

# Robust Accelerated Gradient Method

N.S. Aybat<sup>\*†</sup>, A. Fallah<sup>\*‡</sup>, M. Gürbüzbalaban<sup>\*§</sup>, A. Ozdaglar<sup>\*¶</sup>

May 29, 2018

## Abstract

We study the trade-off between rate of convergence and robustness to gradient errors in designing a first-order algorithm. In particular, we focus on gradient descent (GD) and Nesterov’s accelerated gradient (AG) method for strongly convex quadratic objectives when the gradient has random errors in the form of additive white noise. To characterize robustness, we consider the asymptotic normalized variance of the centered iterate sequence which measures the asymptotic accuracy of the iterates. Using tools from robust control theory, we develop a tractable algorithm that allows us to set the parameters of each algorithm to achieve a particular trade-off between these two performance objectives. Our results show that there is a fundamental lower bound on the robustness level of an algorithm for any achievable rate. For the same achievable rate, we show that AG with tuned parameters is always more robust than GD to gradient errors. Similarly, for the same robustness level, we show that AG can be tuned to be always faster than GD. Our results show that AG can achieve acceleration while being more robust to random gradient errors. This behavior is quite different than previously reported in the deterministic gradient noise setting.

## 1 Introduction

For many large-scale convex optimization and machine learning problems that involve large amounts of data, first-order methods have been the leading computational approach because of their cheap iterations and mild dependence on the problem dimension and data size. The typical analysis of first-order methods assumes the availability of exact gradient information and focuses on the rate of convergence to the optimal solution as the main performance criterion. However, in many applications, the gradient contains deterministic or stochastic errors either because the gradient is computed by solving an auxiliary optimization problem [9], or the method itself involves errors with respect to full gradient as in standard incremental gradient, the classical Robbins-Monro, stochastic gradient, and stochastic approximation methods [22, 25, 2, 3]. When there are persistent errors in gradients, the iterates do not converge and could oscillate in a neighborhood of the optimal solution or even diverge [2, 3, 9, 14]. This makes robustness of the algorithms to gradient errors (in terms of solution accuracy) another important performance objective [9, 16]. In particular, even though accelerated gradient methods proposed by Nesterov converges faster than gradient descent (GD) in the absence of noise for convex problems [20], it was shown that they are less robust to errors, i.e., accelerated methods require higher precision gradient information than GD to achieve the same solution accuracy [9, 7, 14, 26].

In this paper, we study the trade-off between rate of convergence and robustness to gradient errors in designing a first-order algorithm. In particular, we focus on GD and Nesterov’s accelerated gradient (AG) method for quadratic objectives when the gradient has stochastic errors and investigate how to set the parameters of each algorithm to achieve a particular trade-off between these two performance objectives. To study this question systematically, we employ tools from

---

\*The authors are in alphabetical order.

†Department of Industrial and Manufacturing Engineering, Pennsylvania State University, University Park, PA, USA. nsa10@psu.edu.

‡Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Cambridge, MA, USA. afallah@mit.edu.

§Department of Management Science and Information Systems, Rutgers University, Piscataway, NJ, USA. mg1366@rutgers.edu.

¶Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Cambridge, MA, USA. asuman@mit.edu.

control theory whereby we represent each of the algorithms as dynamical systems. This approach has attracted recent attention and has already led to a number of insights for the design and analysis of optimization algorithms [19, 17, 28, 6, 13, 28, 18]. The novelty of our work is to use this approach to provide explicit characterizations of robustness, which can then be placed in a computationally tractable optimization problem to select the algorithm parameters in order to systematically achieve a desired trade-off.

We focus on problems with a strongly convex quadratic objective function. For this case, the rate of convergence of any of the two algorithms we study is given by the spectral radius of the "state-transition" matrix in the dynamical system representation. To characterize robustness, we consider the asymptotic normalized variance of the centered iterate sequence (output vector of the dynamical system) which measures the asymptotic accuracy of the iterates. We show that this limit is well-defined and is given by the  $H_2$  norm of the corresponding linear dynamical system.  $H_2$  norm is a fundamental measure for quantifying robustness of a linear system to noise [29] and admits various definitions and characterizations. We focus on a particular representation of the  $H_2$  norm in terms of the solution of the discrete Lyapunov equation which leads to explicit expressions for GD and AG.

Our first set of results focuses on the GD method. The spectral radius of the state-transition matrix corresponding to GD dynamics, hence the rate of convergence for GD, can be expressed in terms of the smallest and largest eigenvalues of the matrix  $Q$ . We show that the  $H_2$  norm of the linear system, our robustness measure for GD, admits a tractable characterization in terms of the spectrum of  $Q$ . We use these explicit expressions within an optimization problem for selecting the stepsize to minimize  $H_2$  norm subject to a given upper bound on the convergence rate. We also show that there is a fundamental lower bound on the the robustness level of an algorithm for any achievable rate  $\rho$ .

We next consider Nesterov's AG method defined by two parameters: stepsize  $\alpha$  and momentum parameter  $\beta$ . We first characterize the *stability region* of the method, i.e., the values of nonnegative  $(\alpha, \beta)$  for which the spectral radius of the state-transition matrix is less than or equal to one. Similar to GD, we then provide an explicit characterization of the  $H_2$  norm of a relevant dynamical system. Our results show that AG is superior to the GD in the sense that AG can always achieve the same rate with GD while being more robust to noise; similarly, AG can be tuned to be faster than GD while achieving the same robustness level. This behavior is contrary to the AG and GD comparison in the deterministic gradient error setting of [9] when the superiority of one method against the other is not clear. These results show that the random and deterministic noise settings have different behavior.

In addition to the above cited papers, our paper is related to the analysis of the first order methods with *inexact gradients* subject to random noise. Hu *et. al.* [18] analyzes the stochastic gradient method under deterministic noise and study the effect of the stepsize on the convergence rate and the asymptotic neighborhood of convergence. Devolder *et al.* [9] considers a general inexact first-order oracle and provides complexity estimates for GD and AG methods, and highlights the interesting phenomenon that superiority of accelerated methods over the classical ones may not hold with inexact oracles. In another manuscript [8], the authors design a method with intermediate speed and intermediate sensitivity to this inexact oracle which can model deterministic gradient errors, however they are not applicable to random additive gradient errors that is more suitable in a number of settings in machine learning [16, 1, 24].

**Notation.** The  $d \times d$  identity matrix and zero matrix are denoted by  $I_d$  and  $0_d$ , respectively. We define  $\text{diag}(a_1, \dots, a_d)$  as the diagonal matrix with diagonal entries  $a_1, \dots, a_d$ . For matrix  $A \in \mathbb{R}^{d \times d}$ ,  $\text{Tr}(A)$  denotes the trace of  $A$ . The spectral radius of  $A$  is defined as the largest absolute value of its eigenvalues and is denoted by  $\rho(A)$ . We say that a square matrix  $A$  is discrete-time stable, if all of its eigenvalues lie strictly inside the unit disc, i.e., if  $\rho(A) < 1$ . We let  $\delta[k]$  denote the Kronecker delta function, i.e.,  $\delta[0] = 1$  and  $\delta[k] = 0$  for any positive integer  $k$ . For scalars  $0 < \mu \leq L$ , we define  $S_{\mu,L}(\mathbb{R}^d)$  as the set of continuously differentiable functions  $f : \mathbb{R}^d \rightarrow \mathbb{R}$  that are strongly convex with modulus  $\mu$  and having Lipschitz-continuous gradients with constant  $L$ . We use the superscript  $\top$  to denote the transpose of a vector or a matrix depending on the context. The ratio  $\kappa = \frac{L}{\mu}$  is called the *condition number* of  $f$ . The gradient  $\nabla f$  is represented as a column vector.

## 2 Optimization Algorithms as Dynamical Systems

A discrete-time linear dynamical system with a feedback rule  $\phi$  can be expressed as

$$x_{k+1} = Ax_k + Bu_k, \quad y_k = Cx_k + Du_k, \quad u_k = \phi(y_k), \quad (1a)$$

for  $k \geq 0$ , where  $x_k \in \mathbb{R}^m$  is the *state*,  $u_k \in \mathbb{R}^d$  is the *input*,  $y_k \in \mathbb{R}^d$  is the *output* – all represented by column vectors. The matrices  $A, B, C$ , and  $D$  are called the *system matrices*; they are fixed matrices with appropriate dimensions. The function  $\phi : \mathbb{R}^d \rightarrow \mathbb{R}^d$  defines the feedback rule that relates the output of this system to its input.

