

Preconditioned Nesterov's Accelerated Gradient Descent Method and Its Applications to Nonlinear PDE

A Dissertation Presented for the
Doctor of Philosophy
Degree

The University of Tennessee, Knoxville

Jea Hyun Park

August 2021

Acknowledgments

The content of this dissertation is due to collaborative work with my advisors, Dr. Steven Wise and Dr. Abner Salgado. Thus, herein, the pronoun *we/our/us* will be used as I present our findings.

It is obvious that these acknowledgments would be endless if I were to list all people and things I wish to express gratitude to. After reviewing my draft of more than three pages for acknowledgments and still feeling it was missing many others, I have decided to make it brief and simple by focusing on people who are directly related to this work. I never knew that writing the acknowledgment portion was like numerical implementation: our lives are truly infinite dimensional, but data to record them must be of finite dimension. First of all, I thank my advisors, Dr. Wise and Dr. Salgado. They have been good friends as well as academic advisors. I also thank my committee members, Dr. Xiaobing Feng and Dr. Arun Padakandla. They were willing to spend time to make my doctoral work more polished. Of course, I would not be able to finish, let alone enjoy, my graduate study without my family.

Abstract

We develop a theoretical foundation for the application of Nesterov’s accelerated gradient descent method (AGD) to the approximation of solutions of a wide class of partial differential equations (PDEs). This is achieved by proving the existence of an *invariant set* and exponential convergence rates when its preconditioned version (PAGD) is applied to minimize *locally Lipschitz* smooth, strongly convex objective functionals. We introduce a second-order ordinary differential equation (ODE) with a preconditioner built-in and show that PAGD is an explicit time-discretization of this ODE, which requires a natural time step restriction for energy stability. At the continuous time level, we show an exponential convergence of the ODE solution to its steady state using a simple energy argument. At the discrete level, assuming the aforementioned step size restriction, the existence of an invariant set is proved and a matching exponential rate of convergence of the PAGD scheme is derived by mimicking the energy argument and the convergence at the continuous level. Applications of the PAGD method to numerical PDEs are demonstrated with certain nonlinear elliptic PDEs using pseudo-spectral methods for spatial discretization, and several numerical experiments are conducted. The results confirm the global geometric and mesh size-independent convergence of the PAGD method, with an accelerated rate that is improved over the preconditioned gradient descent (PGD) method.

Table of Contents

1	Introduction	1
1.1	Problem of interest	1
1.2	Motivations	1
1.3	Related works	4
1.4	Contributions	7
1.5	Organization	9
2	Preliminaries	11
2.1	Preconditioning	11
3	Optimization Schemes	18
3.1	Preconditioned gradient descent method (PGD)	18
3.2	Preconditioned Nesterov’s accelerated gradient descent method (PAGD)	19
4	ODE Model	21
4.1	The limiting ODE	21
4.2	Analysis of the IVP	22
4.3	Convergence to equilibrium	25
4.4	Derivation of the ODE	28
4.5	Derivation of the initial conditions	29
4.6	PAGD as a discretization of the IVP	30
5	An Energy Approach to Convergence of PAGD	32
5.1	Existence of an invariant set	32

5.2	Convergence of PAGD	38
6	Numerical Experiments	46
6.1	The continuous problem	47
6.2	Discretization	51
6.3	A problem with a manufactured solution	56
6.4	A problem where the solution is unknown	58
6.5	A comparison between PGD and PAGD	63
	Bibliography	65
	Appendix	72
A	Literature comparison	73
	Vita	74

List of Tables

6.1	Performance comparison between PGD and PAGD	62
2	A comparison of recent works from a numerical PDE point of view.	73

List of Figures

4.1	A rolling ball system.	23
6.1	Objective, error, and energy decay plots for GD, AGD, PGD, and PAGD. . .	57
6.2	∞ -norm plots of the search directions for GD, AGD, PGD, and PAGD. . . .	60
6.3	Number of iterations for ∞ -norm of the search directions to reach the tolerance 10^{-3} for GD, AGD, PGD, and PAGD.	61

