui and Michael W Mahoney. Fast randomized kernel ridge regression with statistical guarantees. In *Advances in Neural Information Processing Systems*, pages 775–783, 2015.

- [4] Zeyuan Allen-Zhu. Katyusha: The first direct acceleration of stochastic gradient methods. In *Proceedings of the Symposium on Theory of Computing*, pages 1200–1205, 2017.
- [5] Nachman Aronszajn. Theory of reproducing kernels. *Transactions of the American Mathematical Society*, 68(3):337–404, 1950.
- [6] Francis Bach. Self-concordant analysis for logistic regression. *Electronic Journal of Statistics*, 4:384–414, 2010.
- [7] 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.
- [8] Raghu Bollapragada, Richard H. Byrd, and Jorge Nocedal. Exact and inexact subsampled newton methods for optimization. *IMA Journal of Numerical Analysis*, 39(2):545–578, 2018.
- [9] Léon Bottou and Olivier Bousquet. The tradeoffs of large scale learning. In *Advances in Neural Information Processing Systems*, pages 161–168, 2008.
- [10] Léon Bottou, Frank E. Curtis, and Jorge Nocedal. Optimization methods for large-scale machine learning. *Siam Review*, 60(2):223–311, 2018.
- [11] Christos Boutsidis and Alex Gittens. Improved matrix algorithms via the subsampled randomized hadamard transform. *SIAM Journal on Matrix Analysis and Applications*, 34(3):1301–1340, 2013.
- [12] Stephen Boyd and Lieven Vandenberghe. *Convex optimization*. Cambridge university press, 2004.
- [13] A. Caponnetto and E. De Vito. Optimal rates for the regularized least-squares algorithm. *Found. Comput. Math.*, 7(3):331–368, July 2007.
- [14] Aaron Defazio. A simple practical accelerated method for finite sums. In *Advances in Neural Information Processing Systems*, pages 676–684, 2016.
- [15] Aaron Defazio, Francis Bach, and Simon Lacoste-Julien. SAGA: A fast incremental gradient method with support for non-strongly convex composite objectives. In *Advances in Neural Information Processing systems*, pages 1646–1654, 2014.
- [16] Peter Deufhard. *Newton Methods for Nonlinear Problems: Affine Invariance and Adaptive Algorithms*. Springer, 2011.
- [17] Petros Drineas, Michael W Mahoney, Shan Muthukrishnan, and Tamás Sarlós. Faster least squares approximation. *Numerische mathematik*, 117(2):219–249, 2011.
- [18] Petros Drineas, Malik Magdon-Ismail, Michael W Mahoney, and David P Woodruff. Fast approximation of matrix coherence and statistical leverage. *Journal of Machine Learning Research*, 13(Dec):3475–3506, 2012.
- [19] Gene H. Golub and Charles F. Van Loan. *Matrix Computations*, volume 3. JHU Press, 2012.
- [20] Robert Gower, Filip Hanzely, Peter Richtárik, and Sebastian U. Stich. Accelerated stochastic matrix inversion: general theory and speeding up BFGS rules for faster second-order optimization. In *Advances in Neural Information Processing Systems*, pages 1619–1629, 2018.
- [21] Sai Praneeth Karimireddy, Sebastian U. Stich, and Martin Jaggi. Global linear convergence of newton’s method without strong-convexity or lipschitz gradients. *CoRR*, abs/1806.00413, 2018. URL <http://arxiv.org/abs/1806.00413>.
- [22] Hongzhou Lin, Julien Mairal, and Zaid Harchaoui. A universal catalyst for first-order optimization. In *Advances in Neural Information Processing Systems*, pages 3384–3392, 2015.
- [23] Ulysse Marteau-Ferey, Dmitrii Ostrovskii, Francis Bach, and Alessandro Rudi. Beyond least-squares: Fast rates for regularized empirical risk minimization through self-concordance. In *Proceedings of the Conference on Computational Learning Theory*, 2019.

- [24] Arkadii Nemirovskii and Yurii Nesterov. Interior-point polynomial algorithms in convex programming. *Society for Industrial and Applied Mathematics*, 1994.
- [25] Mert Pilanci and Martin J Wainwright. Newton sketch: A near linear-time optimization algorithm with linear-quadratic convergence. *SIAM Journal on Optimization*, 27(1):205–245, 2017.
- [26] Ali Rahimi and Benjamin Recht. Random features for large-scale kernel machines. In *Advances in Neural Information Processing Systems*, pages 1177–1184, 2008.
- [27] Farbod Roosta-Khorasani and Michael W. Mahoney. Sub-sampled Newton methods. *Math. Program.*, 174(1-2):293–326, 2019.
- [28] Alessandro Rudi, Raffaello Camoriano, and Lorenzo Rosasco. Less is more: Nyström computational regularization. In *Advances in Neural Information Processing Systems 28*, pages 1657–1665. 2015.
- [29] Alessandro Rudi, Luigi Carratino, and Lorenzo Rosasco. FALKON: An optimal large scale kernel method. In *Advances in Neural Information Processing Systems 30*, pages 3888–3898. 2017.
- [30] Alessandro Rudi, Daniele Calandriello, Luigi Carratino, and Lorenzo Rosasco. On fast leverage score sampling and optimal learning. In *Advances in Neural Information Processing Systems*, pages 5672–5682, 2018.
- [31] Y. Saad. *Iterative Methods for Sparse Linear Systems*. Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 2nd edition, 2003.
- [32] John Shawe-Taylor and Nello Cristianini. *Kernel Methods for Pattern Analysis*. Cambridge University Press, 2004.
- [33] Charles Sutton and Andrew McCallum. An introduction to conditional random fields. *Foundations and Trends® in Machine Learning*, 4(4):267–373, 2012.
- [34] Christopher K. I. Williams and Matthias Seeger. Using the Nyström method to speed up kernel machines. In *Advances in Neural Information Processing Systems*, pages 682–688, 2001.

Organization of the Appendix

A. Main results on generalized self-concordant functions

Notations, definitions and basic results concerning generalized self-concordant functions.

B. Results on approximate Newton methods

In this section, the interaction between the notion of Dikin ellipsoid, approximate Newton methods and generalized self-concordant functions is studied. The results needed in the main paper are all concentrated in Appendix B.3. In particular the results in Lemma 2 are proven in a more general form in Lemma 11.

C. Proof of bounds for the globalization scheme

In this section, we leverage the results of the previous two sections to analyze the globalization scheme.

C.1. Main technical lemmas

We start by proving the result on the inclusion of Dikin ellipsoids (Lemma 3).

C.2. Proof of main theorems

In particular, a general version of Thm. 1 is proven. Moreover Remark 1 is proven in Proposition 7, while the fixed scheme to choose $(q_k)_{k \in \mathbb{N}}$ is proven in Proposition 8.

C.3. Proof of Thm. 1

Finally, we prove the properties of the globalization schemes presented in Thm. 1.

D. Non-parametric learning with generalized self-concordant functions

In this section, some basic results about non-parametric learning with generalized self-concordant functions are recalled and the main results of Sec. 4 are proven.

D.1. General setting and assumptions, statistical result for regularized ERM.

More details about the generalization properties of empirical risk minimization as well as the optimal rates in Thm. 2 are recalled.

D.2. Reducing the dimension: projecting on a subspace using Nyström sub-sampling.

D.3. Sub-sampling techniques.

The basics of uniform sub-sampling and sub-sampling with approximate leverage scores are recalled.

D.4. Selecting the M Nyström points

Thm. 3 is proven in a more general version in Thm. 6.

D.5 Performing the globalization scheme to approximate $\beta_{M,\lambda}$

A general scheme is proposed to solve the projected problem approximately using the globalization scheme.

D.5.1. Performing approximate Newton steps

We start by describing the way of computing approximate Newton steps. A generalized version of Lemma 4 is proven in Proposition 12.

D.5.2. Applying the globalization scheme to control $\widehat{v}_{M,\lambda}(\beta)$

We then completely analyse the approximating of $\beta_{M,\lambda}$ from an optimization point of view (see Proposition 13).

D.6. Final algorithm and results

Finally, the proof of Thm. 4 is provided, using the results of the previous subsections.

E. Algorithm

In this section, the pseudocode for the algorithm presented in Sec. 4 and analyzed in Thm. 7 is provided.

F. Experiments

In this section, more details about the experiments are provided.

G. Solving a projected problem to reduce dimension

In this section, more details about the problem of randomized projections are provided.

G.2. Relating the projected to the original problem

In particular, results to relate the ERM with the projected ERM in terms of excess risk are provided for generalized self-concordant functions.

H. Relations between statistical problems and empirical problem.

In this section, we provide results to relate excess expected risk with excess empirical risk for generalized self-concordant functions.

I. Multiplicative approximations for Hermitian operators

In this section, some general analytic results on multiplicative approximations for Hermitian operators are derived. Moreover they are used to provide a simplified proof for the results in Lemma 1. See in particular Lemmas 28 and 29 and [25], Lemma 2.

A Main results on generalized self-concordant functions

In this section, we start by introducing a few notations. We define the key notion of generalized self-concordance in Appendix A.1, and present the main results concerning generalized self-concordant functions. In Appendix A.2, we describe how generalized self-concordance behaves with respect to an expectation or to certain relaxations.

Notations Let $\lambda \geq 0$ and \mathbf{A} be a bounded positive semidefinite Hermitian operator on \mathcal{H} . We denote with \mathbf{I} the identity operator, and

$$\|x\|_{\mathbf{A}} := \|\mathbf{A}^{1/2}x\|, \quad (11)$$

$$\mathbf{A}_{\lambda} := \mathbf{A} + \lambda\mathbf{I}. \quad (12)$$

Let f be a twice differentiable convex function on a Hilbert space \mathcal{H} . We adopt the following notation for the Hessian of f :

$$\forall x \in \mathcal{H}, \mathbf{H}_f(x) := \nabla^2 f(x) \in \mathcal{L}(\mathcal{H}).$$

For any $\lambda > 0$, we define the λ -regularization of f :

$$f_{\lambda} := f + \frac{\lambda}{2} \|\cdot\|^2.$$

f_{λ} is λ -strongly convex and has a unique minimizer which we denote with $x_{*}^{f,\lambda}$. Moreover, define

$$\forall x \in \mathcal{H}, \mathbf{H}_{f,\lambda}(x) := \nabla^2 f_{\lambda}(x) = \mathbf{H}_f(x) + \lambda\mathbf{I}, \quad \nu_{f,\lambda}(x) := \|\nabla f_{\lambda}(x)\|_{\mathbf{H}_{f,\lambda}^{-1}(x)}.$$

The quantity $\nu_{f,\lambda}(x)$ is called the **Newton decrement** at point x and will play a significant role.

When the function f is clear from the context, we will omit the subscripts with f and use $\mathbf{H}, \mathbf{H}_{\lambda}, \nu_{\lambda}, \dots$

A.1 Definitions and results on generalized self-concordant functions

In this section, we introduce the main definitions and results for self-concordant functions. These results are mainly the same as in appendix B of [23].

Definition 3 (generalized self-concordant function). *Let \mathcal{H} be a Hilbert space. Formally, a generalized self-concordant function on \mathcal{H} is a couple (f, \mathcal{G}) where:*

- i \mathcal{G} is a bounded subset of \mathcal{H} ; we will usually denote $\|\mathcal{G}\|$ or R the quantity $\sup_{g \in \mathcal{G}} \|g\| < \infty$;*
- ii f is a convex and three times differentiable mapping on \mathcal{H} such that*

$$\forall x \in \mathcal{H}, \forall h, k \in \mathcal{H}, \nabla^{(3)} f(x)[h, k, k] \leq \sup_{g \in \mathcal{G}} |g \cdot h| \nabla^2 f(x)[k, k].$$

To make notations lighter, we will often omit \mathcal{G} from the notations and simply say that f stands both for the mapping and the couple (f, \mathcal{G}) .

Definition 4 (Definitions). *Let f be a generalized self-concordant function. We define the following quantities.*

- $\forall h \in \mathcal{H}, \mathfrak{t}_f(h) := \sup_{g \in \mathcal{G}} |h \cdot g|;$
- $\forall x \in \mathcal{H}, \forall \lambda > 0, \mathfrak{r}_{f,\lambda}(x) := \frac{1}{\sup_{g \in \mathcal{G}} \|g\|_{\mathbf{H}_{f,\lambda}^{-1}(x)}};$
- $\forall c \geq 0, \forall \lambda > 0, \mathfrak{D}_{f,\lambda}(c) := \{x : \nu_{f,\lambda}(x) \leq c \mathfrak{r}_{f,\lambda}(x)\}.$

We also define the following functions:

$$\psi(t) = \frac{e^t - t - 1}{t^2}, \quad \underline{\phi}(t) = \frac{1 - e^{-t}}{t}, \quad \bar{\phi}(t) = \frac{e^t - 1}{t}. \quad (13)$$

Note that $\psi, \bar{\phi}$ are increasing functions and that $\underline{\phi}$ is a decreasing function. Moreover, $\frac{\bar{\phi}(t)}{\underline{\phi}(t)} = e^t$. Once again, if f is clear, we will often omit the reference to f in the quantities above, keeping only $\mathfrak{t}, \mathfrak{r}_\lambda, \mathfrak{D}_\lambda \dots$

We condense results obtained in [23] under a slightly different form. The proofs, however, are exactly the same.

While in [23], only the regularized case is dealt with, the proof techniques are exactly the same to obtain Proposition 1. Proposition 2 is proved explicitly in Proposition 4 of [23] and Lemma 5 is proved in Proposition 5.

Omitting the subscript f , we get the following results.

Proposition 1 (Bounds for the non-regularized function f). *Let f be a generalized self-concordant function. Then the following bounds hold (we omit f in the subscripts):*

$$\forall x \in \mathcal{H}, \forall h \in \mathcal{H}, e^{-\mathfrak{t}(h)} \mathbf{H}(x) \preceq \mathbf{H}(x+h) \preceq e^{\mathfrak{t}(h)} \mathbf{H}(x), \quad (14)$$

$$\forall x, h \in \mathcal{H}, \forall \lambda > 0, \|\nabla f(x+h) - \nabla f(x)\|_{\mathbf{H}_\lambda^{-1}(x)} \leq \bar{\phi}(\mathfrak{t}(h)) \|h\|_{\mathbf{H}_\lambda(x)}, \quad (15)$$

$$\forall x, h \in \mathcal{H}, \psi(-\mathfrak{t}(h)) \|h\|_{\mathbf{H}(x)}^2 \leq f(x+h) - f(x) - \nabla f(x) \cdot h \leq \psi(\mathfrak{t}(h)) \|h\|_{\mathbf{H}(x)}^2. \quad (16)$$

We get the analogous bounds in the regularized case.

Proposition 2 (Bounds for the regularized function f_λ). *Let f be a generalized self-concordant function and $\lambda > 0$ be a regularizer. Then the following bounds hold:*

$$\forall x, h \in \mathcal{H}, e^{-\mathfrak{t}(h)} \mathbf{H}_\lambda(x) \preceq \mathbf{H}_\lambda(x+h) \preceq e^{\mathfrak{t}(h)} \mathbf{H}_\lambda(x), \quad (17)$$

$$\forall x, h \in \mathcal{H}, \underline{\phi}(\mathfrak{t}(h)) \|h\|_{\mathbf{H}_\lambda(x)} \leq \|\nabla f_\lambda(x+h) - \nabla f_\lambda(x)\|_{\mathbf{H}_\lambda^{-1}(x)} \leq \bar{\phi}(\mathfrak{t}(h)) \|h\|_{\mathbf{H}_\lambda(x)}, \quad (18)$$

$$\forall x, h \in \mathcal{H}, \psi(-\mathfrak{t}(h)) \|h\|_{\mathbf{H}_\lambda(x)}^2 \leq f_\lambda(x+h) - f_\lambda(x) - \nabla f_\lambda(x) \cdot h \leq \psi(\mathfrak{t}(h)) \|h\|_{\mathbf{H}_\lambda(x)}^2. \quad (19)$$

Corollary 1. *Let f be a \mathcal{G} generalized self-concordant function and $\lambda > 0$ be a regularizer, and x_λ^* the unique minimizer of f_λ . Then the following bounds hold for any $x \in \mathcal{H}$:*

$$\underline{\phi}(\mathfrak{t}(x - x_\lambda^*)) \|x - x_\lambda^*\|_{\mathbf{H}_\lambda(x)} \leq \underbrace{\|\nabla f_\lambda(x)\|_{\mathbf{H}_\lambda^{-1}(x)}}_{\nu_\lambda(x)} \leq \bar{\phi}(\mathfrak{t}(x - x_\lambda^*)) \|x - x_\lambda^*\|_{\mathbf{H}_\lambda(x)}, \quad (20)$$

$$\psi(-\mathfrak{t}(x - x_\lambda^*)) \|x - x_\lambda^*\|_{\mathbf{H}_\lambda(x_\lambda^*)}^2 \leq f_\lambda(x) - f_\lambda(x_\lambda^*) \leq \psi(\mathfrak{t}(x - x_\lambda^*)) \|x - x_\lambda^*\|_{\mathbf{H}_\lambda(x_\lambda^*)}^2. \quad (21)$$

Moreover, the following localization lemma holds.

Lemma 5 (localization). *Let $\lambda > 0$ be fixed. If $\frac{\nu_\lambda(x)}{r_\lambda(x)} < 1$, then*

$$\mathfrak{t}(x - x_\lambda^*) \leq -\log \left(1 - \frac{\nu_\lambda(x)}{r_\lambda(x)} \right). \quad (22)$$

In particular, this shows:

$$\forall c < 1, \forall \lambda > 0, x \in D_\lambda(c) \implies \mathfrak{t}(x - x_\lambda^*) \leq -\log(1 - c).$$

We now state a Lemma which shows that the difference to the optimum in function values is equivalent to the squared newton decrement in a small Dikin ellipsoid. We will use this result in the main paper.

Lemma 6 (Equivalence of norms). *Let $\lambda > 0$ and $x \in D_\lambda(\frac{1}{7})$. Then the following holds:*

$$\frac{1}{4}\nu_\lambda(x)^2 \leq f_\lambda(x) - f_\lambda(x_\lambda^*) \leq \nu_\lambda(x)^2.$$

Proof. Apply Lemma 5 knowing $x \in D_\lambda(\frac{1}{7})$ to get $\mathfrak{t}(x - x_\lambda^*) \leq \log(7/6)$. Then apply Eq. (19) and Eq. (18) to get:

$$\begin{aligned} f_\lambda(x) - f_\lambda(x_\lambda^*) &\leq \psi(\mathfrak{t}(x - x_\lambda^*)) \|x - x_\lambda^*\|_{\mathbf{H}_\lambda(x_\lambda^*)}^2 \\ &\leq e^{\mathfrak{t}(x - x_\lambda^*)} \psi(\mathfrak{t}(x - x_\lambda^*)) \|x - x_\lambda^*\|_{\mathbf{H}_\lambda(x)}^2 \\ &\leq \frac{e^{\mathfrak{t}(x - x_\lambda^*)} \psi(\mathfrak{t}(x - x_\lambda^*))}{\underline{\phi}(\mathfrak{t}(x - x_\lambda^*))^2} \nu_\lambda(x)^2. \end{aligned}$$

Replacing with the bound above, we get

$$\forall \lambda > 0, \forall x \in D_\lambda(\frac{1}{7}), f_\lambda(x) - f_\lambda(x_\lambda^*) \leq \nu_\lambda(x)^2.$$

For the lower bound, proceed in exactly the same way. \square

A.2 Comparison between generalized self-concordant functions

The following result is straightforward.

Lemma 7 (Comparison between generalized self-concordant functions). *Let $\mathcal{G}_1 \subset \mathcal{G}_2 \subset \mathcal{H}$ be two bounded subsets. If (f, \mathcal{G}_1) is generalized self-concordant, then (f, \mathcal{G}_2) is also generalized self-concordant. Moreover,*

$$\forall x \in \mathcal{H}, \forall \lambda > 0, r_{(f, \mathcal{G}_1), \lambda}(x) \geq r_{(f, \mathcal{G}_2), \lambda}(x).$$

In particular, we will often use the following fact. If (f, \mathcal{G}) is generalized self-concordant, and \mathcal{G} is bounded by R , then $(f, B_{\mathcal{H}}(R))$ is also generalized self-concordant. Moreover,

$$r_{(f, B_{\mathcal{H}}(R)), \lambda}(x) = \frac{\sqrt{\lambda + \lambda_{\min}(\mathbf{H}_f(x))}}{R} \geq \frac{\sqrt{\lambda}}{R}.$$

We now state a result which shows that, given a family of generalized self-concordant functions, the expectancy of that family is also generalized self-concordant. This can be seen as a reformulation of Proposition 2 of [23].

Proposition 3 (Expectation). *Let \mathcal{Z} be a polish space equipped with its Borel sigma-algebra, and \mathcal{H} be a Hilbert space. Let $((f_z, \mathcal{G}_z))_{z \in \mathcal{Z}}$ be a family of generalized self-concordant functions such that the mapping $(z, x) \mapsto f_z(x)$ is measurable.*

Assume we are given a random variable Z on \mathcal{Z} , whose support we denote with $\text{supp}(Z)$, such that

- *the random variables $\|f_Z(0)\|, \|\nabla f_Z(0)\|, \text{Tr}(\nabla^2 f_Z(0))$ are bounded;*
- *$\mathcal{G} := \bigcup_{z \in \text{supp}(Z)} \mathcal{G}_z$ is a bounded subset of \mathcal{H} .*

Then the mapping $f : x \in \mathcal{H} \mapsto \mathbb{E}[f_Z(x)]$ is well defined, (f, \mathcal{G}) is generalized self-concordant, and we can differentiate under the expectation.

Corollary 2. Let $n \in \mathbb{N}$ and $(f_i, \mathcal{G}_i)_{1 \leq i \leq n}$ be a family of generalized self-concordant functions. Define

$$f(x) = \frac{1}{n} \sum_{i=1}^n f_i(x), \quad \mathcal{G} = \bigcup_{i=1}^n \mathcal{G}_i.$$

Then (f, \mathcal{G}) is generalized self-concordant.

B Results on approximate Newton methods

In this section, we assume we are given a generalized self-concordant function f in the sense of Appendix A. As f will be fixed throughout this part, we will omit it from the notations. Recall the definitions from Definition 4:

$$\nu_\lambda(x) := \|\nabla f_\lambda(x)\|_{\mathbf{H}_\lambda^{-1}(x)}, \quad \frac{1}{r_\lambda(x)} := \sup_{g \in \mathcal{G}} \|g\|_{\mathbf{H}_\lambda^{-1}(x)}, \quad D_\lambda(c) := \left\{ x : \frac{\nu_\lambda(x)}{r_\lambda(x)} \leq c \right\}.$$

Define the following quantities:

- the true Newton step at point x for the λ -regularized problem:

$$\Delta_\lambda(x) := \mathbf{H}_\lambda^{-1}(x) \nabla f_\lambda(x).$$

- the renormalized Newton decrement $\tilde{\nu}_\lambda(x)$:

$$\tilde{\nu}_\lambda(x) := \frac{\nu_\lambda(x)}{r_\lambda(x)}.$$

Moreover, note that a direct application of Eq. (17) yields the following equation which relates the radii at different points:

$$\forall \lambda > 0, \forall x \in \mathcal{H}, \forall h \in \mathcal{H}, e^{-t(h)} r_\lambda(x) \leq r_\lambda(x+h) \leq e^{t(h)} r_\lambda(x). \quad (23)$$

In this appendix, we develop a complete analysis of so-called approximate Newton methods in the case of generalized self-concordant losses. By "approximate Newton method", we mean that instead of performing the classical update $x_{t+1} = x_t - \Delta_\lambda(x_t)$, we perform an update of the form $x_{t+1} = x_t - \tilde{\Delta}_t$ where $\tilde{\Delta}_t$ is an approximation of the real Newton step. We will characterize this approximation by measuring its distance to the real Newton step using two parameters ρ and ϵ_0 :

$$\|\tilde{\Delta}_t - \Delta_\lambda(x_t)\| \leq \rho \nu_\lambda(x_t) + \epsilon_0.$$

We start by presenting a few technical results in Appendix B.1. We continue by proving that an approximate Newton method has linear convergence guarantees in the right Dikin ellipsoid in Appendix B.2. In Appendix B.3, we adapt these results to a certain way of computing approximate Newton steps, which will be the one we use in the core of the paper. In Appendix B.4, we mention ways to reduce the computational burden of these methods by showing that since all Hessians are equivalent in Dikin ellipsoids, one can actually sketch the Hessian at one given point in that ellipsoid instead of re-sketching it at each Newton step. For the sake of simplicity, this is not mentioned in the core paper, but works very well in practice.

B.1 Main technical results

We start with a technical decomposition of the Newton decrement at point $x - \tilde{\Delta}$ for a given $\tilde{\Delta} \in \mathcal{H}$.

Lemma 8 (Technical decomposition). *Let $\lambda > 0$, $x \in \mathcal{H}$ be fixed. Assume we perform a step of the form $x - \tilde{\Delta}$ for a certain $\tilde{\Delta} \in \mathcal{H}$. Define*

$$\delta := \|\tilde{\Delta} - \Delta_\lambda(x)\|_{\mathbf{H}_\lambda(x)}, \quad \tilde{\delta} := \frac{\delta}{r_\lambda(x)}.$$

The following holds:

$$\tilde{\nu}_\lambda(x - \tilde{\Delta}) \leq e^{\tilde{\nu}_\lambda(x) + \tilde{\delta}} \left[\psi(\tilde{\nu}_\lambda(x) + \tilde{\delta})(\tilde{\nu}_\lambda(x) + \tilde{\delta})^2 + \tilde{\delta} \right]; \quad (24)$$

$$\nu_\lambda(x - \tilde{\Delta}_\lambda(x)) \leq e^{\tilde{\nu}_\lambda(x) + \tilde{\delta}} \left[\psi(\tilde{\nu}_\lambda(x) + \tilde{\delta})(\tilde{\nu}_\lambda(x) + \tilde{\delta})(\nu_\lambda(x) + \delta) + \delta \right]. \quad (25)$$

Proof. Note that by definition, $\nabla f_\lambda(x) = \mathbf{H}_\lambda(x) \Delta_\lambda(x)$. Hence

$$\begin{aligned}
\|\nabla f^\lambda(x - \tilde{\Delta})\|_{\mathbf{H}_\lambda^{-1}(x)} &= \|\nabla f^\lambda(x - \tilde{\Delta}) - \nabla f^\lambda(x) + \mathbf{H}_\lambda(x)\Delta_\lambda(x)\|_{\mathbf{H}_\lambda^{-1}(x)} \\
&\leq \|\nabla f^\lambda(x - \tilde{\Delta}) - \nabla f^\lambda(x) + \mathbf{H}_\lambda(x)\tilde{\Delta}\|_{\mathbf{H}_\lambda^{-1}(x)} \\
&\quad + \|\mathbf{H}_\lambda(x)(\Delta_\lambda(x) - \tilde{\Delta})\|_{\mathbf{H}_\lambda^{-1}(x)} \\
&= \left\| \int_0^1 [\mathbf{H}_\lambda(x - s\tilde{\Delta}) - \mathbf{H}_\lambda(x)]\tilde{\Delta} ds \right\|_{\mathbf{H}_\lambda^{-1}(x)} + \delta \\
&\leq \int_0^1 \|\mathbf{H}_\lambda^{-1/2}(x)\mathbf{H}_\lambda(x - s\tilde{\Delta})\mathbf{H}_\lambda^{-1/2}(x) - \mathbf{I}\| ds \|\tilde{\Delta}\|_{\mathbf{H}_\lambda(x)} + \delta.
\end{aligned}$$

Now using Eq. (17), one has $\|\mathbf{H}_\lambda^{-1/2}(x)\mathbf{H}_\lambda(x - s\tilde{\Delta})\mathbf{H}_\lambda^{-1/2}(x) - \mathbf{I}\| \leq e^{st(\tilde{\Delta})} - 1$, whose integral on s is $\psi(\mathbf{t}(\tilde{\Delta}))\mathbf{t}(\tilde{\Delta})$ where ψ is defined in Definition 4. Moreover, bounding

$$\|\tilde{\Delta}\|_{\mathbf{H}_\lambda(x)} \leq \|\tilde{\Delta} - \Delta_\lambda(x)\|_{\mathbf{H}_\lambda(x)} + \|\Delta_\lambda(x)\|_{\mathbf{H}_\lambda(x)} = \delta + \nu_\lambda(x),$$

it holds

$$\|\nabla f^\lambda(x - \tilde{\Delta})\|_{\mathbf{H}_\lambda^{-1}(x)} \leq \psi(\mathbf{t}(\tilde{\Delta}))\mathbf{t}(\tilde{\Delta}) (\nu_\lambda(x) + \delta) + \delta.$$

1. Now note that using Eq. (17), it holds: $\nu_\lambda(x - \tilde{\Delta}) \leq e^{t(\tilde{\Delta})/2} \|\nabla f^\lambda(x - \tilde{\Delta})\|_{\mathbf{H}_\lambda^{-1}(x)}$ and hence:

$$\nu_\lambda(x - \tilde{\Delta}) \leq e^{t(\tilde{\Delta})/2} \left(\psi(\mathbf{t}(\tilde{\Delta}))\mathbf{t}(\tilde{\Delta}) (\nu_\lambda(x) + \delta) + \delta \right). \quad (26)$$

2. Moreover, using Eq. (23),

$$\tilde{\nu}_\lambda(x - \tilde{\Delta}) \leq e^{t(\tilde{\Delta})} \left(\psi(\mathbf{t}(\tilde{\Delta}))\mathbf{t}(\tilde{\Delta}) (\tilde{\nu}_\lambda(x) + \tilde{\delta}) + \tilde{\delta} \right). \quad (27)$$

Noting that

$$\mathbf{t}(\tilde{\Delta}) \leq \frac{\|\tilde{\Delta}\|_{\mathbf{H}_\lambda(x)}}{r_\lambda(x)} \leq \tilde{\nu}_\lambda(x) + \tilde{\delta},$$

and bounding Eq. (26) simply by taking $e^{t(\tilde{\Delta})/2} \leq e^{t(\tilde{\Delta})}$, we get the two bounds in the lemma. \square

We now place ourselves in the case where we are given an approximation of the Newton step of the following form. Assume λ and x are fixed, and that we approximate $\Delta_\lambda(x)$ with $\tilde{\Delta}$ such that there exists $\rho \geq 0$ and $\epsilon_0 \geq 0$ such that it holds:

$$\|\tilde{\Delta} - \Delta_\lambda(x)\|_{\mathbf{H}_\lambda(x)} \leq \rho\nu_\lambda(x) + \epsilon_0.$$

We define/prove the three different following regimes.

Lemma 9 (3 regimes). *Let $x \in \mathcal{D}_\lambda(\frac{1}{7})$ and $\lambda > 0$ be fixed. Let*

$$0 \leq \rho \leq \frac{1}{7}, \quad \epsilon_0 \geq 0 \text{ s.t. } \tilde{\epsilon}_0 := \frac{\epsilon_0}{r_\lambda(x)} \leq \frac{1}{21}.$$

Let $\tilde{\Delta}$ be an approximation of the Newton steps satisfying $\|\tilde{\Delta} - \Delta_\lambda(x)\|_{\mathbf{H}_\lambda(x)} \leq \rho\nu_\lambda(x) + \epsilon_0$. The three following regimes appear.