Iterations of many standard optimization algorithms for minimizing a function  $f \in S_{\mu,L}(\mathbb{R}^d)$  such as GD and AG and the heavy-ball methods admit such a dynamical system representation with  $\phi(y) = \nabla f(y)$ . Consider the GD method for a smooth optimization problem  $\min_{y \in \mathbb{R}^d} f(y)$  where  $f \in S_{\mu,L}(\mathbb{R})$ . The GD iterations with a constant stepsize  $\alpha > 0$  take the following form:

$$x_{k+1} = x_k - \alpha \nabla f(x_k) \quad (2)$$

which can be cast as (1) by setting  $u_k = \nabla f(x_k)$  and letting

$$A = I_d, \quad B = -\alpha I_d, \quad C = I_d, \quad D = 0_d. \quad (3)$$

Another example would be Nesterov's AG method, which computes the iterates as follows:

$$v_{k+1} = y_k - \alpha \nabla f(y_k), \quad y_k = (1 + \beta)v_k - \beta v_{k-1}. \quad (4)$$

Substituting  $y_k$  from the second equation into the first in (4) and defining the state vector  $x_k$  to be

$$x_k = [v_k^\top \quad v_{k-1}^\top]^\top, \quad (5)$$

AG iterations can be rewritten as in (1), where

$$A = \begin{bmatrix} (1 + \beta)I_d & -\beta I_d \\ I_d & 0_d \end{bmatrix}, \quad B = \begin{bmatrix} -\alpha I_d \\ 0_d \end{bmatrix}, \quad C = [(1 + \beta)I_d \quad -\beta I_d], \quad D = 0_d. \quad (6)$$

In this work, we study the performance of optimization algorithms under *additive gradient error* which has been commonly used in the literature (see e.g. [23, 9, 4]). Instead of the actual gradient  $\nabla f(x^k)$ , the assumption is that we have access to a noisy version  $\nabla f(x^k) + w_k$  where  $w_k \in \mathbb{R}^d$  represents the additive noise. In the language of dynamical systems, the noisy iterations of the GD and AG algorithms could be stated as

$$x_{k+1} = Ax_k + B(u_k + w_k), \quad y_k = Cx_k, \quad u_k = \nabla f(y_k), \quad (7)$$

where  $A, B$  and  $C$  are selected according to (3) for GD or (6) for AG, and  $D = 0$ . Although our focus in this paper will be primarily on GD and AG dynamics under noise, it will be clear from our discussion that our ideas naturally extends to many other algorithms that admit such a dynamical system representation including the heavy-ball and the robust momentum methods [6, 19, 17].

In the rest of the paper, we focus on problems with a strongly convex quadratic function. That said, in work subsequent to this paper, we extend our framework to the strongly convex functions which will be the topic of the follow-up work [12].

## 3 Quadratic Functions

Let  $f \in S_{\mu,L}(\mathbb{R}^d)$  be a quadratic function given by  $f(y) = \frac{1}{2}y^\top Qy - p^\top y + r$  where  $Q$  is symmetric and positive definite with eigenvalues  $\{\lambda_i\}_{i=1}^d$  listed in increasing order satisfying  $\mu = \lambda_1 \leq \lambda_2 \leq \dots \leq \lambda_d = L$ . The gradient of  $f$  is given by  $\nabla f(y) = Qy - p = Q(y - y^*)$ , where  $y^* = Q^{-1}p$  is the optimal solution to  $\min_{y \in \mathbb{R}^d} f(y)$  (which is unique due to the strong convexity of  $f$ ). Plugging the formula for the gradient  $\nabla f(y_k)$  into (7), we obtain

$$x_{k+1} = (A + BQC)x_k - BQy^* + Bw_k, \quad y_k = Cx_k. \quad (8a)$$

Given  $A$ ,  $B$ , and  $C$  as in (3) for GD and as in (6) for AG, both GD and AG admit a fixed-point  $x^*$  satisfying  $x^* = Ax^*$ ,  $y^* = Cx^*$ . By defining  $\tilde{y}_k := y_k - y^*$  and  $\xi_k := x_k - x^*$  and substituting these two variables in (8) yields

$$\xi_{k+1} = A_Q \xi_k + B w_k, \quad \tilde{y}_k = C \xi_k. \quad (9)$$

where  $A_Q$  is the state-transition matrix given by

$$A_Q = A + BQC. \quad (10)$$

In the absence of noise (when  $w_k \equiv 0$ ), if the spectral radius of the  $A_Q$  matrix is less than one, then we clearly have  $\xi_k \rightarrow 0$  and  $\tilde{y}_k \rightarrow 0$  linearly.

### 3.1 Performance metric under gradient noise

Our goal is to introduce a performance measure for evaluating the asymptotic behavior of the output  $\tilde{y}^k$  of the dynamical system given in (9) (and hence that of the output of the corresponding optimization algorithm modeled by the dynamical system) under random gradient noise. In particular, assume the additive noise  $\{w_k\}$  is a sequence of zero-mean independent and identically distributed random variables with covariance matrix  $\mathbb{E}(w_k w_k^T) = \sigma^2 I_d$  for some scalar  $\sigma > 0$ . Such noise structure is a natural one arising in the context of stochastic optimization problems [22], for instance in large-scale least-square problems when  $w_k$  has structured randomness [14]. In the special case, when  $w_k$  is Gaussian, the state and the output is also Gaussian. This setting arises in the context of differentially private gradient descent algorithms for convex empirical risk minimization [1] as well as the Euler-Maruyama discretization of the overdamped and underdamped Langevin dynamics [5, 10] in the context of Monte Carlo Markov Chain (MCMC) methods and generalization properties of machine learning algorithms (see e.g. [24, 27]).

As the additive noise  $w_k$  is centered with zero mean in our setting, by taking expectation of both sides of (9), we obtain:  $\mathbb{E}[\xi_k] = A_Q^k \xi_0$ ,  $\mathbb{E}[\tilde{y}_k] = C \mathbb{E}[\xi_k] = C A_Q^k \xi_0$ . Therefore, both sequences  $\{\mathbb{E}[\xi_k]\}_k$  and  $\{\mathbb{E}[\tilde{y}_k]\}_k$  converge to zero with a linear rate  $\rho(A_Q)$  that is determined by the spectral radius of  $A_Q$ . However, due to the noise injected at each step, the limit of the output sequence  $\{\tilde{y}^k\}$  will oscillate around the optimal solution with a non-zero variance. A natural performance measure is then to study the *asymptotic normalized variance* by considering the following limit:

$$\mathcal{J} := \lim_{k \rightarrow \infty} \frac{1}{\sigma^2} \mathbb{E}(\tilde{y}_k^\top \tilde{y}_k) \quad (11)$$

which (if it exists) measures the asymptotic accuracy of the iterations. This quantity can be viewed as the *robustness to noise* of the system because it is equal to the ratio of the power of the output to the power of the input noise, measuring how much a system amplifies input noise. In particular, the smaller this measure is, the more *robust* a system is under additive random noise.

The limit in (11) is studied in *standard  $H_2$  theory* arising in the robust control of dynamical systems (see e.g. [15]). It is a well-known fundamental measure for quantifying the robustness of a linear system to noise in control engineering and has been a key metric to design the parameters of control systems subject to noise.<sup>1</sup> It can be shown that the limit  $\mathcal{J}$  in (11) exists and is exactly equal to the square of the  $H_2$  norm of the linear dynamical system (9). We give a proof of this known fact in Lemma A.1 in the appendix for the sake of completeness.

This motivates us to design the parameters of the GD and AG methods in a way to find a trade-off between the  $H_2$  norm (which measures the robustness to random noise) and the convergence rate. We next introduce preliminary material on the  $H_2$  norm and its computation upon which our framework builds. We will denote the  $H_2$  norm of a system, defined by the system matrices  $(A_Q, B, C)$ , as  $H_2(A_Q, B, C)$  (with dependence on the matrices  $A_Q, B, C$  removed when clear). It is known that the square of the  $H_2$  norm can be computed by the following formula

$$\mathcal{J} = H_2^2 = \text{Tr}(CXC^\top) \quad (12)$$

where  $X$  solves the *discrete Lyapunov equation*<sup>2</sup>

$$A_Q X A_Q^\top - X + B B^\top = 0. \quad (13)$$

<sup>1</sup>The  $H_2$  norm can also be defined in the frequency domain, see e.g. [29] for more information.

<sup>2</sup>The value of  $\mathcal{J} = H_2^2$  can also be computed as  $\text{Tr}(B^\top \tilde{X} B)$  where  $\tilde{X}$  solves  $A_Q^\top \tilde{X} A_Q - \tilde{X} + C^\top C = 0$ .

(see e.g. [15, 29]). Moreover, if  $BB^\top$  is positive definite and  $A_Q$  is *discrete-time stable* (when  $\rho(A_Q) < 1$ ), the solution admits the formula below (see e.g. [29])<sup>3</sup>

$$X = \sum_{k=0}^{\infty} (A_Q^\top)^k B^\top B A_Q^k. \quad (14)$$

Next we focus on the GD and AG algorithms, discuss the dependence of their convergence rate and robustness levels ( $H_2$  norm) on the parameters (stepsize  $\alpha$  and momentum  $\beta$ ) and show how to formulate an optimization problem that takes both robustness and convergence rate into account simultaneously.

### 3.2 Gradient descent (GD) method

For the dynamical system representation of GD, choosing the  $A, B, C$  as in (3) which leads to

$$\mathbb{E}[\xi_{k+1}] = A_Q \mathbb{E}[\xi_k], \quad \text{with } A_Q = I_d - \alpha Q$$

where  $A_Q$  is given by (10). Using this representation, it is easy to see that the GD iterations converge linearly in expectation, i.e.,  $\|\mathbb{E}[\xi_k]\| \leq \rho^k \xi_0$ , where

$$\rho(\alpha) = \rho(A_Q) = \|A_Q\| = \max\{|1 - \alpha\mu|, |1 - \alpha L|\} \quad (15)$$

is the convergence rate of the deterministic GD algorithm [2]. We remind that  $\alpha \in (0, 2/L)$  is a necessary condition for global linear convergence (otherwise  $\rho \geq 1$ ). In particular, it is well-known that the fastest rate is achieved for the stepsize

$$\bar{\alpha} := \arg \min_{\alpha \geq 0} \rho(A_Q) = \frac{2}{\mu + L} \quad (16)$$

which leads to a convergence rate of  $\bar{\rho} = 1 - \frac{2}{\kappa+1}$ . The choice of the stepsize not only affects the rate (see (15)) but also the robustness of the GD algorithm to gradient noise ( $H_2$  norm). The following theorem provides an analytical formula for this dependency as a function of the stepsize.