Chapter 1

Introduction

1.1 Problem of interest

The purpose of this work is to broaden the context in which a well-known and efficient algorithm for unconstrained convex minimization, the so-called *Nesterov's accelerated gradient descent* (AGD) scheme, can be utilized and, further, to shed some light on its convergence properties. This method iteratively finds approximations to the solution of the following optimization problem: given $G : \mathbb{H} \rightarrow \mathbb{R}$, find

$$x^* = \operatorname{argmin} \{G(x) \mid x \in \mathbb{H}\}.$$

Here, and in what follows, $\mathbb{H}$ is a real, separable Hilbert space with inner product $(\cdot, \cdot)_{\mathbb{H}}$ and the so-called *objective* functional, G , is assumed to be strongly convex and *locally Lipschitz* smooth; see Chapter 2 for definitions and notation. We immediately comment that the assumptions on the objective guarantee the existence and uniqueness of a minimizer (e.g., [15, Theorem 7.4-4, Theorem 8.2-2]).

1.2 Motivations

Convex minimization is ubiquitous, and our main interest in this problem comes from the fact that many important nonlinear partial differential equations (PDEs) can be viewed as

the Euler equations of certain convex objective functions. For example, the classical minimal surface problem (see [18]) and the p -Laplacian equation (see [5]) have this structure, just to name a few. However, the method discussed in this work is particularly powerful for semilinear PDEs. For example, time discretizations of many important models in materials science often end up involving such problems, e.g., via the convex splitting technique (see [19]). In a related context, the current explosion of interest in statistical learning has drawn the attention of practitioners to so-called first order schemes, i.e., those that only require knowledge of first order derivatives, which is suitable in dealing with large data sets. These considerations are important for solving nonlinear PDE as well. One of the main thrusts of this research is to show that Nesterov's accelerated schemes, which are popular in statistical learning, can be utilized as fast solvers for nonlinear PDE once we resolve some nontrivial, technical difficulties specific to such problems.

The first and most naïve approach to find x^* would be to appeal directly to the first order necessary (and, in this context, sufficient) optimality condition, namely, the Euler equation:

$$G'(x^*) = 0, \tag{1.2.1}$$

where G' denotes the Fréchet derivative of G . In the examples that we have in mind, however, this requires the simultaneous solution to a very large number of nonlinear equations, and the direct solution of the system is not feasible in practice. Other approaches better suited for minimization must be constructed. According to [8], iterative methods for minimizing functionals date back in 1847 when Cauchy proposed the so-called gradient descent method (GD). The solution to (1.2.1) can be seen as the steady state of the *gradient flow*

$$\begin{cases} X(0) = x_0, \\ \dot{X}(t) = -G'(X(t)), \quad (t > 0). \end{cases} \tag{1.2.2}$$

Here $x_0 \in \mathbb{H}$ is arbitrary and, in the second equation, we are implicitly identifying the dual space of $\mathbb{H}$, denoted by $\mathbb{H}'$, with $\mathbb{H}$ itself. Under the assumptions we have imposed on the objective G , it is possible to show that this flow satisfies $X(t) \rightarrow x^*$ as $t \rightarrow \infty$, see [48, Theorem 2.4]. The idea of GD is to approximate the solution to this flow via a *forward Euler*

time discretization with a fixed step size s : given $x_0 \in \mathbb{H}$, for $k \geq 0$, find x_{k+1} satisfying