- *If $\tilde{\nu}_\lambda(x) \geq \rho$ and $\tilde{\nu}_\lambda(x)^2 \geq \tilde{\epsilon}_0$, then we are in the **quadratic regime**, i.e.*

$$\frac{10\tilde{\nu}_\lambda(x - \tilde{\Delta}_\lambda(x))}{3} \leq \left(\frac{10\tilde{\nu}_\lambda(x)}{3} \right)^2, \quad \nu_\lambda(x - \tilde{\Delta}_\lambda(x)) \leq \frac{10}{3}\tilde{\nu}_\lambda(x)\nu_\lambda(x).$$

- If $\rho \geq \tilde{\nu}_\lambda(x)$ and $\rho \tilde{\nu}_\lambda(x) \geq \tilde{\epsilon}_0$, then we are in the **linear regime**, i.e.

$$\frac{10}{3} \tilde{\nu}_\lambda(x - \tilde{\Delta}_\lambda(x)) \leq \left(\frac{10\rho}{3}\right) \left(\frac{10}{3} \tilde{\nu}_\lambda(x)\right), \quad \nu_\lambda(x - \tilde{\Delta}_\lambda(x)) \leq \frac{10}{3} \tilde{\nu}_\lambda(x) \nu_\lambda(x).$$

- If $\tilde{\epsilon}_0 \geq \tilde{\nu}_\lambda(x)^2$, $\rho \tilde{\nu}_\lambda(x)$, then the **maximal precision** of the approximation is reached, and it holds:

$$\tilde{\nu}_\lambda(x - \tilde{\Delta}_\lambda(x)) \leq 3\tilde{\epsilon}_0 \leq \frac{1}{7}, \quad \nu_\lambda(x - \tilde{\Delta}_\lambda(x)) \leq 3\epsilon_0.$$

Proof. Using the previous lemma,

$$\begin{aligned} \tilde{\nu}_\lambda(x - \tilde{\Delta}_\lambda(x)) &\leq e^{(1+\rho)\tilde{\nu}_\lambda(x) + \tilde{\epsilon}_0} [\psi((1+\rho)\tilde{\nu}_\lambda(x) + \tilde{\epsilon}_0)((1+\rho)\tilde{\nu}_\lambda(x) + \tilde{\epsilon}_0)^2 + \rho\tilde{\nu}_\lambda(x) + \tilde{\epsilon}_0] \\ &\leq \square_1(\tilde{\nu}_\lambda(x), \rho, \tilde{\epsilon}_0) \tilde{\nu}_\lambda(x)^2 + \square_2(\tilde{\nu}_\lambda(x), \rho, \tilde{\epsilon}_0) \rho\tilde{\nu}_\lambda(x) + \square_3(\tilde{\nu}_\lambda(x), \rho, \tilde{\epsilon}_0) \tilde{\epsilon}_0, \end{aligned}$$

and

$$\nu_\lambda(x - \tilde{\Delta}_\lambda(x)) \leq \square_1(\tilde{\nu}_\lambda(x), \rho, \tilde{\epsilon}_0) \tilde{\nu}_\lambda(x) \nu_\lambda(x) + \square_2(\tilde{\nu}_\lambda(x), \rho, \tilde{\epsilon}_0) \rho \nu_\lambda(x) + \square_3(\tilde{\nu}_\lambda(x), \rho, \tilde{\epsilon}_0) \epsilon_0,$$

where the following definitions are used:

$$\begin{aligned} \square_1(\tilde{\nu}, \rho, \tilde{\epsilon}_0) &:= e^{(1+\rho)\tilde{\nu} + \tilde{\epsilon}_0} \psi((1+\rho)\tilde{\nu} + \tilde{\epsilon}_0)(1+\rho)^2, \\ \square_2(\tilde{\nu}, \rho, \tilde{\epsilon}_0) &:= e^{(1+\rho)\tilde{\nu} + \tilde{\epsilon}_0}, \\ \square_3(\tilde{\nu}, \rho, \tilde{\epsilon}_0) &:= e^{(1+\rho)\tilde{\nu} + \tilde{\epsilon}_0} [2\psi((1+\rho)\tilde{\nu} + \tilde{\epsilon}_0)(1+\rho)\tilde{\nu} + 1]. \end{aligned}$$

Now assume $\tilde{\epsilon}_0 \leq \frac{1}{21}$, $\tilde{\nu}_\lambda(x), \rho \leq \frac{1}{7}$. Replacing these values in the functions above bounds \square_1 , \square_2 and \square_3 , and using the case distinction, we get the result. \square

B.2 General analysis of an approximate Newton method

The following proposition describes the behavior of an approximate newton method where ρ and ϵ_0 are fixed a priori.

Proposition 4 (General approximate Newton scheme results). *Let $c \leq \frac{1}{7}$ be fixed and $x_0 \in D_\lambda(c)$ be a given starting point.*

Let $\rho \leq \frac{1}{7}$ and ϵ_0 such that $\epsilon_0 \leq \frac{c}{4} r_\lambda(x_0)$.

Define the following approximate Newton scheme:

$$\forall t \geq 0, \quad x_{t+1} = x_t - \tilde{\Delta}_t, \quad \|\tilde{\Delta}_t - \Delta_\lambda(x_t)\|_{\mathbf{H}_\lambda(x_t)} \leq \rho \nu_\lambda(x_t) + \epsilon_0.$$

The following guarantees hold.

- $\forall t \geq 0, \quad x_t \in D_\lambda(c)$.
- Let $t_c = \left\lceil \log_2 \log_2 \frac{3}{10\rho} \right\rceil + 1$.

$$\forall t \leq t_c, \quad \frac{10\tilde{\nu}_\lambda(x_t)}{3} \leq \max\left(\frac{12\epsilon_0}{r_\lambda(x_0)}, 2^{-2^t}\right),$$

$$\forall t \geq t_c, \quad \frac{10\tilde{\nu}_\lambda(x_t)}{3} \leq \max\left(\frac{12\epsilon_0}{r_\lambda(x_0)}, \left(\frac{10\rho}{3}\right)^{t-t_c+1}\right).$$

- We can bound the relative decrease for both the Newton decrement and the renormalized Newton decrement:

$$\begin{aligned}
\forall t \leq t_c, \quad \nu_\lambda(x_t) &\leq \max\left(3\epsilon_0, \left(\frac{1}{2}\right)^{2^t-1} \nu_\lambda(x_0)\right), \\
\tilde{\nu}_\lambda(x_t) &\leq \max\left(\frac{18\epsilon_0}{5r_\lambda(x_0)}, \left(\frac{1}{2}\right)^{2^t-1} \tilde{\nu}_\lambda(x_0)\right). \\
\forall t \geq t_c, \quad \nu_\lambda(x_t) &\leq \max\left(3\epsilon_0, \left(\frac{10\rho}{3}\right)^{t-t_c+1} \nu_\lambda(x_0)\right), \\
\tilde{\nu}_\lambda(x_t) &\leq \max\left(\frac{18\epsilon_0}{5r_\lambda(x_0)}, \left(\frac{10\rho}{3}\right)^{t-t_c+1} \tilde{\nu}_\lambda(x_0)\right).
\end{aligned}$$

Proof. Start by noting, using Eq. (23),

$$\forall x \in D_\lambda\left(\frac{1}{7}\right), \quad \epsilon \leq \frac{r_\lambda(x)}{21}, \quad \frac{6}{7}r_\lambda(x_0) \leq r_\lambda(x) \leq \frac{7}{6}r_\lambda(x_0). \quad (28)$$

In particular, this holds for any $x \in D_\lambda(c)$, $c \leq \frac{1}{7}$. Thus,

$$\forall c \leq \frac{1}{7}, \quad \forall x_0 \in D_\lambda(c), \quad \frac{\epsilon_0}{r_\lambda(x_0)} \leq \frac{c}{4} \implies \forall x \in D_\lambda(c), \quad \frac{\epsilon_0}{r_\lambda(x)} \leq \frac{c}{3}.$$

1. Proving the first point is simple by induction. Indeed, assume $\tilde{\nu}_\lambda(x_t) \leq c$. We can apply Lemma 9 since the conditions on ϵ and ρ guarantee that the conditions of this lemma are satisfied.

If we are in either the linear or quadratic regime, the fact that $\frac{10\rho}{3}, \frac{10\tilde{\nu}_\lambda(x_t)}{3} \leq \frac{10}{21}$ show that $\tilde{\nu}_\lambda(x_{t+1}) \leq \frac{10}{21}\tilde{\nu}_\lambda(x_t) \leq c$.

If we are in the last case, $\tilde{\nu}_\lambda(x_{t+1}) \leq \frac{3\epsilon_0}{r_\lambda(x_t)} \leq c$.

2. Let us prove the second bullet point by induction. Start by assuming the property holds at t . By the previous point, the hypothesis of Lemma 9 are satisfied at x_t with ρ and ϵ . Assume we are in the limiting case; we easily show that in this case,

$$\frac{10\tilde{\nu}_\lambda(x_{t+1})}{3} \leq \frac{10}{3} \frac{3\epsilon_0}{r_\lambda(x_t)} \leq \frac{35\epsilon_0}{3r_\lambda(x_0)}.$$

Here, the last inequality comes from Eq. (28). If we are not in the limiting case, let us distinguish between the two following cases.

If $t \leq t_c - 1$,

$$\begin{aligned}
\frac{10\tilde{\nu}_\lambda(x_{t+1})}{3} &\leq \frac{10\tilde{\nu}_\lambda(x_t)}{3} \max\left(\frac{10\tilde{\nu}_\lambda(x_t)}{3}, \frac{10\rho}{3}\right) \\
&\leq \max\left(\frac{35\epsilon_0}{3r_\lambda(x_0)}, \frac{10\tilde{\nu}_\lambda(x_t)}{3} \max\left(\left(\frac{1}{2}\right)^{2^t}, \frac{10\rho}{3}\right)\right),
\end{aligned}$$

where the last inequality comes from using the induction hypothesis and the fact that $\frac{10\tilde{\nu}_\lambda(x_t)}{3} \leq 1$.

Using once again the induction hypotheses and the fact that $t \leq \left\lfloor \log_2 \log_2 \frac{3}{10\rho} \right\rfloor$ which implies

$\frac{10\rho}{3} \leq \left(\frac{1}{2}\right)^{2^t}$, we finally get

$$\frac{10\tilde{\nu}_\lambda(x_{t+1})}{3} \leq \max\left(\frac{35\epsilon_0}{3r_\lambda(x_0)}, \left(\frac{1}{2}\right)^{2^{t+1}}\right).$$

The fact that the second property holds for $t = t_c$ is trivial. Now consider the case where $t \geq t_c$. Using the same technique as before but noting that in this case

$$\frac{10\tilde{\nu}_\lambda(x_t)}{3} \leq \max\left(\frac{35\epsilon_0}{3r_\lambda(x_0)}, \left(\frac{10\rho}{3}\right)^{t-t_c+1}\right) \leq \max\left(\frac{35\epsilon_0}{3r_\lambda(x_0)}, \frac{10\rho}{3}\right),$$

We easily use Lemma 9 to reach the desired conclusion.

3. Let $t < t_c$. Then using Lemma 9:

$$\forall s \leq t, \nu_\lambda(x_{s+1}) \leq \max\left(3\epsilon_0, \max\left(\frac{10\rho}{3}, \frac{10\tilde{\nu}_\lambda(x_s)}{3}\right)\nu_\lambda(x_s)\right).$$

Using the fact that for any $s \leq t$, $\frac{10\tilde{\nu}_\lambda(x_s)}{3} \leq \max\left(\frac{35\epsilon_0}{3r_\lambda(x_0)}, \left(\frac{1}{2}\right)^{2^s}\right)$:

$$\forall s \leq t, \nu_\lambda(x_{s+1}) \leq \max\left(3\epsilon_0, \frac{35\epsilon_0}{3} \frac{\nu_\lambda(x_s)}{r_\lambda(x_0)}, \max\left(\frac{10\rho}{3}, \left(\frac{1}{2}\right)^{2^s}\right)\nu_\lambda(x_s)\right).$$

Now using the fact that for any $s \leq t$, $\tilde{\nu}_\lambda(x_s) \leq \frac{1}{7}$, we see that $\frac{\nu_\lambda(x_s)}{r_\lambda(x_0)} \leq \frac{7}{6}\tilde{\nu}_\lambda(x_s) \leq \frac{1}{6}$ and hence $\frac{35\epsilon_0}{3} \frac{\nu_\lambda(x_s)}{r_\lambda(x_0)} \leq 3\epsilon_0$. Moreover, since $s \leq t < t_c$, $\max\left(\frac{10\rho}{3}, \left(\frac{1}{2}\right)^{2^s}\right) = \left(\frac{1}{2}\right)^{2^s}$. Thus:

$$\forall s \leq t, \nu_\lambda(x_{s+1}) \leq \max\left(3\epsilon_0, \left(\frac{1}{2}\right)^{2^s} \nu_\lambda(x_s)\right).$$

Combining these results yields:

$$\nu_\lambda(x_{t+1}) \leq \max\left(3\epsilon_0, \left(\frac{1}{2}\right)^{2^{t+1}-1} \nu_\lambda(x_0)\right).$$

This shows the first equation, that is:

$$\forall t \leq t_c, \nu_\lambda(x_t) \leq \max\left(3\epsilon_0, \left(\frac{1}{2}\right)^{2^t-1} \nu_\lambda(x_0)\right).$$

The case for $t \geq t_c$ is completely analogous. We can also reproduce the same proof to get the same bounds for $\tilde{\nu}$, since the bounds in Lemma 9 are the same for both. □

B.3 Main results in the paper

In the main paper, we mention two types of Newton method. First, we present a result of convergence on the full Newton method:

Lemma 10 (Quadratic convergence of the full Newton method). *Let $c \leq \frac{1}{7}$ and $x_0 \in D_\lambda(c)$. Define*

$$x_{t+1} = x_t - \Delta_\lambda(x_t).$$

Then this scheme converges quadratically, i.e.:

$$\forall t \in \mathbb{N}, \frac{\nu_\lambda(x_t)}{\nu_\lambda(x_0)}, \frac{\tilde{\nu}_\lambda(x_t)}{\tilde{\nu}_\lambda(x_0)} \leq 2^{-(2^t-1)}.$$

Thus :

- $\forall t \in \mathbb{N}, x_t \in D_\lambda(c)$.
- For any $\tilde{c} \leq c$ then $\forall t \geq \lceil \log_2(1 + \log_2 \frac{c}{\tilde{c}}) \rceil$, $x_t \in D_\lambda(\tilde{c})$.
- For any $\epsilon > 0$, $\forall t \geq \lceil \log_2\left(1 + \log_2 \frac{\nu_\lambda(x_0)}{\sqrt{\epsilon}}\right) \rceil$, $\nu_\lambda(x_t) \leq \sqrt{\epsilon}$, $f_\lambda(x) - f_\lambda(x_\lambda^*) \leq \epsilon$.

- If we perform the Newton method and return the first x_t such that $\nu_\lambda(x_t) \leq \sqrt{\varepsilon}$, then the number of Newton steps computations is at most $1 + \left\lceil \log_2 \left(1 + \log_2 \frac{\nu_\lambda(x_0)}{\sqrt{\varepsilon}} \right) \right\rceil$.

Proof. A full Newton method is an approximate Newton method where $\rho, \epsilon_0 = 0$. Thus apply Proposition 4; note that in this case $t_c = +\infty$. The last point shows that if $c \leq \frac{1}{7}$, and if we perform the Newton method with a full Newton step, then

$$\forall t \geq 0, \tilde{\nu}_\lambda(x_t) \leq 2^{-(2^t-1)} \nu_\lambda(x_0), \tilde{\nu}_\lambda(x_t) \leq 2^{-(2^t-1)} \nu_\lambda(x_0).$$

This shows the quadratic convergence, and the first two points directly follow. For the third point, the result for $\nu_\lambda(x_t)$ directly follows from the previous equation, and the one on function values is a direct consequence of Lemma 6 and the fact that $x_t \in D_\lambda(1/7)$.

For the last point, note that $\nu_t(x_t) = \nabla f_\lambda(x_t) \cdot \Delta_\lambda(x_t)$ is accessible. Moreover, the bound on t is given in the point before, and since one has to compute $\Delta_\lambda(x_s)$ for $0 \leq s \leq t$, there are at most $t + 1$ computations. \square

In the main paper, we compute approximate Newton steps by considering methods which naturally yield only a relative error ρ and no absolute error ϵ_0 . Indeed, we take the following notation.

Approximate solutions to linear problems. Let \mathbf{A} be a positive definite Hermitian operator on \mathcal{H} , b in \mathcal{H} , and a wanted relative precision ρ .

We say that x is a ρ -relative approximation to the linear problem $\mathbf{A}x = b$ and write $x \in \text{LinApprox}(\mathbf{A}, b, \rho)$ if the following holds:

$$\|\mathbf{A}^{-1}b - x\|_{\mathbf{A}} \leq \rho \|b\|_{\mathbf{A}^{-1}} = \rho \|\mathbf{A}^{-1}b\|_{\mathbf{A}}.$$

Note that if $x \in \text{LinApprox}(\mathbf{A}, b, \rho)$ for $\rho < 1$, then

$$(1 - \rho) \|b\|_{\mathbf{A}^{-1}} \leq x \cdot b \leq (1 + \rho) \|b\|_{\mathbf{A}^{-1}}.$$

The following lemma shows that if, instead of computing the exact Newton step, we compute a relative approximation of the Newton step belonging to $\text{LinApprox}(\mathbf{H}_\lambda(x), \nabla f_\lambda(x), \rho)$ for a given $\rho < 1$, then one has linear convergence. Moreover, we show that we can still perform a method which automatically stops.

Proposition 5 (relative approximate Newton method). *Let $\lambda > 0$, $\rho \leq \frac{1}{7}$, $c \leq \frac{1}{7}$ and a starting point $x_0 \in D_\lambda(c)$. Assume we perform the following Newton scheme:*

$$\forall t \geq 0, x_{t+1} = x_t - \tilde{\Delta}_t, \quad \tilde{\Delta}_t \in \text{LinApprox}(\mathbf{H}_\lambda(x_t), \nabla f_\lambda(x_t), \rho).$$

Then the scheme converges linearly, i.e.

$$\forall t \in \mathbb{N}, \frac{\nu_\lambda(x_t)}{\nu_\lambda(x_0)}, \frac{\tilde{\nu}_\lambda(x_t)}{\tilde{\nu}_\lambda(x_0)} \leq 2^{-t}.$$

Thus,

- $\forall t \in \mathbb{N}, x_t \in D_\lambda(c)$.
- For any $\tilde{c} \leq c$ then $\forall t \geq \lceil \log_2 \frac{c}{\tilde{c}} \rceil$, $x_t \in D_\lambda(\tilde{c})$.
- For any $\varepsilon > 0$, $\forall t \geq \left\lceil \log_2 \frac{\nu_\lambda(x_0)}{\sqrt{\varepsilon}} \right\rceil$, $\nu_\lambda(x_t) \leq \sqrt{\varepsilon}$, $f_\lambda(x) - f_\lambda(x_\lambda^*) \leq \varepsilon$
- If the method is performed and returns the first x_t such that $x_t \cdot \tilde{\Delta}_t \leq \frac{6}{7}\varepsilon$, then at most $2 + \left\lceil \log_2 \left(\sqrt{\frac{4}{3}} \frac{\nu_\lambda(x_0)}{\sqrt{\varepsilon}} \right) \right\rceil$ approximate Newton steps computations have been performed, and $\nu_\lambda(x_t) \leq \sqrt{\varepsilon}$, $f_\lambda(x) - f_\lambda(x_\lambda^*) \leq \varepsilon$.

Proof. Apply Proposition 4 with $\epsilon_0 = 0$ and $\rho = \frac{1}{7}$, since if $\rho \leq \frac{1}{7}$, then a fortiori the approximation satisfies the condition for $\rho = \frac{1}{7}$. The last point clearly states that

$$\forall t \in \mathbb{N}, \frac{\nu_\lambda(x_t)}{\nu_\lambda(x_0)}, \frac{\tilde{\nu}_\lambda(x_t)}{\tilde{\nu}_\lambda(x_0)} \leq \left(\frac{10}{21}\right)^t \leq 2^{-t}.$$

From this, using Lemma 6 for the third point, the first three points are easily proven.

For the last point, note that since $\tilde{\Delta}_t \in \text{LinApprox}(\mathbf{H}_\lambda(x_t), \nabla f_\lambda(x_t), \rho)$, the following holds: $\nabla f_\lambda(x_t) \cdot \tilde{\Delta}_t = \nu_\lambda(x_t)^2 + \nabla f_\lambda(x_t) \cdot (\tilde{\Delta}_t - \mathbf{H}_\lambda^{-1}(x_t) \nabla f_\lambda(x_t))$. Now bound

$$|\nabla f_\lambda(x_t) \cdot (\tilde{\Delta}_t - \mathbf{H}_\lambda^{-1}(x_t) \nabla f_\lambda(x_t))| \leq \nu_\lambda(x_t) \|\tilde{\Delta}_t - \mathbf{H}_\lambda^{-1}(x_t) \nabla f_\lambda(x_t)\|_{\mathbf{H}_\lambda(x_t)} \leq \rho \nu_\lambda(x_t)^2.$$

Thus:

$$(1 - \rho) \nu_\lambda(x_t)^2 \leq \nabla f_\lambda(x_t) \cdot \tilde{\Delta}_t \leq (1 + \rho) \nu_\lambda(x_t)^2.$$

Since $\rho = \frac{1}{7}$, we see that if $\nabla f_\lambda(x_t) \cdot \tilde{\Delta}_t \leq \frac{6}{7}\epsilon$, then $\nu_\lambda(x_t)^2 \leq \epsilon$. Moreover, since we stop at the first t where $\nabla f_\lambda(x_t) \cdot \tilde{\Delta}_t \leq \frac{6}{7}\epsilon$, then if t denotes the time at which we stop,

$$\frac{6}{7}\epsilon < \nabla f_\lambda(x_{t-1}) \cdot \tilde{\Delta}_{t-1} \leq \frac{8}{7} \nu_\lambda(x_{t-1})^2 \implies \nu_\lambda(x_{t-1})^2 \geq \frac{3}{4}\epsilon.$$

Since $\nu_\lambda(x_{t-1})^2 \leq 2^{-2(t-1)} \nu_\lambda(x_0)^2$, this implies in turn that $t - 1 \leq \log_2 \left(\sqrt{\frac{4}{3}} \frac{\nu_\lambda(x_0)}{\sqrt{\epsilon}} \right)$. Thus, necessarily, $t \leq 1 + \left\lceil \log_2 \left(\sqrt{\frac{4}{3}} \frac{\nu_\lambda(x_0)}{\sqrt{\epsilon}} \right) \right\rceil$, and since we compute approximate Newton steps for $s = 0, \dots, t$, we finally have that the number of approximate Newton steps is bounded by

$$2 + \left\lceil \log_2 \left(\sqrt{\frac{4}{3}} \frac{\nu_\lambda(x_0)}{\sqrt{\epsilon}} \right) \right\rceil.$$

□

Last but not least, we summarize all these theorem in the following simple result.

Lemma 11. *Let $\lambda > 0, c \leq 1/7$, let f_λ be generalized self-concordant and $x \in D_\lambda(c)$. It holds: $\frac{1}{4}\nu_\lambda(x)^2 \leq f_\lambda(x) - f_\lambda(x_\lambda^*) \leq \nu_\lambda(x)^2$. Moreover, the full Newton method starting from x_0 has quadratic convergence, i.e. if x_t is obtained via $t \in \mathbb{N}$ steps of the Newton method Eq. (2), then $\nu_\lambda(x_t) \leq 2^{-(2^t-1)}\nu_\lambda(x_0)$. Finally, the approximate Newton method starting from x_0 has linear convergence, i.e. if x_t is obtained via $t \in \mathbb{N}$ steps of Eq. (3), with $\tilde{\Delta}_t \in \text{LinApprox}(\mathbf{H}_\lambda(x_t), \nabla f_\lambda(x_t), \rho)$ and $\rho \leq 1/7$, then $\nu_\lambda(x_t) \leq 2^{-t}\nu_\lambda(x_0)$.*

Proof. The three points are obtained in the following lemmas, assuming $x \in D_\lambda(1/7)$.

- For $\frac{1}{4}\nu_\lambda(x)^2 \leq f_\lambda(x) - f_\lambda(x_\lambda^*) \leq \nu_\lambda(x)^2$, see Lemma 6 in Appendix A.1.
- The convergence rate of the full Newton method starting in $D_\lambda(1/7)$ is obtained in Lemma 10.
- The convergence rate of the approximate Newton method starting in $D_\lambda(1/7)$ is obtained in Proposition 5.

□

B.4 Sketching the Hessian only once in each Dikin ellipsoid

In this section, we provide a lemma which shows in essence that if we are in a small Dikin ellipsoid, then we can keep the Hessian of the starting point and compute approximations of $\mathbf{H}_\lambda^{-1}(x_0) \nabla f_\lambda(x_t)$; they will be good approximations to $\mathbf{H}_\lambda^{-1}(x_t) \nabla f_\lambda(x_t)$ as well.

Lemma 12. Let $c < 1$ and $x_0 \in D_\lambda(c)$ be fixed.

Let $\tilde{\mathbf{H}}$ be an approximation of the Hessian at x_0 , approximation which we quantify with

$$t := \|\mathbf{H}_\lambda^{-1/2}(x_0) (\mathbf{H}_\lambda(x_0) - \tilde{\mathbf{H}}) \mathbf{H}_\lambda^{-1/2}(x_0)\|.$$

Assume

$$1 + t < 2(1 - c)^2.$$

Let $b \in \mathcal{H}$. If $\tilde{\Delta} \in \text{LinApprox}(\tilde{\mathbf{H}}_\lambda, b, \tilde{\rho})$, then

$$\forall x \in D_\lambda(c), \tilde{\Delta} \in \text{LinApprox}(\mathbf{H}_\lambda(x), b, \rho), \rho = \frac{(\tilde{\rho} - 1)(1 - c)^2 + (1 + t)}{2(1 - c)^2 - (1 + t)}.$$

In particular, if $c \leq \frac{1}{30}$, $x_0 \in D_\lambda(c)$,

$$\forall x \in D_\lambda(c), \forall b \in \mathcal{H}, \tilde{\Delta} \in \text{LinApprox}(\mathbf{H}_\lambda(x_0), b, \frac{1}{20}) \implies \tilde{\Delta} \in \text{LinApprox}(\mathbf{H}_\lambda(x), b, \frac{1}{7}).$$

Proof. First, start with a general theoretical result.

1. Let \mathbf{A} and \mathbf{B} be two positive semi-definite hermitian operators. Let $\lambda > 0$, $b \in \mathcal{H}$ and $\tilde{\Delta} \in \text{LinApprox}(\mathbf{B}_\lambda, b, \tilde{\rho})$. Decompose

$$\begin{aligned} \|\mathbf{A}_\lambda^{-1}b - \tilde{\Delta}\|_{\mathbf{A}_\lambda} &\leq \|\mathbf{A}_\lambda^{-1}b - \mathbf{B}_\lambda^{-1}b\|_{\mathbf{A}_\lambda} + \|\mathbf{B}_\lambda^{-1}b - \tilde{\Delta}\|_{\mathbf{A}_\lambda} \\ &\leq \|\mathbf{A}_\lambda^{1/2}(\mathbf{A}_\lambda^{-1} - \mathbf{B}_\lambda^{-1})\mathbf{A}_\lambda^{1/2}\| \|b\|_{\mathbf{A}_\lambda^{-1}} + \|\mathbf{A}_\lambda^{1/2}\mathbf{B}_\lambda^{-1/2}\| \|\mathbf{B}_\lambda^{-1}b - \tilde{\Delta}\|_{\mathbf{B}_\lambda}. \end{aligned}$$

Now using the fact that $\mathbf{A}_\lambda^{-1} - \mathbf{B}_\lambda^{-1} = \mathbf{B}_\lambda^{-1}(\mathbf{B} - \mathbf{A})\mathbf{A}_\lambda^{-1}$,

$$\begin{aligned} \|\mathbf{A}_\lambda^{1/2}(\mathbf{A}_\lambda^{-1} - \mathbf{B}_\lambda^{-1})\mathbf{A}_\lambda^{1/2}\| &\leq \|\mathbf{A}_\lambda^{-1/2}(\mathbf{B} - \mathbf{A})\mathbf{A}_\lambda^{-1/2}\| \|\mathbf{A}_\lambda^{1/2}\mathbf{B}_\lambda^{-1}\mathbf{A}_\lambda^{1/2}\| \\ &= \|\mathbf{A}_\lambda^{-1/2}(\mathbf{B} - \mathbf{A})\mathbf{A}_\lambda^{-1/2}\| \|\mathbf{A}_\lambda^{1/2}\mathbf{B}_\lambda^{-1/2}\|^2. \end{aligned}$$

Moreover,

$$\|\mathbf{B}_\lambda^{-1}b - \tilde{\Delta}\|_{\mathbf{B}_\lambda} \leq \tilde{\rho} \|b\|_{\mathbf{B}_\lambda^{-1}} \leq \|\mathbf{A}^{1/2}\mathbf{B}^{-1/2}\| \|b\|_{\mathbf{A}_\lambda^{-1}}.$$

Putting things together, and noting that from Lemma 21, $\|\mathbf{A}^{1/2}\mathbf{B}^{-1/2}\|^2 \leq \frac{1}{1 - \|\mathbf{A}_\lambda^{-1/2}(\mathbf{B} - \mathbf{A})\mathbf{A}_\lambda^{-1/2}\|}$

as soon as $\|\mathbf{A}_\lambda^{-1/2}(\mathbf{B} - \mathbf{A})\mathbf{A}_\lambda^{-1/2}\| < 1$, it holds:

$$\tilde{\Delta} \in \text{LinApprox}(\mathbf{A}_\lambda, b, \rho), \rho = \frac{\tilde{\rho} + \|\mathbf{A}_\lambda^{-1/2}(\mathbf{B} - \mathbf{A})\mathbf{A}_\lambda^{-1/2}\|}{1 - \|\mathbf{A}_\lambda^{-1/2}(\mathbf{B} - \mathbf{A})\mathbf{A}_\lambda^{-1/2}\|}.$$

The aim is now to apply this lemma to $\mathbf{A} = \mathbf{H}(x)$ and $\mathbf{B} = \tilde{\mathbf{H}}$.

2. Let $x, x_0 \in D_\lambda(c)$. Using Lemma 22, we see that

$$1 + \|\mathbf{H}_\lambda^{-1/2}(x)(\tilde{\mathbf{H}} - \mathbf{H}(x))\mathbf{H}_\lambda^{-1/2}(x)\| \leq (1 + t)(1 + \|\mathbf{H}_\lambda^{-1/2}(x)(\mathbf{H}(x_0) - \mathbf{H}(x))\mathbf{H}_\lambda^{-1/2}(x)\|).$$

Using Eq. (17), it holds:

$$(e^{-t(x-x_0)} - 1)\mathbf{I} \preceq \mathbf{H}_\lambda^{-1/2}(x)(\mathbf{H}(x_0) - \mathbf{H}(x))\mathbf{H}_\lambda^{-1/2}(x) \preceq (e^{t(x_0-x)} - 1)\mathbf{I}.$$

Thus,

$$\|\mathbf{H}_\lambda^{-1/2}(x)(\mathbf{H}(x_0) - \mathbf{H}(x))\mathbf{H}_\lambda^{-1/2}(x)\| \leq \max(1 - e^{-t(x-x_0)}, e^{t(x-x_0)} - 1) = e^{t(x-x_0)} - 1.$$

Finally, using the fact that $x_0, x \in D_\lambda(c)$ for $c < 1$ yields $t(x - x_0) \leq 2 \log \frac{1}{1-c}$. Hence

$$1 + \|\mathbf{H}_\lambda^{-1/2}(x)(\mathbf{H}(x_0) - \mathbf{H}(x))\mathbf{H}_\lambda^{-1/2}(x)\| \leq \frac{1}{(1 - c)^2}.$$

Thus,

$$\|\mathbf{H}_\lambda^{-1/2}(x)(\tilde{\mathbf{H}} - \mathbf{H}(x))\mathbf{H}_\lambda^{-1/2}(x)\| \leq \frac{1 + t}{(1 - c)^2} - 1.$$

The result then follows. \square

C Proof of bounds for the globalization scheme

In this section, we prove that the scheme of decreasing μ towards λ converges.

C.1 Main technical lemmas

Lemma 13 (Next μ). *Let $\mu > 0$, $c < 1$.*

$$\nu_\mu(x) \leq \frac{c}{3} \frac{\sqrt{\mu}}{R} \implies \nu_{\tilde{\mu}}(x) \leq c \frac{\sqrt{\tilde{\mu}}}{R}, \quad \tilde{\mu} := q\mu, \quad q \geq \frac{\frac{1}{3} + \frac{R\sqrt{\mu}\|x\|_{\mathbf{H}_\mu^{-1}(x)}}{c}}{1 + \frac{R\sqrt{\mu}\|x\|_{\mathbf{H}_\mu^{-1}(x)}}{c}}.$$

$$x \in \mathbf{D}_\mu\left(\frac{c}{3}\right) \implies x \in \mathbf{D}_{\tilde{\mu}}(c), \quad \tilde{\mu} := q\mu, \quad q \geq \frac{\frac{1}{3} + \frac{\mu\|x\|_{\mathbf{H}_\mu^{-1}(x)}}{c r_\mu(x)}}{1 + \frac{\mu\|x\|_{\mathbf{H}_\mu^{-1}(x)}}{c r_\mu(x)}}.$$

Proof. For any $\tilde{\mu} < \mu$, note that

$$\forall x \in \mathcal{H}, \|\mathbf{H}_{\tilde{\mu}}^{-1/2}(x)\mathbf{H}_\mu^{1/2}(x)\| = \sqrt{\frac{\lambda_{\min}(\mathbf{H}(x)) + \mu}{\lambda_{\min}(\mathbf{H}(x)) + \tilde{\mu}}} \leq \sqrt{\mu/\tilde{\mu}}.$$

This shows that $\|\cdot\|_{\mathbf{H}_{\tilde{\mu}}^{-1}(x)} \leq \sqrt{\frac{\mu}{\tilde{\mu}}}\|\cdot\|_{\mathbf{H}_\mu^{-1}(x)}$, and in particular that $\frac{1}{r_{\tilde{\mu}}(x)} \leq \sqrt{\mu/\tilde{\mu}}\frac{1}{r_\mu(x)}$.