**Proposition 3.1.** *Let  $f$  be a quadratic function of the form  $f(y) = \frac{1}{2}y^\top Qy - p^\top y + r$ . Consider the GD iterations given by (2) with constant stepsize  $\alpha \in (0, 2/L)$ . Then the square of the  $H_2$  norm of the corresponding linear system (9) with system matrices selected according to (3) is given by*

$$H_2^2(\alpha) = \alpha^2 \sum_{i=1}^d \frac{1}{1 - (1 - \alpha\lambda_i)^2} \quad (17)$$

where  $0 < \mu = \lambda_1 \leq \lambda_2 \leq \dots \leq \lambda_d = L$  are the eigenvalues of  $Q$ .

Having explicit expressions for both convergence rate and robustness to noise ( $H_2$  norm) for GD (see (15) and (17)), given a tolerance level  $\epsilon > 0$  on the optimal convergence rate  $\bar{\rho} = 1 - \frac{2}{\kappa+1}$ , a natural approach to account for the trade-off between these two measures is to choose the stepsize  $\alpha$  that results in the most robust algorithm satisfying the rate constraints, i.e. optimizing

$$\min_{\alpha \in (0, 2/L)} H_2(\alpha) \quad \text{subject to } \rho(\alpha) \leq (1 + \epsilon)\bar{\rho}. \quad (18)$$

This problem is equivalent to the following convex problem for  $\epsilon \in [0, \frac{2}{\kappa-1})$  (which ensures that the upper bound on the rate is less than one and the optimization problem (18) admits a solution):

$$\min_{\alpha \in (0, 2/L)} H_2^2(\alpha) \quad \text{subject to } \frac{1}{1 - \rho^2(\alpha)} \leq \frac{1}{1 - (1 + \epsilon)^2 \bar{\rho}^2}. \quad (19)$$

Indeed,  $1/(1 - \rho^2)$  is a nondecreasing convex function for  $\rho \in (0, 1)$  and  $\rho(\alpha)$  is convex in  $\alpha$ ; therefore, both  $1/(1 - \rho(\alpha)^2)$  and  $H_2^2(\alpha)$  in (17) is convex for  $\alpha \in (0, \frac{2}{L})$  and is increasing in  $\alpha$ .

<sup>3</sup>In this paper we focus on the *unstructured noise*, i.e., when the noise covariance matrix is given by  $\mathbb{E}(w_k w_k^\top) = \sigma^2 I_d$ . However, if the noise covariance is structured and is known, i.e.  $\mathbb{E}(w_k w_k^\top) = C$  for some given positive definite  $C$ , then this can also be handled similarly. We simply consider the Cholesky decomposition of  $C = LL^\top$  and substitute  $w_k = L\tilde{w}_k$  into (9), then the  $H_2$  norm of this system can be investigated by replacing  $B$  with  $BL$  in the above formulas.

Moreover, (19) satisfies the Slater condition. Thus, duality implies that there exists  $\tau(\epsilon)$  such that the above minimization problem is equivalent to the following unconstrained problem:

$$\alpha_*(\tau) := \arg \min_{\alpha \in (0, 2/L)} F_\tau(\alpha) := H_2^2(\alpha) + \tau \frac{1}{1 - \rho^2(\alpha)}. \quad (20)$$

The parameter  $\tau > 0$  determines the trade-off between rate and the stability. For small  $\tau$ , the dominant term in the cost would be  $H_2^2$  so that we expect the optimal stepsize to be small. On the other hand, for large enough  $\tau$ , the convergence rate is the dominant term in the cost; therefore, one would expect the optimal stepsize (that solves the problem (20)) to be close to  $\bar{\alpha}$  which corresponds to the fastest achievable rate  $\bar{\rho}$  (see (16)). In order to get more intuition about the effect of the choice of the stepsize parameter, we next give an illustrative example in dimension  $d = 2$  to show the behavior of the optimal  $\alpha_*(\tau)$  as the tradeoff parameter  $\tau$  is varied from zero to infinity. For computational tractability, we consider the unconstrained version of the problem in (20).<sup>4</sup>

**Example 3.2.** In dimension  $d = 2$ , let  $\tau = 5$  and consider the parameters

$$\mu = \lambda_1 = 0.1 \quad \text{and} \quad L = \lambda_2 = 1 \quad \text{with} \quad \kappa = \frac{L}{\mu} = 10. \quad (21)$$

Writing the terms appearing in the non-linear equations in Proposition C.1 under a common denominator, both of the equality conditions can be reformulated as a polynomial root finding problem in  $\alpha$  for a polynomial with degree 4. The roots of polynomials can easily be found efficiently up to arbitrary accuracy by calculating the eigenvalues of the corresponding companion matrix [11], for instance using the `roots` function in MATLAB. After a careful examination of all the roots, we conclude that the optimal stepsize  $\alpha_*$  that minimizes the cost  $F_\tau(\alpha)$  is  $\alpha_* \approx 1.43$  with rate  $\rho(\alpha_*) \approx 0.86$  and robustness  $H_2^2(\alpha_*) \approx 10.3$ . This point is marked Figure 1 below which shows the robustness level  $H_2^2$  as a function of the optimal convergence rate  $\rho$  when we change  $\tau$  from zero to infinity (corresponds to the rightmost point in the curve) to infinity (corresponds to the uppermost point in the curve) for the parameters in (21).

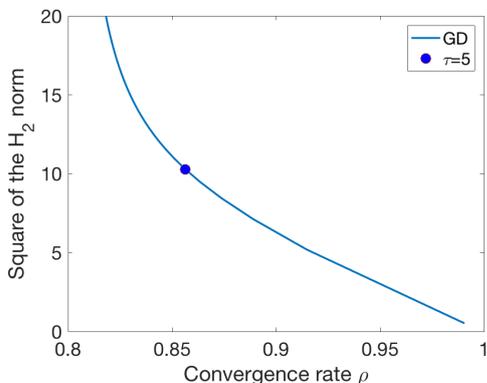


Figure 1: Robustness to gradient noise in terms of  $H_2$  norm as a function of  $\rho$  for GD.

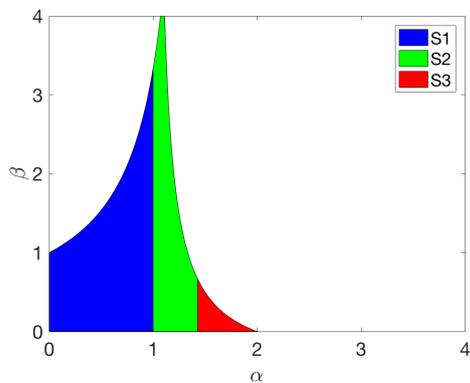


Figure 2: The stability region  $\mathcal{S} = \mathcal{S}_1 \cup \mathcal{S}_2 \cup \mathcal{S}_3$  with parameters  $\mu = 0.7$  and  $L = 1$ .

Figure 1 illustrate the trade-off between the rate and robustness. When  $\tau$  is very small, the robustness is dominant in the cost function  $F$ , which leads to a smaller stepsize with improved robustness but slower convergence. The fact that smaller stepsizes lead to improved robustness has also been observed in the literature for both additive and multiplicative deterministic noise [19, 13]. As  $\tau$  grows, we achieve faster rate at the expense of a larger robustness level. In particular, as  $\tau$  goes to infinity; the optimal stepsize  $\alpha_*(\tau) \rightarrow \bar{\alpha} = 2/(L + \mu)$  which leads to the fastest rate  $\bar{\rho} = 1 - \frac{2}{\kappa+1} \approx 0.8182$ . We see from Figure 1 that smaller values of  $\rho$  (or equivalently smaller values of  $\frac{1}{1-\rho^2}$ ) are accompanied by larger values of  $H_2^2$ . This suggests that the product  $H_2^2 \frac{1}{1-\rho^2}$  cannot be too small for any choice of the stepsize  $\alpha$ . The next lemma shows that for a given achievable rate  $\rho$ , there are some fundamental limits (lower bounds) to how robust the GD can be.

<sup>4</sup>In Proposition C.1 of the appendix, we derive the first-order conditions for  $\alpha_*(\tau)$  that allows it to be computed up to arbitrary accuracy.

**Proposition 3.3.** *Under the assumptions of Proposition 3.1, the following inequality holds*

$$H_2^2(\alpha) \geq (1 - \rho^2(\alpha)) \sum_{i=1}^d \frac{1}{4\lambda_i^2} \quad (22)$$

for any choice of the stepsize  $\alpha > 0$  where  $\rho(\alpha)$  is given by (15) and  $H_2^2(\alpha)$  is given by (17).