$$x_{k+1} = x_k - sG'(x_k). \quad (1.2.3)$$

While this idea seems straightforward, more in-depth discussions on this method started only in the 1960s, where some practical step size rules and convergence analyses were established. It was shown that if the objective functional is convex and Lipschitz smooth, then GD converges to the minimizer, x^* , and it exhibits a first order rate of convergence in the objective. Here, and in what follows, by an n^{th} order (*algebraic*) convergence in the objective, we mean that $G(x_k) - G^* \leq O(1/k^n)$, as $k \rightarrow \infty$, where $G^* = G(x^*)$ is the minimum of G . By an *exponential* or a *geometric convergence* in the objective we mean that $G(x_k) - G^* \leq O(r^k)$, as $k \rightarrow \infty$, for some $r \in (0, 1)$. In the latter case, we call r the rate of (exponential) convergence. It can further be shown that, if the objective is, in addition, strongly convex, then the rate of convergence is exponential, and that it matches the rate of convergence of the solution of (1.2.2) to x^* . (See [33, Theorem 2.1.15] or Remark 5.2.7). Some physical intuition for the evolution of the solution to (1.2.2) is provided in Chapter 4. See, in particular, Remark 4.1.1.

Evidently, all considerations regarding convergence are subject to the norm $\|\cdot\|_{\mathbb{H}}$. It is possible to improve the convergence rate by using an equivalent norm, through which the level sets of the objective G look “more circular.” In the numerical linear algebra and numerical PDE communities, this is commonly known as *preconditioning*. In the context of (1.2.2) and GD, this is achieved by introducing an operator $\mathcal{L} : \mathbb{H} \rightarrow \mathbb{H}'$ and considering the evolution of

$$\dot{X}(t) = -\mathcal{L}^{-1}G'(X(t)).$$

Notice that we no longer implicitly identify $\mathbb{H}'$ with $\mathbb{H}$. The time-discrete counterpart of (1.2.3) is known as the *preconditioned gradient descent method* (PGD) and is as follows: given $x_0 \in \mathbb{H}$, for $k \geq 0$, find x_{k+1} such that

$$x_{k+1} = x_k - s\mathcal{L}^{-1}G'(x_k).$$

If the preconditioner is suitably chosen, then the convergence rate of GD can be substantially improved (see [20]). Note that we will tacitly assume in the sequel that $\mathcal{L}$ is independent of the iteration index k . We remark that Newton’s method may be viewed as a kind of *generalized* preconditioned gradient descent method if we assume that G is twice Fréchet differentiable and allow for the possibility that the preconditioner can change at each iteration. In particular, Newton’s method is expressed as

$$G''(x_k)(x_{k+1} - x_k) = -G'(x_k) =: r_k,$$

where $G''(x_k)$ is the second Fréchet derivative of G , and r_k is the so-called residual. Then, Newton’s method is a generalized preconditioned gradient descent method for which the preconditioner satisfies $s\mathcal{L}_k = G''(x_k)$. One of the difficulties with Newton’s method is that the preconditioner constantly changes, in general, and must be recomputed and re-inverted at each iteration step, which can prove quite costly. Furthermore, G'' may not exist in all applications of interest. Indeed, in the sequel, we will not assume that G'' exists.

1.3 Related works

To improve the convergence rate of GD, Nesterov [32] suggested a scheme that accelerates the GD method. For convex and Lipschitz smooth objectives the Nesterov’s accelerated gradient descent (AGD) scheme achieves a second order convergence rate. Later, he showed that if the objective is, in addition, strongly convex, then AGD achieves a faster exponential convergence rate than GD (see [33, Theorem 2.1.15]).

However, while the GD scheme has a strong physical intuition behind it, it is not completely clear what mechanism is at play to provide an acceleration in the AGD scheme. Some attempts have been made to understand this in the literature. Attouch et al. [4] studied the asymptotic behavior of the solution to a heavy ball system (similar to (4.1.1)) and showed their convergence to minimizers (if they exist) of locally Lipschitz objectives that are bounded below at the continuous time level. Goudou and Munier [24] looked into a similar system with quasiconvex, locally Lipschitz objectives at the continuous time

level and the convergence of its implicit discretization, what they call *proximal inertial algorithm*. Apparently, [46] is the first work that explains the acceleration happening in AGD both quantitatively and intuitively and inspired many researchers including us. For convex, Lipschitz smooth objectives, they were able to show that the solutions to a second order ODE