Using this fact, it holds:

$$\begin{aligned} \tilde{\nu}_{\tilde{\mu}}(x) &= \frac{\|\nabla f_{\tilde{\mu}}(x)\|_{\mathbf{H}_{\tilde{\mu}}^{-1}(x)}}{r_{\tilde{\mu}}(x)} \\ &= \frac{\|\nabla f_\mu(x) - (\mu - \tilde{\mu})x\|_{\mathbf{H}_{\tilde{\mu}}^{-1}(x)}}{r_{\tilde{\mu}}(x)} \\ &\leq \frac{\mu}{\tilde{\mu}} \frac{\|\nabla f_\mu(x)\|_{\mathbf{H}_\mu^{-1}(x)}}{r_\mu(x)} + \left(\frac{\mu}{\tilde{\mu}} - 1\right) \frac{\|\mu x\|_{\mathbf{H}_\mu^{-1}(x)}}{r_\mu(x)}. \end{aligned}$$

Hence, if $\tilde{\nu}_\mu(x) \leq \frac{c}{3}$, a condition to obtain $\tilde{\nu}_{\tilde{\mu}}(x) \leq c$ is the following:

$$\frac{\mu}{\tilde{\mu}} \left(\frac{c}{3} + t\right) \leq c + t \Leftrightarrow \tilde{\mu} \geq \mu \frac{c/3 + t}{c + t} \quad t = \frac{\|\mu x\|_{\mathbf{H}_\mu^{-1}(x)}}{r_\mu(x)}.$$

This yields the second point of the lemma. The analysis is completely analogous for the first. \square

Lemma 14 (Useful bounds for q). *Let $\mu > 0$. Then the following hold:*

$$\forall x \in \mathcal{H}, \frac{\mu\|x\|_{\mathbf{H}_\mu^{-1}(x)}}{r_\mu(x)} \leq R\sqrt{\mu}\|x\|_{\mathbf{H}_\mu^{-1}(x)} \leq R\|x\|.$$

Moreover, we can bound all of these quantities using x_μ^* :

- For any $c < 1$, $x \in \mathcal{H}$, if $x \in \mathbf{D}_\mu(c/3)$, then the following holds:

$$\frac{\mu\|x\|_{\mathbf{H}_\mu^{-1}(x)}}{c r_\mu(x)} \leq \frac{1}{3} \left(1 + \frac{1}{1 - c/3}\right) + \frac{1}{1 - c/3} \frac{\|\mu x_\mu^*\|_{\mathbf{H}_\mu^{-1}(x_\mu^*)}}{c r_\mu(x_\mu^*)}.$$

- For any $c < 1$, $x \in \mathcal{H}$, if $\frac{R\sqrt{\mu}\|x\|_{\mathbf{H}_\mu^{-1}(x)}}{c} \leq \frac{c}{3}$, then the following holds:

$$\frac{R\sqrt{\mu}\|x\|_{\mathbf{H}_\mu^{-1}(x)}}{c} \leq \left(1 + \frac{1}{1 - c/3}\right) \frac{1}{3} + \sqrt{\frac{1}{1 - c/3}} \frac{R\sqrt{\mu}\|x_\mu^*\|_{\mathbf{H}_\mu^{-1}(x_\mu^*)}}{c}.$$

Likewise, it can be shown that under the same conditions:

$$\frac{R\|x\|}{c} \leq \frac{R\|x_\mu^*\|}{c} + \frac{1}{3}\bar{\phi}(-\log(1 - c/3)).$$

Proof. The first bound is obvious. Moreover, the fact that $\tilde{\nu}_\mu(x) \leq \frac{c}{3}$ implies that $\mathfrak{t}(x - x_\mu^*) \leq \log \frac{1}{1-c/3}$. Thus, we get the classical bounds on the Hessian using Eq. (14):

$$e^{-\mathfrak{t}(x-x_\mu^*)} \mathbf{H}(x) \preceq \mathbf{H}(x_\mu^*) \preceq e^{\mathfrak{t}(x-x_\mu^*)} \mathbf{H}(x).$$

1. Bound on $\mu\|x\|_{\mathbf{H}_\mu^{-1}(x)}$. Using Eqs. (17) and (18),

$$\begin{aligned} \mu\|x\|_{\mathbf{H}_\mu^{-1}(x)} &= \|\nabla f_\mu(x) - \nabla f(x) + \nabla f(x_\mu^*) - \nabla f(x_\mu^*)\|_{\mathbf{H}_\mu^{-1}(x)} \\ &\leq \nu_\mu(x) + \int_0^1 \|\mathbf{H}_\mu(x)^{-1/2} \mathbf{H}(x_t)(x - x_\mu^*)\| dt + \|\nabla f(x_\mu^*)\|_{\mathbf{H}_\mu(x)}, \quad x_t = tx + (1-t)x_\mu^*. \end{aligned}$$

Now bound $\|\mathbf{H}_\mu(x)^{-1/2} \mathbf{H}(x_t)(x - x_\mu^*)\| \leq \|\mathbf{H}_\mu(x)^{-1/2} \mathbf{H}_\mu(x_t)^{1/2}\| \|x - x_\mu^*\|_{\mathbf{H}(x_t)}$ and use Eq. (17) and Eq. (14) to get:

$$\|\mathbf{H}_\mu(x)^{-1/2} \mathbf{H}(x_t)(x - x_\mu^*)\| \leq e^{\mathfrak{t}(x-x_\mu^*)} \|x - x_\mu^*\|_{\mathbf{H}(x)}.$$

Integrating this yields:

$$\int_0^1 \|\mathbf{H}_\mu(x)^{-1/2} \mathbf{H}(x_t)(x - x_\mu^*)\| dt \leq \bar{\phi}(\mathfrak{t}(x - x_\mu^*)) \|x - x_\mu^*\|_{\mathbf{H}(x)} \leq e^{\mathfrak{t}(x-x_\mu^*)} \nu_\mu(x).$$

Where the last inequality is obtained using the bounds between gradient and hessian distance Eq. (18). Finally, using the bound on $\mathfrak{t}(x - x_\mu^*)$,

$$\mu\|x\|_{\mathbf{H}_\mu^{-1}(x)} \leq \left(1 + \frac{1}{1-c/3}\right) \nu_\mu(x) + \sqrt{\frac{1}{1-c/3}} \|\nabla f(x_\mu^*)\|_{\mathbf{H}_\mu^{-1}(x_\mu^*)}.$$

2. Bound on $R\|x\|$. Start by decomposing

$$R\|x\| \leq R\|x_\mu^*\| + R\|x - x_\mu^*\|.$$

Now bound

$$R\|x - x_\mu^*\| \leq \frac{R}{\sqrt{\mu}} \|x - x_\mu^*\|_{\mathbf{H}_\mu(x)}.$$

Using Eq. (17), $\|x - x_\mu^*\|_{\mathbf{H}_\mu(x)} \leq \bar{\phi}(-\log(1-c/3))\nu_\mu(x)$. Hence:

$$R\|x\| \leq R\|x_\mu^*\| + \bar{\phi}(-\log(1-c/3)) \frac{R\nu_\mu(x)}{\sqrt{\mu}}.$$

3. Now assume $x \in D_\mu(c/3)$. Using the bound on $\mu\|x\|_{\mathbf{H}_\mu^{-1}(x)}$, and noting that

$$\frac{1}{r_\mu(x)} \leq e^{\mathfrak{t}(x-x_\mu^*)/2} \frac{1}{r_\mu(x_\mu^*)},$$

it holds:

$$\frac{\mu\|x\|_{\mathbf{H}_\mu^{-1}(x)}}{c r_\mu(x)} \leq \frac{1}{3} \left(1 + \frac{1}{1-c/3}\right) + \frac{1}{1-c/3} \frac{\|\mu x_\mu^*\|_{\mathbf{H}_\mu^{-1}(x_\mu^*)}}{c r_\mu(x_\mu^*)}.$$

4. Now assume $\frac{R\nu_\mu(x)}{\sqrt{\mu}} \leq \frac{c}{3}$. We know that in particular, $x \in D_\mu(c/3)$ and hence:

$$\begin{aligned} R\sqrt{\mu}\|x\|_{\mathbf{H}_\mu^{-1}(x)} &\leq \left(1 + \frac{1}{1-c/3}\right) \frac{R\nu_\mu(x)}{\sqrt{\mu}} + \sqrt{\frac{1}{1-c/3}} \frac{R\mu\|x_\mu^*\|_{\mathbf{H}_\mu^{-1}(x_\mu^*)}}{\sqrt{\mu}} \\ &\leq \left(1 + \frac{1}{1-c/3}\right) \frac{c}{3} + \sqrt{\frac{1}{1-c/3}} R\sqrt{\mu}\|x_\mu^*\|_{\mathbf{H}_\mu^{-1}(x_\mu^*)}. \end{aligned}$$

Hence

$$\frac{R\sqrt{\mu}\|x\|_{\mathbf{H}_\mu^{-1}(x)}}{c} \leq \left(1 + \frac{1}{1-c/3}\right) \frac{1}{3} + \sqrt{\frac{1}{1-c/3}} \frac{R\sqrt{\mu}\|x_\mu^*\|_{\mathbf{H}_\mu^{-1}(x_\mu^*)}}{c}.$$

Likewise:

$$\frac{R\|x\|}{c} \leq \frac{R\|x_\mu^*\|}{c} + \frac{1}{3}\bar{\phi}(-\log(1-c/3)).$$

□

We can get the following simpler bounds.

Corollary 3 (Application to $c = \frac{1}{7}$). *Applying Lemma 14 to $c = \frac{1}{7}$, we get the following bounds. Let $\mu > 0$.*

- For any $x \in \mathcal{H}$, if $x \in \mathcal{D}_\mu(c/3)$, then the following holds:

$$\frac{7\mu\|x\|_{\mathbf{H}_\mu^{-1}(x)}}{r_\mu(x)} \leq 1 + \frac{8\|\mu x_\mu^*\|_{\mathbf{H}_\mu^{-1}(x_\mu^*)}}{r_\mu(x_\mu^*)}.$$

- For any $c < 1$, $x \in \mathcal{H}$, if $\frac{R\nu_\mu(x)}{\sqrt{\mu}} \leq \frac{c}{3}$, then the following hold:

$$\begin{aligned} 7R\sqrt{\mu}\|x\|_{\mathbf{H}_\mu^{-1}(x)} &\leq 1 + 8R\sqrt{\mu}\|x_\mu^*\|_{\mathbf{H}_\mu^{-1}(x_\mu^*)}. \\ 7R\|x\| &\leq 7R\|x_\mu^*\| + 1. \end{aligned}$$

C.2 Proof of main theorems

In this section, we bound the number of iterations of our scheme in different cases.

Recall the proposed globalization scheme in the paper, where $\text{ANM}_\rho(f, x, t)$ is a method performing t successive ρ -relative approximate Newton steps of f starting at x .

Proposed Globalization Scheme

Phase I: Getting in the Dikin ellipsoid of f_λ

Start with $x_0 \in \mathcal{H}$, $\mu_0 > 0$, $t, T \in \mathbb{N}$ and $(q_k)_{k \in \mathbb{N}} \in (0, 1]$.

For $k \in \mathbb{N}$

$$x_{k+1} \leftarrow \text{ANM}_\rho(f_{\mu_k}, x_k, t)$$

$$\mu_{k+1} \leftarrow q_{k+1}\mu_k$$

Stop when $\mu_{k+1} < \lambda$ and set $x_{last} \leftarrow x_k$. $K \leftarrow k$

Phase II: reach a certain precision starting from inside the Dikin ellipsoid

Return $\hat{x} \leftarrow \text{ANM}_\rho(f_\lambda, x_{last}, T)$

Throughout this section, we will denote with K the value of k when the scheme stops, i.e. the first value of k such that $\mu_{k+1} < \lambda$.

Adaptive methods We start by presenting an adaptive way to select μ_{k+1} from μ_k , with theoretical guarantees. The main result is the following.

Proposition 6 (Adaptive, simple version). *Assume that we perform phase I starting at x_0 such that*

$$\frac{R\nu_{\mu_0}(x_0)}{\sqrt{\mu_0}} \leq \frac{1}{7}.$$

Assume that at each step k , we compute x_{k+1} using $t = 2$ iterations of the ρ -relative approximate Newton method. Then if at each iteration, we set:

$$\mu_{k+1} = q_{k+1} \mu_k, \quad q_{k+1} := \frac{\frac{1}{3} + 7R\|x_{k+1}\|}{1 + 7R\|x_{k+1}\|}.$$

Then the following hold:

$$1. \quad \forall k \leq K + 1, \quad \frac{R\nu_{\mu_k}(x_k)}{\sqrt{\mu_k}} \leq \frac{1}{7}.$$

2. The decreasing parameter q_{k+1} is bounded above before reaching K :

$$\forall k \leq K, \quad q_{k+1} \leq \frac{\frac{4}{3} + 7R\|x_{\mu_k}^*\|}{2 + 7R\|x_{\mu_k}^*\|} \leq \frac{\frac{4}{3} + 7R\|x_\lambda^*\|}{2 + 7R\|x_\lambda^*\|}.$$

3. K is finite,

$$K \leq \left\lceil \frac{\log \frac{\mu_0}{\lambda}}{\log \frac{2+7R\|x_\lambda^*\|}{\frac{4}{3}+7R\|x_\lambda^*\|}} \right\rceil \leq \left\lceil (3 + 11R\|x_\lambda^*\|) \log \frac{\mu_0}{\lambda} \right\rceil,$$

$$\text{and } \frac{R\nu_\lambda(x_{K+1})}{\sqrt{\lambda}} \leq \frac{1}{7}.$$

Proof. Let us prove the three points one by one.

1. This is easily proved by induction, the keys to the induction hypothesis being:

- Using the induction hypothesis, $x_k \in \mathcal{D}_{\mu_k}(c)$ and hence, using Proposition 5 shows that after two iterations of the approximate Newton scheme, $\frac{\nu_{\mu_k}(x_{k+1})}{\nu_{\mu_k}(x_k)} \leq \frac{1}{3}$ which implies $\frac{R\nu_{\mu_k}(x_{k+1})}{\sqrt{\mu_k}} \leq \frac{c}{3}$.
- Now using Lemma 13, we see that that since

$$7R\|x_{k+1}\| = \frac{R\|x_{k+1}\|}{c} \geq \frac{R\sqrt{\mu_k}\|x_{k+1}\|_{\mathbf{H}_{\mu_k}^{-1}(x_{k+1})}}{c},$$

the hypotheses to guarantee the bound for q_{k+1} hold, hence

$$\frac{R\nu_{\mu_{k+1}}(x_{k+1})}{\sqrt{\mu_{k+1}}} \leq c.$$

2. Using the second bullet point of Cor. 3, we see that the previous point implies

$$\forall k \leq K, \quad 7R\|x_{k+1}\| \leq 7R\|x_{\mu_k}^*\| + 1 \implies q_{k+1} \leq \frac{4/3 + 7R\|x_{\mu_k}^*\|}{2 + 7R\|x_{\mu_k}^*\|}.$$

Now using the fact that for any $k \leq K$, $\mu_k > \lambda$, we can use the simple fact that $\|x_\lambda^*\| \geq \|x_{\mu_k}^*\|$ to get the desired bound for q_{k+1} .

3. Using the previous point clearly shows the following bound:

$$\forall k \leq K + 1, \quad \mu_k \leq \left(\frac{\frac{4}{3} + 7R\|x_\lambda^*\|}{2 + 7R\|x_\lambda^*\|} \right)^k \mu_0.$$

As this clearly converges to 0 when k goes to infinity, K is necessarily finite. Applying this for $k = K$, we see that:

$$\lambda \leq \mu_K \leq \left(\frac{\frac{4}{3} + 7R\|x_\lambda^*\|}{2 + 7R\|x_\lambda^*\|} \right)^K \mu_0.$$

This shows that $K \leq \frac{\log \frac{\mu_0}{\lambda}}{\log \frac{2+7R\|x_\lambda^*\|}{\frac{4}{3}+7R\|x_\lambda^*\|}}$.

The final bound is obtained noting that

$$\frac{2 + 7R\|x_\lambda^*\|}{\frac{4}{3} + 7R\|x_\lambda^*\|} = 1 + \frac{1}{t}, \quad t = 2 + \frac{21}{2}R\|x_\lambda^*\|,$$

and using the classical bound:

$$\frac{1}{\log(1 + \frac{1}{t})} \leq t + 1.$$

Finally, the fact that $\frac{R\nu_\lambda(x_{K+1})}{\sqrt{\lambda}} \leq c$ is just a consequence of the fact that $\mu_{K+1} \leq \lambda \leq \mu_K$ and thus that $\lambda = q\mu_K$ with $q \geq q_{K+1}$, which is shown to satisfy the condition in Lemma 13. Hence, the lemma holds not only for μ_{K+1} but also for λ . \square

Remark 2 (μ_0). *In the previous proposition, we assume start at x_0, μ_0 such that*

$$\frac{R\nu_{\mu_0}(x_0)}{\sqrt{\mu_0}} \leq \frac{1}{7}.$$

A simple way to have such a pair is simply to select:

$$x_0 = 0, \mu_0 = 7R\|\nabla f(0)\|,$$

$$\text{since } \frac{R\nu_{\mu_0}(x_0)}{\sqrt{\mu_0}} = \frac{R\|\nabla f(0)\|_{\mathbf{H}_{\mu_0}^{-1}(0)}}{\sqrt{\mu_0}} \leq \frac{R\|\nabla f(0)\|}{\mu_0}.$$

Alternatively, if one can approximately compute $\|x\|_{\mathbf{H}_{\mu}^{-1}(x)}$, one can propose the following variant, whose proof is completely analogous.

Proposition 7 (Adaptive, small variant version). *Assume that we perform phase I starting at x_0 such that*

$$\frac{R\nu_{\mu_0}(x)}{\sqrt{\mu_0}} \leq \frac{1}{7}.$$

Then if at each iteration, we set:

$$t_{k+1} = 7\sqrt{\frac{7}{6}}R\sqrt{\mu_k}\sqrt{x_{k+1} \cdot s_{k+1}}, s_{k+1} \in \text{LinApprox}(\mathbf{H}_{\mu_k}(x_{k+1}), x_{k+1}, \frac{1}{7}),$$

and

$$\mu_{k+1} = q_{k+1} \mu_k, \quad q_{k+1} := \frac{\frac{1}{3} + t_{k+1}}{1 + t_{k+1}}.$$

Then the following hold:

1. $\forall k \leq K, \frac{R\nu_{\mu_k}(x_k)}{\sqrt{\mu_k}} \leq \frac{1}{7}$.
2. *The decreasing parameter q_{k+1} is bounded above before reaching K :*

$$\forall k \leq K, q_{k+1} \leq \sup_{\mu_0 \geq \mu \geq \lambda} \frac{\frac{7}{3} + 10R\sqrt{\mu}\|x_\mu^*\|_{\mathbf{H}_\mu^{-1}(x_\mu^*)}}{3 + 10R\sqrt{\mu}\|x_\mu^*\|_{\mathbf{H}_\mu^{-1}(x_\mu^*)}} \leq \frac{\frac{7}{3} + 10R\|x_\lambda^*\|}{3 + 10R\|x_\lambda^*\|}.$$

3. *K is finite,*

$$K \leq \left(\frac{9}{2} + 15 \sup_{\lambda \leq \mu \leq \mu_0} R\sqrt{\mu}\|x_\mu^*\|_{\mathbf{H}_\mu^{-1}(x_\mu^*)} \right) \log \frac{\mu_0}{\lambda},$$

$$\text{and } \frac{R\nu_\lambda(x_{K+1})}{\sqrt{\lambda}} \leq \frac{1}{7}.$$

Proof. The main thing to note is that because of the properties of $\frac{1}{7}$ -approximations, if $s_{k+1} \in \text{LinApprox}(\mathbf{H}_{\mu_k}(x_{k+1}), x_{k+1}, \frac{1}{7})$,

$$(1 - \frac{1}{7})\|x_{k+1}\|_{\mathbf{H}_{\mu_k}^{-1}(x_{k+1})}^2 \leq x_{k+1} \cdot s_{k+1} \leq (1 + \frac{1}{7})\|x_{k+1}\|_{\mathbf{H}_{\mu_k}^{-1}(x_{k+1})}^2.$$

Hence,

$$\|x_{k+1}\|_{\mathbf{H}_{\mu_k}^{-1}(x_{k+1})} \leq \sqrt{\frac{7}{6}}\sqrt{x_{k+1} \cdot s_{k+1}} \leq \sqrt{\frac{4}{3}}\|x_{k+1}\|_{\mathbf{H}_{\mu_k}^{-1}(x_{k+1})}.$$

Hence, $t_{k+1} \geq 7R\sqrt{\mu_k}\|x_{k+1}\|_{\mathbf{H}_{\mu_k}^{-1}(x_{k+1})}$, and we can apply Lemma 13 to get the first point. To get the second point, we bound t_{k+1} above:

$$t_{k+1} \leq 7\sqrt{\frac{4}{3}}R\sqrt{\mu_k}\|x_{k+1}\|_{\mathbf{H}_{\mu_k}^{-1}(x_{k+1})}.$$

Now use Cor. 3 to find:

$$t_{k+1} \leq \sqrt{\frac{4}{3}} \left(1 + 8R\sqrt{\mu_k}\|x_{\mu_k}^*\|_{\mathbf{H}_{\mu_k}^{-1}(x_{\mu_k}^*)} \right) \leq 2 + 10R\sqrt{\mu_k}\|x_{\mu_k}^*\|_{\mathbf{H}_{\mu_k}^{-1}(x_{\mu_k}^*)}.$$

Thus,

$$q_{k+1} \leq \frac{\frac{7}{3} + 10R\sqrt{\mu_k}\|x_{\mu_k}^*\|_{\mathbf{H}_{\mu_k}^{-1}(x_{\mu_k}^*)}}{3 + 10R\sqrt{\mu_k}\|x_{\mu_k}^*\|_{\mathbf{H}_{\mu_k}^{-1}(x_{\mu_k}^*)}}.$$

Note that as long as $k \geq K$,

$$q_{k+1} \leq \sup_{\mu \geq \lambda} \frac{\frac{7}{3} + 10R\sqrt{\mu}\|x_{\mu}^*\|_{\mathbf{H}_{\mu}^{-1}(x_{\mu}^*)}}{3 + 10R\sqrt{\mu}\|x_{\mu}^*\|_{\mathbf{H}_{\mu}^{-1}(x_{\mu}^*)}} \leq \frac{\frac{7}{3} + 10R\|x_{\lambda}^*\|}{3 + 10R\|x_{\lambda}^*\|}.$$

This guarantees convergence.

For the last point, the proof is exactly the same as in the previous proposition. □

General non-adaptive result. As mentioned in the core of the article, in practice, we do not select q_{k+1} at each iteration using a safe adaptative value, but rather decrease $\mu_{k+1} = q\mu_k$ with a constant q , which we see as a parameter to tune. The following result shows that for q large enough, this is justified, and that the lower bound we get for q depends on the radius of the Dikin ellipsoid $r_{\mu}(x)$, instead of $\frac{\sqrt{\mu}}{R}$ in the previous bounds, which is somewhat finer and shows that if the data is structured such that this radius is very big, then q might actually be very small.

Proposition 8 (Fixed q). *Assume that we perform phase I starting at x_0 such that*

$$x_0 \in \mathbf{D}_{\mu_0}\left(\frac{1}{7}\right).$$

Assume we perform the method with a fixed $q_{k+1} = q$, satisfying

$$q \geq \sup_{\lambda \leq \mu \leq \mu_0} \frac{\frac{4}{3} + 8 \frac{\mu\|x_{\mu}^*\|_{\mathbf{H}_{\mu}^{-1}(x_{\mu}^*)}}{r_{\mu}(x_{\mu}^*)}}{2 + 8 \frac{\mu\|x_{\mu}^*\|_{\mathbf{H}_{\mu}^{-1}(x_{\mu}^*)}}{r_{\mu}(x_{\mu}^*)}}.$$

Then the following hold:

1. $\forall k \leq K + 1, x_k \in \mathbf{D}_{\mu_k}\left(\frac{1}{7}\right).$
2. K is finite,

$$K \leq \frac{1}{1-q} \log \frac{\mu_0}{\lambda},$$

and $x_{K+1} \in \mathbf{D}_{\lambda}\left(\frac{1}{7}\right).$

Proof. Let us prove the two points.

1. Let us prove the result by induction. The initialization is trivial. Now assume $x_k \in D_{\mu_k}(\frac{1}{7})$. Performing two iterations of the approximate Newton method guarantees that

$$x_{k+1} \in D_{\mu_k}(\frac{1}{21}),$$

as show in Proposition 5. Now using Lemma 13, we see that $x_{k+1} \in D_{q\mu_k}(\frac{1}{7})$, provided that

$$q \geq \frac{\frac{1}{3} + \frac{7\mu_k \|x_{k+1}\|_{\mathbf{H}_{\mu_k}^{-1}(x_{k+1})}}{r_{\mu_k}(x_{k+1})}}{1 + \frac{7\mu_k \|x_{k+1}\|_{\mathbf{H}_{\mu_k}^{-1}(x_{k+1})}}{r_{\mu_k}(x_{k+1})}}.$$

Now using Cor. 3, we get that

$$\frac{7\mu_k \|x_{k+1}\|_{\mathbf{H}_{\mu_k}^{-1}(x_{k+1})}}{r_{\mu_k}(x_{k+1})} \leq 1 + \frac{8\mu_k \|x_{\mu_k}^*\|_{\mathbf{H}_{\mu_k}^{-1}(x_{\mu_k}^*)}}{r_{\mu_k}(x_{\mu_k}^*)} \leq 1 + 8 \sup_{\lambda \leq \mu \leq \mu_0} \frac{\mu \|x_{\mu}^*\|_{\mathbf{H}_{\mu}^{-1}(x_{\mu}^*)}}{r_{\mu}(x_{\mu}^*)}.$$

Hence the result.

2. This point just follows, using the bound $\frac{1}{\log \frac{1}{q}} \leq \frac{1}{1-q}$. □

C.3 Proof of Thm. 1

Using Remark 2, the fact that $x_0 = 0$ and $\mu_0 = 7R\|\nabla f(0)\|$, as well as the hypotheses of the theorem, we can apply Proposition 6, and show that the number of steps K performed in the first phase is bounded:

$$K \leq \lfloor (3 + 11R\|x_{\lambda}^*\|) \log(7R\|\nabla f(0)\|/\lambda) \rfloor.$$

Moreover, this proposition also shows that $R\nu_{\lambda}(x_{last})/\sqrt{\lambda} \leq \frac{1}{7}$. Hence, we can use Proposition 5: if

$$t \geq T = \left\lceil \log_2 \sqrt{\frac{\lambda \varepsilon^{-1}}{R^2}} \right\rceil \geq \left\lceil \log_2 \frac{\nu_{\lambda}(x_{last})}{\sqrt{\varepsilon}} \right\rceil,$$

then it holds $\nu_{\lambda}(\hat{x}) \leq \sqrt{\varepsilon}$ and $f_{\lambda}(\hat{x}) - f_{\lambda}(x_{\lambda}^*) \leq \varepsilon$. □

D Non-parametric learning with generalized self-concordant functions

In this section, the aim is to provide a fast algorithm in the case of Kernel methods which achieves the optimal statistical guarantees.

D.1 General setting and assumptions, statistical result for regularized ERM.

In this section, we consider the supervised learning problem of learning a predictor $f : \mathcal{X} \rightarrow \mathcal{Y}$ from training samples $(x_i, y_i)_{1 \leq i \leq n}$ which we assume to be realisations from a certain random variable $Z = (X, Y) \in \mathcal{Z} = \mathcal{X} \times \mathcal{Y}$ whose distribution is ρ . In what follows, for simplification purposes, we assume $\mathcal{Y} = \mathbb{R}$; however, this analysis can easily be adapted (although with heavier notations) to the setting where $\mathcal{Y} = \mathbb{R}^p$. Our aim is to compute the predictor of minimal generalization error

$$\inf_{f \in \mathcal{H}} L(f) := \mathbb{E}_{z \sim \rho} [\ell_z(f(x))], \quad (29)$$

where \mathcal{H} is a space of candidate solutions and $\ell_z : \mathbb{R} \rightarrow \mathbb{R}$ is a loss function comparing the prediction $f(x)$ to the objective y .

Kernel methods. Kernel methods consider a space of functions \mathcal{H}_K implicitly constructed from a symmetric positive semi-definite Kernel $K : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ and whose basic functions are the $K_x : t \in \mathcal{X} \mapsto K(x, t)$ and the linear combinations of such functions $f = \sum_{j=1}^m \alpha_j K_{x_j}$.

It is endowed with a scalar product such that: $\forall x_1, x_2 \in \mathcal{X}, K_{x_1} \cdot K_{x_2} = K(x_1, x_2)$, and as a consequence, \mathcal{H}_K satisfies the self-reproducing property:

$$\forall x \in \mathcal{X}, \forall f \in \mathcal{H}, f(x) = \langle f, K_x \rangle_{\mathcal{H}}.$$

In order to find a good predictor for Eq. (29), the following estimator, called the regularized ERM estimator, is often computed:

$$\hat{f}_\lambda := \arg \min_{f \in \mathcal{H}} \hat{L}_\lambda(f) := \frac{1}{n} \sum_{i=1}^n \ell_{z_i}(f(x_i)) + \frac{\lambda}{2} \|f\|_{\mathcal{H}}^2.$$

The properties of this estimator have been studied in [13] for the square loss and [23] for generalized self-concordant functions. In Appendix H, we recall the full setting of [23], and extend it to include the statistical properties of the projected problem.

Assumptions In this section, we will make the following assumptions, which are reformulations of the assumptions of [23], which we recall in Appendix H, in order to have the statistical properties of the regularized ERM. First, we assume that the (x_i, y_i) are i.i.d. samples.

Assumption 1 (i.i.d. data). *The samples $(z_i)_{1 \leq i \leq n} = (x_i, y_i)_{1 \leq i \leq n} \in \mathcal{Z}^n$ are independently and identically distributed according to ρ .*

In the case where $\mathcal{Y} = \mathbb{R}$, we make the following assumptions on the loss, which leads to the self-concordance of the mappings $f \mapsto \ell_z(f(x))$ and that of L, \hat{L} ...

Assumption 2 (Technical assumptions). *The mapping $(z, t) \in \mathcal{Z} \times \mathbb{R} \mapsto \ell_z(t)$ is measurable. Moreover,*

- *there exists $R_\ell < \infty$ such that for all $z \in \text{supp}(Z)$,*

$$\forall t \in \mathbb{R}, |\ell_z^{(3)}(t)| \leq R_\ell \ell_z''(t),$$

- *the random variables $|\ell_z(0)|, |\ell_z'(0)|, |\ell_z''(0)|$ are bounded;*
- *The kernel is bounded, i.e. $\forall x \in \text{supp}(X), K(x, x) \leq \kappa^2$ for a certain κ .*

Using these assumptions, we see that the following properties are satisfied. Define $L_z(f) := \ell_z(f(x))$. Then the L_z satisfy the following properties:

- For any $z \in \mathcal{Z}$, $(L_z, \{\mathbf{R}_\ell K_x\})$ is a generalized self-concordant function in the sense of Definition 4.
- The mapping $(z, f) \in \mathcal{Z} \times \mathcal{H} \mapsto L_z(f)$ is measurable;
- the random variables $\|L_Z(0)\|$, $\|\nabla L_Z(0)\|$, $\text{Tr}(\nabla^2 L_Z(0))$ are bounded by $|\ell_Z(0)|$, $\kappa|\ell'_Z(0)|$, $\kappa^2|\ell''_Z(0)|$;
- $\mathcal{G} := \{\mathbf{R}_\ell K_x : z \in \text{supp}(Z)\}$ is a bounded subset of \mathcal{H} , bounded by $R = R_\ell \kappa$.

This shows that Assumption 7 and Assumption 8 are satisfied by the L_z and hence, using Proposition 16 in the next appendix, L is well-defined, generalized self-concordant with \mathcal{G} . Moreover, the empirical loss

$$\widehat{L} = \frac{1}{n} \sum_{i=1}^n L_{z_i},$$

is also generalized self-concordant with $\widehat{\mathcal{G}} := \{\mathbf{R}_\ell K_{x_i} : 1 \leq i \leq n\}$.