### 3.3 Nesterov's accelerated gradient (AG) method

As we discussed in Section 2, the dynamical system model of the Nesterov's accelerated gradient method admits the state vector (5) and the system matrices (6). This results in the iterations stated in (9) for the system matrices  $A_Q, B, C$  given by (6) and (10), i.e.,

$$A_Q = \begin{bmatrix} (1 + \beta)(I_d - \alpha Q) & -\beta(I_d - \alpha Q) \\ I_d & 0_d \end{bmatrix}, B = \begin{bmatrix} -\alpha I_d \\ 0_d \end{bmatrix}, C = [(1 + \beta)I_d \quad -\beta I_d]. \quad (23)$$

In this section, we will formulate an analogous problem to (20) for the AG method to design the parameters  $(\alpha, \beta)$  in a way to find a trade-off between the rate and the robustness. Because AG has the pair  $(\alpha, \beta)$  as design parameters, the analogue of (20) is

$$(\alpha_*, \beta_*) := \arg \min_{(\alpha, \beta) \in \mathcal{S}} F_\tau(\alpha, \beta) := H_2^2(\alpha, \beta) + \tau \frac{1}{1 - \rho(\alpha, \beta)^2} \quad (24)$$

where  $H_2^2(\alpha, \beta)$  is the square of the  $H_2$  norm for the system (9),  $\rho(\alpha, \beta)$  is the convergence rate of AG with parameters  $(\alpha, \beta)$  and  $\mathcal{S}$  is the set of all possible choices of the tuple  $(\alpha, \beta)$  so that the AG iterations are globally convergent, i.e.

$$\mathcal{S} = \{(\alpha, \beta) : \rho(A_Q) < 1, \alpha \geq 0, \beta \geq 0\} \subset \mathbb{R}^2. \quad (25)$$

We call the set  $\mathcal{S}$ , the *stability region* of AG, in analogy with the stability region of numerical methods that arise in the discretization of continuous-time differential equations.

The optimization problem (24) finds the parameters that find a trade-off between the convergence rate and the robustness for a given value of the trade-off parameter  $\tau$  among all possible choices of  $(\alpha, \beta)$ . Similar to the GD case, this requires an understanding of the convergence rate as well as the robustness level for given any parameters  $(\alpha, \beta) \in \mathcal{S}$ . The convergence rate  $\rho$  of the AG method as a function of  $\alpha$  and  $\beta$  is well-known. Diagonalizing the  $A_Q$  matrix using the eigenvalue decomposition of  $Q$ , it can be shown after some computations that the rate  $\rho = \rho(\alpha, \beta)$  admits the following formula

$$\rho(\alpha, \beta) = \rho(A_Q) = \max\{\rho_\mu(\alpha, \beta), \rho_L(\alpha, \beta)\} \quad (26)$$

where  $A_Q$  is defined by (23) and  $\rho_\lambda$  is defined for  $\lambda \in \{\mu, L\}$  as follows:

$$\rho_\lambda(\alpha, \beta) = \begin{cases} \frac{1}{2}|(1 + \beta)(1 - \alpha\lambda)| + \frac{1}{2}\sqrt{\Delta_\lambda} & \text{if } \Delta_\lambda \geq 0 \\ \sqrt{\beta(1 - \alpha\lambda)} & \text{otherwise} \end{cases}, \quad \Delta_\lambda = (1 + \beta)^2(1 - \alpha\lambda)^2 - 4\beta(1 - \alpha\lambda) \quad (27)$$

(see e.g. [19, Appendix A], [21, Section 4.3]). The explicit expression (26) for the rate allows us to characterize the set  $\mathcal{S}$  in the next proposition. We illustrate the set  $\mathcal{S}$  in Figure 2.<sup>5</sup>

**Proposition 3.4.** *Let  $\mathcal{S}$  be the stability set of Nesterov's accelerated method defined by (25). Then its closure is given by the union of the following three sets:*

$$\begin{aligned} \mathcal{S}_1 &:= \left\{ (\alpha, \beta) : 0 \leq \alpha \leq \frac{1}{L}, 0 \leq \beta(1 - \alpha\mu) \leq 1 \right\}, \\ \mathcal{S}_2 &:= \left\{ (\alpha, \beta) : \frac{1}{L} < \alpha \leq \min \left\{ \frac{2}{L}, \frac{1}{\mu} \right\}, \alpha L - 1 \leq \frac{1}{2\beta + 1}, \beta(1 - \alpha\mu) \leq 1 \right\}, \\ \mathcal{S}_3 &:= \left\{ (\alpha, \beta) : \frac{1}{\mu} \leq \alpha \leq \frac{2}{L}, \alpha L - 1 \leq \frac{1}{2\beta + 1} \right\}, \end{aligned} \quad (28)$$

with the convention that  $\mathcal{S}_3$  is the empty set if  $\mu < \frac{L}{2}$ .

<sup>5</sup>We note that the stability region of a second-order difference equation that arises in accelerated algorithms that are sublinearly convergent for weakly convex quadratic functions has been studied in [14], however these results do not apply to the set  $\mathcal{S}$  as we do not require the rate to be accelerated (we consider not only accelerated rates but also any rate  $\rho$  less than one) and we consider strongly convex functions instead of the weakly convex.

The following proposition gives a characterization of the  $H_2$  norm of the AG that we will need to study for the optimization problem (24).

**Proposition 3.5.** *Consider the dynamical system model (9) of AG algorithm with system matrices  $A_Q, B, C$  given by (6) and (23), and parameters  $(\alpha, \beta) \in \mathcal{S}$ . Then, the  $H_2$  norm of this system is*

$$H_2^2(\alpha, \beta) = \sum_{i=1}^d u_{\alpha, \beta}(\lambda_i) \quad (29)$$

where  $\mu = \lambda_1 \leq \lambda_2 \leq \dots \leq \lambda_d = L$  are the eigenvalues of  $Q$  and

$$u_{\alpha, \beta}(\lambda) := \alpha \frac{1 + \beta + \alpha\lambda\beta(2\beta + 1)}{\lambda(1 - \beta(1 - \alpha\lambda))(2 + 2\beta - \alpha\lambda(1 + 2\beta))}. \quad (30)$$

In the special case, choosing  $\beta = 0$  reduces to the formula (17) derived for GD.

Since we have an exact characterization of the  $H_2$  norm, we can derive the optimality conditions for the problem (24) by an approach similar to Proposition C.1 where the optimizer can be characterized as a root of a polynomial. In dimension  $d = 2$ , the optimizer is easy to compute given parameters  $\mu$  and  $L$ ; however in high dimensions, this is computationally expensive as it would require determining all the eigenvalues of  $Q$  (there are typically other eigenvalues than  $\mu$  and  $L$ ). Nevertheless, exploiting the convexity properties of the  $H_2^2$  function, we develop a tractable upper bound to the function  $F_\tau(\alpha, \beta)$  that is easy to optimize. We defer the technical details to the appendix, however we present numerical experiments in the next section that illustrate that this approach can lead to good approximations to the optimal parameters.

## 4 Numerical Experiments

Our first set of experiments, illustrated in Figure 3, concern a further study of Example 3.2 for comparing AG and GD in terms of performance. We vary the trade-off parameter from  $\tau = 0$  to  $\tau = \infty$  and plot the robustness level  $H_2^2(\alpha_*(\tau), \beta_*(\tau))$  versus the rate  $\rho(\alpha_*(\tau), \beta_*(\tau))$  corresponding to the optimal parameters  $(\alpha_*(\tau), \beta_*(\tau))$ , we do the same for GD as in Figure 1. We observe that for the same achievable rate, the optimized AG parameters lead to more robust algorithms compared to the optimized GD algorithms as AG have an additional parameter to optimize over. This shows that AG can improve GD in terms of both convergence rate and robustness at the same time when gradients are subject to white noise, which is not the case for deterministic gradient errors in the worst case [9]. It is interesting that the popular choice of parameters (blue and red dots), as well as the parameters that lead to the optimal (fastest) rate (green and purple dots) lie on the curves that trade robustness with rate in an optimal fashion.

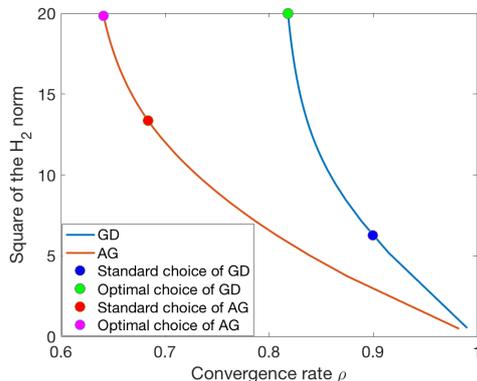


Figure 3: Robustness to gradient noise in terms of  $H_2$  norm as a function of the convergence rate  $\rho$ .

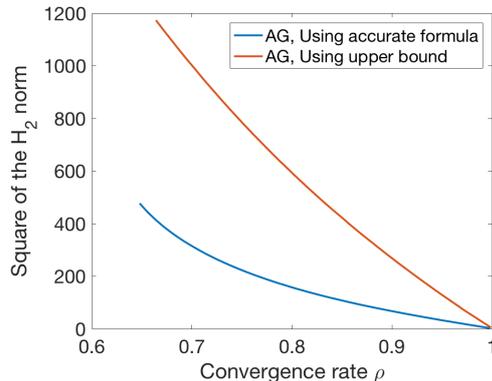


Figure 4:  $H_2^2$ , computed both accurately and by upper bound, as a function of the convergence rate for  $d = 100$ .

In higher dimensions, optimizing the costs (20) and (24) accurately is expensive. Therefore, we choose the parameters based on optimizing the upper bounds developed in Section H which can be done efficiently. Figure (4) plots the convergence rate and the  $H_2$  norm obtained by solving (24) accurately versus solving it approximately based on optimizing an upper bound. The objective is

a random quadratic function in dimension  $d = 100$  with parameters. Our results show that for any trade-off parameter  $\tau$  our upper bound is within a factor of 2.5 of true parameters, illustrating the quality of our approximation to the optimal parameter choice.