$$\ddot{X} + \frac{3}{t}\dot{X} + \nabla G(X) = 0$$

converges to the set of minimizers of G quadratically fast, as $t \rightarrow \infty$, and a matching, discrete convergence rate was established for a version of AGD. Wibisono et al. [48] took a similar approach in more generality in the language of *Bregman Lagrangian flow*. However, these two works did not explain the exponential acceleration for the strongly convex objectives. This limitation was one of the motivations of our work and removing it is one of the goals of this thesis. Recently, there have appeared more works that address the same issue and provide more general or unifying frameworks. For (globally) L -smooth, μ -strongly convex objectives (see Chapter 2 for definitions and notation), the best known convergence rates of AGD and the associated ODE are $G(x_k) - G^* \leq O((1 - \sqrt{\mu/L})^k)$ as $k \rightarrow \infty$ and $G(X(t)) - G^* \leq O(e^{-\sqrt{\mu}t})$ as $t \rightarrow \infty$ respectively. Similarly, for (globally) L -smooth, convex objectives, $G(x_k) - G^* \leq O(1/k^2)$ as $k \rightarrow \infty$ and $G(X(t)) - G^* \leq O(1/t^2)$ as $t \rightarrow \infty$ are the best known convergence rates respectively. Within the same framework as in [48], but using a different Lyapunov function, Wilson et al. [49] showed the best convergence rates for both convex and strongly convex cases at both continuous and discrete level. Shi et al. [44] looked into what they call *high-resolution ODE* and provided a finer understanding about the momentum-type schemes and obtained similar results. The convergence rates for the strongly convex case that they derived were not the best. However, they were able to explain the difference in the performance of *Polyak's momentum method* and AGD at the continuous level. Siegel [45] also analyzed a system of ODEs to study a version of AGD and obtained the best known rate of convergence for the strongly convex case at both continuous and discrete levels. He also studied non-smooth but still strongly convex objectives and stochastic versions. Luo and Chen [31] obtained the same best convergence rates for all the four cases mentioned above using a single ODE system but using a time rescaling argument when

dealing with the convex case. Laborde and Oberman [30] studied perturbed ODE systems and the corresponding version of AGD in the stochastic framework. As a byproduct, they obtained the same best convergence rates for the perturbed version of AGD with strongly convex objectives in the deterministic setting. There are also related works from a PDE point of view rather than convex optimization. Schaeffer and Hou [40] studied accelerated methods for nonlinear elliptic operators, which may not have a variational structure, i.e., the PDE may not have an appropriate objective. They also proposed similar methods for viscosity solutions. Benyamin et al. [7] and Calder and Yezzi [11] studied *PDE accelerations* that are similar to [48] in spirit and applied them to image processing and minimal surface obstacle problems respectively.

It is worth discussing *nonlinear conjugate gradient methods* (nonlinear CG) as these have the same objective as the main discussion of this thesis (unconstrained, non-quadratic minimization). In fact, in certain settings, the conjugate gradient method is equivalent to Polyak’s momentum method (see [9]). The *conjugate gradient method* (CG) was first proposed by Hestenes and Stiefel [28] (1952) as a direct method to solve a system of linear equations, which is, in turn, an explicit and practical version of the *conjugate direction method* proposed by Schmidt [41] (1908) and reinvented by Fox et al. [23] (1948). However, the CG does not solve a system of linear equation in finite steps in practice due to round-off errors. It became popular only after Reid [39] (1971) showed its value as an iterative method for large, sparse matrices. Extensions of the CG to nonlinear systems have been studied by many researchers. Given an objective function $G : \mathbb{R}^d \rightarrow \mathbb{R}$, each of the nonlinear CGs has the form ($k \geq 0$)

$$\begin{aligned} d_{(0)} &= -G'(x_0), & (\text{initial guess and the steepest descent direction}) \\ x_{k+1} &= x_k + s_k d_k, & (\text{exact or some inexact line search}) \\ d_{k+1} &= -G'(x_{k+1}) + \lambda_k d_k, & (\text{new search direction}) \end{aligned}$$

and a different formula for λ_k gives rise to a different nonlinear CG scheme. (λ_k is related to the Gram-Schmidt process to find a new search direction that is conjugate to the previous search directions in the linear case; see [43] for more detail about the CG for linear systems.)