Finally, as in Appendix H, we make an assumption on the regularity of the problem; namely, we assume that a solution to the learning problem exists in \mathcal{H} .

Assumption 3 (Existence of a minimizer). *There exists $f^* \in \mathcal{H}$ such that $L(f^*) = \inf_{f \in \mathcal{H}} L(f)$.*

We adopt all the notations from Appendix H, doing the distinction between expected an empirical problems by adding a $\widehat{\cdot}$ over the quantities related to the empirical problem. We continue using the standard notations for L : for any $f \in \mathcal{H}$ and $\lambda > 0$,

$$\begin{aligned} L_\lambda(f) &= L(f) + \frac{\lambda}{2} \|f\|^2, & \widehat{L}_\lambda(f) &= \widehat{L}(f) + \frac{\lambda}{2} \|f\|^2 \\ \mathbf{H}(f) &= \nabla^2 L(f), & \mathbf{H}_\lambda(f) &= \nabla^2 L_\lambda(f) = \mathbf{H}(f) + \lambda \mathbf{I} \\ \widehat{\mathbf{H}}(f) &= \nabla^2 \widehat{L}(f), & \widehat{\mathbf{H}}_\lambda(f) &= \nabla^2 \widehat{L}_\lambda(f) = \widehat{\mathbf{H}}(f) + \lambda \mathbf{I} \end{aligned}$$

Recall that \widehat{f}_λ is defined as the minimizer of \widehat{L}_λ .

Define the following bounds on the second order derivatives:

$$\forall f \in \mathcal{H}, \mathbf{b}_2(f) = \sup_{z \in \text{supp}(Z)} \ell''_z(f(x)).$$

Statistical properties of the estimator The statistical properties of the estimator \widehat{f}_λ have been studied in [23] in the case of generalized self concordance, an are reported in the main lines in Appendix H. The statistical rates of this estimator and the optimal choice of λ is determined by two parameters, defined in Proposition 17 and which we adapt to the Kernel problem here.

- the *bias* $\mathbf{b}_\lambda = \|\mathbf{H}_\lambda(f^*)^{-1/2} \nabla L_\lambda(f^*)\| = \lambda \|f^*\|_{\mathbf{H}_\lambda^{-1}(f^*)}$, which characterizes the regularity of the optimum. The faster \mathbf{b}_λ decreases to zero, the more regular f^* is.
- the *effective dimension*

$$\text{df}_\lambda = \mathbb{E} \left[\|\mathbf{H}_\lambda(f^*)^{-1/2} \nabla L_Z(f^*)\|^2 \right]. \quad (30)$$

This quantity characterizes the size of the space \mathcal{H} with respect to the problem; the slower it explodes as λ goes to zero, the smaller the size of \mathcal{H} .

For more complete explanations on the meaning of these quantities, we refer to [23].

Moreover, as mentioned in Proposition 17, one can define

$$\mathbf{B}_1^* := \sup_{z \in \text{supp}(Z)} \|\nabla L_z(f^*)\|, \mathbf{B}_2^* := \sup_{z \in \text{supp}(Z)} \text{Tr}(\nabla^2 L_z(f^*)), \mathbf{Q}^* = \frac{\mathbf{B}_1^*}{\sqrt{\mathbf{B}_2^*}}, \mathbf{b}_2^* = \mathbf{b}_2(f^*). \quad (31)$$

We assume the following regularity condition on the minimizer f^* , in order to get statistical bounds.

Assumption 4 (Source condition). *There exists $r > 0$ and $g \in \mathcal{H}$ such that $f^* = \mathbf{H}^r(f^*)g$. This implies the following decrease rate of the bias:*

$$b_\lambda \leq L\lambda^{1/2+r}, \quad L = \|g\|_{\mathcal{H}}.$$

This is a stronger assumption than the existence of the minimizer as $r > 0$ is crucial for our analysis.

We also quantify the effective dimension df_λ : (however, since it always holds for $\alpha = 1$, this is not, strictly speaking, an additional assumption).

Assumption 5 (Effective dimension). *The effective dimension decreases as $\text{df}_\lambda \leq Q\lambda^{-1/\alpha}$.*

If these two assumptions hold, define:

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

Under these assumptions, one can obtain the following statistical rates (which can be found in [23] or in Cor. 4).

Proposition 9. *Let $\delta \in (0, 1/2]$. Under Assumptions 1 to 5, when $n \geq N$ and $\lambda = (C_0/n)^\beta$, then with probability at least $1 - 2\delta$,*

$$L(\widehat{f}_\lambda) - L(f^*) \leq C_1 n^{-\gamma} \log \frac{2}{\delta},$$

with $C_0 = 256(Q/L)^2$, $C_1 = 8(256)^\gamma (Q^\gamma L^{1-\gamma})^2$ and N defined in [23], and satisfying $N = O(\text{poly}(B_1^*, B_2^*, L, Q, R, \log(1/\delta)))$.

D.2 Reducing the dimension: projecting on a subspace using Nyström sub-sampling.

Computations Using a representer theorem, one of the key properties of Kernel spaces is that, owing to the reproducing property,

$$\widehat{f}_\lambda \in \mathcal{H}_n := \left\{ \sum_{i=1}^n \alpha_i K_{x_i} : (\alpha_i) \in \mathbb{R}^n \right\}.$$

This means that solving the regularized empirical problem can be turned into a finite dimensional problem in α . Indeed $\widehat{f}_\lambda = \sum_{i=1}^n \alpha_i K_{x_i}$ where $\alpha = (\alpha_i)_{1 \leq i \leq n}$ is the solution to the following problem:

$$\alpha = \arg \min_{\alpha \in \mathbb{R}^n} \frac{1}{n} \sum_{i=1}^n \ell_{z_i}(\alpha^\top \mathbf{K}_{nn} e_i) + \frac{\lambda}{2} \alpha^\top \mathbf{K}_{nn} \alpha, \quad \mathbf{K}_{nn} = (K(x_i, x_j))_{1 \leq i, j \leq n} \in \mathbb{R}^{n \times n}.$$

The previous problem is usually too costly to solve directly for large values of n , both in time and memory, because of the operations involving \mathbf{K}_{nn} . A solution consists in looking for a solution in a smaller dimensional sub-space \mathcal{H}_M constructed from sub-samples of the data $\{\tilde{x}_1, \dots, \tilde{x}_M\} \subset \{x_1, \dots, x_n\}$:

$$\mathcal{H}_M := \left\{ \sum_{j=1}^M \tilde{\alpha}_j K_{\tilde{x}_j} : \tilde{\alpha} \in \mathbb{R}^M \right\}.$$

In this case, the minimizer $\widehat{f}_{M,\lambda} = \arg \min_{f \in \mathcal{H}_M} \widehat{L}_\lambda(f)$ can be written $\widehat{f}_{M,\lambda} = \sum_{j=1}^M \tilde{\alpha}_j K_{\tilde{x}_j}$, where $\tilde{\alpha}$ is the solution to the following problem:

$$\tilde{\alpha} = \arg \min_{\alpha \in \mathbb{R}^M} \frac{1}{n} \sum_{i=1}^n \ell_{z_i}(\alpha^\top \mathbf{K}_{Mn} e_i) + \frac{\lambda}{2} \alpha^\top \mathbf{K}_{MM} \alpha,$$

where

$$\mathbf{K}_{nM} = (K(x_i, \tilde{x}_j))_{\substack{1 \leq i \leq n \\ 1 \leq j \leq M}}, \quad \mathbf{K}_{Mn} = \mathbf{K}_{nM}^\top, \quad \mathbf{K}_{MM} := (K(\tilde{x}_i, \tilde{x}_j))_{1 \leq i, j \leq M}.$$

Let \mathbf{T} be an upper triangular matrix such that $\mathbf{T}^\top \mathbf{T} = \mathbf{K}_{MM}$. One can re-parametrize the previous problem in the following way. For any $\beta \in \mathbb{R}^M$, define $f_\beta = \sum_{j=1}^M [\mathbf{T}^\dagger \beta]_j K_{\tilde{x}_j}$. This implies in particular that $\|f_\beta\|_{\mathcal{H}} = \|\beta\|_{\mathbb{R}^M}$. Then $\widehat{f}_{M,\lambda} = f_{\beta_{M,\lambda}}$, where

$$\beta_{M,\lambda} = \arg \min_{\beta \in \mathbb{R}^M} \widehat{L}_{M,\lambda}(\beta) := \frac{1}{n} \sum_{i=1}^n \ell_{z_i}(e_i^\top \mathbf{K}_{nM} \mathbf{T}^\dagger \beta) + \frac{\lambda}{2} \|\beta\|^2.$$

Using the properties the ℓ_{z_i} , one easily shows that $\beta \mapsto \ell_{z_i}(e_i^\top \mathbf{K}_{nM} \mathbf{T}^\dagger \beta)$ is $\{\mathbf{R}_\ell \mathbf{T}^{-\top} \mathbf{K}_{Mn} e_i\}$ generalized self-concordant, and $\|\mathbf{R}_\ell \mathbf{T}^{-\top} \mathbf{K}_{Mn} e_i\| \leq \mathbf{R}_\ell \sqrt{K(x_i, x_i)}$. Thus, \widehat{L}_M is also generalized self-concordant, and the associated $\widehat{\mathcal{G}}_M$ is bounded by $R = \mathbf{R}_\ell k$. It will therefore be possible to apply the second order scheme presented in this paper to approximately compute $\beta_{M,\lambda}$.

Statistics Let $\widehat{\nu}_{\lambda,M}(\beta)$ denote the Newton decrement of $\widehat{L}_{\lambda,M}$ at point β and \mathbf{P}_M denote the orthogonal projection on \mathcal{H}_M . Then the following statistical result shows that provided β is a good enough approximation of the optimum, and provided \mathcal{H}_M is large enough, then f_β has the same generalization error as the empirical risk minimizer \widehat{f}_λ .

Recall the following result proved in Proposition 19 in Appendix H.3.

Proposition 10 (Behavior of an approximation to the projected problem). *Suppose that Assumptions 1 to 3 are satisfied. Let $n \in \mathbb{N}$, $\delta \in (0, 1/2]$, $0 < \lambda \leq \mathbf{B}_2^*$. Whenever*

$$n \geq \Delta_1 \frac{\mathbf{B}_2^*}{\lambda} \log \frac{8\mathbf{B}_1^2 \mathbf{B}_2^*}{\lambda \delta}, \quad \mathbf{C}_1 \sqrt{\frac{\text{df}_\lambda \vee (\mathbf{Q}^*)^2}{n}} \log \frac{2}{\delta} \leq \frac{\lambda^{1/2}}{R}, \quad \mathbf{C}_1 \mathbf{b}_\lambda \leq \frac{\lambda^{1/2}}{R},$$

if

$$\|\mathbf{H}^{1/2}(f^*)(\mathbf{I} - \mathbf{P}_M)\|^2 \leq \lambda \frac{\sqrt{2}}{480}, \quad 126 \widehat{\nu}_{M,\lambda}(\beta) \leq \frac{\lambda^{1/2}}{R},$$

the following holds, with probability at least $1 - 2\delta$.

$$L(f_\beta) - L(f^*) \leq \mathbf{K}_1 \mathbf{b}_\lambda^2 + \mathbf{K}_2 \frac{\text{df}_\lambda \vee (\mathbf{Q}^*)^2}{n} \log \frac{2}{\delta} + \mathbf{K}_3 \widehat{\nu}_{M,\lambda}^2(\beta), \quad R \|f_\beta - f^*\|_{\mathcal{H}} \leq 10,$$

where $\mathbf{K}_1 \leq 6.0\text{e}4$, $\mathbf{K}_2 \leq 6.0\text{e}6$ and $\mathbf{K}_3 \leq 810$, \mathbf{C}_1 is defined in Lemma 19, and the other constants are defined in Thm. 8.

In particular, if we apply the previous result for a fixed λ , the following theorem holds (for a proof, see Appendix H.4).

Theorem 5 (Quantitative result with source $r > 0$). *Suppose that Assumptions 1 to 5 are satisfied.*

Let $n \geq N$ and $\delta \in (0, \frac{1}{2}]$. If $\lambda = \left(\left(\frac{\mathbf{Q}}{\mathbf{L}}\right)^2 \frac{1}{n}\right)^{\frac{\alpha}{\alpha(1+2r)+1}}$, and if

$$\|\mathbf{H}^{1/2}(f^*)(\mathbf{I} - \mathbf{P}_M)\|^2 \leq \lambda \frac{\sqrt{2}}{480}, \quad \widehat{\nu}_{M,\lambda}(\beta) \leq \mathbf{Q}^\gamma \mathbf{L}^{1-\gamma} n^{-\gamma/2},$$

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

$$L(f_\beta) - L(f^*) \leq \mathbf{K} (\mathbf{Q}^\gamma \mathbf{L}^{1-\gamma})^2 \frac{1}{n^\gamma} \log \frac{2}{\delta}, \quad R \|f_\beta - f^*\| \leq 10,$$

where N is defined in Eq. (42) and $\mathbf{K} \leq 7.0\text{e}6$.

The proof of the previous result is quite technical and can be found in Appendix H, in Thm. 9.

D.3 A note on sub-sampling techniques

Let Z be a random variable on a Polish space \mathcal{Z} and $(v_z)_{z \in \mathcal{Z}}$ be a family of vectors in \mathcal{H} such that $\|v\|_{L^\infty(Z)} := \sup_{z \in \text{supp}(Z)} \|v_z\| < \infty$ is bounded. Assume that z_1, \dots, z_n are i.i.d. samples from Z .

Define the following trace class Hermitian operators:

$$\mathbf{A} = \mathbb{E}[v_Z \otimes v_Z], \quad \widehat{\mathbf{A}} = \frac{1}{n} \sum_{i=1}^n v_{z_i} \otimes v_{z_i}.$$

Define

$$\mathcal{N}^{\mathbf{A}}(\lambda) := \text{Tr}(\mathbf{A}_\lambda^{-1} \mathbf{A}), \quad \mathcal{N}_\infty^{\mathbf{A}}(\lambda) := \sup_{z \in \text{supp}(Z)} \|\mathbf{A}_\lambda^{-1/2} v_z\|^2. \quad (32)$$

We typically have:

$$\mathcal{N}^{\mathbf{A}}(\lambda) \leq \mathcal{N}_\infty^{\mathbf{A}}(\lambda) \leq \frac{\|v\|_{L^\infty(Z)}^2}{\lambda}.$$

We define the leverage scores associated to the points z_i and \mathbf{A} :

$$\forall 1 \leq i \leq n, \forall t > 0, l_i^{\mathbf{A}}(t) = \|\widehat{\mathbf{A}}_t^{-1/2} v_{z_i}\|^2 = n \left((\mathbf{G}_{nn} + t n \mathbf{I})^{-1} \mathbf{G}_{nn} \right)_{ii}, \quad (33)$$

where $\mathbf{G}_{nn} = (v_{z_i} \cdot v_{z_j})_{1 \leq i, j \leq n}$ denotes the Gram matrix associated to the family v_{z_i} .

As in [28], definition 1, we give the following definition for leverage scores.

Definition 5 (q -approximate leverage scores). *given t_0 , a family $(\tilde{l}_i^{\mathbf{A}}(t))_{1 \leq i \leq n}$ is said to be a family of q -approximate leverage scores with respect to \mathbf{A} if*

$$\forall 1 \leq i \leq n, \forall t \geq t_0, \frac{1}{q} l_i^{\mathbf{A}}(t) \leq \tilde{l}_i^{\mathbf{A}}(t) \leq q l_i^{\mathbf{A}}(t).$$

We say that a subset of m points $\{\tilde{z}_1, \dots, \tilde{z}_m\} \subset \{z_i : 1 \leq i \leq n\}$ is:

- **Sampled using q -approximate leverage scores for t** if the $\tilde{z}_j = z_{i_j}$ where the i_j are m i.i.d. samples from $\{1, \dots, n\}$ using the probability vector $p_i = \frac{\tilde{l}_i^{\mathbf{A}}(t)}{\sum_{i=1}^n \tilde{l}_i^{\mathbf{A}}(t)}$. In that case, we define $\widehat{\mathbf{A}}_m := \frac{1}{m} \sum_{j=1}^m \frac{1}{np_{i_j}} v_{\tilde{z}_j} \otimes v_{\tilde{z}_j}$.
- **Sampled uniformly** if the $\{i_j : 1 \leq j \leq m\}$ is a uniformly chosen subset of $\{1, \dots, n\}$ of size m . In this case, we define $\widehat{\mathbf{A}}_m := \frac{1}{m} \sum_{j=1}^m v_{\tilde{z}_j} \otimes v_{\tilde{z}_j}$.

In Appendix I.1, we present technical lemmas which allow us to show that if m is large enough, the following hold:

- $\|\mathbf{A}_\eta(\mathbf{I} - \mathbf{P}_m)\|^2 \leq 3\eta$, where \mathbf{P}_m is the orthogonal projection on the subspace induced by the $v_{\tilde{z}_j}$;
- $\widehat{\mathbf{A}}_{m,\lambda}$ is equivalent to $\widehat{\mathbf{A}}_\lambda$.

Remark 3 (cost of computing q -approximate leverage scores). *In [30], one can show that the complexity of computing q -approximate leverage scores can be achieved in: $c_{\text{samp}} = O(q^2 \mathcal{N}^{\mathbf{A}}(\lambda)^2 \min(n, 1/\lambda))$ time (where a unit of time is a scalar product evaluation) and $O(\mathcal{N}^{\mathbf{A}}(\lambda)^2 + n)$ in memory.*

D.4 Selecting the M Nyström points

In order for Thm. 5 to hold, we must subsample the M points such as to guarantee $\|\mathbf{H}^{1/2}(f^*)(\mathbf{I} - \mathbf{P}_M)\|^2 \leq \frac{\sqrt{2}\lambda}{480}$.

Since we must sub-sample the M points a priori, i.e. before performing the method, it is necessary to have sub-sampling schemes which do not depend heavily on the point. Define the covariance operator:

$$\Sigma = \mathbb{E}[K_X \otimes K_X].$$

Since $\mathbf{H}(f^*) = \mathbb{E}[\ell_Z''(f(X)) K_X \otimes K_X]$, it is easy to see that $\mathbf{H}(f^*) \preceq b_2^* \Sigma$. Note that for Σ , since $\widehat{\Sigma} = \frac{1}{n} \sum_{i=1}^n K_{x_i} \otimes K_{x_i}$, the leverage scores have the following form:

$$\forall 1 \leq i \leq n, l_i^\Sigma(t) = n \left((\mathbf{K}_{nn} + \lambda n \mathbf{I})^{-1} \mathbf{K}_{nn} \right)_{ii}.$$

Proposition 11 (Selecting Nyström points). *Let $\delta > 0$. Let $\eta = \min(\|\Sigma\|, \frac{\lambda\sqrt{2}}{1440(b_2^* \vee 1)})$. Assume the samples $\{\tilde{x}_1, \dots, \tilde{x}_M\}$ are obtained with one of the following.*

1. $n \geq M \geq (10 + 160\mathcal{N}_\infty^\Sigma(\eta)) \log \frac{8\kappa^2}{\eta\delta}$ using uniform sampling;
 2. $M \geq (6 + 486q^2\mathcal{N}^\Sigma(\eta)) \log \frac{8\kappa^2}{\eta\delta}$ using q -approximate leverage scores with respect to Σ for $t = \eta$, $t_0 \vee \frac{19\kappa^2}{n} \log \frac{n}{2\delta} < \eta$, $n \geq 405\kappa^2 \vee 67\kappa^2 \log \frac{12\kappa^2}{\delta}$.
- Then it holds, with probability at least $1 - \delta$:*

$$\|\Sigma_\eta^{1/2}(\mathbf{I} - \mathbf{P}_M)\| \leq 3\eta \implies \|\mathbf{H}^{1/2}(f^*)(\mathbf{I} - \mathbf{P}_M)\|^2 \leq \lambda \frac{\sqrt{2}}{480}.$$

Proof. The proof is a direct application of the lemmas in Appendix I.1. Indeed, note that since $\Sigma = \mathbb{E}[K_X \otimes K_X]$, then the results can be applied with $Z \leftarrow X$ and $v_z \leftarrow K_x$. Indeed, from Assumption 2, it holds:

$$\sup_{x \in \text{supp}(X)} \|K_x\|^2 \leq \kappa^2.$$

□

We can now combine Proposition 11 and Proposition 10 to obtain the following statistical bounds for the optimizer of the projected Nyström problem $\beta_{M,\lambda}$.

Theorem 6. *Suppose that Assumptions 1 to 3 are satisfied. Let $n \in \mathbb{N}$, $\delta \in (0, 1/2]$, $0 < \lambda \leq B_2^* \wedge 720\sqrt{2}(b_2^* \vee 1)\|\Sigma\|$. Assume*

$$n \geq \Delta_1 \frac{B_2^*}{\lambda} \log \frac{8\Box_1^2 B_2^*}{\lambda\delta}, \quad C_1 \sqrt{\frac{\text{df}_\lambda \vee (\mathbf{Q}^*)^2}{n} \log \frac{2}{\delta}} \leq \frac{\lambda^{1/2}}{R}, \quad C_1 b_\lambda \leq \frac{\lambda^{1/2}}{R},$$

Let $\eta = \frac{\lambda\sqrt{2}}{1440(b_2^ \vee 1)}$. Assume the samples $\{\tilde{x}_1, \dots, \tilde{x}_M\}$ are obtained with one of the following.*

1. $n \geq M \geq (10 + 160\mathcal{N}_\infty^\Sigma(\eta)) \log \frac{8\kappa^2}{\eta\delta}$ using uniform sampling;
 2. $M \geq (6 + 486q^2\mathcal{N}^\Sigma(\eta)) \log \frac{8\kappa^2}{\eta\delta}$ using q -approximate leverage scores with respect to Σ for $t = \eta$, $t_0 \vee \frac{19\kappa^2}{n} \log \frac{n}{2\delta} < \eta$, $n \geq 405\kappa^2 \vee 67\kappa^2 \log \frac{12\kappa^2}{\delta}$.
- The following holds, with probability at least $1 - 3\delta$.*

$$L(f_{\beta_{M,\lambda}}) - L(f^*) \leq K_1 b_\lambda^2 + K_2 \frac{\text{df}_\lambda \vee (\mathbf{Q}^*)^2}{n} \log \frac{2}{\delta}, \quad R\|\beta_{M,\lambda}\| \leq R\|f^*\| + 10,$$

where $K_1 \leq 6.0\text{e}4$, $K_2 \leq 6.0\text{e}6$ and $K_3 \leq 810$, C_1 is defined in Lemma 19, and the other constants are defined in Thm. 8.

Proof. This is simply a reformulation of Proposition 10, noting that $\hat{v}_{M,\lambda}(\beta_{M,\lambda}) = 0$ and that Proposition 11 implies the condition on the Hessian at the optimum. □

Provided source condition holds with $r > 0$, the conditions of this theorem are not void.

D.5 Performing the globalization scheme to approximate $\beta_{M,\lambda}$

In order to apply Proposition 10, one needs to control $\hat{v}_{M,\lambda}(\beta)$.

We will apply our general scheme to $\hat{L}_{M,\lambda}$ in order to obtain such a control.

D.5.1 Performing approximate Newton steps

The key element in the globalization scheme is to be able to compute $\frac{1}{7}$ -approximate Newton steps.

Note that at a given point β and for a given $\mu > 0$ the Hessian is of the form:

$$\hat{\mathbf{H}}_{M,\mu}(\beta) = \frac{1}{n} \mathbf{T}^{-\top} \mathbf{K}_{Mn} \mathbf{D}_n(\beta) \mathbf{K}_{nM} \mathbf{T}^{-1} + \mu \mathbf{I}_M,$$

where $\mathbf{D}_n(\beta) = \text{diag}((d_i(\beta))_{1 \leq i \leq n})$ is a diagonal matrix whose elements are given by $d_i(\beta) = \ell''_{z_i}(e_i^\top \mathbf{K}_{nM} \mathbf{T}^{-1} \beta)$.

Note that we can always write

$$\widehat{\mathbf{H}}_{M,\mu}(\beta) = \frac{1}{n} \sum_{i=1}^n u_i(\beta) u_i(\beta)^\top + \mu \mathbf{I}, \quad u_i(\beta) = \sqrt{d_i(\beta)} \mathbf{T}^{-\top} \mathbf{K}_{Mn} e_i$$

The gradient can be put in the following form:

$$\nabla \widehat{L}_{M,\mu}(\beta) = \frac{1}{n} \mathbf{T}^{-\top} \mathbf{K}_{Mn} v + \mu \beta, \quad v = (\ell'_{z_i}(e_i^\top \mathbf{K}_{nM} \mathbf{T}^{-1} \beta))_{1 \leq i \leq n}.$$

Computing the gradient at one point therefore costs $O(nM + M^2)$, this being the cost of computing \mathbf{K}_{nM} times a vector costs $O(nM)$ and computing \mathbf{T}^{-1} times a vector takes $O(M^2)$ since \mathbf{T} is triangular. Moreover, the cost in memory is $O(M^2 + n)$, M^2 being needed for the saving of \mathbf{T} and n for the saving of the gradient; \mathbf{K}_{nM} times a vector can also be done in $O(n)$ memory, provided we compute it by blocks.

On the other hand, computing the full Hessian matrix would cost nM^2 operations, which is untractable. However, computing a Hessian vector product can be done in $O(nM + M^2)$ time, as for the gradient, which suggest using an iterative solver with preconditioning.

Computing $x \in \text{LinApprox}(\mathbf{A}, b, \rho)$ through pre-conditioned conjugate gradient descent. Assume we wish to solve the problem $\mathbf{A}x = b$ where $\mathbf{A} \in \mathbb{R}^{M \times M}$ is a positive definite matrix and b is a vector of \mathbb{R}^M . If one uses the conjugate gradient method starting from zero, then if x_k denotes the k -th iterate of the conjugate gradient algorithm, Theorem 6.6 in [31] shows that

$$x_k \in \text{LinApprox}(\mathbf{A}, b, \rho), \quad \rho = 2 \left(\frac{\sqrt{\text{Cond}(\mathbf{A})} - 1}{\sqrt{\text{Cond}(\mathbf{A})} + 1} \right)^k.$$

where $\text{Cond}(\mathbf{A})$ is the condition number of the matrix \mathbf{A} , namely the ratio $\frac{\lambda_{\max}(\mathbf{A})}{\lambda_{\min}(\mathbf{A})}$. If $\text{Cond}(\mathbf{A})$ is large, this convergence can be very slow. The idea of preconditioning is to compute an approximation matrix $\widetilde{\mathbf{A}}$ such that

$$\frac{1}{2} \widetilde{\mathbf{A}} \preceq \mathbf{A} \preceq \frac{3}{2} \widetilde{\mathbf{A}}. \quad (34)$$

We then compute \mathbf{B} a triangular matrix such that $\mathbf{B}^\top \mathbf{B} = \widetilde{\mathbf{A}}$ using a cholesky decomposition, which can be done in $O(M^3)$, and note that $\mathbf{B}^{-\top} \mathbf{A} \mathbf{B}^{-1}$ is very well conditioned; indeed, its condition number is bounded by 3.

Perform a conjugate gradient method to solve the pre-conditioned problem $\mathbf{B}^{-\top} \mathbf{A} \mathbf{B}^{-1} z = \mathbf{B}^{-\top} b$, and denote with z_τ the τ -th iteration of this method. Then using the bound on the condition number, we find

$$z_\tau \in \text{LinApprox}(\mathbf{B}^{-\top} \mathbf{A} \mathbf{B}^{-1}, \mathbf{B}^{-\top} b, \rho), \quad \rho = 2 \left(\frac{\sqrt{3} - 1}{\sqrt{3} + 1} \right)^\tau,$$

which in turn implies that by setting $x_\tau := \mathbf{B}^{-1} z_\tau$,

$$x_\tau \in \text{LinApprox}(\mathbf{A}, b, \rho), \quad \rho = 2 \left(\frac{\sqrt{3} - 1}{\sqrt{3} + 1} \right)^\tau.$$

This shows that after at most $\tau = 3$ iterations, provided $\widetilde{\mathbf{A}}$ satisfies Eq. (34), $x_\tau \in \text{LinApprox}(\mathbf{A}, b, \frac{1}{7})$. The cost of this method is therefore $O(M^3 + nM)$ in time, and $O(n + M^2)$ due to the computing of the preconditioner and computing matrix vector products by block. This does not include the cost of finding a suitable $\widetilde{\mathbf{A}}$.

Computing a suitable approximation of $\widehat{\mathbf{H}}_{M,\mu}(\beta)$ To compute a good pre-conditioner, we will subsample Q points i_1, \dots, i_Q points from $\{1, \dots, n\}$, and sketch the Hessian using these Q points.

Proposition 12 (Computing approximate newton steps). *Let $\delta > 0$. Let $\beta \in \mathbb{R}^M$ and $\mu \geq \lambda$, and assume $\frac{19b_2(f_\beta)\kappa^2}{n} \log \frac{n}{2\delta} < \lambda$ and $n \geq 405b_2(f_\beta)\kappa^2 \vee 67b_2(f_\beta)\kappa^2 \log \frac{12b_2(f_\beta)\kappa^2}{\delta}$. Let $\tilde{\mu} = \min(\mu, \|\mathbf{H}(f_\beta)\|)$. Assume one of the following properties is satisfied*

1. $Q \geq \left(10 + 160\mathcal{N}_\infty^{\mathbf{H}(f_\beta)}(\tilde{\mu})\right) \log \frac{8b_2(f_\beta)\kappa^2}{\tilde{\mu}\delta}$ with uniform sampling of the $\{i_1, \dots, i_Q\}$. We set $\mathbf{D}_Q = \text{diag}(\ell''_{z_{i_j}}(f_\beta(x_{i_j})))_{1 \leq j \leq Q}$
2. $Q \geq (6 + 486q^2\mathcal{N}^{\mathbf{H}(f_\beta)}(\tilde{\mu})) \log \frac{8b_2(f_\beta)\kappa^2}{\tilde{\mu}\delta}$ using q -approximate leverage scores associated to $\mathbf{H}(f_\beta)$ for $t = \tilde{\mu}$. We set $\mathbf{D}_Q = \text{diag}\left(\frac{\ell''_{z_{i_j}}(f_\beta(x_{i_j}))}{p_{i_j}}\right)$, where the p_{i_j} are the probabilities computed from the leverage scores.

Assume we use a pre-conditioner \mathbf{B} such that

$$\mathbf{B}^\top \mathbf{B} = \frac{1}{Q} \mathbf{T}^{-\top} \mathbf{K}_{MQ} \mathbf{D}_Q \mathbf{K}_{QM} \mathbf{T}^{-1} + \mu \mathbf{I}_M, \quad \mathbf{K}_{QM} = (K(x_{i_j}, \tilde{x}_k))_{\substack{1 \leq j \leq Q \\ 1 \leq k \leq M}}$$

If we perform $\tau = \log(\rho/2)/\log((\sqrt{3}+1)/\sqrt{3}-1)$ iterations of the conjugate gradient descent on the pre-conditioned Newton system using \mathbf{B} as a preconditioner, then with probability at least $1 - \delta$, this procedure returns $\tilde{\Delta} \in \text{LinApprox}(\widehat{\mathbf{H}}_{M,\lambda}(\beta), \nabla \widehat{L}_{M,\lambda}(\beta), \rho)$, and the computational time is of order $O(\tau(Mn + M^2Q + M^3 + c_{\text{samp}}))$, and the memory requirements can be reduced to $O(M^2 + n)$. Here c_{samp} stands for the complexity of computing Nystrom leverage scores, and using Remark 3 or [30], $c_{\text{samp}} = O(1)$ if uniform sampling is used, and $c_{\text{samp}} = O(\mathcal{N}^{\mathbf{H}(f_\beta)}(\tilde{\mu})^2/\lambda)$ if Nystrom sub-sampling is used. Note that for $\tau = 3$, $\rho = \frac{1}{7}$.

Proof. Start by defining the following operators:

- $K_n : f \in \mathcal{H} \rightarrow (f(x_i))_{1 \leq i \leq n} \in \mathbb{R}^n$;
- $K_M : f \in \mathcal{H} \rightarrow (f(\tilde{x}_j))_{1 \leq j \leq M} \in \mathbb{R}^M$;
- $V = K_M^* \mathbf{T}^{-1}$, where \mathbf{T} is an upper triangular matrix such that $\mathbf{T}^\top \mathbf{T} = \mathbf{K}_{MM} = K_M K_M^*$.

Note that $K_n V = \mathbf{K}_{nM} \mathbf{T}^{-1}$.