## 5 Conclusion

We consider the gradient descent (GD) and accelerated gradient (AG) methods for optimizing strongly convex quadratic functions. We developed a computationally tractable framework to design their parameters in a way to trade between two conflicting performance measures: the convergence rate and the robustness to additive white noise in the gradient computations measured in terms of final asymptotic variance of the algorithm output. We show that this robustness measure is equal to the  $H_2$  norm of a dynamical system associated to the optimization algorithm and give an explicit characterization of this quantity. Our results show that for the same achievable rate, AG can always be tuned to be more robust. Similarly, for the same robustness level, we show that AG can be tuned to be always faster than GD. We also give fundamental lower bounds on the achievable robustness level for gradient descent for a given achievable rate.

## References

- [1] R. Bassily, A. Smith, and A. Thakurta. Private empirical risk minimization: Efficient algorithms and tight error bounds. In *Foundations of Computer Science (FOCS), 2014 IEEE 55th Annual Symposium on*, pages 464–473. IEEE, 2014.
- [2] D.P Bertsekas. *Nonlinear programming*. Athena scientific Belmont, 1999.
- [3] D.P Bertsekas. Incremental gradient, subgradient, and proximal methods for convex optimization: A survey. *Optimization for Machine Learning*, 2010(1-38):3, 2011.
- [4] D.P. Bertsekas and J.N. Tsitsiklis. Gradient convergence in gradient methods with errors. *SIAM Journal on Optimization*, 10(3):627–642, 2000.
- [5] X. Cheng, N. S. Chatterji, P. L. Bartlett, and M. I. Jordan. Underdamped Langevin MCMC: A non-asymptotic analysis. *ArXiv e-prints*, July 2017.
- [6] S. Cyrus, B. Hu, B. Van Scoy, and L. Lessard. A robust accelerated optimization algorithm for strongly convex functions. *arXiv preprint arXiv:1710.04753*, 2017.
- [7] A. d’Aspremont. Smooth optimization with approximate gradient. *SIAM Journal on Optimization*, 19(3):1171–1183, 2008.
- [8] O. Devolder, F. Glineur, and Y. Nesterov. Intermediate gradient methods for smooth convex problems with inexact oracle. Technical report, Université catholique de Louvain, Center for Operations Research and Econometrics (CORE), 2013.
- [9] O. Devolder, F. Glineur, and Y. Nesterov. First-order methods of smooth convex optimization with inexact oracle. *Mathematical Programming*, 146(1-2):37–75, 2014.
- [10] A. Eberle, A. Guillin, and R. Zimmer. Couplings and quantitative contraction rates for Langevin dynamics. *ArXiv e-prints*, March 2017.
- [11] A. Edelman and H. Murakami. Polynomial roots from companion matrix eigenvalues. *Mathematics of Computation*, 64(210):763–776, 1995.
- [12] A. Fallah, M. Gürbüzbalaban, and A. Ozdaglar. Robust accelerated gradient methods for strongly convex functions. In *In preparation*, 2018.
- [13] M. Fazlyab, A. Ribeiro, M. Morari, and V.M. Preciado. Analysis of optimization algorithms via integral quadratic constraints: Nonstrongly convex problems. *arXiv preprint arXiv:1705.03615*, 2017.
- [14] N. Flammarion and F. Bach. From averaging to acceleration, there is only a step-size. In *Conference on Learning Theory*, pages 658–695, 2015.

- [15] W.M. Haddad and D.S. Bernstein. Parameter-dependent Lyapunov functions and the discrete-time popov criterion for robust analysis. *Automatica*, 30(6):1015–1021, 1994.
- [16] M. Hardt. Robustness versus acceleration. August 18th, 2014. <http://blog.mrtz.org/2014/08/18/robustness-versus-acceleration.htm>.
- [17] B. Hu and L. Lessard. Dissipativity theory for nesterov’s accelerated method. *arXiv preprint arXiv:1706.04381*, 2017.
- [18] B. Hu, P. Seiler, and L. Lessard. Analysis of approximate stochastic gradient using quadratic constraints and sequential semidefinite programs. *arXiv preprint arXiv:1711.00987*, 2017.
- [19] Laurent Lessard, Benjamin Recht, and Andrew Packard. Analysis and design of optimization algorithms via integral quadratic constraints. *SIAM Journal on Optimization*, 26(1):57–95, 2016.
- [20] Y. Nesterov. A method of solving a convex programming problem with convergence rate  $o(1/k^2)$ . In *Soviet Mathematics Doklady*, volume 27, pages 372–376, 1983.
- [21] B. O’Donoghue and E. Candès. Adaptive restart for accelerated gradient schemes. *Foundations of Computational Mathematics*, 15(3):715–732, Jun 2015.
- [22] B. T. Polyak and A. B. Juditsky. Acceleration of stochastic approximation by averaging. *SIAM Journal on Control and Optimization*, 30(4):838–855, 1992.
- [23] B.T. Polyak. Introduction to optimization. translations series in mathematics and engineering. *Optimization Software*, 1987.
- [24] M. Raginsky, A. Rakhlin, and M. Telgarsky. Non-convex learning via stochastic gradient langevin dynamics: a nonasymptotic analysis. *arXiv preprint arXiv:1702.03849*, 2017.
- [25] H. Robbins and S. Monro. A stochastic approximation method. *The annals of mathematical statistics*, pages 400–407, 1951.
- [26] M. Schmidt, Nicolas L.R., and F.R. Bach. Convergence rates of inexact proximal-gradient methods for convex optimization. In J. Shawe-Taylor, R. S. Zemel, P. L. Bartlett, F. Pereira, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 24*, pages 1458–1466. Curran Associates, Inc., 2011.
- [27] B. Tzen, T.n Liang, and M. Raginsky. Local optimality and generalization guarantees for the langevin algorithm via empirical metastability. *arXiv preprint arXiv:1802.06439*, 2018.
- [28] A.C. Wilson, B. Recht, and M.I. Jordan. A lyapunov analysis of momentum methods in optimization. *arXiv preprint arXiv:1611.02635*, 2016.
- [29] Kemin Zhou, John Comstock Doyle, Keith Glover, et al. *Robust and optimal control*, volume 40. Prentice hall New Jersey, 1996.

## A Equivalence between the performance measure $\mathcal{J}$ and the $H_2$ norm

**Lemma A.1.** Consider the linear dynamical system (9) defined by the  $(A_Q, B, C)$  matrices and assume that the additive noise  $\{w_k\}$  is a sequence of zero-mean, independent and identically distributed (i.i.d.) random variables with a covariance matrix  $\mathbb{E}(w_k w_k^T) = \sigma^2 I_d$  for some  $\sigma > 0$ . If the spectral radius of  $A_Q$  is less than one; then the limit (11) exists and is equal to the  $H_2$  norm of the system, i.e. we have

$$\mathcal{J} = H_2^2 = \lim_{k \rightarrow \infty} \frac{1}{\sigma^2} \mathbb{E}(\tilde{y}_k^T \tilde{y}_k)$$

where  $H_2$  satisfies (12)–(14).

*Proof.* Let  $X$  be the solution of Lyapunov equation (13). Using (12) and (9), we obtain

$$\begin{aligned} \mathbb{E}[\tilde{y}_k^T \tilde{y}_k] - \sigma^2 H_2^2 &= \text{Tr}(\mathbb{E}[\tilde{y}_k \tilde{y}_k^T]) - \sigma^2 \text{Tr}(CXC^T) \\ &= \text{Tr}(C\mathbb{E}[\xi_k \xi_k^T]C^T) - \sigma^2 \text{Tr}(CXC^T) \\ &= \text{Tr}(C(V_k - \sigma^2 X)C^T) \end{aligned} \quad (31)$$

where  $V_k$  is defined as  $V_k = \mathbb{E}[\xi_k \xi_k^T]$ . Hence, it suffices to show that

$$\lim_{k \rightarrow \infty} \text{Tr}(C(V_k - \sigma^2 X)C^T) = 0. \quad (32)$$

To show this, first note that the following recursive equation holds using (9):

$$\begin{aligned} V_{k+1} &= \mathbb{E}[\xi_{k+1} \xi_{k+1}^T] \\ &= \mathbb{E}[(A_Q \xi_k + B w_k)(A_Q \xi_k + B w_k)^T] \\ &= A_Q \mathbb{E}[\xi_k \xi_k^T] A_Q^T + A_Q \mathbb{E}[\xi_k w_k^T] B^T + B \mathbb{E}[w_k \xi_k^T] A_Q^T + B \mathbb{E}[w_k w_k^T] B^T \\ &= A_Q V_k A_Q^T + \sigma^2 B B^T \end{aligned} \quad (33)$$

where in the last equality we used the fact that  $w_k$  are zero-mean with covariance matrix  $\mathbb{E}[w_k w_k^T] = \sigma^2 I_d$ , and it is independent from  $\xi_k$ . Moreover, by (13) we have

$$X = A_Q X A_Q^T + B B^T. \quad (34)$$

Subtracting  $\sigma^2 X$  from both sides of (33) and replacing  $X$  by (34), we obtain

$$\begin{aligned} V_{k+1} - \sigma^2 X &= A_Q (V_k - \sigma^2 X) A_Q^T \\ &= A_Q^k (V_1 - \sigma^2 X) (A_Q^T)^k \end{aligned} \quad (35)$$

where the last equality comes from recursively using the preceding equality. Since spectral radius of  $A_Q$  is less than one, if we consider the Jordan decomposition of  $A_Q$ , it follows that  $A_Q^k \rightarrow 0$  as well as  $(A_Q^T)^k \rightarrow 0$  as  $k \rightarrow \infty$ . This proves (32) as desired.  $\square$