Now note that

$$\forall f \in \mathcal{H}, \mathbf{H}(f) = \mathbb{E}[v_z \otimes v_z], \quad \widehat{\mathbf{H}}(f) = \frac{1}{n} \sum_{i=1}^n v_{z_i} \otimes v_{z_i}, \quad v_z = \sqrt{\ell''_z(f(x))} K_x.$$

Since for any $f \in \mathcal{H}$, $\widehat{\mathbf{H}}(f) = \frac{1}{n} K_n^* \mathbf{D}_n(f) K_n$, where $\mathbf{D}_n(f) = \text{diag}(\ell''_{z_i}(f(x_i)))$, we see that

$$\widehat{\mathbf{H}}_{M,\mu}(\beta) = V^* \widehat{\mathbf{H}}(f_\beta) V + \mu \mathbf{I}_M.$$

Thus, the last lemma of Appendix I.1 can be applied, using the fact that $\|v_z\|^2 \leq b_2(f)\kappa^2$, to get that in both cases of the proposition, under the corresponding assumptions:

$$\frac{1}{2} \left(\frac{1}{Q} \mathbf{T}^{-\top} \mathbf{K}_{MQ} \mathbf{D}_Q \mathbf{K}_{QM} \mathbf{T}^{-1} + \mu \mathbf{I}_M \right) \preceq \widehat{\mathbf{H}}_{M,\mu}(\beta) \preceq \frac{3}{2} \left(\frac{1}{Q} \mathbf{T}^{-\top} \mathbf{K}_{MQ} \mathbf{D}_Q \mathbf{K}_{QM} \mathbf{T}^{-1} + \mu \mathbf{I}_M \right).$$

The rest of the proposition follows from the previous discussion. \square

D.5.2 Applying the globalization scheme to control $\widehat{v}_{M,\lambda}(\beta)$

In order to apply Proposition 12 to each point β in our method, we need to have a globalized version of the condition of this proposition.

First, we start by localizing the different values of β we will visit throughout the algorithm.

Definition 6 (path of regularized solutions). *Let $\lambda > 0$, $\varepsilon > 0$. Define the path of regularized solutions*

$$\widehat{\Gamma}_\lambda^M := \{\beta_{M,\mu} : \mu \geq \lambda\}. \quad (35)$$

And the ε approximation of this path:

$$\widehat{\Gamma}_{\lambda,\varepsilon}^M := \left\{ \beta \in \mathbb{R}^M : d(\beta, \widehat{\Gamma}_\lambda^M) \leq \varepsilon \right\}. \quad (36)$$

Note that we always have $\widehat{\Gamma}_\lambda^M \subset \mathcal{B}_{\mathbb{R}^M}(\|\beta_{M,\lambda}\|)$. We now state a lemma proving that all the values visited during the algorithm will lie in an approximation of this path.

Lemma 15. *Define Let $\beta \in \mathbb{R}^M$ such that $\widehat{v}_{M,\mu}(\beta) \leq \frac{\mu^{1/2}}{7R}$ for some $\mu \geq \lambda$. Then the following holds:*

$$\beta \in \widehat{\Gamma}_{\lambda, \frac{1}{6R}}^M.$$

Proof. Bound

$$R\|\beta - \beta_{M,\mu}\| \leq \frac{R}{\mu^{1/2}} \|\beta - \beta_{M,\mu}\|_{\widehat{\mathbf{H}}_{M,\mu}(\beta)} \leq \frac{1}{\underline{\phi}(\mathbf{t}_M(\beta - \beta_{M,\mu}))} \frac{R\widehat{v}_{M,\mu}(\beta)}{\mu^{1/2}}.$$

Just apply Eq. (18) to obtain $R\|\beta - \beta_{M,\mu}\| \leq \frac{1}{6}$. \square

We now introduce the following quantities which will allow to control the number of sub-samples throughout the whole algorithm.

Definition 7. *Define*

- $\bar{b}_2 := \sup_{\beta \in \widehat{\Gamma}_{\lambda, 1/6R}^M} b_2(f_\beta)$.
- $\bar{\mathcal{N}}^{\mathbf{H}}(\lambda) = \sup_{\beta \in \widehat{\Gamma}_{\lambda, 1/6R}^M} \mathcal{N}^{\mathbf{H}(f_\beta)}(\lambda)$.
- $\bar{\mathcal{N}}_\infty^{\mathbf{H}}(\lambda) = \sup_{\beta \in \widehat{\Gamma}_{\lambda, 1/6R}^M} \mathcal{N}_\infty^{\mathbf{H}(f_\beta)}(\lambda)$.
- $\|\bar{\mathbf{H}}\| = \min_{\beta \in \widehat{\Gamma}_{\lambda, 1/6R}^M} \|\mathbf{H}(f_\beta)\|$.

Proposition 13 (Performance of the globalization scheme). *Let $\varepsilon > 0$, $\delta > 0$, $\tilde{\lambda} = \min(\lambda, \|\bar{\mathbf{H}}\|)$. Assume $\frac{19\bar{b}_2\kappa^2}{n} \log \frac{n}{2\delta} < \tilde{\lambda}$ and $n \geq 405\bar{b}_2\kappa^2 \vee 67\bar{b}_2\kappa^2 \log \frac{12\bar{b}_2\kappa^2}{\delta}$.*

Assume we perform the globalization scheme with the parameters in Thm. 1, where in order to compute any ρ approximation of a regularized Newton step, we use a conjugate gradient descent on the pre-conditioned system, where the pre-conditioner is computed as in Proposition 12 using

1. $Q \geq \left(10 + 160\bar{\mathcal{N}}_\infty^{\mathbf{H}}(\tilde{\lambda})\right) \log \frac{8\bar{b}_2\kappa^2}{\lambda\delta}$ if using uniform sampling
2. $Q \geq \left(6 + 486q^2\bar{\mathcal{N}}^{\mathbf{H}}(\tilde{\lambda})\right) \log \frac{8\bar{b}_2\kappa^2}{\lambda\delta}$ if using Nyström leverage scores

Recall that t denotes the number of approximate Newton steps performed at for each μ in Phase I and T denotes the number of approximate Newton steps performed in Phase II, and that using Thm. 1, $t = 2$ and $T = \lceil \log_2 \sqrt{1 \vee (\lambda\varepsilon^{-1}/R^2)} \rceil$. Moreover, recall that K denotes the number of steps performed in Phase I. Define

$$N_{ns} := 2 \left[\left(3 + 11R\|\beta_{M,\lambda}\|\right) \log_2(7R\|\nabla \widehat{L}_M(0)\|/\lambda) \right] + \lceil \log_2 \sqrt{1 \vee (\lambda\varepsilon^{-1}/R^2)} \rceil.$$

Then with probability at least $(1 - \delta)^{N_{ns}}$:

- The method presented in Proposition 12 returns a 1/7- approximate Newton step at each time it is called in the algorithm.
- If β denotes the result of the method, $\widehat{v}_{M,\lambda}(\beta) \leq \sqrt{\varepsilon}$.
- The number of approximate Newton steps computed during the algorithm is bounded by $N_{n,s}$; the complexity of the method is therefore of order $O(N_{n,s}(M^2 \max(M, Q) + nM + c_{\text{samp}}(\lambda)))$ in time and $O(MQ + M^2 + n)$ in memory, where $c_{\text{samp}}(\lambda)$ is a bound on the complexity associated to the computing of leverage scores (see [30] for details).

The algorithm is detailed in Appendix E, in algorithm 1. Note however that the notations are those of the main paper, which are slightly different from the ones used here.

Proof. If we take the globalization scheme, using the parameters of Thm. 1. Assume that all previous approximate Newton steps have been computed in a good way. Then the β at which we are belongs to $\widehat{\Gamma}_{\lambda,1/6R}^M$. Thus, the hypotheses of this proposition imply that the hypothesis of Proposition 12 are satisfied; and hence, up to a $(1 - \delta)$ probability factor, we can assume that the next approximate Newton step is performed correctly, continuing the globalization scheme in the right way. Thus, the globalization scheme converges as in Thm. 1. \square

D.6 Statistical properties of the algorithm

The following theorem describes the computational and statistical behavior of our algorithm.

Proposition 14 (Behavior of an approximation to the projected problem). *Suppose that Assumptions 1 to 3 are satisfied.*

Let $n \in \mathbb{N}$, $\varepsilon > 0$, $\delta \in (0, 1/2]$, $0 < \lambda \leq \mathbf{B}_2^*$.

Define $\tilde{\lambda} = \min(\lambda, \|\mathbf{H}\|)$ and assume $\frac{19\bar{\mathbf{b}}_2\kappa^2}{n} \log \frac{n}{2\delta} < \tilde{\lambda}$, $n \geq 405\bar{\mathbf{b}}_2\kappa^2 \vee 67\bar{\mathbf{b}}_2\kappa^2 \log \frac{12\bar{\mathbf{b}}_2\kappa^2}{\delta}$, and $n \geq \Delta_1 \frac{\mathbf{B}_2^*}{\lambda} \log \frac{8\mathbf{Q}_1^2\mathbf{B}_2^*}{\lambda\delta}$. Assume

$$C_1 \sqrt{\frac{\text{df}_\lambda \vee (\mathbf{Q}^*)^2}{n} \log \frac{2}{\delta}} \leq \frac{\lambda^{1/2}}{R}, \quad C_1 \mathbf{b}_\lambda \leq \frac{\lambda^{1/2}}{R}, \quad 126\sqrt{\varepsilon} \leq \frac{\lambda^{1/2}}{R}.$$

Assume that the M points $\tilde{x}_1, \dots, \tilde{x}_M$ are obtained through Nyström sub-sampling using $\eta = \|\Sigma\| \wedge \frac{\lambda\sqrt{2}}{1440(\bar{\mathbf{b}}_2^*\sqrt{1})}$, with either

1. $M \geq (10 + 160\mathcal{N}_\infty^\Sigma(\eta)) \log \frac{8\kappa^2}{\eta\delta}$ if using uniform sampling;
2. $M \geq (6 + 486q^2\mathcal{N}^\Sigma(\eta)) \log \frac{8\kappa^2}{\eta\delta}$ if using q -approximate leverage scores for η , associated to the co-variance operator Σ .

Assume we perform the globalization scheme as in Proposition 13, i.e. with the parameters in Thm. 1, where in order to compute any ρ approximation of a regularized Newton step, we use a conjugate gradient descent on the pre-conditioned system, where the pre-conditioner is computed as in Proposition 12 using

1. $Q \geq \left(10 + 160\bar{\mathcal{N}}_\infty^{\mathbf{H}}(\tilde{\lambda})\right) \log \frac{8\bar{\mathbf{b}}_2\kappa^2}{\lambda\delta}$ if using uniform sampling
2. $Q \geq \left(6 + 486q^2\bar{\mathcal{N}}^{\mathbf{H}}(\tilde{\lambda})\right) \log \frac{8\bar{\mathbf{b}}_2\kappa^2}{\lambda\delta}$ if using Nyström leverage scores

Let $N_{n,s}$ be defined as in Proposition 13. Recall $N_{n,s}$ is an upper bound for the number of approximate Newton steps performed in the algorithm. One can bound

$$N_{n,s} \leq 2 \left[(113 + 11R\|f^*\|) \log_2 \frac{7R\|\nabla\widehat{L}_M(0)\|}{\lambda} \right] + \left\lceil \log_2 \frac{\lambda^{1/2}}{R\varepsilon} \right\rceil.$$

Moreover, with probability at least $1 - (N_{n,s} + 2)\delta$, the following holds:

$$L(f_\beta) - L(f^*) \leq K_1 \mathbf{b}_\lambda^2 + K_2 \frac{\text{df}_\lambda \vee (\mathbf{Q}^*)^2}{n} \log \frac{2}{\delta} + K_3 \varepsilon.$$

where $K_1 \leq 6.0\text{e}4$, $K_2 \leq 6.0\text{e}6$ and $K_3 \leq 810$, C_1 is defined in Lemma 19, and the other constants are defined in Thm. 8.

Proof. This is a simple combination between Propositions 10, 11 and 13. To bound the number of Newton steps N_{ns} , one simply uses the fact that under the conditions of the theorem, $R\|\beta_{M,\lambda}\| \leq 10 + R\|f^*\|$. \square

Remark 4 (Complexity). *Let $L = \bar{b}_2\kappa^2$. The complexity of the previous method using leverage scores computed for Σ for the Nystrom projections and for $\mathbf{H}(f_\beta)$ for choosing the Q points at the different stages is the following. The total complexity in time will be of order:*

$$O\left(N_{ns} \left(n\bar{\mathbf{H}}(\lambda) \log(L\lambda^{-1}\delta^{-1}) + \bar{b}_2^3 \mathcal{N}^\Sigma(\lambda)^3 \log^3(L\lambda^{-1}\delta^{-1}) + L/\lambda \bar{b}_2 \mathcal{N}^\Sigma(\lambda)^2 \right)\right).$$

The memory complexity can be bounded by

$$O(\bar{b}_2^2 \mathcal{N}^\Sigma(\lambda)^2 \log^2(L\lambda^{-1}\delta^{-1}) + n).$$

Here, we use the fact that $\mathbf{H} \leq \bar{b}_2 \Sigma$.

We can now write down the previous proposition by classifying problems using Assumptions 4 and 5 and in order to get optimal rates.

Theorem 7 (Performance of the scheme using pre-conditioning). *Let $\delta > 0$. Assume Assumptions 1 to 5 are satisfied. Let $n \geq \tilde{N}$, where \tilde{N} is characterized in the proof, $\lambda = \left(\frac{Q}{L}\right)^2 \frac{1}{n} \frac{\alpha}{\alpha(1+2r)+1}$. Assume that the M points $\tilde{x}_1, \dots, \tilde{x}_M$ are obtained through Nyström sub-sampling using $\eta = \frac{\lambda\sqrt{2}}{1440(\bar{b}_2^2\sqrt{1})}$, with either*

1. $M \geq (10 + 160\bar{\mathcal{N}}_\infty^\Sigma(\eta)) \log \frac{8\kappa^2}{\eta\delta}$ if using uniform sampling;
2. $M \geq (6 + 486q^2\bar{\mathcal{N}}^\Sigma(\eta)) \log \frac{8\kappa^2}{\eta\delta}$ if using q -approximate leverage scores for η , associated to the co-variance operator Σ .

Assume we perform the globalization scheme as in Proposition 13, i.e. with the parameters in Thm. 1, where in order to compute any ρ approximation of a regularized Newton step, we use a conjugate gradient descent on the pre-conditioned system, where the pre-conditioner is computed as in Proposition 12 using

1. $Q \geq (10 + 160\bar{\mathcal{N}}_\infty^{\mathbf{H}}(\lambda)) \log \frac{8\bar{b}_2\kappa^2}{\lambda\delta}$ if using uniform sampling
2. $Q \geq (6 + 486q^2\bar{\mathcal{N}}^{\mathbf{H}}(\lambda)) \log \frac{8\bar{b}_2\kappa^2}{\lambda\delta}$ if using Nyström leverage scores

Let N_{ns} be defined as in Proposition 13. Recall N_{ns} is an upper bound for the number of approximate Newton steps performed in the algorithm. One can bound

$$N_{ns} \leq (227 + 22R\|f^*\|) \left(\left\lceil \log_2 \left(7R\|\nabla\widehat{L}_M(0)\| \right) \right\rceil + \left\lceil \log_2 \frac{nL^2}{Q^2} \right\rceil + \left\lceil \log_2 \frac{1}{RL} \right\rceil \right).$$

Moreover, with probability at least $1 - (N_{ns} + 2)\delta$, the following holds:

- all of the approximate Newton methods yield $\frac{1}{7}$ -approximate Newton steps
- The scheme finishes, and the number of approximate Newton steps is bounded by N_{ns} . The total complexity of the method is therefore

$$O((nM + M^3 + M^2Q + c_{s\text{amp}})N_{ns}) \text{ in time , } \quad O(n + M^2) \text{ in memory.}$$

- The returned β is statistically optimal:

$$L(f_\beta) - L(f^*) \leq K (Q^\gamma L^{1-\gamma})^2 \frac{1}{n^\gamma} \log \frac{2}{\delta},$$

where K is defined in Thm. 5.

Proof. The proof consists mainly of combining Propositions 11 and 13 and Thm. 5.

Recall that we set $\lambda = \left(\frac{Q^2}{L^2} \frac{1}{n}\right) \frac{\alpha}{\alpha(1+2r)+1}$.

1. Start by defining \tilde{N} such that:

- $\tilde{N} \geq N$ where N is defined in Thm. 5;
- $\forall n \geq \tilde{N}$, $\lambda \leq \overline{\|\mathbf{H}\|}$. This is possible as $\frac{\alpha}{\alpha(1+2r)+1}$ is a strictly positive exponent.
- $\forall n \geq \tilde{N}$, $\frac{19\bar{b}_2 \vee 1}{n} \kappa^2 \log \frac{n}{2\delta} < \lambda$; this is possible as soon as $\frac{\alpha}{\alpha(1+2r)+1} < 1$, i.e. this is satisfied since $r > 0$;
- $\tilde{N} \geq 405\bar{b}_2 \vee 1 \kappa^2 \vee 67\bar{b}_2 \vee 1 \kappa^2 \log \frac{12\bar{b}_2 \vee 1}{\delta} \kappa^2$;
- $\forall n \geq \tilde{N}$, $\frac{\lambda\sqrt{2}}{1440(\bar{b}_2^2 \vee 1)} \leq \|\Sigma\|$.

We see that such a \tilde{N} can be defined explicitly.

2. Combining the assumptions on \tilde{N} with the ones on M , we see that all the assumptions of Proposition 11 are satisfied and thus that with probability at least $1 - \delta$, all the hypotheses for Thm. 5 are satisfied except the bound on $\hat{\nu}_{M,\lambda}(\beta)$.

3. Applying Proposition 13, taking $\sqrt{\varepsilon} = Q^\gamma L^{1-\gamma} n^{-\gamma/2}$ and $\lambda = \left(\frac{Q^2}{L^2} \frac{1}{n}\right)^{\frac{\alpha}{\alpha(1+2r)+1}}$, we see that under these hypotheses,

$$N_{ns} := 2 \left[(3 + 11R\|\beta_{M,\lambda}\|) \log_2 \left(7R\|\nabla\hat{L}_M(0)\| \left(\frac{nL^2}{Q^2}\right)^{\frac{\alpha}{\alpha(1+2r)+1}} \right) \right] \\ + \left[\log_2 \left(\frac{1}{RL} \left(\frac{nL^2}{Q^2}\right)^{\frac{r\alpha}{\alpha(1+2r)+1}} \right) \right].$$

Now we can bound this harshly:

$$N_{ns} \leq (7 + 22R\|\beta_{M,\lambda}\|) \left(\left\lceil \log_2 \left(7R\|\nabla\hat{L}_M(0)\| \right) \right\rceil + \left\lceil \log_2 \frac{nL^2}{Q^2} \right\rceil + \left\lceil \log_2 \frac{1}{RL} \right\rceil \right).$$

Now bounding $R\|\beta_{M,\lambda}\| \leq 10 + R\|f^*\|$, we get

$$N_{ns} \leq (227 + 22R\|f^*\|) \left(\left\lceil \log_2 \left(7R\|\nabla\hat{L}_M(0)\| \right) \right\rceil + \left\lceil \log_2 \frac{nL^2}{Q^2} \right\rceil + \left\lceil \log_2 \frac{1}{RL} \right\rceil \right).$$

4. Finally, we use a union bound to conclude. □

E Algorithm

Algorithm 1 Algorithm efficient non-parametric learning for generalized self-concordant losses with optimal statistical guarantees discussed in Sec. 4 of the main paper.

<p>Input: $(x_i, y_i)_{i=1}^n$, $n \in \mathbb{N}$, ℓ loss function, k kernel function and $\lambda > 0$. Return: estimated function $\hat{g} : \mathcal{X} \rightarrow \mathbb{R}$ Parameters: $Q, M, T \in \mathbb{N}$, $\mu_0 > 0$, $(q_k)_{k \in \mathbb{N}}$. Fixed parameters: $t = 2$ from Thm. 1, $\tau = 3$ from Proposition 12 in Appendix D.5.1. $(\bar{x}_j)_{j=1}^M \leftarrow \text{leverage-scores-sampling}((x_i)_{i=1}^n, M, \lambda, k)$ $\mathbf{K} \leftarrow \text{kernel-matrix}((\bar{x}_j)_{j=1}^M, (\bar{x}_j)_{j=1}^M)$ $\mathbf{T} \leftarrow \text{cholesky-upper-triangular}(\mathbf{K})$ define the function $v(\cdot) = (k(\bar{x}_1, \cdot), \dots, k(\bar{x}_M, \cdot)) \in \mathbb{R}^M$</p>
<p>define compute-preconditioner: Input: $\alpha \in \mathbb{R}^M$, $\lambda > 0$ $c_i \leftarrow \sqrt{\ell^{(2)}(v(x_i)^\top \mathbf{T}^{-1} \alpha, y_i)}$ for all $i = 1, \dots, n$ define the function $k'(\circ, \bullet)$ as $k'(\circ, \bullet) := c_\circ \times c_\bullet \times k(x_\circ, x_\bullet)$ for $\circ, \bullet \in \{1, \dots, n\}$ $(h_s)_{s=1}^Q \leftarrow \text{leverage-scores-sampling}((i)_{i=1}^n, Q, \lambda, k')$ $\mathbf{G} \leftarrow \text{kernel-matrix}((\bar{x}_j)_{j=1}^M, (x_{h_s})_{s=1}^Q, k)$ $\mathbf{H} \leftarrow \mathbf{T}^{-\top} \times \mathbf{G} \times \text{diag}((c_{i_h}^2)_{h=1}^Q) \times \mathbf{G}^\top \times \mathbf{T}^{-1}$ $\mathbf{B} \leftarrow \text{cholesky-upper-triangular}(\frac{1}{Q} \mathbf{H} + \lambda \mathbf{I})$ return \mathbf{B}</p>
<p>define preconditioned-conj-grad: Input: $\alpha \in \mathbb{R}^M$, $\mu > 0$, $r \in \mathbb{R}^M$, $\tau \in \mathbb{N}$, $\mathbf{B} \in \mathbb{R}^{M \times M}$ $p \leftarrow r$, $s_0 \leftarrow \ r\ ^2$, $\beta \leftarrow 0$ For $i = 1, \dots, \tau$ $z \leftarrow \mu \mathbf{B}^{-\top} \mathbf{B}^{-1} p + \frac{1}{n} \sum_{i=1}^n \ell^{(2)}(v(x_i)^\top \mathbf{T}^{-1} \alpha, y_i) (v(x_i)^\top \mathbf{T}^{-1} \mathbf{B}^{-1} p) \mathbf{B}^{-\top} \mathbf{T}^{-\top} v(x_i)$ $a \leftarrow s_0 / (p^\top z)$ $\beta \leftarrow \beta + ap$ $r \leftarrow r - az$, $s_1 \leftarrow \ r\ ^2$ $p \leftarrow r + (s_1 / s_0) p$ $s_0 \leftarrow s_1$ return β</p>
<p>define appr-linear-solver: Input: $\alpha \in \mathbb{R}^M$, $\mu > 0$, $g \in \mathbb{R}^M$ $\mathbf{B} \leftarrow \text{compute-preconditioner}(\alpha, \mu)$ $u \leftarrow \text{preconditioned-conjugate-gradient}(\alpha, \mu, \mathbf{B}^{-\top} g, \tau = 3, \mathbf{B})$ return $\mathbf{B}^{-1} u$</p>
<p>define approximate-Newton: Input: $\alpha_0 \in \mathbb{R}^M$, $\mu > 0$, $t \in \mathbb{N}$ For $j = 1, \dots, t$ $g \leftarrow \mu \alpha_{j-1} + \frac{1}{n} \sum_{i=1}^n \ell^{(1)}(v(x_i)^\top \mathbf{T}^{-1} \alpha_{j-1}, y_i) \mathbf{T}^{-\top} v(x_i)$ $\alpha_j \leftarrow \alpha_{j-1} - \text{appr-linear-solver}(\alpha_{j-1}, \mu, g)$ return α_t</p>
<p>$\alpha_0 \leftarrow 0$ For $k \in \mathbb{N}$ $\alpha_{k+1} \leftarrow \text{approximate-Newton}(\alpha_k, \mu_k, t = 2)$ $\mu_{k+1} \leftarrow q_{k+1} \mu_k$ Stop when $\mu_{k+1} < \lambda$ and set $\alpha_{last} \leftarrow \alpha_k$ $\hat{\alpha} \leftarrow \text{approximate-Newton}(\alpha_{last}, \lambda, T)$ return $\hat{g}(\cdot) := v(\cdot)^\top \mathbf{T}^{-1} \hat{\alpha}$</p>

Let $N, M \in \mathbb{N}$ with $M \leq N$. In Alg. 1, `leverage-scores-sampling` $((z_i)_{i=1}^N, M, k, \lambda)$ returns a subset of $(z_i)_{i=1}^N$ of cardinality M sampled by using (approximate) leverage scores at scale $\lambda > 0$ and computed using the kernel k . An explicit example of an algorithm computing `leverage-scores-sampling` is in [30]. Moreover `kernel-matrix` $((x_i)_{i=1}^N, (x'_i)_{i=1}^M, k)$ computes the kernel matrix $K \in \mathbb{R}^{N \times M}$ where $K_{ij} = k(x_i, x'_j)$, with $N, M \in \mathbb{N}$.

F Experiments

We present our algorithm’s performance for logistic regression on two large scale data sets: Higgs and Susy. We have implemented our method using pytorch, and performed computations on one node of a Tesla P100-PCIE-16GB GPU. Recall that in the case of logistic regression, $\ell_{(x,y)}(t) = \log(1 + e^{-yt})$.

In what follows, denote with n the cardinality of the data set and d the number of features of this data set. The error is measured in terms of classification error for both data sets. In both cases, we pre-process the data by subtracting the mean and dividing by the standard deviation for each feature. The data sets are the following.

Susy ($n = 5 \times 10^6$, $d = 18$, binary classification). We always use a Gaussian Kernel with $\sigma = 5$ for logistic loss (obtained through a grid search; note that in [29], $\sigma = 4$ is used for the square loss), and will always use 10^4 Nystrom points.

Higgs ($n = 1.1 \times 10^7$, $d = 28$, binary classification). We then apply a Gaussian Kernel with $\sigma = 5$, as in [29] (we have also performed a grid search).

For these data sets, we do not have a fixed test set, and thus set apart 20% of the data set at random to be the test set, and use the rest of the 80% to train the classifier.

In practice, we perform our globally convergent scheme with the following parameters.

- We use $Q = M$ uniform random features to compute the pre-conditioner for each approximate Newton step;
- In the first phase, we decrease μ in a very fast way to λ by starting at $\mu = 1$ and dividing μ by 1000 after performing only a single approximate Newton step (using 2 iterations of conjugate gradient descent);
- In the second phase, we perform 10 approximate Newton steps (each ANS is computed using 8 iterations of conjugate gradient descent).

Selection of λ In the introduction, we claim that in many a learning problem, the parameter λ obtained through cross validation is often much smaller than the ones obtained in statistical bounds which are usually of order $\frac{1}{\sqrt{n}}$. This leads to very ill conditioned problems.

For both data sets, we select λ (and σ , but we omit the double tables from this paper) by computing the test loss and classification errors for different values of λ , and report the evolution of these losses as a function of the parameter λ in Fig. 2 for the Higgs data set, and Fig. 3 for the Susy data set. We see that the optimal λ yield strongly ill-conditioned problems.




Figure 2: **(Left)** Classification error as a function of the regularization parameter and **(Right)** test loss as a function of the regularization parameter, when performing a logistic regression with $M = 2 \times 10^4$ Nyström features on the entire Higgs data set; we select $\lambda = 10^{-9}$.




Figure 3: **(Left)** Classification error as a function of the regularization parameter and **(Right)** test loss as a function of the regularization parameter, when performing a logistic regression with $M = 10^4$ Nyström features on the entire Susy data set; we select $\lambda = 10^{-10}$.

Comparison with accelerated methods Given the M Nystrom points, our aim is to minimize $\widehat{L}_{M,\lambda}$. From an optimization point of view, i.e. from a point of view where the aim is to minimize $\widehat{L}_{M,\lambda}$, we compare our method with a large mini-batch version of Katyusha accelerated SVRG (see [4]).

Indeed, we perform this method using batch sizes of size M ; the theoretical bounds provided in [4] show that the algorithm has linear convergence, with a time complexity of order $O(nM + M^3 + M^2\sqrt{\frac{L}{\lambda}}) \log \frac{1}{\varepsilon}$ to reach precision ε . In the following plots, we compare both methods in terms of passes and time.

By pass, we mean the following.

- In the case of our second-order scheme, we define a pass on the data to be one step of the conjugate gradient descent used to compute approximate newton steps.
- In the case of Katyusha SVRG, we define a pass on the data to be either a full gradient computation or n/M computations of the type $K_{\tau M} T^{-1} \beta$ where T is an upper triangular matrix, and $K_{\tau M}$ is a $M \times M$ kernel matrix, associated to one batch gradient.

We use this notion to measure the speed of our method as they both correspond to natural $O(nM)$ operations, and incorporate the essential of the computing time. However, the second point is often much slower to compute than the first, due to the solving of the triangular system. Thus, the notion of passes is to take with precaution, as a pass for the accelerated SVRG algorithm takes much longer to run than a pass for our method. This is confirmed by the time plots (see Fig. 5 for instance).

Comparison between the two methods - Due to the running time of K-SVRG, we compare both methods for $M = 10000$ Nyström points for both data sets. We compare the performance of these two algorithms with respect to the distance to the optimum in function values as well as classification error Fig. 4 for the Higgs data set, and in Fig. 5 for the Susy data set.

Note on the need for precise optimization - As noted in the introduction, we see in both Fig. 5 and Fig. 4 that precise optimization of the objective function is needed in order to get a good classification error. This justifies a posteriori the use of a second order method. In particular, in Fig. 5, one notes the difference in behavior between the two methods : the second order method converges linearly in a fast way while the first order method slows down because of the condition number.

Note on ill-conditioning - First note that in order to optimize test error, one gets very poorly conditioned problems. As predicted by the rates, we observe that K-SVRG is more sensible to ill-conditioning than our second order scheme. Indeed, in Fig. 6, we have plotted the results for Susy for a smaller condition number with $\lambda = 10^{-8}$, compared to $\lambda = 10^{-10}$ to get optimal test error in Fig. 5. We see that the difference in number of passes needed to reach a certain precision is much lower when $\lambda = 10^{-8}$ in Fig. 6, confirming that K-SVRG behaves better when the condition number is smaller.




Figure 4: **(Left)** Distance to optimum as a function of time and **(Right)** distance to optimum and classification error as a function of the number of passes on the data when performing our second order scheme and K-SVRG to minimize the train loss on Higgs, with 1.0×10^4 Nyström points and $\lambda = 10^{-9}$.




Figure 5: **(Left)** Distance to optimum as a function of time and **(Right)** distance to optimum and classification error as a function of the number of passes on the data when performing our second order scheme and K-SVRG to minimize the train loss on Susy, with 1.0×10^4 Nyström points and $\lambda = 10^{-10}$.

Performance of our method. In Table 1, we record the performance of the following methods, taking the λ values we have obtained previously for the different data sets.

For FALKON (see [29]), we take the parameters suggested in the paper (except for the number of Nyström points needed for Higgs, as our computational capacity is limited).

Method	Susy			Higgs		
	c-error	M	time(m)	c-error	M	time(m)
Logistic regression with K-SVRG	19.64%	10^4	230	27.82 %	10^4	500
Logistic regression with our scheme	19.5%	10^4	15	26.9 %	2.5×10^4	65
Ridge Regression with FALKON ([29])	19.7%	10^4	5	27.16 %	2.5×10^4	60

Table 1: Classification error of different methods




Figure 6: **(Left)** Distance to optimum as a function of time and **(Right)** distance to optimum and classification error as a function of the number of passes on the data when performing our second order scheme and K-SVRG to minimize the train loss on Susy, with 1.0×10^4 Nyström points and $\lambda = 10^{-8}$.

G Solving a projected problem to reduce dimension

G.1 Introduction and notations

In this section, we give ourselves a generalized self-concordant function f whose associated subset we denote with \mathcal{G} . Once again, we will always omit the subscript f in the notations associated to f .

The aim of this section is the following. Given f and $\lambda > 0$, computing an approximate solution to

$$x_\lambda^* = \arg \min_{x \in \mathcal{H}} f_\lambda(x),$$

is often too costly. Instead, we look for a solution in a small subset of \mathcal{H} which we see as the image of a certain orthogonal projector \mathbf{P} and which we denote $\mathcal{H}_\mathbf{P}$. Usually, this subset will be finite dimensional and admit an easy parametrization. Thus we will compare an approximation of x_λ^* to an approximation of

$$x_{\mathbf{P},\lambda}^* = \arg \min_{x \in \mathcal{H}_\mathbf{P}} f_\lambda(x) = \arg \min_{x \in \mathcal{H}} f(\mathbf{P}x) + \frac{\lambda}{2} \|x\|^2.$$

Denote with $f_\mathbf{P}$ the mapping $x \in \mathcal{H} \mapsto f(\mathbf{P}x)$. It is easy to see that, as f is a generalized self-concordant function with \mathcal{G} , $f_\mathbf{P}$ is naturally a generalized self-concordant with $\mathcal{G}_\mathbf{P} := \mathbf{P}\mathcal{G} = \{\mathbf{P}g : g \in \mathcal{G}\}$. Moreover, $x_{\mathbf{P},\lambda}^* = x_{f_\mathbf{P},\lambda}^*$.

We will adopt the following notations for the quantities related to the generalized self-concordant function $f_\mathbf{P}$. Essentially, we always replace $f_\mathbf{P}$ simply by \mathbf{P} from our definitions in appendix.

- For the regularized function :

$$\forall x \in \mathcal{H}, \forall \lambda > 0, f_{\mathbf{P},\lambda}(x) = f_\mathbf{P}(x) + \frac{\lambda}{2} \|x\|^2.$$

- For the Hessians

$$\forall x \in \mathcal{H}, \lambda > 0, \mathbf{H}_{\mathbf{P},\lambda}(x) = \mathbf{H}_{f_{\mathbf{P},\lambda}}(x) = \mathbf{P}\mathbf{H}(\mathbf{P}x)\mathbf{P} + \lambda\mathbf{I}.$$

- $\forall h \in \mathcal{H}, \mathbf{t}_\mathbf{P}(h) := \mathbf{t}_{f_\mathbf{P}}(h) = \mathbf{t}(\mathbf{P}h)$.

- For the Newton decrement:

$$\forall x \in \mathcal{H}, \lambda > 0, \nu_{\mathbf{P},\lambda}(x) = \nu_{f_{\mathbf{P},\lambda}}(x) = \|\nabla f_{\mathbf{P},\lambda}\|_{\mathbf{H}_{\mathbf{P},\lambda}^{-1}(x)} = \|\mathbf{P}\nabla f(\mathbf{P}x) + \lambda x\|_{\mathbf{H}_{\mathbf{P},\lambda}^{-1}(x)}.$$

- For the Dikin ellipsoid radius:

$$\forall \lambda > 0, \forall x \in \mathcal{H}, r_{\mathbf{P},\lambda}(x) := r_{f_{\mathbf{P},\lambda}}(x) = \frac{1}{\sup_{g \in \mathcal{G}} \|\mathbf{P}g\|_{\mathbf{H}_{\lambda,\mathbf{P}}^{-1}(x)}};$$

- For the Dikin ellipsoid:

$$\forall \lambda > 0, \forall c \geq 0, D_{\mathbf{P},\lambda}(c) := D_{f_{\mathbf{P}},\lambda}(c).$$

Note that for any $x \in \mathcal{H}_{\mathbf{P}}$, $r_{\mathbf{P},\lambda}(x) \geq r_{\lambda}(x)$.

We will now introduce the key quantities in order to compare an approximation of $x_{\mathbf{P},\lambda}^*$ to an approximation of x_{λ}^* .

Definition 8 (key quantities). *Define the following quantities*

- For any $\lambda > 0$, the source term $s_{\lambda} := \lambda \|x_{\lambda}^*\|_{\mathbf{H}_{\lambda}^{-1}(x_{\lambda}^*)} = \|\nabla f(x_{\lambda}^*)\|_{\mathbf{H}_{\lambda}^{-1}(x_{\lambda}^*)}$;
- Given an orthogonal projector \mathbf{P} , $\lambda > 0$, and $x \in \mathcal{H}$, the capacity of the projector $C_{\mathbf{P}}(x, \lambda) := \frac{\|\mathbf{H}(x)^{1/2}(\mathbf{I}-\mathbf{P})\|^2}{\lambda}$.

G.2 Relating the projected to the original problem

Given $x \in \mathcal{H}_{\mathbf{P}}$, our aim is to bound $\nu_{\lambda}(x)$ given $\nu_{\lambda,\mathbf{P}}(x)$ and s_{λ} .

Proposition 15. *Let $x \in \mathcal{H}_{\mathbf{P}}$. If*

$$\frac{s_{\lambda}}{r_{\lambda}(x_{\lambda}^*)} \leq \frac{1}{4}, \quad C_{\mathbf{P}}(x_{\lambda}^*, \lambda) \leq \frac{1}{120}, \quad \nu_{\mathbf{P},\lambda}(x) \leq \frac{r_{\mathbf{P},\lambda}(x)}{2},$$

Then it holds:

$$\nu_{\lambda}(x) \leq 3(\nu_{\mathbf{P},\lambda}(x) + s_{\lambda}).$$

Moreover, under these conditions,

- $\|x - x_{\lambda}^*\| \leq 7\lambda^{-1/2}(\nu_{\mathbf{P},\lambda}(x) + s_{\lambda})$;
- $\lambda \|x\|_{\mathbf{H}_{\mathbf{P},\lambda}^{-1}(x)} \leq 7\nu_{\mathbf{P},\lambda}(x) + 9s_{\lambda}$.

Proof. In this proof, introduce the following auxiliary quantity:

$$\gamma_{\lambda} := \frac{s_{\lambda}}{r_{\lambda}(x_{\lambda}^*)}.$$

1) Start by bounding $t(\mathbf{P}x_{\lambda}^* - x_{\lambda}^*)$. It holds:

$$\begin{aligned} t(\mathbf{P}x - x_{\lambda}^*) &= \sup_{g \in \mathcal{G}} |g \cdot (\mathbf{I} - \mathbf{P})x_{\lambda}^*| \\ &\leq \frac{1}{r_{\lambda}(x_{\lambda}^*)} \|(\mathbf{I} - \mathbf{P})x_{\lambda}^*\|_{\mathbf{H}_{\lambda}(x_{\lambda}^*)} \\ &\leq \frac{1}{r_{\lambda}(x_{\lambda}^*)} \|\mathbf{H}_{\lambda}(x_{\lambda}^*)^{1/2}(\mathbf{I} - \mathbf{P})\mathbf{H}_{\lambda}(x_{\lambda}^*)^{1/2}\| \|\mathbf{H}_{\lambda}^{-1/2}(x_{\lambda}^*)x_{\lambda}^*\| \\ &= (1 + C_{\mathbf{P}}(x_{\lambda}^*, \lambda)) \frac{\lambda \|\mathbf{H}_{\lambda}^{-1/2}(x_{\lambda}^*)x_{\lambda}^*\|}{r_{\lambda}(x_{\lambda}^*)} \\ &= (1 + C_{\mathbf{P}}(x_{\lambda}^*, \lambda)) \gamma_{\lambda}. \end{aligned}$$

2) Then bound $t(x_{\mathbf{P},\lambda}^* - \mathbf{P}x_{\lambda}^*)$ First, bound $\nu_{\mathbf{P},\lambda}(\mathbf{P}x_{\lambda}^*)$:

$$\begin{aligned} \nu_{\mathbf{P},\lambda}(\mathbf{P}x_{\lambda}^*) &= \|\mathbf{P}\nabla f_{\lambda}(\mathbf{P}x_{\lambda}^*)\|_{\mathbf{H}_{\lambda,\mathbf{P}}(\mathbf{P}x_{\lambda}^*)^{-1}} \\ &\leq \|\nabla f_{\lambda}(\mathbf{P}x_{\lambda}^*)\|_{\mathbf{H}_{\lambda}(\mathbf{P}x_{\lambda}^*)^{-1}}. \end{aligned}$$

Using Eq. (17), we get $\|\nabla f_{\lambda}(\mathbf{P}x_{\lambda}^*)\|_{\mathbf{H}_{\lambda}(\mathbf{P}x_{\lambda}^*)^{-1}} \leq e^{t((\mathbf{I}-\mathbf{P})x_{\lambda}^*)/2} \nu_{\lambda}(\mathbf{P}x_{\lambda}^*)$. Using Eq. (20), we can bound

$$\nu_{\lambda}(\mathbf{P}x_{\lambda}^*) \leq \bar{\phi}(t((\mathbf{I} - \mathbf{P})x_{\lambda}^*)) \|(\mathbf{I} - \mathbf{P})x_{\lambda}^*\|_{\mathbf{H}_{\lambda}(x_{\lambda}^*)} \leq \bar{\phi}(t((\mathbf{I} - \mathbf{P})x_{\lambda}^*)) (1 + C_{\mathbf{P}}(x_{\lambda}^*, \lambda))s_{\lambda}.$$

Putting things together,

$$\nu_{\mathbf{P},\lambda}(\mathbf{P}x_\lambda^*) \leq e^{t((\mathbf{I}-\mathbf{P})x_\lambda^*)/2} \bar{\phi}(t((\mathbf{I}-\mathbf{P})x_\lambda^*)) (1 + \mathbf{C}_{\mathbf{P}}(x_\lambda^*, \lambda)) s_\lambda.$$

Now

$$\frac{1}{r_{\mathbf{P},\lambda}(\mathbf{P}x_\lambda^*)} \leq \frac{1}{r_\lambda(\mathbf{P}x_\lambda^*)} \leq e^{t((\mathbf{I}-\mathbf{P})x_\lambda^*)/2} \frac{1}{r_\lambda(x_\lambda^*)}.$$

Hence,

$$\frac{\nu_{\mathbf{P},\lambda}(\mathbf{P}x_\lambda^*)}{r_{\mathbf{P},\lambda}(\mathbf{P}x_\lambda^*)} \leq e^{\tilde{t}\lambda} \bar{\phi}(\tilde{t}\lambda) \tilde{t}\lambda, \quad \tilde{t}\lambda = (1 + \mathbf{C}_{\mathbf{P}}(x_\lambda^*, \lambda))\gamma_\lambda.$$

Since $t \mapsto e^t \bar{\phi}(t)$ t is an increasing function whose value in 0 is 0, we find numerically that for $t = \frac{3}{10}$, $e^t \bar{\phi}(t) t \leq \frac{1}{2}$. Hence, if $(1 + \mathbf{C}_{\mathbf{P}}(x_\lambda^*, \lambda))\gamma_\lambda \leq \frac{3}{10}$, then $\frac{\nu_{\mathbf{P},\lambda}(\mathbf{P}x_\lambda^*)}{r_{\mathbf{P},\lambda}(\mathbf{P}x_\lambda^*)} \leq \frac{1}{2}$. Using Lemma 5, this shows that

$$\mathbf{t}_{\mathbf{P}}(\mathbf{P}x_\lambda^* - x_{\mathbf{P},\lambda}^*) = \mathbf{t}(\mathbf{P}x_\lambda^* - x_{\mathbf{P},\lambda}^*) \leq \log 2.$$

3) Getting a bound for $\mathbf{t}(x - x_\lambda^*)$. To do so, combine the two previous bounds with the fact that if $\nu_{\mathbf{P},\lambda}(x) \leq \frac{r_{\mathbf{P},\lambda}(x)}{2}$, then using Lemma 5 with $f_{\mathbf{P}}$, $\mathbf{t}_{\mathbf{P}}(x - x_{\mathbf{P},\lambda}^*) = \mathbf{t}(x - x_{\mathbf{P},\lambda}^*) \leq \log 2$. Thus, if

$$(1 + \mathbf{C}_{\mathbf{P}}(x_\lambda^*, \lambda))\gamma_\lambda \leq \frac{3}{10}, \quad \nu_{\mathbf{P},\lambda}(x) \leq \frac{r_{\mathbf{P},\lambda}(x)}{2},$$

then it holds

$$\mathbf{t}(x - x_\lambda^*) \leq \frac{3}{10} + 2 \log 2.$$

4) A technical result to bound $\|\mathbf{H}_\lambda(x)^{-1/2} \mathbf{H}_{\mathbf{P},\lambda}(x)^{1/2}\|$. Using the fact that $\mathbf{P}x = x$, and Lemma 23, applied to $\mathbf{A} = \mathbf{H}(x)$, we get

$$\|\mathbf{H}_\lambda(x)^{-1/2} \mathbf{H}_{\mathbf{P},\lambda}(x)^{1/2}\| \leq 1 + \sqrt{\mathbf{C}_{\mathbf{P}}(x, \lambda)}.$$

Then, one can easily bound $\mathbf{C}_{\mathbf{P}}(x, \lambda) \leq e^{t(x-x_\lambda^*)} \mathbf{C}_{\mathbf{P}}(x_\lambda^*, \lambda)$.

5) Let us now bound $\nu_\lambda(x)$. First, decompose the term

$$\nu_\lambda(x) = \|\nabla f_\lambda(x)\|_{\mathbf{H}_\lambda^{-1}(x)} \leq \|\mathbf{P}\nabla f_\lambda(x)\|_{\mathbf{H}_\lambda^{-1}(x)} + \|(\mathbf{I} - \mathbf{P})\nabla f(x)\|_{\mathbf{H}_\lambda^{-1}(x)}.$$

Since $x \in \mathcal{H}_{\mathbf{P}}$, $\|\mathbf{P}\nabla f_\lambda(x)\|_{\mathbf{H}_\lambda^{-1}(x)} = \|\nabla f_{\mathbf{P},\lambda}(x)\|_{\mathbf{H}_\lambda^{-1}(x)}$, and using the previous point, we get

$$\|\mathbf{P}\nabla f_\lambda(x)\|_{\mathbf{H}_\lambda^{-1}(x)} \leq \left(1 + e^{t(x-x_\lambda^*)/2} \sqrt{\mathbf{C}_{\mathbf{P}}(x_\lambda^*, \lambda)}\right) \nu_{\mathbf{P},\lambda}(x).$$

Let us now bound the second term. We divide it into two terms:

$$\|(\mathbf{I} - \mathbf{P})\nabla f(x)\|_{\mathbf{H}_\lambda^{-1}(x)} \leq \|(\mathbf{I} - \mathbf{P})(\nabla f(x) - \nabla f(x_\lambda^*))\|_{\mathbf{H}_\lambda^{-1}(x)} + \|(\mathbf{I} - \mathbf{P})\nabla f(x_\lambda^*)\|_{\mathbf{H}_\lambda^{-1}(x)}.$$

The second term can be bounded in the following way:

$$\|(\mathbf{I} - \mathbf{P})\nabla f(x_\lambda^*)\|_{\mathbf{H}_\lambda^{-1}(x)} \leq \frac{1}{\sqrt{\lambda}} \|(\mathbf{I} - \mathbf{P})\mathbf{H}_\lambda^{1/2}(x_\lambda^*)\| \|\nabla f(x_\lambda^*)\|_{\mathbf{H}_\lambda^{-1}(x_\lambda^*)} \leq \sqrt{1 + \mathbf{C}_{\mathbf{P}}(x_\lambda^*, \lambda)} s_\lambda.$$

For the first term, we proceed in the following way.

$$\begin{aligned} \|(\mathbf{I} - \mathbf{P})(\nabla f(x) - \nabla f(x_\lambda^*))\|_{\mathbf{H}_\lambda^{-1}(x)} &= \left\| \int_0^1 \mathbf{H}_\lambda^{-1/2}(x) (\mathbf{I} - \mathbf{P}) \mathbf{H}(x_t) (x - x_\lambda^*) dt \right\| \\ &\leq \frac{1}{\sqrt{\lambda}} \int_0^1 \|(\mathbf{I} - \mathbf{P})\mathbf{H}^{1/2}(x_t)\| \|\mathbf{H}^{1/2}(x_t)(x - x_\lambda^*)\| dt \\ &\leq \sqrt{\mathbf{C}_{\mathbf{P}}(x_\lambda^*, \lambda)} \bar{\phi}(\mathbf{t}(x - x_\lambda^*)) \|x - x_\lambda^*\|_{\mathbf{H}(x_\lambda^*)} \\ &\leq \sqrt{\mathbf{C}_{\mathbf{P}}(x_\lambda^*, \lambda)} e^{t(x-x_\lambda^*)} \nu_\lambda(x). \end{aligned}$$

Hence the final bound:

$$\left(1 - \sqrt{\mathbf{C}_{\mathbf{P}}(x_{\lambda}^*, \lambda)} e^{\mathbf{t}(x-x_{\lambda}^*)}\right) \nu_{\lambda}(x) \leq \left(1 + e^{\mathbf{t}(x-x_{\lambda}^*)/2} \sqrt{\mathbf{C}_{\mathbf{P}}(x_{\lambda}^*, \lambda)}\right) \nu_{\mathbf{P},\lambda}(x) + \sqrt{1 + \mathbf{C}_{\mathbf{P}}(x_{\lambda}^*, \lambda)} s_{\lambda}.$$

Now if $\mathbf{C}_{\mathbf{P}}(x_{\lambda}^*, \lambda) \leq \frac{1}{120}$, we see that $\sqrt{\mathbf{C}_{\mathbf{P}}(x_{\lambda}^*, \lambda)} e^{\mathbf{t}(x-x_{\lambda}^*)} \leq \frac{1}{2}$, and hence, using the bound on $\mathbf{t}(x - x_{\lambda}^*)$,

$$\nu_{\lambda}(x) \leq 3(\nu_{\mathbf{P},\lambda}(x) + s_{\lambda}).$$

6) Showing the last two points . We leverage the fact that $\nu_{\lambda}(x) \leq 3(\nu_{\mathbf{P},\lambda}(x) + s_{\lambda})$ and $\mathbf{t}(x - x_{\lambda}^*) \leq \frac{3}{10} + 2 \log 2$.

To show the first bound, we plug in the previous results in the following equation:

$$\|x - x_{\lambda}^*\| \leq \lambda^{-1/2} \|x - x_{\lambda}^*\|_{\mathbf{H}_{\lambda}(x)} \leq \frac{1}{\underline{\phi}(\mathbf{t}(x - x_{\lambda}^*))} \lambda^{-1/2} \nu_{\lambda}(x).$$

The last inequality is obtained using Eq. (18).

To show the second point, we use the fact that $x \in \mathcal{H}_{\mathbf{P}}$ to show that

$$\lambda \|x\|_{\mathbf{H}_{\mathbf{P},\lambda}^{-1}(x)} \leq \lambda \|x\|_{\mathbf{H}_{\lambda}^{-1}(x)} \leq \lambda \|x - x_{\lambda}^*\|_{\mathbf{H}_{\lambda}(x)} + \lambda \|x_{\lambda}^*\|_{\mathbf{H}_{\lambda}^{-1}(x)}.$$

Then applying Eq. (17) and Eq. (18):

$$\lambda \|x\|_{\mathbf{H}_{\mathbf{P},\lambda}^{-1}(x)} \leq \frac{1}{\underline{\phi}(\mathbf{t}(x - x_{\lambda}^*))} \nu_{\lambda}(x) + e^{\mathbf{t}(x-x_{\lambda}^*)/2} s_{\lambda}.$$

We then use the previous results to conclude. □

G.3 Finding a good projector

Lemma 16. *If for a certain $\eta \leq \lambda$ and for a certain constant C , $\|\mathbf{H}_{\eta}^{1/2}(x)(\mathbf{I} - \mathbf{P})\|^2 \leq C\eta$, then*

$$\mathbf{C}_{\mathbf{P}}(x, \lambda) \leq \frac{C\eta}{\lambda}.$$

Proof. This is completely direct, using the fact that $\mathbf{H}^{1/2}(x) \preceq \mathbf{H}_{\eta}^{1/2}(x)$. □

H Relations between statistical problems and empirical problem.

In this section, we recall and reformulate the framework from [23].

H.1 Statistical problem and ERM estimator

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 Hilbert space, with norm $\|\cdot\|$, and let $(f_z)_{z \in \mathcal{Z}}$ be a family of functions on \mathcal{H} . Our goal is to minimize the *expected risk* with respect to $x \in \mathcal{H}$:

$$\inf_{x \in \mathcal{H}} f(x) := \mathbb{E} [f_Z(x)].$$

Given $(z_i)_{i=1}^n \in \mathcal{Z}^n$, we define the *empirical risk*:

$$\widehat{f}(x) := \frac{1}{n} \sum_{i=1}^n f_{z_i}(x),$$

and consider the following estimator based on regularized empirical risk minimization given $\lambda > 0$ (note that the minimizer is unique in this case):

$$\widehat{x}_\lambda^* = \arg \min_{x \in \mathcal{H}} \widehat{f}_\lambda(x) := \widehat{f}(x) + \frac{\lambda}{2} \|x\|^2,$$

where we assume the following.

Assumption 6 (i.i.d. data). *The samples $(z_i)_{1 \leq i \leq n}$ are independently and identically distributed according to ρ .*

We make the following assumption on the family (f_z) (this is a reformulation of Assumption 8 in [23])

Assumption 7 (Generalized self-concordance). *For any $z \in \mathcal{Z}$, there exists an associated subset $\mathcal{G}_z \subset \mathcal{H}$ such that (f_z, \mathcal{G}_z) is generalized self-concordant in the sense of Definition 3.*

Moreover we require the following technical assumption to guarantee that f and its derivatives are well defined for any $x \in \mathcal{H}$ (this is a reformulation of Assumptions 3 and 4 in [23], and the necessary conditions to obtain Proposition 3).

Assumption 8 (Technical assumptions). *The mapping $(z, x) \in \mathcal{Z} \times \mathcal{H} \mapsto f_z(x)$ is measurable. Moreover,*

- *the random variables $\|f_Z(0)\|, \|\nabla f_Z(0)\|, \text{Tr}(\nabla^2 f_Z(0))$ are bounded;*
- *$\mathcal{G} := \bigcup_{z \in \text{supp}(Z)} \mathcal{G}_z$ is a bounded subset of \mathcal{H} .*

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 $z \mapsto \mathcal{G}_z$ is uniformly bounded on bounded sets, then they hold.

Proposition 16. *Under Assumptions 7 and 8, the function (f, \mathcal{G}) (or simply f) is generalized self-concordant.*

Moreover, under Assumption 6, define

$$\widehat{\mathcal{G}} := \bigcup_{i=1}^n \mathcal{G}_{z_i}.$$

Then $(\widehat{f}, \widehat{\mathcal{G}})$ (or simply \widehat{f}) is generalized self-concordant. Moreover, note that $\widehat{\mathcal{G}} \subset \mathcal{G}$.

The main regularity assumption we make on our statistical problems follows (see Assumption 5 in [23]).

Assumption 9 (Existence of a minimizer). *There exists $x^* \in \mathcal{H}$ such that $f(x^*) = \inf_{x \in \mathcal{H}} f(x)$.*

Notations We adopt all the notations from Appendix A for f and \widehat{f} , which are generalized self-concordant functions with associated subsets given in Proposition 16 with the following conventions:

- For all quantities relating to f , we omit the subscript f as usual;
- For all quantities relating to \widehat{f} , we omit the subscript \widehat{f} and instead put a hat over all these quantities. For instance:

$$\widehat{\mathbf{H}}(x) := \mathbf{H}_{\widehat{f}}(x) = \frac{1}{n} \sum_{i=1}^n \nabla^2 f_{z_i}(x), \quad \widehat{r}_\lambda(x) := r_{\widehat{f}, \lambda}(x) = \frac{1}{\sup_{g \in \widehat{\mathcal{G}}} \|g\|_{\widehat{\mathbf{H}}_\lambda^{-1}(x)}}, \text{ etc...}$$

Recall the two main quantities introduced in [23] to establish the quality of our estimator \widehat{x}_λ^* (in [23], this is a mix between Proposition 2 and Definition 3).

Proposition 17 (Bias, degrees of freedom). *Suppose Assumptions 7 to 9 are satisfied. The following key quantities are well defined:*

- the bias $\mathbf{b}_\lambda = \|\mathbf{H}_\lambda(x^*)^{-1/2} \nabla f_\lambda(x^*)\|$;
- the effective dimension $\text{df}_\lambda = \mathbb{E} [\|\mathbf{H}_\lambda(x^*)^{-1/2} \nabla f_Z(x^*)\|^2]$.

Moreover, we also introduce the following quantities:

$$\mathbf{B}_1^* := \sup_{z \in \text{supp}(Z)} \|\nabla f_z(x^*)\|, \quad \mathbf{B}_2^* := \sup_{z \in \text{supp}(Z)} \text{Tr}(\nabla^2 f_z(x^*)), \quad \mathbf{Q}^* = \frac{\mathbf{B}_1^*}{\sqrt{\mathbf{B}_2^*}}.$$

We can now recall the main theorem of [23] (Theorem 4), which quantifies the behavior of the ERM estimator:

Theorem 8 (Bound for the ERM estimator). *Let $n \in \mathbb{N}$, $\delta \in (0, 1/2]$, $0 < \lambda \leq \mathbf{B}_2^*$. Whenever*

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

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

$$f(\widehat{x}_\lambda^*) - f(x^*) \leq C_{\text{bias}} \mathbf{b}_\lambda^2 + C_{\text{var}} \frac{\text{df}_\lambda \vee (\mathbf{Q}^*)^2}{n} \log \frac{2}{\delta}, \quad (37)$$

where $C_{\text{bias}}, C_{\text{var}}, \Box_1 \leq 414$, $\Delta_1, \Delta_2 \leq 5184$.

H.2 Link between a good approximation of \widehat{x}_λ^* and x^*

In this paper, we provide an algorithm which can effectively compute a good approximation of \widehat{x}_λ^* (as it is a finite sum problem which can be solved). This algorithm will return a certain $x \in \mathcal{H}$, whose precision with respect to the empirical problem will be characterized by $\widehat{v}_\lambda(x)$. The aim of the following lemma is to see how this approximation x behaves with respect to the statistical problem.

Lemma 17. *Suppose the conditions for Thm. 8 are satisfied, i.e. let $n \in \mathbb{N}$, $\delta \in (0, 1/2]$, $0 < \lambda \leq \mathbf{B}_2^*$ and suppose*

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

Let x be an approximation of \widehat{x}_λ^* characterized by its Newton decrement $\widehat{v}_\lambda(x)$. If

$$\widehat{v}_\lambda(x) \leq \frac{\widehat{r}_\lambda(x)}{2}, \quad \widehat{v}_\lambda(x) \leq \frac{r_\lambda(x^*)}{2},$$

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

$$f(x) - f(x^*) \leq 14(f(\widehat{x}_\lambda^*) - f(x^*)) + 30\widehat{v}_\lambda(x)^2.$$

Proof. Using Eq. (16),

$$\begin{aligned} f(x) - f(\widehat{x}_\lambda^*) &\leq \langle \nabla f(\widehat{x}_\lambda^*), x - \widehat{x}_\lambda^* \rangle_{\mathcal{H}} + \psi(\mathbf{t}(x - \widehat{x}_\lambda^*)) \|x - \widehat{x}_\lambda^*\|_{\mathbf{H}_\lambda(\widehat{x}_\lambda^*)}^2 \\ &\leq \frac{1}{2} \|\nabla f(\widehat{x}_\lambda^*)\|_{\mathbf{H}_\lambda^{-1}(\widehat{x}_\lambda^*)}^2 + \left(\psi(\mathbf{t}(x - \widehat{x}_\lambda^*)) + \frac{1}{2} \right) \|x - \widehat{x}_\lambda^*\|_{\mathbf{H}_\lambda(\widehat{x}_\lambda^*)}^2. \end{aligned}$$

1. Let us bound $\|\nabla f(\widehat{x}_\lambda^*)\|_{\mathbf{H}_\lambda^{-1}(\widehat{x}_\lambda^*)}$

$$\begin{aligned} \|\nabla f(\widehat{x}_\lambda^*)\|_{\mathbf{H}_\lambda^{-1}(\widehat{x}_\lambda^*)} &\leq \int_0^1 \|\mathbf{H}_\lambda^{-1/2}(\widehat{x}_\lambda^*)\mathbf{H}(x_t)(\widehat{x}_\lambda^* - x^*)\| dt, & x_t &= (1-t)\widehat{x}_\lambda^* + tx^* \\ &\leq \int_0^1 \|\mathbf{H}_\lambda^{-1/2}(\widehat{x}_\lambda^*)\mathbf{H}^{1/2}(x_t)\| \|\mathbf{H}^{1/2}(x_t)(\widehat{x}_\lambda^* - x^*)\| dt. \end{aligned}$$

Now using equation Eq. (14)

$$\mathbf{H}(x_t) \preceq e^{t\mathfrak{t}(\widehat{x}_\lambda^* - x^*)} \mathbf{H}(\widehat{x}_\lambda^*), \quad \mathbf{H}(x_t) \preceq e^{(1-t)\mathfrak{t}(\widehat{x}_\lambda^* - x^*)}.$$

Thus:

$$\|\nabla f(\widehat{x}_\lambda^*)\|_{\mathbf{H}_\lambda^{-1}(\widehat{x}_\lambda^*)} \leq e^{\mathfrak{t}(\widehat{x}_\lambda^* - x^*)/2} \|\widehat{x}_\lambda^* - x^*\|_{\mathbf{H}(x^*)}.$$

Finally, using equation Eq. (16)

$$\|\nabla f(\widehat{x}_\lambda^*)\|_{\mathbf{H}_\lambda^{-1}(\widehat{x}_\lambda^*)} \leq \frac{e^{\mathfrak{t}(\widehat{x}_\lambda^* - x^*)/2}}{\psi(-\mathfrak{t}(\widehat{x}_\lambda^* - x^*))^{1/2}} (f(\widehat{x}_\lambda^*) - f(x^*))^{1/2}.$$

2. Let us bound the terms involving $\|x - \widehat{x}_\lambda^*\|_{\mathbf{H}_\lambda(\widehat{x}_\lambda^*)}$ Note that using Eq. (18) and Eq. (17) applied to \widehat{f} ,

$$\|x - \widehat{x}_\lambda^*\|_{\mathbf{H}_\lambda(\widehat{x}_\lambda^*)} \leq \|\mathbf{H}_\lambda^{1/2}(\widehat{x}_\lambda^*)\widehat{\mathbf{H}}_\lambda^{-1/2}(\widehat{x}_\lambda^*)\| \frac{e^{\widehat{\mathfrak{t}}(x - \widehat{x}_\lambda^*)/2}}{\underline{\phi}(\widehat{\mathfrak{t}}(x - \widehat{x}_\lambda^*))} \widehat{\nu}_\lambda(x).$$

This also leads to:

$$\begin{aligned} \mathfrak{t}(x - \widehat{x}_\lambda^*) &\leq \frac{1}{r_\lambda(\widehat{x}_\lambda^*)} \|\mathbf{H}_\lambda^{1/2}(\widehat{x}_\lambda^*)\widehat{\mathbf{H}}_\lambda^{-1/2}(\widehat{x}_\lambda^*)\| \|x - \widehat{x}_\lambda^*\|_{\widehat{\mathbf{H}}_\lambda(\widehat{x}_\lambda^*)} \\ &\leq \frac{1}{r_\lambda(\widehat{x}_\lambda^*)} \|\mathbf{H}_\lambda^{1/2}(\widehat{x}_\lambda^*)\widehat{\mathbf{H}}_\lambda^{-1/2}(\widehat{x}_\lambda^*)\| \frac{e^{\widehat{\mathfrak{t}}(x - \widehat{x}_\lambda^*)/2}}{\underline{\phi}(\widehat{\mathfrak{t}}(x - \widehat{x}_\lambda^*))} \widehat{\nu}_\lambda(x). \end{aligned}$$

3. Putting things together In the end, we get

$$\begin{aligned} f(x) - f(x^*) &\leq \left(1 + \frac{e^{\mathfrak{t}(\widehat{x}_\lambda^* - x^*)}}{\psi(-\mathfrak{t}(\widehat{x}_\lambda^* - x^*))}\right) (f(\widehat{x}_\lambda^*) - f(x^*)) \\ &\quad + \left(\psi(\mathfrak{t}(x - \widehat{x}_\lambda^*)) + \frac{1}{2}\right) \left(e^{\mathfrak{t}(\widehat{x}_\lambda^* - x^*)/2} \|\mathbf{H}_\lambda^{1/2}(x^*)\widehat{\mathbf{H}}_\lambda^{-1/2}(x^*)\| \frac{e^{\widehat{\mathfrak{t}}(x - \widehat{x}_\lambda^*)/2}}{\underline{\phi}(\widehat{\mathfrak{t}}(x - \widehat{x}_\lambda^*))}\right) \widehat{\nu}_\lambda(x)^2. \end{aligned}$$

Moreover, we bound

$$\mathfrak{t}(x - \widehat{x}_\lambda^*) \leq e^{(\mathfrak{t}(x^* - \widehat{x}_\lambda^*) + \mathfrak{t}(\widehat{x}_\lambda^* - x^*)) / 2} \|\mathbf{H}_\lambda^{1/2}(x^*)\widehat{\mathbf{H}}_\lambda^{-1/2}(x^*)\| \frac{e^{\widehat{\mathfrak{t}}(x - \widehat{x}_\lambda^*)/2}}{\underline{\phi}(\widehat{\mathfrak{t}}(x - \widehat{x}_\lambda^*))} \widehat{\nu}_\lambda(x).$$

4. Plugging in previous results Under the assumptions of this lemma, which include the assumptions of Theorem 4. in [23], we get the following bounds.