## B Proof of Theorem 3.1

We first show that without loss of generality we can assume  $Q$  is a diagonal matrix, this simplifies the proof argument. To show this, let  $Q = U \Lambda U^T$  be the eigenvalue decomposition of  $Q$  where  $U$  is a unitary matrix and  $\Lambda = \text{diag}(\lambda_1, \dots, \lambda_d)$  is a diagonal matrix containing the eigenvalues of  $Q$ . We also define  $A_\Lambda = I_d - \alpha \Lambda$  which is a diagonal matrix. Multiplying  $A_Q$  by  $U^T$  and  $U$  from left and right respectively yields to

$$U^T A_Q U = U^T (I_d - \alpha Q) U = I_d - \alpha \Lambda = A_\Lambda. \quad (36)$$

Similarly, we multiply the Lyapunov equation (13) from left and right by  $U^T$  and  $U$  respectively which implies that

$$U^T A_Q X A_Q^T U - U^T X U + \alpha^2 I_d = 0$$

where we used the fact that  $B = -\alpha I_d$  for the dynamical system representation of the gradient descent (see (3)). It follows from (36) that  $A_Q = U(I_d - \alpha\Lambda)U^T$ , plugging this into the Lyapunov equation above, we obtain

$$(I_d - \alpha\Lambda)U^T XU(I_d - \alpha\Lambda) - U^T XU + \alpha^2 I_d = 0.$$

This means that the matrix  $U^T XU$  solves the Lyapunov equation obtained by replacing  $A_Q$  by  $A_\Lambda$  in (13). Furthermore, the  $H_2$  norm after this substitution (corresponding to the system matrices  $(A_\Lambda, B, C)$ ) would be equal to

$$\sqrt{\text{Tr}(CU^T XUC)} = \sqrt{\text{Tr}(U^T XU)} = \sqrt{\text{Tr}(X)} = \sqrt{\text{Tr}(CXCT)}$$

where we used the fact that  $C = I_d$  in the dynamical system representation of the gradient descent (see (3)). Therefore, the value of the  $H_2$  norm would be invariant if we were to replace  $Q$  by  $\Lambda$  and solve the Lyapunov equations for the triplet  $(A_\Lambda, B, C)$ . It is straightforward to verify that

$$X = \text{diag}\left(\frac{\alpha^2}{1 - (1 - \alpha\lambda_1)^2}, \dots, \frac{\alpha^2}{1 - (1 - \alpha\lambda_d)^2}\right)$$

satisfies the Lyapunov equation for  $(A_\Lambda, B, C)$ , and plugging it into (12) completes the proof.

## C First-order optimality conditions for the objective $F_\tau(\alpha)$

**Proposition C.1.** *There exists an optimizer  $\alpha_*(\tau)$  to the minimization problem (20). Furthermore, any optimizer is either  $\alpha_*(\tau) = 2/(\mu + L)$  or it satisfies one of the following two conditions:*

$$\left\{ \begin{array}{l} \alpha^2 \sum_{i=1}^d \frac{1}{\lambda_i(2 - \alpha\lambda_i)^2} + \frac{\tau(\alpha\mu - 1)}{\mu(2 - \alpha\mu)^2} = 0 \quad \text{and} \quad |1 - \alpha\mu| > |1 - \alpha L|. \end{array} \right. \quad (37)$$

$$\left\{ \begin{array}{l} \alpha^2 \sum_{i=1}^d \frac{1}{\lambda_i(2 - \alpha\lambda_i)^2} + \frac{\tau(\alpha L - 1)}{L(2 - \alpha L)^2} = 0 \quad \text{and} \quad |1 - \alpha\mu| < |1 - \alpha L|. \end{array} \right. \quad (38)$$

Therefore, by examining the values of  $F$  at the points that satisfy this equality and inequality constraints, we can determine the optimal stepsize  $\alpha^*(\tau)$ .

*Proof.* To solve (20) explicitly note that

$$\frac{1}{1 - \rho^2} = \begin{cases} \frac{1}{1 - (1 - \alpha\mu)^2} & \text{if } |1 - \alpha\mu| > |1 - \alpha L| \\ \frac{1}{1 - (1 - \alpha L)^2} & \text{if } |1 - \alpha\mu| \leq |1 - \alpha L|. \end{cases}$$

The optimal  $\alpha^*$  cannot be attained on the boundary points of the interval  $[0, 2/L]$  as the objective  $F$  is infinite at these points. Therefore, the optimum is attained at an interior point of the interval  $[0, 2/L]$  where  $F$  is differentiable with respect to  $\alpha$  except when  $|1 - \alpha\mu| = |1 - \alpha L|$ , i.e. when  $\alpha = 2/(\mu + L)$ . For  $\alpha^* \neq 2/(\mu + L)$ , we can write-down the first-order conditions of optimality

$$\frac{\partial F}{\partial \alpha} = 0$$

which leads to (37) and (38). □

## D Proof of Proposition 3.3

It follows from (15) that for every  $i \in \{1, \dots, d\}$ , we have  $\rho \geq |1 - \alpha\lambda_i|$ . This implies that

$$\frac{1}{1 - \rho^2} \geq \frac{1}{1 - (1 - \alpha\lambda_i)^2}.$$

Multiplying both sides by  $\frac{\alpha^2}{1 - (1 - \alpha\lambda_i)^2}$  and summing over all  $i$  yields

$$\frac{1}{1 - \rho^2} \sum_{i=1}^d \frac{\alpha^2}{1 - (1 - \alpha\lambda_i)^2} \geq \sum_{i=1}^d \frac{\alpha^2}{(1 - (1 - \alpha\lambda_i)^2)^2}.$$

Then, it follows from (17) that

$$\frac{1}{1-\rho^2} H_2^2 \geq \sum_{i=1}^d \frac{\alpha^2}{(1-(1-\alpha\lambda_i)^2)^2}. \quad (39)$$

The right hand side of (39) admits a lower bound as follows

$$\begin{aligned} \sum_{i=1}^d \frac{\alpha^2}{(1-(1-\alpha\lambda_i)^2)^2} &= \sum_{i=1}^d \frac{\alpha^2}{(\alpha\lambda_i(2-\alpha\lambda_i))^2} \\ &= \sum_{i=1}^d \frac{1}{\lambda_i^2(2-\alpha\lambda_i)^2} \geq \sum_{i=1}^d \frac{1}{4\lambda_i^2}. \end{aligned} \quad (40)$$

where the last inequality stems from the fact that  $|2-\alpha\lambda_i| \leq 2$ . Using the lower bound (40) along with (39) completes the proof.

## E Proof of Proposition 3.4

In the light of the formula (26) that characterizes  $\rho(A_Q)$ , the closure of the stability set  $\mathcal{S}$  admits the representation

$$\mathcal{S} = \mathcal{S}_\mu \cap \mathcal{S}_L \quad (41)$$

where for  $\lambda \in \{\mu, L\}$  we define

$$\mathcal{S}_\lambda = \{(\alpha, \beta) : \rho_\lambda(\alpha, \beta) \leq 1, \alpha \geq 0, \beta \geq 0\} \subset \mathbb{R}^2. \quad (42)$$

We first write  $\mathcal{S}_\lambda$  as a union of two disjoint sets depending on the signature of  $\Delta_\lambda$ :

$$\mathcal{S}_\lambda = \mathcal{S}_{\lambda,1} \cup \mathcal{S}_{\lambda,2}$$

where

$$\mathcal{S}_{\lambda,1} = \mathcal{S}_\lambda \cap \{(\alpha, \beta) : \Delta_\lambda \leq 0\}, \quad \mathcal{S}_{\lambda,2} = \mathcal{S}_\lambda \cap \{(\alpha, \beta) : \Delta_\lambda > 0\}. \quad (43)$$

It follows from the definition of  $\Delta_\lambda$  in (27) that  $\Delta_\lambda \leq 0$  if and only if

$$0 \leq 1 - \alpha\lambda \leq \frac{4\beta}{(1+\beta)^2} \quad (44)$$

and similarly, we have  $\Delta_\lambda \geq 0$  if and only if

$$1 - \alpha\lambda \leq 0 \quad \text{or} \quad 1 - \alpha\lambda \geq \frac{4\beta}{(1+\beta)^2}. \quad (45)$$

If (44) holds then,  $\rho_\lambda = \sqrt{\beta(1-\alpha\lambda)} \leq 1$  if and only if

$$0 \leq 1 - \alpha\lambda \leq \frac{1}{\beta}.$$

Therefore,

$$\mathcal{S}_{\lambda,1} = \{(\alpha, \beta) : 0 \leq 1 - \alpha\lambda \leq \min\{\frac{1}{\beta}, \frac{4\beta}{(1+\beta)^2}\}\}. \quad (46)$$

We next focus on  $\mathcal{S}_{\lambda,2}$ . If (45) is satisfied, then  $\rho_\lambda \leq 1$  if and only if

$$\frac{1}{2}(1+\beta)(1-\alpha\lambda)\text{sign}(1-\alpha\lambda) + \frac{1}{2}\sqrt{\Delta_\lambda} \leq 1. \quad (47)$$

There are two cases:

- $\Delta_\lambda > 0$  and  $1 - \alpha\lambda < 0$ : In this case,  $\rho_\lambda \leq 1$  if and only if

$$\sqrt{\Delta_\lambda} \leq 2 - (1 + \beta)c_\lambda$$

where  $c_\lambda = -(1 - \alpha\lambda) > 0$ . By squaring both sides, this is if and only if,

$$\Delta_\lambda \leq (2 - (1 + \beta)c_\lambda)^2 \quad \text{and} \quad 2 - (1 + \beta)c_\lambda \geq 0$$