- In [23], the assumptions of Theorem 4 imply that we can use Lemma 9, which uses Lemma 8 in which we show that with probability at least $1 - \delta$,

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

- Still using the assumptions of Theorem 4, we see in the proof of this theorem that the assumptions of Theorem 7 of [23] are satisfied in the case where $b_\lambda \leq \frac{r_\lambda(x^*)}{2}$, and thus that

$$\mathfrak{t}(\hat{x}_\lambda^* - x_\lambda^*) \leq \log 2, \quad \mathfrak{t}(x_\lambda^* - x^*) \leq \log 2.$$

Plugging in all these bounds, we get

$$\begin{aligned} & \left(1 + \frac{e^{\mathfrak{t}(\hat{x}_\lambda^* - x^*)}}{\psi(-\mathfrak{t}(\hat{x}_\lambda^* - x^*))} \right) \leq 14, \quad \mathfrak{t}(x - \hat{x}_\lambda^*) \leq 6, \\ & \left(\psi(\mathfrak{t}(x - \hat{x}_\lambda^*)) + \frac{1}{2} \right) \left(e^{\mathfrak{t}(\hat{x}_\lambda^* - x_\lambda^*)/2} \|\mathbf{H}_\lambda^{1/2}(x_\lambda^*) \widehat{\mathbf{H}}_\lambda^{-1/2}(x_\lambda^*)\| \frac{e^{\widehat{\mathfrak{t}}(x - \hat{x}_\lambda^*)/2}}{\underline{\phi}(\widehat{\mathfrak{t}}(x - \hat{x}_\lambda^*))} \right) \leq 30. \end{aligned}$$

□

H.3 Bounds when we solve a projected empirical problem

In this section, we place ourselves in the setting of Appendix G. In this section, we had argued that for computational purposes, it was less costly to compute an approximate solution to a projected problem.

In this section, we assume that we are going to project the regularized empirical problem, that is solve approximately

$$x \approx \arg \min_{x \in \mathcal{H}} \widehat{f}_{\mathbf{P}, \lambda}(x) = \widehat{f}(\mathbf{P}x) + \frac{\lambda}{2} \|x\|^2.$$

for a given orthogonal projection \mathbf{P} . Recall from Appendix G that there is a natural way of seeing $\widehat{f}_{\mathbf{P}}$ as a generalized self-concordant function. We import all the notations from this section, keeping a $\widehat{\cdot}$ over all notations to mark the fact that we are projecting \widehat{f} and not f .

To quantify the quality of the approximation x , we will use the Newton decrement for the empirical projected problem $\widehat{\nu}_{\mathbf{P}, \lambda}(x) := \nu_{\widehat{f}_{\mathbf{P}, \lambda}}(x)$.

As we see in Proposition 15, under certain conditions, bounding $\widehat{\nu}_\lambda(x)$ amounts to bounding two terms:

- The empirical source $\widehat{s}_\lambda := \lambda \|\widehat{x}_\lambda^*\|_{\widehat{\mathbf{H}}_\lambda^{-1}(\widehat{x}_\lambda^*)}$,
- The projected empirical Newton decrement $\widehat{\nu}_{\mathbf{P}, \lambda}(x)$.

1. Bounding the empirical source term \widehat{s}_λ Start by bounding the source empirical source term using quantities we know.

Lemma 18 (Empirical source). *Let $n \in \mathbb{N}$, $\delta \in (0, 1/2]$, $0 < \lambda \leq \mathbf{B}_2^*$. Whenever*

$$n \geq \Delta_1 \frac{\mathbf{B}_2^*}{\lambda} \log \frac{8 \square_1^2 \mathbf{B}_2^*}{\lambda \delta}, \quad \sqrt{\Delta_2 \frac{\text{df}_\lambda \vee (\mathbf{Q}^*)^2}{n} \log \frac{2}{\delta}} \leq r_\lambda(x^*), \quad 2b_\lambda \leq r_\lambda(x^*).$$

The following holds, with probability at least $1 - 2\delta$.

$$\widehat{s}_\lambda \leq 8 b_\lambda + 80 \sqrt{\frac{\text{df}_\lambda \vee (\mathbf{Q}^*)^2 \log \frac{2}{\delta}}{n}}.$$

Moreover, we also have the following bound :

$$\|\widehat{x}_\lambda^* - x^*\| \leq 3 \lambda^{-1/2} b_\lambda + 8 \lambda^{-1/2} \sqrt{\frac{\text{df}_\lambda \vee (\mathbf{Q}^*)^2 \log \frac{2}{\delta}}{n}}.$$

Proof. We first decompose the source term into two terms, and then apply different bounds from [23] to effectively bound it. We will use the following quantity:

$$\widehat{v}_\lambda := \|\mathbf{H}_\lambda^{1/2}(x_\lambda^*) \widehat{\mathbf{H}}_\lambda^{-1/2}(x_\lambda^*)\|^2 \|\nabla \widehat{f}_\lambda(x_\lambda^*)\|_{\mathbf{H}_\lambda^{-1}(x_\lambda^*)}.$$

It is also defined in equation (23) in [23].

1. Dividing \widehat{s}_λ into two controllable terms . Decompose

$$\begin{aligned} \widehat{s}_\lambda &= \|\lambda \widehat{x}_\lambda^*\|_{\widehat{\mathbf{H}}_\lambda^{-1}(\widehat{x}_\lambda^*)} \leq \|\widehat{\mathbf{H}}_\lambda^{-1/2}(\widehat{x}_\lambda^*) \mathbf{H}_\lambda^{1/2}(x_\lambda^*)\| \|\lambda \widehat{x}_\lambda^*\|_{\mathbf{H}_\lambda^{-1}(x_\lambda^*)} \\ &\leq \|\widehat{\mathbf{H}}_\lambda^{-1/2}(\widehat{x}_\lambda^*) \mathbf{H}_\lambda^{1/2}(x_\lambda^*)\| \left(\|\nabla f_\lambda(x_\lambda^*)\|_{\mathbf{H}_\lambda^{-1}(x_\lambda^*)} + \|\nabla f(x_\lambda^*)\|_{\mathbf{H}_\lambda^{-1}(x_\lambda^*)} \right). \end{aligned}$$

On the one hand, from the previous proof, we get

$$\begin{aligned} \|\nabla f(x_\lambda^*)\|_{\mathbf{H}_\lambda^{-1}(x_\lambda^*)} &\leq e^{t(\widehat{x}_\lambda^* - x^*)/2} \|\widehat{x}_\lambda^* - x^*\|_{\mathbf{H}(x^*)} \\ &\leq e^{t(\widehat{x}_\lambda^* - x^*)/2} \left(e^{t(x_\lambda^* - x^*)} \|\widehat{x}_\lambda^* - x_\lambda^*\|_{\mathbf{H}_\lambda(x_\lambda^*)} + \|x_\lambda^* - x^*\|_{\mathbf{H}_\lambda(x^*)} \right) \\ &\leq e^{t(\widehat{x}_\lambda^* - x^*)/2} \left(\frac{e^{t(x_\lambda^* - x^*)}}{\underline{\phi}(t(\widehat{x}_\lambda^* - x_\lambda^*))} \widehat{v}_\lambda + \frac{1}{\underline{\phi}(t(x_\lambda^* - x^*))} \mathbf{b}_\lambda \right). \end{aligned}$$

In the last line, we use the fact that $\|\widehat{x}_\lambda^* - x_\lambda^*\|_{\mathbf{H}_\lambda(x_\lambda^*)} \leq \|\mathbf{H}_\lambda^{1/2}(x_\lambda^*) \widehat{\mathbf{H}}_\lambda^{-1/2}(x_\lambda^*)\| \|\widehat{x}_\lambda^* - x_\lambda^*\|_{\widehat{\mathbf{H}}_\lambda(x_\lambda^*)}$ and then bound it using Eq. (18) applied to \widehat{f} to get

$$\begin{aligned} \|\widehat{x}_\lambda^* - x_\lambda^*\|_{\widehat{\mathbf{H}}_\lambda(x_\lambda^*)} &\leq \frac{1}{\underline{\phi}(\widehat{t}(x_\lambda^* - \widehat{x}_\lambda^*))} \|\nabla \widehat{f}_\lambda(x_\lambda^*)\|_{\widehat{\mathbf{H}}_\lambda^{-1}(x_\lambda^*)} \\ &\leq \frac{1}{\underline{\phi}(t(x_\lambda^* - \widehat{x}_\lambda^*))} \|\mathbf{H}_\lambda^{1/2}(x_\lambda^*) \widehat{\mathbf{H}}_\lambda^{-1/2}(x_\lambda^*)\| \|\nabla \widehat{f}_\lambda(x_\lambda^*)\|_{\mathbf{H}_\lambda^{-1}(x_\lambda^*)}. \end{aligned}$$

On the other hand, apply successively Eq. (18) to f and \widehat{f} using the fact that $\widehat{t} \leq t$ to get

$$\begin{aligned} \|\nabla f_\lambda(x_\lambda^*)\|_{\mathbf{H}_\lambda^{-1}(x_\lambda^*)} &= \|\nabla f_\lambda(\widehat{x}_\lambda^*) - \nabla f_\lambda(x_\lambda^*)\|_{\mathbf{H}_\lambda^{-1}(x_\lambda^*)} \\ &\leq e^{t(\widehat{x}_\lambda^* - x_\lambda^*)/2} \overline{\phi}(t(\widehat{x}_\lambda^* - x_\lambda^*)) \|\widehat{x}_\lambda^* - x_\lambda^*\|_{\mathbf{H}_\lambda(x_\lambda^*)} \\ &\leq e^{t(\widehat{x}_\lambda^* - x_\lambda^*)/2} \overline{\phi}(t(\widehat{x}_\lambda^* - x_\lambda^*)) \|\mathbf{H}_\lambda^{1/2}(x_\lambda^*) \widehat{\mathbf{H}}_\lambda^{-1/2}(x_\lambda^*)\| \|\widehat{x}_\lambda^* - x_\lambda^*\|_{\widehat{\mathbf{H}}_\lambda(x_\lambda^*)} \\ &\leq \frac{e^{t(\widehat{x}_\lambda^* - x_\lambda^*)/2} \overline{\phi}(t(\widehat{x}_\lambda^* - x_\lambda^*))}{\underline{\phi}(t(\widehat{x}_\lambda^* - x_\lambda^*))} \|\mathbf{H}_\lambda^{1/2}(x_\lambda^*) \widehat{\mathbf{H}}_\lambda^{-1/2}(x_\lambda^*)\|^2 \|\nabla \widehat{f}_\lambda(x_\lambda^*)\|_{\mathbf{H}_\lambda(x_\lambda^*)} \\ &= e^{3t(\widehat{x}_\lambda^* - x_\lambda^*)/2} \widehat{v}_\lambda. \end{aligned}$$

Putting things together:

$$\widehat{s}_\lambda \leq \|\widehat{\mathbf{H}}_\lambda^{-1/2}(\widehat{x}_\lambda^*) \mathbf{H}_\lambda^{1/2}(x_\lambda^*)\| \left(e^{3t(x_\lambda^* - \widehat{x}_\lambda^*)/2} \left(1 + \frac{1}{\underline{\phi}(t(x_\lambda^* - \widehat{x}_\lambda^*))} \right) \widehat{v}_\lambda + \frac{e^{t(x_\lambda^* - \widehat{x}_\lambda^*)/2}}{\underline{\phi}(t(x_\lambda^* - x^*))} \mathbf{b}_\lambda \right).$$

2. We now import the results from [23] . The following hypotheses imply those of Thms 4 and 7 in [23]:

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

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

In particular, they imply that with probability at least $1 - 2\delta$:

- $\widehat{v}_\lambda \leq \frac{1}{2}b_\lambda + 4\Box_1 \sqrt{\frac{\text{df}_\lambda \vee (\mathbf{Q}^*)^2 \log \frac{2}{\delta}}{n}}$;
- $\|\mathbf{H}_\lambda^{1/2}(x_\lambda^*) \widehat{\mathbf{H}}_\lambda^{-1/2}(x_\lambda^*)\| \leq \sqrt{2}$;
- $\mathfrak{t}(x^* - x_\lambda^*) \leq \log 2$;
- $\mathfrak{t}(\widehat{x}_\lambda^* - x_\lambda^*) \leq \log 2$.

Hence, plugging these bounds in the previous equation, we get

$$\widehat{s}_\lambda \leq 8b_\lambda + 80 \sqrt{\frac{\text{df}_\lambda \vee (\mathbf{Q}^*)^2 \log \frac{2}{\delta}}{n}}.$$

3. Note that in what has been done previously, we can bound:

$$\|\widehat{x}_\lambda^* - x_\lambda^*\|_{\mathbf{H}_\lambda(x_\lambda^*)} \leq \frac{1}{\underline{\phi}(\mathfrak{t}(x_\lambda^* - \widehat{x}_\lambda^*))} \widehat{v}_\lambda \leq b_\lambda + 8 \sqrt{\frac{\text{df}_\lambda \vee (\mathbf{Q}^*)^2 \log \frac{2}{\delta}}{n}}.$$

Moreover,

$$\|x_\lambda^* - x^*\|_{\mathbf{H}_\lambda(x^*)} \leq \frac{1}{\underline{\phi}(\mathfrak{t}(x_\lambda^* - x^*))} \|\nabla f_\lambda(x^*)\|_{\mathbf{H}_\lambda^{-1}(x^*)} \leq 2b_\lambda.$$

Hence:

$$\|\widehat{x}_\lambda^* - x^*\| \leq 3 \lambda^{-1/2} b_\lambda + 8 \lambda^{-1/2} \sqrt{\frac{\text{df}_\lambda \vee (\mathbf{Q}^*)^2 \log \frac{2}{\delta}}{n}}.$$

□

2. Final bound for the projected ERM approximation In this paragraph, denote with $C_{\mathbf{P}}(x, \lambda)$ the quantity $\frac{\|\mathbf{H}^{1/2}(x)(\mathbf{I}-\mathbf{P})\|^2}{\lambda}$ and $\widehat{C}_{\mathbf{P}}(x, \lambda)$ the quantity $\frac{\|\widehat{\mathbf{H}}^{1/2}(x)(\mathbf{I}-\mathbf{P})\|^2}{\lambda}$

Lemma 19. Let $n \in \mathbb{N}$, $\delta \in (0, 1/2]$, $0 < \lambda \leq B_2^*$. Whenever

$$n \geq \Delta_1 \frac{B_2^*}{\lambda} \log \frac{8\Box_1^2 B_2^*}{\lambda \delta}, \quad C_1 \sqrt{\frac{\text{df}_\lambda \vee (\mathbf{Q}^*)^2 \log \frac{2}{\delta}}{n}} \leq r_\lambda(x^*), \quad C_1 b_\lambda \leq r_\lambda(x^*),$$

if

$$C_{\mathbf{P}}(x^*, \lambda) \leq \frac{\sqrt{2}}{480}, \quad \widehat{v}_{\mathbf{P}, \lambda}(x) \leq \frac{\widehat{r}_{\mathbf{P}, \lambda}(x)}{2} \wedge \frac{r_\lambda(x^*)}{126},$$

the following holds, with probability at least $1 - 2\delta$.

$$\widehat{v}_\lambda(x) \leq \frac{\widehat{r}_\lambda(x)}{2}, \quad \widehat{v}_\lambda(x) \leq \frac{r_\lambda(x^*)}{2}.$$

Here, $C_1 = 1008$.

Proof. Proceed in the following way.

1. It is easy to see that the conditions of this lemma imply the conditions of Thm. 8. Hence, as in the previous proofs, the following hold:

- $\|\mathbf{H}_\lambda^{1/2}(x_\lambda^*) \widehat{\mathbf{H}}_\lambda^{-1/2}(x_\lambda^*)\| \leq \sqrt{2}$;
- $\mathfrak{t}(x^* - x_\lambda^*) \leq \log 2$;
- $\mathfrak{t}(\widehat{x}_\lambda^* - x_\lambda^*) \leq \log 2$.

2. Let us now apply Proposition 15 to \widehat{f} . If

$$\frac{\widehat{s}_\lambda}{\widehat{r}_\lambda(\widehat{x}_\lambda^*)} \leq \frac{1}{4}, \quad \widehat{\mathbf{C}}_{\mathbf{P}}(\widehat{x}_\lambda^*, \lambda) \leq \frac{1}{120}, \quad \widehat{\nu}_{\mathbf{P},\lambda}(x) \leq \frac{\widehat{r}_{\mathbf{P},\lambda}(x)}{2},$$

Then it holds:

$$\widehat{\nu}_\lambda(x) \leq 3(\widehat{\nu}_{\mathbf{P},\lambda}(x) + \widehat{s}_\lambda), \quad \widehat{\mathbf{t}}(x - \widehat{x}_\lambda^*) \leq \frac{3}{10} + 2 \log 2. \quad (38)$$

where the second bound is obtained in the proof of this proposition. Now since

$$\begin{aligned} \frac{1}{\widehat{r}_\lambda(\widehat{x}_\lambda^*)} &\leq e^{\widehat{\mathbf{t}}(\widehat{x}_\lambda^* - x_\lambda^*)/2} \frac{1}{\widehat{r}_\lambda(x_\lambda^*)} && \text{Eq. (17)} \\ &\leq e^{\widehat{\mathbf{t}}(\widehat{x}_\lambda^* - x_\lambda^*)/2} \|\mathbf{H}_\lambda^{1/2}(x_\lambda^*) \widehat{\mathbf{H}}_\lambda^{-1/2}(x_\lambda^*)\| \sup_{g \in \widehat{\mathcal{G}}} \|g\|_{\mathbf{H}_\lambda^{-1}(x_\lambda^*)} && \text{Def} \\ &\leq e^{\mathbf{t}(\widehat{x}_\lambda^* - x_\lambda^*)/2} \|\mathbf{H}_\lambda^{1/2}(x_\lambda^*) \widehat{\mathbf{H}}_\lambda^{-1/2}(x_\lambda^*)\| \sup_{g \in \mathcal{G}} \|g\|_{\mathbf{H}_\lambda^{-1}(x_\lambda^*)} && \widehat{\mathcal{G}} \subset \mathcal{G} \\ &= e^{\mathbf{t}(\widehat{x}_\lambda^* - x_\lambda^*)/2} \|\mathbf{H}_\lambda^{1/2}(x_\lambda^*) \widehat{\mathbf{H}}_\lambda^{-1/2}(x_\lambda^*)\| \frac{1}{r_\lambda(x_\lambda^*)} && \text{Def} \\ &\leq e^{(\mathbf{t}(\widehat{x}_\lambda^* - x_\lambda^*) + \mathbf{t}(x_\lambda^* - x^*)) / 2} \|\mathbf{H}_\lambda^{1/2}(x_\lambda^*) \widehat{\mathbf{H}}_\lambda^{-1/2}(x_\lambda^*)\| \frac{1}{r_\lambda(x^*)} && \text{Eq. (17)} \\ &\leq \frac{2\sqrt{2}}{r_\lambda(x^*)}. && \text{previous bounds} \end{aligned}$$

In a similar way, we get $\widehat{\mathbf{C}}_{\mathbf{P}}(\widehat{x}_\lambda^*, \lambda) \leq 2\sqrt{2}\mathbf{C}_{\mathbf{P}}(x^*, \lambda)$. Thus, the conditions above are satisfied if the following conditions are satisfied:

$$\frac{\widehat{s}_\lambda}{r_\lambda(x^*)} \leq \frac{\sqrt{2}}{16}, \quad \mathbf{C}_{\mathbf{P}}(x^*, \lambda) \leq \frac{\sqrt{2}}{480}, \quad \widehat{\nu}_{\mathbf{P},\lambda}(x) \leq \frac{\widehat{r}_{\mathbf{P},\lambda}(x)}{2}.$$

Finally, note that under these conditions,

$$\frac{1}{\widehat{r}_\lambda(x)} \leq \frac{e^{\widehat{\mathbf{t}}(x - \widehat{x}_\lambda^*)/2}}{\widehat{r}_\lambda(x)} \leq \frac{7}{r_\lambda(x^*)}. \quad (39)$$

using the previous bound and the bound on $\widehat{\mathbf{t}}(x - \widehat{x}_\lambda^*)$.

3. Let us assume

$$\frac{\widehat{s}_\lambda}{r_\lambda(x^*)} \leq \frac{\sqrt{2}}{16}, \quad \mathbf{C}_{\mathbf{P}}(x^*, \lambda) \leq \frac{\sqrt{2}}{480}, \quad \widehat{\nu}_{\mathbf{P},\lambda}(x) \leq \frac{\widehat{r}_{\mathbf{P},\lambda}(x)}{2}.$$

According to Eq. (39), and to Eq. (38), if

$$\widehat{\nu}_{\mathbf{P},\lambda}(x) + \widehat{s}_\lambda \leq \frac{r_\lambda(x^*)}{42},$$

then it holds

$$\widehat{\nu}_\lambda(x) \leq \frac{\widehat{r}_\lambda(x)}{2}, \quad \widehat{\nu}_\lambda(x) \leq \frac{r_\lambda(x^*)}{2}.$$

We simplify this condition as:

$$\widehat{\nu}_{\mathbf{P},\lambda}(x) \leq \frac{r_\lambda(x^*)}{126}, \quad \widehat{s}_\lambda \leq \frac{2r_\lambda(x^*)}{126}.$$

4. Now using the fact that under the conditions of this lemma, those of Lemma 18 are satisfied:

$$\widehat{s}_\lambda \leq 8 \mathbf{b}_\lambda + 80 \sqrt{\frac{\text{df}_\lambda \vee (\mathbf{Q}^*)^2 \log \frac{2}{\delta}}{n}}.$$

Thus, $\widehat{s}_\lambda \leq \frac{2r_\lambda(x^*)}{126}$ holds, provided

$$\mathbf{b}_\lambda \leq \frac{r_\lambda(x^*)}{\mathbf{C}_1}, \quad n \geq \mathbf{C}_1^2 \frac{\text{df}_\lambda \vee (\mathbf{Q}^*)^2 \log \frac{2}{\delta}}{r_\lambda(x^*)^2},$$

where $\mathbf{C}_1 = 1008$. □

Proposition 18 (Behavior of an approximation to the projected problem). *Let $n \in \mathbb{N}$, $\delta \in (0, 1/2]$, $0 < \lambda \leq B_2^*$. Let $x \in \mathcal{H}_{\mathbf{P}}$. Whenever*

$$n \geq \Delta_1 \frac{B_2^*}{\lambda} \log \frac{8\Box_1^2 B_2^*}{\lambda \delta}, \quad C_1 \sqrt{\frac{\text{df}_\lambda \vee (Q^*)^2}{n} \log \frac{2}{\delta}} \leq r_\lambda(x^*), \quad C_1 b_\lambda \leq r_\lambda(x^*),$$

if

$$C_{\mathbf{P}}(x^*, \lambda) \leq \frac{\sqrt{2}}{480}, \quad \widehat{\nu}_{\mathbf{P}, \lambda}(x) \leq \frac{\widehat{r}_{\mathbf{P}, \lambda}(x)}{2} \wedge \frac{r_\lambda(x^*)}{126}.$$

The following holds, with probability at least $1 - 2\delta$.

$$f(x) - f(x^*) \leq K_1 b_\lambda^2 + K_2 \frac{\text{df}_\lambda \vee (Q^*)^2}{n} \log \frac{2}{\delta} + K_3 \widehat{\nu}_{\mathbf{P}, \lambda}^2(x),$$

where $K_1 \leq 6.0e4$, $K_2 \leq 6.0e6$ and $K_3 \leq 810$, C_1 are defined in Lemma 19, and the other constants are defined in Thm. 8.

Remark 5 (Constants). *In this result, absolutely huge constants are obtained. They are (of course) totally sub-optimal. Indeed, this analysis has been simplified by dividing the bound into blocks: error of the empirical risk minimization with regularization, error of the projection compared to this empirical risk minimizer. Going back and forth from empirical to statistical, from projected to non projected induces exponential explosion of the constants. There is a way of doing the analysis directly by projecting the statistical problem. However, in order to relate to our previous work [23] and avoid re-doing all of our work we discarded this. If we were to perform this more direct analysis, we could keep the constants to a reasonable level, of order 10^2 .*

Proof. We apply Lemma 17, using the previous lemma to guarantee the conditions.

1. Under the conditions of this proposition, applying Lemma 19, the conditions of Lemma 17 are satisfied. Thus,

$$f(x) - f(x^*) \leq 14(f(\widehat{x}_\lambda^*) - f(x^*)) + 30\widehat{\nu}_\lambda(x)^2.$$

Moreover, from the previous proof,

$$\widehat{\nu}_\lambda(x) \leq 3(\widehat{\nu}_{\mathbf{P}, \lambda}(x) + \widehat{s}_\lambda),$$

and seeing as Lemma 18 is satisfied,

$$\widehat{s}_\lambda \leq 8 b_\lambda + 80 \sqrt{\frac{\text{df}_\lambda \vee (Q^*)^2 \log \frac{2}{\delta}}{n}}.$$

This therefore yields:

$$\widehat{\nu}_\lambda(x)^2 \leq 27\widehat{\nu}_{\mathbf{P}, \lambda}(x)^2 + 1726b_\lambda^2 + 172600 \frac{\text{df}_\lambda \vee (Q^*)^2 \log \frac{2}{\delta}}{n}.$$

2. Moreover, from Thm. 8, it holds:

$$f(\widehat{x}_\lambda^*) - f(x^*) \leq 414 b_\lambda^2 + 414 \frac{\text{df}_\lambda \vee (Q^*)^2}{n} \log \frac{2}{\delta}.$$

3. Putting things together:

$$f(x) - f(x^*) \leq K_1 b_\lambda^2 + K_2 \frac{\text{df}_\lambda \vee (Q^*)^2}{n} \log \frac{2}{\delta} + K_3 \widehat{\nu}_{\mathbf{P}, \lambda}^2(x).$$

We bound the constants in the theorem. □

Lemma 20. *Under the conditions of the previous theorem, the following hold:*

- $\frac{1}{\widehat{r}_{\mathbf{P}, \lambda}(x)} \leq \frac{8}{r_\lambda(x^*)}$;

- $\lambda^{1/2}\|x - x^*\| \leq 7\widehat{\nu}_{\mathbf{P},\lambda}(x) + 59\mathbf{b}_\lambda + 568\sqrt{\frac{\mathbf{d}f_\lambda \vee (\mathbf{Q}^*)^2 \log \frac{2}{\delta}}{n}}$;
- $\lambda\|x\|_{\widehat{\mathbf{H}}_{\mathbf{P},\lambda}^{-1}(x)} \leq 7\widehat{\nu}_{\mathbf{P},\lambda}(x) + 72\mathbf{b}_\lambda + 720\sqrt{\frac{\mathbf{d}f_\lambda \vee (\mathbf{Q}^*)^2 \log \frac{2}{\delta}}{n}}$.

In particular, $\frac{\lambda\|x\|_{\widehat{\mathbf{H}}_{\mathbf{P},\lambda}^{-1}(x)}}{\widehat{r}_{\mathbf{P},\lambda}(x)} \leq 11$.

Proof. Let us prove the three statements.

1. Write $\frac{1}{\widehat{r}_{\mathbf{P},\lambda}(x)} = \sup_{g \in \widehat{\mathcal{G}}} \|\mathbf{P}g\|_{\widehat{\mathbf{H}}_{\mathbf{P},\lambda}^{-1}(x)}$. Now

$$\sup_{g \in \widehat{\mathcal{G}}} \|\mathbf{P}g\|_{\widehat{\mathbf{H}}_{\mathbf{P},\lambda}^{-1}(x)} \leq \sup_{g \in \widehat{\mathcal{G}}} \|g\|_{\widehat{\mathbf{H}}_\lambda^{-1}(x)} \leq e^{\widehat{\mathbf{t}}(x - \widehat{x}_\lambda^*)/2} \sup_{g \in \widehat{\mathcal{G}}} \|g\|_{\widehat{\mathbf{H}}_\lambda^{-1}(\widehat{x}_\lambda^*)}.$$

Now bound

$$\sup_{g \in \widehat{\mathcal{G}}} \|g\|_{\widehat{\mathbf{H}}_\lambda^{-1}(\widehat{x}_\lambda^*)} \leq e^{\widehat{\mathbf{t}}(x_\lambda^* - \widehat{x}_\lambda^*)/2} \sup_{g \in \widehat{\mathcal{G}}} \|g\|_{\widehat{\mathbf{H}}_\lambda^{-1}(x_\lambda^*)} \leq e^{\widehat{\mathbf{t}}(x_\lambda^* - \widehat{x}_\lambda^*)/2} \|\mathbf{H}_\lambda^{1/2}(x_\lambda^*) \widehat{\mathbf{H}}_\lambda^{-1/2}(x_\lambda^*)\| \sup_{g \in \widehat{\mathcal{G}}} \|g\|_{\mathbf{H}_\lambda^{-1}(x_\lambda^*)}.$$

Finally bound

$$\sup_{g \in \widehat{\mathcal{G}}} \|g\|_{\mathbf{H}_\lambda^{-1}(x_\lambda^*)} \leq e^{\mathbf{t}(x^* - x_\lambda^*)/2} \frac{1}{r_\lambda(x^*)}.$$

Now using the fact that under the previous assumptions $\mathbf{t}(x^* - x_\lambda^*), \mathbf{t}(x_\lambda^* - \widehat{x}_\lambda^*) \leq \log 2$, $\widehat{\mathbf{t}}(x - \widehat{x}_\lambda^*) \leq \frac{3}{10} + 2 \log 2$ and $\|\mathbf{H}_\lambda^{1/2}(x_\lambda^*) \widehat{\mathbf{H}}_\lambda^{-1/2}(x_\lambda^*)\| \leq \sqrt{2}$, we get the first equation.

2. In order to bound $\lambda^{1/2}\|x - x^*\|$, decompose

$$\lambda^{1/2}\|x - x^*\| \leq \lambda^{1/2}\|x - \widehat{x}_\lambda^*\| + \lambda^{1/2}\|\widehat{x}_\lambda^* - x^*\|.$$

Now use Proposition 15 to bound $\lambda^{1/2}\|x - \widehat{x}_\lambda^*\| \leq 7(\widehat{\nu}_{\mathbf{P},\lambda}(x) + \widehat{s}_\lambda)$. Using Lemma 18, under the conditions above,

$$\widehat{s}_\lambda \leq 8\mathbf{b}_\lambda + 80\sqrt{\frac{\mathbf{d}f_\lambda \vee (\mathbf{Q}^*)^2 \log \frac{2}{\delta}}{n}}.$$

Hence

$$\lambda^{1/2}\|x - \widehat{x}_\lambda^*\| \leq 7\widehat{\nu}_{\mathbf{P},\lambda}(x) + 56\mathbf{b}_\lambda + 560\sqrt{\frac{\mathbf{d}f_\lambda \vee (\mathbf{Q}^*)^2 \log \frac{2}{\delta}}{n}}.$$

Moreover, using again Lemma 18

$$\lambda^{1/2}\|\widehat{x}_\lambda^* - x^*\| \leq 3\mathbf{b}_\lambda + 8\sqrt{\frac{\mathbf{d}f_\lambda \vee (\mathbf{Q}^*)^2 \log \frac{2}{\delta}}{n}}.$$

Combining these two inequalities, we get:

$$\lambda^{1/2}\|x - x^*\| \leq 7\widehat{\nu}_{\mathbf{P},\lambda}(x) + 59\mathbf{b}_\lambda + 568\sqrt{\frac{\mathbf{d}f_\lambda \vee (\mathbf{Q}^*)^2 \log \frac{2}{\delta}}{n}}.$$

3. In order to bound $\lambda \|x\|_{\widehat{\mathbf{H}}_{\mathbf{P},\lambda}^{-1}(x)}$, use Proposition 15 to get $\lambda \|x\|_{\widehat{\mathbf{H}}_{\mathbf{P},\lambda}^{-1}(x)} \leq 7\widehat{\nu}_{\mathbf{P},\lambda}(x) + 9\widehat{s}_\lambda$.

Now using Lemma 18, the following bound holds:

$$\lambda \|x\|_{\widehat{\mathbf{H}}_{\mathbf{P},\lambda}^{-1}(x)} \leq 7\widehat{\nu}_{\mathbf{P},\lambda}(x) + 72b_\lambda + 720\sqrt{\frac{\text{df}_\lambda \vee (\mathbf{Q}^*)^2 \log \frac{2}{\delta}}{n}}.$$

□

Proposition 19 (Simplification). *Let $n \in \mathbb{N}$, $\delta \in (0, 1/2]$, $0 < \lambda \leq \mathbf{B}_2^*$. Let $x \in \mathcal{H}_{\mathbf{P}}$. Whenever*

$$n \geq \Delta_1 \frac{\mathbf{B}_2^*}{\lambda} \log \frac{8\Box_1^2 \mathbf{B}_2^*}{\lambda \delta}, \quad \mathbf{C}_1 \sqrt{\frac{\text{df}_\lambda \vee (\mathbf{Q}^*)^2}{n} \log \frac{2}{\delta}} \leq \frac{\sqrt{\lambda}}{R}, \quad \mathbf{C}_1 b_\lambda \leq \frac{\sqrt{\lambda}}{R},$$

if

$$\mathbf{C}_{\mathbf{P}}(x^*, \lambda) \leq \frac{\sqrt{2}}{480}, \quad \widehat{\nu}_{\mathbf{P},\lambda}(x) \leq \frac{\sqrt{\lambda}}{126R},$$

then the following holds, with probability at least $1 - 2\delta$.

$$f(x) - f(x^*) \leq \mathbf{K}_1 b_\lambda^2 + \mathbf{K}_2 \frac{\text{df}_\lambda \vee (\mathbf{Q}^*)^2}{n} \log \frac{2}{\delta} + \mathbf{K}_3 \widehat{\nu}_{\mathbf{P},\lambda}^2(x),$$

where $\mathbf{K}_1 \leq 6.0e4$, $\mathbf{K}_2 \leq 6.0e6$ and $\mathbf{K}_3 \leq 810$, \mathbf{C}_1 are defined in Lemma 19, and the other constants are defined in Thm. 8.

Moreover, in that case, $R\|x - x^*\| \leq 10$.

H.4 Optimal choice of λ , specific source conditions

In this part, we continue to assume Assumptions 6 to 9. We present a classification of distributions ρ and show that we can achieve better rates than the classical slow rates, as presented in Appendix F of [23].

H.4.1 Classification of distributions and statistical bounds for the ERM

We use the following classification for distributions.

Definition 9 (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 $L, Q \geq 0$,

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

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

Note that given our assumptions, we always have

$$\rho \in \mathcal{P}_{1,0}, \quad L = \|x^*\|, \quad Q = \mathbf{B}_1^*. \quad (40)$$

We also define

$$\lambda_1 = \left(\frac{Q}{\mathbf{Q}^*} \right)^{2\alpha} \wedge 1, \quad (41)$$

such that

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

Interpretation of the classes

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

Remark 6 (source condition). *Assume there exists $0 \leq r \leq 1/2$ and $v \in \mathcal{H}$ such that*

$$\mathbf{P}_{\mathbf{H}(x^*)} x^* = \mathbf{H}(x^*)^r v.$$

Then it holds:

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

- 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, $\alpha = +\infty$.

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 [23] (see corollary 3), explicit bounds are given for the performance of the regularized expected risk minimizer \hat{x}_λ^* depending on which class ρ belongs to, i.e., as a function of α, r .

Corollary 4. *Let $\delta \in (0, 1/2]$. Under Assumptions 6 to 9, if $\rho \in \mathcal{P}_{\alpha, r}$ with $r > 0$, when $n \geq N$ and $\lambda = (C_0/n)^\beta$, then with probability at least $1 - 2\delta$,*

$$f(\hat{x}_\lambda^*) - f(x^*) \leq C_1 n^{-\gamma} \log \frac{2}{\delta},$$

with $C_0 = 256(Q/L)^2$, $C_1 = 8(256)^\gamma (Q^\gamma L^{1-\gamma})^2$ and N defined in [23], and satisfying $N = O(\text{poly}(B_1^, B_2^*, L, Q, R, \log(1/\delta)))$.*

H.4.2 Quantitative bounds for the projected problem

In this part, the aim is to show that if we approximately solve the projected problem up to a certain precision, then this approximation has the same statistical rates as the regularized ERM with the good choice of λ . For the sake of simplicity, we will assume that $r > 0$.

In what follows, we define

$$N = \frac{Q^2}{L^2} (B_2^* \wedge \lambda_0 \wedge \lambda_1)^{-1/\beta} \vee \left(2.1e4 \frac{1}{1-\beta} A \log \left(1.4e6 \frac{1}{1-\beta} A^2 \frac{1}{\delta} \right) \right)^{1/(1-\beta)}, \quad (42)$$

where $A = \frac{B_2^* L^{2\beta}}{Q^{2\beta}}$, $\lambda_0 = (C_1 L R \log \frac{2}{\delta})^{-1/r} \wedge 1$ and $\lambda_1 = \frac{Q^{2\alpha}}{(Q^*)^{2\alpha}}$.

Theorem 9 (Quantitative result with source $r > 0$). *Let $\rho \in \mathcal{P}_{\alpha, r}$ and assume $r > 0$. Let $\delta \in (0, \frac{1}{2}]$. Let \mathbf{P} be an orthogonal projection, $x \in \mathcal{H}$. If*

$$n \geq N, \quad \lambda = \left(\left(\frac{Q}{L} \right)^2 \frac{1}{n} \right)^\beta, \quad \mathbf{C}_{\mathbf{P}}(x^*, \lambda) \leq \frac{\sqrt{2}}{480}, \quad \hat{v}_{\mathbf{P}, \lambda}(x) \leq Q^\gamma L^{1-\gamma} n^{-\gamma/2}$$

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

$$f(x) - f(x^*) \leq K (Q^\gamma L^{1-\gamma})^2 \frac{1}{n^\gamma} \log \frac{2}{\delta},$$

where N is defined in Eq. (42) and $K \leq 7.0e6$. Moreover, $R\|x - x^\| \leq 10$.*

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

Let us formulate Proposition 19 using the fact that $\rho \in \mathcal{P}_{\alpha, r}$.

Let $n \in \mathbb{N}$, $\delta \in (0, 1/2]$, $0 < \lambda \leq B_2^*$, $x \in \mathcal{H}_{\mathbf{P}}$. Whenever

$$n \geq \Delta_1 \frac{B_2^*}{\lambda} \log \frac{8\Box_1^2 B_2^*}{\lambda \delta}, \quad C_1 \sqrt{\frac{Q^2}{\lambda^{1/\alpha n}} \log \frac{2}{\delta}} \leq \frac{\lambda^{1/2}}{R}, \quad C_1 L \lambda^{1/2+r} \leq \frac{\lambda^{1/2}}{R},$$

if

$$C_{\mathbf{P}}(x^*, \lambda) \leq \frac{\sqrt{2}}{480}, \quad \widehat{v}_{\mathbf{P}, \lambda}(x) \leq L \lambda^{1/2+r},$$

The following holds, with probability at least $1 - 2\delta$.

$$f(x) - f(x^*) \leq (K_1 + K_3) L^2 \lambda^{1+2r} + K_2 \frac{Q^2}{\lambda^{1/\alpha n}} \log \frac{2}{\delta}, \quad R \|x - x^*\| \leq 10,$$

where all constants are defined in Proposition 19.

Assume that $r > 0$. Define

$$\lambda_0 = (C_1 L R \log \frac{2}{\delta})^{-1/r} \wedge 1.$$

Then for any $\lambda \leq \lambda_0$:

$$L \lambda^{1/2+r} \leq \frac{1}{C_1} \frac{\sqrt{\lambda}}{R}.$$

1) First, we find a simple condition to guarantee

$$r_\lambda(x^*)^2 \lambda^{1/\alpha} \geq C_2 Q^2 \frac{1}{n} \log \frac{2}{\delta}.$$

We see that if $\lambda \leq \lambda_0$, then $r_\lambda \geq C_1 L \lambda^{1/2+r} \log \frac{2}{\delta}$. Hence, this condition is satisfied if

$$\lambda \leq \lambda_0, \quad C_1^2 L^2 \lambda^{1+2r+1/\alpha} \geq C_2 Q^2 \frac{1}{n}.$$

Using the fact that $C_2 = C_1^2$, we reformulate:

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

2) Now fix

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

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

Using our restatement of Proposition 18, with probability at least $1 - 2\delta$,

$$L(x) - L(x^*) \leq \left(K_1 + K_3 + K_2 \log \frac{2}{\delta} \right) L^2 \lambda^{1+2r} \leq K \log \frac{2}{\delta} L^2 \lambda^{1+2r},$$

where $K = K_1 + K_3 + K_2 \leq 7.0e6$ (see Proposition 18).

This result holds provided

$$0 < \lambda \leq B_2^* \wedge \lambda_0 \wedge \lambda_1, \quad n \geq \Delta_1 \frac{B_2^*}{\lambda} \log \frac{8\Box_1^2 B_2^*}{\lambda \delta}. \quad (43)$$

Indeed, it is shown in the previous point that the other conditions are satisfied.

3) Let us now work to guarantee the conditions in Eq. (43).

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

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

Then apply lemma 15 from [23] with $a_1 = 2\Delta_1$, $a_2 = 16\Box_1^2$, $A = \frac{B_2^* L^{2\beta}}{Q^{2\beta}}$. Since $\beta \geq 1/2$, using the bounds in Thm. 8, we find $a_1 \leq 10400$ and $a_2 \leq 64$, hence the following sufficient condition:

$$n \geq \left(2.1e4 \frac{1}{1-\beta} A \log \left(1.4e6 \frac{1}{1-\beta} A^2 \frac{1}{\delta} \right) \right)^{1/(1-\beta)}.$$

Then, to guarantee the condition

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

we simply need

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

Hence, defining

$$N = \frac{Q^2}{L^2} (B_2^* \wedge \lambda_0 \wedge \lambda_1)^{-1/\beta} \vee \left(2.1e4 \frac{1}{1-\beta} A \log \left(1.4e6 \frac{1}{1-\beta} A^2 \frac{1}{\delta} \right) \right)^{1/(1-\beta)},$$

we see that as soon as $n \geq N$, Eq. (43) holds.

□

I Multiplicative approximations for Hermitian operators

In this section, we put together useful tools for approximating linear operators and solving linear systems with regularization.

In this section, \mathbf{A} and \mathbf{B} will always denote positive semi-definite Hermitian operators on a Hilbert space \mathcal{H} , and \mathbf{P} will denote an orthogonal projection operator. Moreover, given a positive semi-definite operator \mathbf{A} , and $\lambda > 0$, \mathbf{A}_λ will stand for the regularized operator $\mathbf{A} + \lambda\mathbf{I}$.

Lemma 21 (Equivalence of Hermitian operators). *Let \mathbf{A} and \mathbf{B} be two semi-definite Hermitian operators. Let $\lambda > 0$. Assume you have access to*

$$t := \|\mathbf{A}_\lambda^{-1/2}(\mathbf{B} - \mathbf{A})\mathbf{A}_\lambda^{-1/2}\|.$$

It holds:

$$\|\mathbf{A}_\lambda^{-1/2}\mathbf{B}_\lambda^{1/2}\|^2 \leq 1 + t \Leftrightarrow \mathbf{B}_\lambda \preceq (1 + t)\mathbf{A}_\lambda.$$

Moreover, if $t < 1$,

$$\|\mathbf{B}_\lambda^{-1/2}\mathbf{A}_\lambda^{1/2}\|^2 \leq \frac{1}{1-t} \Leftrightarrow (1-t)\mathbf{A}_\lambda \preceq \mathbf{B}_\lambda.$$

Proof. For the first point, simply note that:

$$\|\mathbf{A}_\lambda^{-1/2}\mathbf{B}_\lambda^{1/2}\|^2 = \|\mathbf{A}_\lambda^{-1/2}\mathbf{B}_\lambda\mathbf{A}_\lambda^{-1/2}\| = \|\mathbf{I} + \mathbf{A}_\lambda^{-1/2}(\mathbf{B} - \mathbf{A})\mathbf{A}_\lambda^{-1/2}\| \leq 1 + t.$$

For the second point,

$$\|\mathbf{B}_\lambda^{-1/2}\mathbf{A}_\lambda^{1/2}\|^2 = \|\left(\mathbf{A}_\lambda^{-1/2}\mathbf{B}_\lambda\mathbf{A}_\lambda^{-1/2}\right)^{-1}\| = \|\left(\mathbf{I} + \mathbf{A}_\lambda^{-1/2}(\mathbf{B} - \mathbf{A})\mathbf{A}_\lambda^{-1/2}\right)^{-1}\|.$$

Moreover, we know that if $\|\mathbf{H}\| < 1$ with \mathbf{H} a Hermitian operator, then $\|(\mathbf{I} + \mathbf{H})^{-1}\| \leq \frac{1}{1 - \|\mathbf{H}\|}$. The result follows. \square

We will now state a technical lemma which describes how combining approximation behaves.

Lemma 22 (Combination of approximations). *Let $N \geq 1$. Let $(\mathbf{A}_i)_{1 \leq i \leq N+1}$ be a sequence of positive semi-definite Hermitian operators. Define*

$$t_i := \|\mathbf{A}_{i,\lambda}^{-1/2}(\mathbf{A}_{i+1} - \mathbf{A}_i)\mathbf{A}_{i,\lambda}^{-1/2}\|.$$

For any $1 \leq i, j \leq N + 1$, define

$$t_{i:j} := \|\mathbf{A}_{i,\lambda}^{-1/2}(\mathbf{A}_j - \mathbf{A}_i)\mathbf{A}_{i,\lambda}^{-1/2}\|.$$

In particular, $t_i = t_{i:i+1}$. Then the following holds:

$$\forall 1 \leq i \leq j \leq N, 1 + t_{i:j} \leq \prod_{k=i}^{j-1} (1 + t_k)$$

Moreover, if $t_i < 1$, then it holds:

$$\|\mathbf{A}_{i+1,\lambda}^{-1/2}(\mathbf{A}_i - \mathbf{A}_{i+1})\mathbf{A}_{i+1,\lambda}^{-1/2}\| \leq \frac{t_i}{1 - t_i}$$

Hence, in that case

$$\forall 1 \leq j \leq i \leq N, 1 + t_{j:i} \leq \prod_{k=i}^{j-1} \frac{1}{1 - t_k}$$

Proof. Let us prove everything for a sequence of three operators; the rest follows by induction. Let $\mathbf{A}_1, \mathbf{A}_2, \mathbf{A}_3$ be three positive semi-definite operators.

1. Bound

$$\begin{aligned}
t_{1:3} &= \|\mathbf{A}_{1,\lambda}^{-1/2} (\mathbf{A}_1 - \mathbf{A}_3) \mathbf{A}_{1,\lambda}^{-1/2}\| \\
&\leq \|\mathbf{A}_{1,\lambda}^{-1/2} (\mathbf{A}_1 - \mathbf{A}_2) \mathbf{A}_{1,\lambda}^{-1/2}\| + \|\mathbf{A}_{1,\lambda}^{-1/2} (\mathbf{A}_2 - \mathbf{A}_3) \mathbf{A}_{1,\lambda}^{-1/2}\| \\
&\leq t_{1:2} + \|\mathbf{A}_{1,\lambda}^{-1/2} \mathbf{A}_{2,\lambda}^{1/2}\|^2 t_{2:3} \\
&\leq t_{1:2} + (1 + t_{1:2}) t_{2:3}.
\end{aligned}$$

The last line comes from Lemma 21. Thus

$$1 + t_{1:3} \leq 1 + t_{1:2} + t_{2:3} + t_{1:2} t_{2:3} = (1 + t_{1:2})(1 + t_{2:3}).$$

2. Let us now bound $t_{2:1}$ knowing $t_{1:2}$. This will imply the rest of the lemma.

Indeed, simply note that

$$t_{2:1} = \|\mathbf{A}_{2,\lambda}^{-1/2} (\mathbf{A}_2 - \mathbf{A}_1) \mathbf{A}_{2,\lambda}^{-1/2}\| \leq \|\mathbf{A}_{2,\lambda}^{-1/2} \mathbf{A}_{1,\lambda}^{1/2}\|^2 t_{1:2}.$$

Using Lemma 21, if $t_{1:2} < 1$, $\|\mathbf{A}_{2,\lambda}^{-1/2} \mathbf{A}_{1,\lambda}^{1/2}\|^2 \leq \frac{1}{1-t_{1:2}}$, hence

$$t_{2:1} \leq \frac{t_{1:2}}{1-t_{1:2}}.$$

□

Lemma 23 (Projection of Hermitian operators). *For any Hermitian operator \mathbf{A} and orthogonal projection \mathbf{P} , the following holds:*

$$\|\mathbf{A}_\lambda^{-1/2} (\mathbf{A} - \mathbf{PAP}) \mathbf{A}_\lambda^{-1/2}\| \leq \left(1 + \frac{\|\mathbf{A}^{1/2}(\mathbf{I} - \mathbf{P})\|}{\sqrt{\lambda}}\right)^2 - 1.$$

In particular,

$$\|\mathbf{A}_\lambda^{-1/2} (\mathbf{PAP} + \lambda\mathbf{I})^{1/2}\| \leq 1 + \frac{\|\mathbf{A}^{1/2}(\mathbf{I} - \mathbf{P})\|}{\sqrt{\lambda}}.$$

Moreover, if

$$\frac{\|\mathbf{A}^{1/2}(\mathbf{I} - \mathbf{P})\|}{\sqrt{\lambda}} < \sqrt{2} - 1,$$

then it holds

$$\|\mathbf{A}_\lambda^{1/2} (\mathbf{PAP} + \lambda\mathbf{I})^{-1/2}\|^2 \leq \frac{1}{2 - \left(1 + \frac{\|\mathbf{A}^{1/2}(\mathbf{I} - \mathbf{P})\|}{\sqrt{\lambda}}\right)^2}.$$

We also always have:

$$\|(\mathbf{PAP} + \lambda\mathbf{I})^{-1/2} \mathbf{PA}_\lambda^{1/2}\|^2 \leq 1.$$

Proof. For any Hermitian operator \mathbf{A} , the following computation holds:

$$\begin{aligned}
\|\mathbf{A}_\lambda^{-1/2} (\mathbf{A} - \mathbf{PAP}) \mathbf{A}_\lambda^{-1/2}\| &= \|\mathbf{A}_\lambda^{-1/2} (\mathbf{A} - (\mathbf{I} - (\mathbf{I} - \mathbf{P}))\mathbf{A}(\mathbf{I} - (\mathbf{I} - \mathbf{P}))) \mathbf{A}_\lambda^{-1/2}\| \\
&\leq 2\|\mathbf{A}_\lambda^{-1/2} (\mathbf{I} - \mathbf{P})\mathbf{A} \mathbf{A}_\lambda^{-1/2}\| + \|\mathbf{A}_\lambda^{-1/2} (\mathbf{I} - \mathbf{P})\mathbf{A}(\mathbf{I} - \mathbf{P}) \mathbf{A}_\lambda^{-1/2}\| \\
&\leq \frac{2\|\mathbf{A}^{1/2}(\mathbf{I} - \mathbf{P})\|}{\sqrt{\lambda}} + \frac{\|\mathbf{A}^{1/2}(\mathbf{I} - \mathbf{P})\|^2}{\lambda} \\
&= \left(1 + \frac{\|\mathbf{A}^{1/2}(\mathbf{I} - \mathbf{P})\|}{\sqrt{\lambda}}\right)^2 - 1.
\end{aligned}$$

□

Lemma 24 (Relationship between approximations). *Let \mathbf{A} and \mathbf{B} be two positive semi-definite hermitian operators. Let $\lambda > 0$, $b \in \mathcal{H}$ and $\rho > 0$. If*

$$\|\mathbf{A}_\lambda^{-1/2} (\mathbf{B} - \mathbf{A}) \mathbf{A}_\lambda^{-1/2}\| \leq \frac{1}{2} \wedge \frac{\rho}{4}, \quad \tilde{\Delta} \in \text{LinApprox}(\mathbf{B}_\lambda, b, \rho/4),$$

then it holds:

$$\tilde{\Delta} \in \text{LinApprox}(\mathbf{A}_\lambda, b, \rho).$$

Proof. Assume $\tilde{\Delta} \in \text{LinApprox}(\mathbf{B}_\lambda, b, \tilde{\rho}/4)$ for a certain $\tilde{\rho}$. Decompose

$$\begin{aligned} \|\mathbf{A}_\lambda^{-1}b - \tilde{\Delta}\|_{\mathbf{A}_\lambda} &\leq \|\mathbf{A}_\lambda^{-1}b - \mathbf{B}_\lambda^{-1}b\|_{\mathbf{A}_\lambda} + \|\mathbf{B}_\lambda^{-1}b - \tilde{\Delta}\|_{\mathbf{A}_\lambda} \\ &\leq \|\mathbf{A}_\lambda^{1/2}(\mathbf{A}_\lambda^{-1} - \mathbf{B}_\lambda^{-1})\mathbf{A}_\lambda^{1/2}\| \|b\|_{\mathbf{A}_\lambda^{-1}} + \|\mathbf{A}_\lambda^{1/2}\mathbf{B}_\lambda^{-1/2}\| \|\mathbf{B}_\lambda^{-1}b - \tilde{\Delta}\|_{\mathbf{B}_\lambda}. \end{aligned}$$

Now using the fact that $\mathbf{A}_\lambda^{-1} - \mathbf{B}_\lambda^{-1} = \mathbf{B}_\lambda^{-1}(\mathbf{B} - \mathbf{A})\mathbf{A}_\lambda^{-1}$,

$$\begin{aligned} \|\mathbf{A}_\lambda^{1/2}(\mathbf{A}_\lambda^{-1} - \mathbf{B}_\lambda^{-1})\mathbf{A}_\lambda^{1/2}\| &\leq \|\mathbf{A}_\lambda^{-1/2}(\mathbf{B} - \mathbf{A})\mathbf{A}_\lambda^{-1/2}\| \|\mathbf{A}_\lambda^{1/2}\mathbf{B}_\lambda^{-1}\mathbf{A}_\lambda^{1/2}\| \\ &= \|\mathbf{A}_\lambda^{-1/2}(\mathbf{B} - \mathbf{A})\mathbf{A}_\lambda^{-1/2}\| \|\mathbf{A}_\lambda^{1/2}\mathbf{B}_\lambda^{-1/2}\|^2. \end{aligned}$$

Moreover,

$$\|\mathbf{B}_\lambda^{-1}b - \tilde{\Delta}\|_{\mathbf{B}_\lambda} \leq \tilde{\rho}\|b\|_{\mathbf{B}_\lambda^{-1}} \leq \|\mathbf{A}^{1/2}\mathbf{B}^{-1/2}\| \|b\|_{\mathbf{A}_\lambda^{-1}}.$$

Putting things together, and noting that from Lemma 21, $\|\mathbf{A}^{1/2}\mathbf{B}^{-1/2}\|^2 \leq \frac{1}{1 - \|\mathbf{A}_\lambda^{-1/2}(\mathbf{B} - \mathbf{A})\mathbf{A}_\lambda^{-1/2}\|}$

as soon as $\|\mathbf{A}_\lambda^{-1/2}(\mathbf{B} - \mathbf{A})\mathbf{A}_\lambda^{-1/2}\| < 1$, it holds:

$$\tilde{\Delta} \in \text{LinApprox}(\mathbf{A}_\lambda, b, \rho), \quad \rho = \frac{\tilde{\rho} + \|\mathbf{A}_\lambda^{-1/2}(\mathbf{B} - \mathbf{A})\mathbf{A}_\lambda^{-1/2}\|}{1 - \|\mathbf{A}_\lambda^{-1/2}(\mathbf{B} - \mathbf{A})\mathbf{A}_\lambda^{-1/2}\|}.$$

Choosing the right values for $\tilde{\rho}$ and $\|\mathbf{A}_\lambda^{-1/2}(\mathbf{B} - \mathbf{A})\mathbf{A}_\lambda^{-1/2}\|$ yields the result. \square

I.1 Results for Nystrom sub-sampling

Recall the notations from Appendix D.

We write without proof the following lemmas, which are just restatements of lemmas 9 and 10 of [29].

Lemma 25 (Uniform sampling). *Let $\delta > 0$. If $\{\tilde{z}_1, \dots, \tilde{z}_m\}$ are sampled uniformly, then if $0 < \lambda < \|\mathbf{A}\|$, $m \leq n$ and*

$$m \geq (10 + 160\mathcal{N}_\infty^{\mathbf{A}}(\lambda)) \log \frac{8\|v\|_{L^\infty(Z)}^2}{\lambda\delta}.$$

Then it holds, with probability at least $1 - \delta$:

$$\|\mathbf{A}_\lambda^{-1/2}(\hat{\mathbf{A}} - \mathbf{A})\mathbf{A}_\lambda^{-1/2}\| \leq \frac{1}{2}, \quad \|\hat{\mathbf{A}}_{m,\lambda}^{-1/2}(\hat{\mathbf{A}} - \hat{\mathbf{A}}_m)\hat{\mathbf{A}}_{m,\lambda}^{-1/2}\| \leq \frac{1}{2}.$$

Lemma 26 (Nystrom sampling). *Let $\delta > 0$. If $\{\tilde{z}_1, \dots, \tilde{z}_m\}$ are sampled using q -approximate leverage scores for $t = \lambda$, then if $t_0 \vee \frac{19\|v\|_{L^\infty(Z)}^2}{n} \log \frac{n}{2\delta} < \lambda < \|\mathbf{A}\|$, and $n \geq 405\|v\|_{L^\infty(Z)}^2 \vee 67\|v\|_{L^\infty(Z)}^2 \log \frac{12\|v\|_{L^\infty(Z)}^2}{\delta}$, if*

$$m \geq (6 + 486q^2\mathcal{N}^{\mathbf{A}}(\lambda)) \log \frac{8\|v\|_{L^\infty(Z)}^2}{\lambda\delta}.$$

Then it holds, with probability at least $1 - \delta$:

$$\|\mathbf{A}_\lambda^{-1/2}(\hat{\mathbf{A}} - \mathbf{A})\mathbf{A}_\lambda^{-1/2}\| \leq \frac{1}{2}, \quad \|\hat{\mathbf{A}}_{m,\lambda}^{-1/2}(\hat{\mathbf{A}} - \hat{\mathbf{A}}_m)\hat{\mathbf{A}}_{m,\lambda}^{-1/2}\| \leq \frac{1}{2}.$$

Lemma 27. *Let $\lambda > 0$. Assume:*

$$\|\mathbf{A}_\lambda^{-1/2}(\hat{\mathbf{A}} - \mathbf{A})\mathbf{A}_\lambda^{-1/2}\| \leq \frac{1}{2}, \quad \|\hat{\mathbf{A}}_{m,\lambda}^{-1/2}(\hat{\mathbf{A}} - \hat{\mathbf{A}}_m)\hat{\mathbf{A}}_{m,\lambda}^{-1/2}\| \leq \frac{1}{2}.$$

Denote with P_m the projection on $\text{span}(v_{\tilde{z}_j})_{1 \leq j \leq m}$. Then the following holds:

$$\|\mathbf{A}_\lambda^{1/2}(\mathbf{I} - P_m)\|^2 \leq 3\lambda,$$

and for any partial isometry V ,

$$\frac{1}{2} \left(V^* \hat{\mathbf{A}}_m V + \lambda \mathbf{I} \right) \preceq V^* \hat{\mathbf{A}} V + \lambda \mathbf{I} \preceq \frac{3}{2} \left(V^* \hat{\mathbf{A}}_m V + \lambda \mathbf{I} \right).$$

Proof. For the first point, use the well known fact that

$$\mathbf{I} - \mathbf{P}_m \leq \lambda \widehat{\mathbf{A}}_{m,\lambda}^{-1},$$

since the range of \mathbf{P}_m contains that of $\widehat{\mathbf{A}}_m$. Thus,

$$\|\mathbf{A}_\lambda^{1/2}(\mathbf{I} - \mathbf{P}_m)\|^2 \leq \lambda \|\mathbf{A}_\lambda^{1/2} \widehat{\mathbf{A}}_{m,\lambda}^{-1/2}\|^2.$$

Now using Lemma 22,

$$\|\mathbf{A}_\lambda^{-1/2}(\widehat{\mathbf{A}} - \mathbf{A})\mathbf{A}_\lambda^{-1/2}\| \leq \frac{1}{2} \implies \|\widehat{\mathbf{A}}_\lambda^{-1/2}(\widehat{\mathbf{A}} - \mathbf{A})\widehat{\mathbf{A}}_\lambda^{-1/2}\| \leq 1.$$

Hence, again using Lemma 22,

$$\|\widehat{\mathbf{A}}_{m,\lambda}^{-1/2}(\widehat{\mathbf{A}}_m - \mathbf{A})\widehat{\mathbf{A}}_{m,\lambda}^{-1/2}\| \leq 2,$$

and therefore, using Lemma 21,

$$\|\mathbf{A}_\lambda^{1/2} \widehat{\mathbf{A}}_{m,\lambda}^{-1/2}\|^2 \leq 3.$$

For the second point, this is only a consequence of Lemma 21. \square

Now state two results which show that

Lemma 28 (Uniform sampling yielding ρ -approximation). *Let $0 < \rho \leq 1$ and $\delta > 0$. Let $b \in \mathcal{H}$. If $\{\tilde{z}_1, \dots, \tilde{z}_m\}$ are sampled uniformly, $0 < \lambda < \|\mathbf{A}\|$, $m \leq n$ and*

$$m \geq \left(2 + \frac{48}{\rho} + \frac{5000}{\rho^2} \mathcal{N}_\infty^{\mathbf{A}}(\lambda)\right) \log \frac{8\|v\|_{L^\infty(Z)}^2}{\lambda\delta}.$$

Then it holds, with probability at least $1 - \delta$:

$$x \in \text{LinApprox}(\widehat{\mathbf{A}}_{m,\lambda}, b, \rho/4) \implies x \in \text{LinApprox}(\mathbf{A}_\lambda, b, \rho).$$

In particular, with probability $1 - \delta$,

$$\widehat{\mathbf{A}}_{m,\lambda}^{-1} b \in \text{LinApprox}(\mathbf{A}_\lambda, b, \rho).$$

Proof. Apply Lemma 9 from [29] with $\eta = \frac{\rho}{12} < \frac{1}{2}$. We find that under the conditions above, with probability at least $1 - \delta$,

$$\|\mathbf{A}_\lambda^{-1/2}(\widehat{\mathbf{A}} - \mathbf{A})\mathbf{A}_\lambda^{-1/2}\| \leq \eta, \quad \|\widehat{\mathbf{A}}_{m,\lambda}^{-1/2}(\widehat{\mathbf{A}} - \widehat{\mathbf{A}}_m)\widehat{\mathbf{A}}_{m,\lambda}^{-1/2}\| \leq \eta.$$

Now use Lemma 22 to see that

$$\|\mathbf{A}_\lambda^{-1/2}(\widehat{\mathbf{A}}_m - \mathbf{A})\mathbf{A}_\lambda^{-1/2}\| \leq (1 + \eta^2) - 1 \leq 3\eta \leq \rho/4.$$

Thus, we can apply Lemma 24 to get the desired result. \square

Lemma 29 (Leverage scores Nystrom sampling yielding ρ -approximation). *Let $\delta > 0$. If $\{\tilde{z}_1, \dots, \tilde{z}_m\}$ are sampled using q -approximate leverage scores for $t = \lambda$, then if $t_0 \vee \frac{19\|v\|_{L^\infty(Z)}^2}{n} \log \frac{n}{2\delta} < \lambda < \|\mathbf{A}\|$, and $n \geq 405\|v\|_{L^\infty(Z)}^2 \vee 67\|v\|_{L^\infty(Z)}^2 \log \frac{12\|v\|_{L^\infty(Z)}^2}{\delta}$, if*

$$m \geq \left(2 + \frac{24}{\rho} + \frac{13000q^2}{\rho^2} \mathcal{N}_\infty^{\mathbf{A}}(\lambda)\right) \log \frac{8\|v\|_{L^\infty(Z)}^2}{\lambda\delta}.$$

Then it holds, with probability at least $1 - \delta$:

$$x \in \text{LinApprox}(\widehat{\mathbf{A}}_{m,\lambda}, b, \rho/4) \implies x \in \text{LinApprox}(\mathbf{A}_\lambda, b, \rho).$$

In particular, with probability $1 - \delta$,

$$\widehat{\mathbf{A}}_{m,\lambda}^{-1} b \in \text{LinApprox}(\mathbf{A}_\lambda, b, \rho).$$

Proof. The proof is exactly the same as that of the previous lemma, using Lemma 10 instead of Lemma 9 in [29]. \square