The first inequality holds if

$$c_\lambda = -(1 - \alpha\lambda) \leq \frac{1}{2\beta + 1}$$

whereas the second inequality holds if

$$c_\lambda = -(1 - \alpha\lambda) \leq \frac{2}{\beta + 1}.$$

The first inequality is more binding, if it holds the second inequality holds too. Therefore,

$$\{(\alpha, \beta) : 0 \leq -(1 - \alpha\lambda) \leq \frac{1}{2\beta + 1}\} \subset \mathcal{S}_{\lambda,2}. \quad (48)$$

- $\Delta_\lambda > 0$  and  $1 - \alpha\lambda > 0$ : In this case,  $\rho_\lambda \leq 1$  if and only if

$$\sqrt{\Delta_\lambda} \leq 2 - (1 + \beta)d_\lambda$$

where  $d_\lambda := -c_\lambda = (1 - \alpha\lambda) > 0$ . After squaring both sides, this is if and only if

$$\Delta_\lambda \leq (2 - (1 + \beta)d_\lambda)^2 \quad \text{and} \quad 2 - (1 + \beta)d_\lambda \geq 0 \quad (49)$$

where the first inequality simplifies to  $1 \geq d_\lambda$ . (49) along with (45) means

$$\frac{4\beta}{(1 + \beta)^2} \leq 1 - \alpha\lambda \leq \min\left\{1, \frac{2}{1 + \beta}\right\}$$

which implies  $\beta \leq 1$ , and therefore

$$\{(\alpha, \beta) : \frac{4\beta}{(1 + \beta)^2} \leq 1 - \alpha\lambda \leq \frac{2}{1 + \beta}\} \subset \mathcal{S}_{\lambda,2}. \quad (50)$$

Merging (46), (48), and (50) yields

$$\mathcal{S}_\lambda = \left\{(\alpha, \beta) : 1 - \alpha\lambda \in \left[-\frac{1}{1 + 2\beta}, \min\left\{\frac{1}{\beta}, \frac{2}{1 + \beta}\right\}\right]\right\}. \quad (51)$$

In order to complete the proof, due to the representation (42), it suffices to compute the intersection  $\mathcal{S}_\mu \cap \mathcal{S}_L$ . There are several cases to consider depending on the value of  $\alpha$ :

- First, consider  $\alpha \in [0, \frac{1}{L}]$ . In this case  $1 - \alpha\mu \geq 1 - \alpha L \geq 0$ , and hence (51) implies

$$1 - \alpha\mu \leq \begin{cases} \frac{2}{1 + \beta} & \text{if } \beta \leq 1 \\ \frac{1}{\beta} & \text{if } \beta \geq 1. \end{cases} \quad (52)$$

$$\quad (53)$$

Nevertheless, if  $\beta \leq 1$  then  $\frac{2}{1 + \beta} \geq 1$ , so the first case always holds. As a result, the above inequality simplifies to

$$(1 - \alpha\mu)\beta \leq 1.$$

- Now, assume  $\alpha \in [\frac{1}{L}, \min\{\frac{2}{L}, \frac{1}{\mu}\}]$ . Then  $1 - \alpha\mu \geq 0 \geq 1 - \alpha L$ , and thus (51) yields

$$1 - \alpha L \geq -\frac{1}{1 + 2\beta}, \quad 1 - \alpha\mu \leq \min\left\{\frac{1}{\beta}, \frac{2}{1 + \beta}\right\}$$

where the second inequality again simplifies to  $(1 - \alpha\mu)\beta \leq 1$ .

- The last possible case happens when  $\mu \geq \frac{L}{2}$ , and so  $\alpha \in [\frac{1}{\mu}, \frac{2}{L}]$  is possible. In this case  $1 - \alpha L \leq 1 - \alpha\mu \leq 0$ , and so using (51), we just need to check

$$1 - \alpha L \geq -\frac{1}{1 + 2\beta}$$

Considering all these cases along with the fact that (51) shows  $\alpha$  cannot be greater than  $\frac{2}{L}$  completes the proof.

## F Proof of Proposition 3.5

*Proof.* Similar to the analysis for the gradient descent, we can assume without loss of generality that  $Q$  is diagonal. The proof is also similar. Consider  $U\Lambda U^T$  be the eigenvalue decomposition of  $Q$ . Then  $A_Q$  in (23) can be written as

$$A_Q = \tilde{U} A_\Lambda \tilde{U}^T \quad (54)$$

where

$$\tilde{U} = \begin{bmatrix} U & 0_d \\ 0_d & U \end{bmatrix}, A_\Lambda = \begin{bmatrix} (1 + \beta)(I_d - \alpha\Lambda) & -\beta(I_d - \alpha\Lambda) \\ I_d & 0_d \end{bmatrix} \quad (55)$$

Replacing  $A_Q$  from (54) in Lyapunov equation (13) implies

$$\tilde{U} A_\Lambda \tilde{U}^T X \tilde{U} A_\Lambda^T \tilde{U}^T - X + BB^T = 0. \quad (56)$$

Multiplying by  $\tilde{U}$  and  $\tilde{U}^T$  from right and left respectively yields

$$A_\Lambda \tilde{U}^T X \tilde{U} A_\Lambda^T - \tilde{U}^T X \tilde{U} + BB^T = 0 \quad (57)$$

where we used the fact that  $B$  in (6) has the property that  $\tilde{U}^T BB^T \tilde{U} = BB^T$ . Equation (57) shows that  $\tilde{U}^T X \tilde{U}$  satisfies the Lyapunov equation when  $A_Q$  is replaced by  $A_\Lambda$ , and therefore the  $H_2$  norm after substituting  $Q$  by  $\Lambda$  would be equal to

$$\begin{aligned} \sqrt{\text{Tr}(C\tilde{U}^T X \tilde{U} C^T)} &= \sqrt{\text{Tr}(\tilde{U} C^T C \tilde{U}^T X)} \\ &= \sqrt{\text{Tr}(C^T C X)} \\ &= \sqrt{\text{Tr}(C X C^T)} \end{aligned}$$

where the first and last equality come from cyclic property of the trace, and the second equality is true for  $C$  in (6). This result basically completes the proof of our claim that we can assume  $Q$  is diagonal. For simplicity we will continue our analysis with  $A_Q$ , assuming its a diagonal matrix.

Let  $P_\pi$  be the permutation matrix associated with the permutation  $\pi$  over  $\{1, 2, \dots, 2d\}$  that satisfies

$$\pi(i) = \begin{cases} 2i - 1 & \text{if } 1 \leq i \leq d \\ 2(i - d) & \text{if } d + 1 \leq i \leq 2d. \end{cases} \quad (58)$$

$$(59)$$

Also, it is well-known that permutation matrices satisfy

$$P_\pi^{-1} = P_\pi^T = P_{\pi^{-1}}. \quad (60)$$

Multiplying Lyapunov equation (13) by  $P_\pi$  and  $P_\pi^T$  from left and right respectively, along with using (60) leads to

$$(P_\pi A_Q P_\pi^T) Y (P_\pi A_Q^T P_\pi^T) - Y + P_\pi BB^T P_\pi^T = 0 \quad (61)$$

where  $Y = P_\pi X P_\pi^T$ . Considering (23),  $P_\pi A_Q P_\pi^T$  is a block diagonal matrix in the following form

$$\begin{bmatrix} T_1 & 0_d & \dots & 0_d \\ 0_d & T_2 & \dots & 0_d \\ \vdots & \vdots & \ddots & \vdots \\ 0_d & 0_d & \dots & T_d \end{bmatrix}$$

where

$$T_i = \begin{bmatrix} (1 + \beta)(1 - \alpha\lambda_i) & -\beta(1 - \alpha\lambda_i) \\ 1 & 0 \end{bmatrix}$$

where  $0 < \mu = \lambda_1 \leq \lambda_2 \leq \dots \leq \lambda_d = L$  are eigenvalues of  $Q$ . Also,

$$BB^T = \begin{bmatrix} \alpha^2 I_d & 0_d \\ 0_d & 0_d \end{bmatrix}$$

, and so  $P_\pi B B^T P_\pi^T$  is a  $2d$  times  $2d$  diagonal matrix with  $\alpha^2$  on entries  $(1, 1), (3, 3), \dots, (2d-1, 2d-1)$  and zero elsewhere. Hence,  $Y$  as the solution of (61), would be a block diagonal matrix in the form of

$$\begin{bmatrix} Y_1 & 0_d & \dots & 0_d \\ 0_d & Y_2 & \dots & 0_d \\ \vdots & \vdots & \ddots & \vdots \\ 0_d & 0_d & \dots & Y_d \end{bmatrix}$$

where

$$Y_i = \begin{bmatrix} y_i^u & y_i^o \\ y_i^o & y_i^d \end{bmatrix}$$

satisfies the equality

$$\begin{bmatrix} (1+\beta)(1-\alpha\lambda_i) & -\beta(1-\alpha\lambda_i) \\ 1 & 0 \end{bmatrix} \begin{bmatrix} y_i^u & y_i^o \\ y_i^o & y_i^d \end{bmatrix} \begin{bmatrix} (1+\beta)(1-\alpha\lambda_i) & 1 \\ -\beta(1-\alpha\lambda_i) & 0 \end{bmatrix} - \begin{bmatrix} y_i^u & y_i^o \\ y_i^o & y_i^d \end{bmatrix} + \begin{bmatrix} \alpha^2 & 0 \\ 0 & 0 \end{bmatrix} = 0$$

We observe that this leads to linear equations to solve for  $y_i^u, y_i^o$  and  $y_i^d$ :

$$\begin{bmatrix} (1+\beta)^2(1-\alpha\lambda_i)^2 - 1 & -2\beta(1+\beta)(1-\alpha\lambda_i)^2 & \beta^2(1-\alpha\lambda_i)^2 \\ (1+\beta)(1-\alpha\lambda_i) & -1-\beta(1-\alpha\lambda_i) & 0 \\ 1 & 0 & -1 \end{bmatrix} \begin{bmatrix} y_i^u \\ y_i^o \\ y_i^d \end{bmatrix} = \begin{bmatrix} -\alpha^2 \\ 0 \\ 0 \end{bmatrix}.$$

Solving this system of equations, we obtain:

$$\begin{aligned} y_i^u &= y_i^d = \alpha^2 \frac{1 + \beta(1 - \alpha\lambda_i)}{\alpha\lambda_i(1 - \beta(1 - \alpha\lambda_i))(2 + 2\beta - \alpha\lambda_i(1 + 2\beta))} \\ y_i^o &= \frac{\alpha^2(1 + \beta)(1 - \alpha\lambda_i)}{\alpha\lambda_i(1 - \beta(1 - \alpha\lambda_i))(2 + 2\beta - \alpha\lambda_i(1 + 2\beta))}. \end{aligned} \quad (62)$$

And now, the  $H_2$  norm can be computed using

$$\begin{aligned} \text{Tr}(CXC^T) &= \text{Tr}(CP_\pi^T Y P_\pi C^T) \\ &= \text{Tr}(P_\pi C^T C P_\pi^T Y). \end{aligned} \quad (63)$$

The matrix  $P_\pi C^T C P_\pi^T$  is block diagonal with two times two matrices

$$\begin{bmatrix} (1+\beta)^2 & -\beta(1+\beta) \\ -\beta(1+\beta) & \beta^2 \end{bmatrix}$$

on its diagonal. Therefore, using (62), the  $H_2$  norm is equal to

$$\begin{aligned} &\sum_{i=1}^d \alpha^2 \frac{(\beta^2 + (1+\beta)^2)(1+\beta(1-\alpha\lambda_i)) - 2\beta(1+\beta)^2(1-\alpha\lambda_i)}{\alpha\lambda_i(1-\beta(1-\alpha\lambda_i))(2+2\beta-\alpha\lambda_i(1+2\beta))} \\ &= \sum_{i=1}^d \alpha^2 \frac{1+\beta+\alpha\lambda_i\beta(2\beta+1)}{\alpha\lambda_i(1-\beta(1-\alpha\lambda_i))(2+2\beta-\alpha\lambda_i(1+2\beta))} \end{aligned} \quad (64)$$

which completes the proof.  $\square$

## G Convexity of $u_{\alpha,\beta}(\lambda)$

We next show that the function  $u_{\alpha,\beta}(\lambda)$  appearing in the definition of the  $H_2$  norm of the dynamical system representation of the AG algorithm is convex with respect to  $\lambda$ .

**Lemma G.1.** *Let  $(\alpha, \beta) \in \mathcal{S}$  where  $\mathcal{S}$  is the stability region of the dynamical system representation of AG given by (25). The function  $u_{\alpha,\beta}(\lambda)$  defined by (30) is convex on the interval  $[\mu, L]$ .*

*Proof.* Since the  $H_2$  norm is finite on  $\mathcal{S}$ , the term  $u_{\alpha,\beta}(\lambda)$  appearing in the characterization of the  $H_2$  norm is finite for every  $\lambda \in [\mu, L]$  as well. In particular the function  $u_{\alpha,\beta}$  is differentiable up to arbitrary order with respect to both  $\alpha$  and  $\beta$  for any  $(\alpha, \beta) \in \mathcal{S}$ . Let  $(\alpha, \beta) \in \mathcal{S}$  is given. We

necessarily have  $\alpha > 0$ , otherwise if the stepsize  $\alpha = 0$ , the AG iterations would clearly not be convergent. Assume that  $\beta > 0$ , the case  $\beta = 0$  can be treated similarly. The function  $u_{\alpha,\beta}$  is of the form

$$u_{\alpha,\beta}(\lambda) = u_{\alpha,\beta}^{(1)}(\lambda) + u_{\alpha,\beta}^{(2)}(\lambda)$$

where

$$u_{\alpha,\beta}^{(1)}(\lambda) = \frac{\alpha(1+\beta)}{\lambda(1-\beta(1-\alpha\lambda))(2+2\beta-\alpha\lambda(1+2\beta))},$$

and

$$u_{\alpha,\beta}^{(2)}(\lambda) = \frac{\alpha^2\beta(2\beta+1)}{(1-\beta(1-\alpha\lambda))(2+2\beta-\alpha\lambda(1+2\beta))}.$$

It suffices to show that  $u_{\alpha,\beta}^{(1)}(\lambda)$  and  $u_{\alpha,\beta}^{(2)}(\lambda)$  are convex over  $[\mu, L]$  as a function of  $\lambda$  which will imply that the sum  $u_{\alpha,\beta}(\lambda)$  is convex. For this purpose,

$$u_{\alpha,\beta}^{(1)}(\lambda) = \frac{a_1}{\lambda(\lambda+b_1)(c_1-\lambda)}, \quad u_{\alpha,\beta}^{(2)}(\lambda) = \frac{a_2}{(\lambda+b_2)(c_2-\lambda)},$$

for some constants  $a_i, b_i$  and  $c_i$  that depend only on  $\alpha$  and  $\beta$  for  $i = 1, 2$ . In particular,  $a_1, a_2, c_1, c_2 > 0$  and due to the positivity of the  $H_2$  norm, the quadratic terms in the denominators satisfy

$$q_i(\lambda) = (\lambda + b_i)(c_i - \lambda) > 0 \quad \text{for any } \lambda \in [\mu, L], \quad \text{and } i = 1, 2.$$

where  $q_i$  is concave. We also have

$$\log u_{\alpha,\beta}^{(1)}(\lambda) = \log(a_1) - (\log \lambda + \log q_1(\lambda)) \tag{65}$$

$$\log u_{\alpha,\beta}^{(2)}(\lambda) = \log(a_2) - \log q_2(\lambda). \tag{66}$$

The function  $x \mapsto \log(x)$  is concave and is increasing over  $(0, \infty)$ . Since the composition of a concave, non-decreasing function with a concave function is concave in dimension one; the functions  $\log \lambda + \log q_1(\lambda)$  and  $\log q_2(\lambda)$  are both concave. Then, it from the above formulas (65)-(66) that both  $u_{\alpha,\beta}^{(1)}(\lambda)$  and  $u_{\alpha,\beta}^{(2)}(\lambda)$  are log-concave. This implies that both functions are convex and their sum  $u_{\alpha,\beta}(\lambda)$  is convex. This completes the proof.  $\square$

## H Tractable upper bounds for the objective $F_\tau(\alpha, \beta)$ in arbitrary dimensions

Although computing the optimizer to (24) is easy to do in dimension  $d = 2$ , for higher dimensions, this is computationally challenging; because the optimizer depends on the eigenvalues  $\lambda_i$  of  $Q$  that are not known in general. Computing all the eigenvalues can be as expensive as optimizing the objective function  $f$ . Therefore, we next develop computationally tractable upper bounds to the objective  $F_\tau(\alpha, \beta)$  that depends only on  $\mu$  and  $L$  instead of all the eigenvalues  $\{\lambda_i\}_{i=1}^d$ .

First, we show in Lemma G.1 that the function  $u_{\alpha,\beta}(\lambda)$  defined in (30) is convex in  $\lambda \in [\mu, L]$  for fixed  $(\alpha, \beta) \in \mathcal{S}$ . Therefore, its maximum is attained at one of the ends of this interval, i.e.

$$u_{\alpha,\beta}(\lambda) \leq \bar{u}_{\alpha,\beta} := \max[u_{\alpha,\beta}(\mu), u_{\alpha,\beta}(L)] \quad \text{for } \lambda \in [\mu, L].$$

Substituting this upper bound in (29) and (24) leads to the following upper bound on the  $H_2$  norm

$$H_2^2(\alpha, \beta) \leq \bar{H}_2^2(\alpha, \beta) := d\bar{u}_{\alpha,\beta} \tag{67}$$

and the relaxed optimization problem

$$(\alpha_*, \beta_*) := \arg \min_{(\alpha,\beta) \in \mathcal{S}} \bar{F}_\tau(\alpha, \beta) := \bar{H}_2^2(\alpha, \beta) + \tau \frac{1}{1 - \rho(\alpha, \beta)^2}. \tag{68}$$

This objective is differentiable everywhere in the interior of the stability region  $\mathcal{S}$  except when the first term is not differentiable when  $u_{\alpha,\beta}(\mu) = u_{\alpha,\beta}(L)$  and the second term is not differentiable when  $\rho_\mu = \rho_L$  or  $\Delta_\mu = 0$  or  $\Delta_L = 0$ . Furthermore, following a similar approach to Example 3.2, the first order optimality conditions with respect to  $\alpha$  and  $\beta$  results in low-order polynomials (that are

independent from the dimension  $d$ ) which can be solved efficiently up to any accuracy. Then, after checking a finite set of non-differentiable points; the optimization problem has the computational complexity equivalent to computing the roots of a polynomial of a small order which is cheap to compute in high dimensions. Our results reported in the numerical experiments section show that this upper bound results in good performance in terms of trading the speed and the robustness of an algorithm.