They are all equivalent to each other when applied to a strongly convex, quadratic objective function. Fletcher and Reeves [22] (1964; $\lambda_k = \frac{\|G'(x_{k+1})\|_{\ell^2}}{\|G'(x_k)\|_{\ell^2}}$) are considered to have carried out such an attempt for the first time, whose method is called the *Fletcher-Reeves CG* in the rest of the section. Its convergence result is established in [51, 38] (under the exact line search) and [2] (under an inexact line search). In the literature, the scheme proposed by Polak and Ribiere [35] (1969; $\lambda_k = \frac{G'(x_{k+1})^T(G'(x_{k+1}) - G'(x_k))}{\|G'(x_k)\|_{\ell^2}}$), which we call the *Polak-Ribiere CG*, is said to be the best option in most cases among the simple versions of the CG excluding more fine-tuned modifications. They also showed the convergence of the method for strongly convex functions when the exact line search is used. Later, more convergence proofs followed under different assumptions such as inexact line searches and various versions of so-called *Wolfe conditions* (see, e.g. [50, 37, 25]). Though there are many other versions of the CG, I conclude this section by briefly mentioning efforts to take advantage of the previous two types of the CG, which are called the *hybrid CG* methods. According to [27], the Fletcher-Reeves CG often slows down as the iteration proceeds as pointed out in [37], but it has a better convergence theory than the Polak-Ribiere CG. On the other hand, while the Polak-Ribiere CG requires stronger assumptions for convergence, it performs better than the Fletcher-Reeves CG in practice. Intuitively, the Polak-Ribiere CG bears a “restarting feature” in itself when the progress gets slow (see [21, p. 86] and [27]). As an example of the hybrid CGs, in [47], it is proposed to take $\lambda_k = \lambda_K^{PR}$ if $0 \leq \lambda_k^{PR} \leq \lambda_k^{FR}$ or $\lambda_k = \lambda_K^{FR}$ otherwise, where λ_K^{FR} and λ_K^{PR} are the λ_k computed by the Fletcher-Reeves CG and Polak-Ribiere CG, respectively. In [27], it is reported that today’s best performing CGs are among the hybrid CGs. For example, a version of the hybrid CG proposed by Hager and Zhang [26] and that by Dai and Yuan [17] showed top performances in a certain numerical experiment. For more information about many more versions of the CG, we refer the readers to [27].

1.4 Contributions

The work contained herein includes the following important contributions, in particular, from a numerical PDE point of view, which can be seen more clearly from the literature comparison of Appendix A. See Table 2.

1. We prove all of our results under the more general assumption that the objective functional is *locally Lipschitz* smooth. Almost all earlier works assume that the objective is *globally Lipschitz* smooth (see e.g., [33, 48, 46, 3, 49, 30, 40, 31, 45, 44]). This is too restrictive to approximate solutions of nonlinear PDEs. If the objective functional associated with the PDE of interest grows just a bit faster than quadratic functionals (i.e., those of very mild nonlinearity), it violates the global Lipschitz condition and is beyond the theoretical guarantee. On the other hand, the local condition does not require anything outside of a certain bounded set so that many more nonlinear PDEs can be dealt with. Only a few works from a dynamical system point of view (e.g., [4, 24]) assume local Lipschitz condition. However, those works address only convergence itself at the continuous time level and do not discuss discrete level analysis. To the best of our knowledge, this is the first work that provides convergence rates under the local Lipschitz smoothness assumption at the continuous or discrete level.
2. We prove the existence of an invariant set $\mathfrak{B}$ of the PAGD method. That is, every sequence generated by the scheme stays in a certain bounded set. The local Lipschitz assumption is meaningful when it is furnished with an invariant set so that we have no restriction in exploiting the Lipschitz condition. We emphasize that this is not a trivial technicality. Unlike the gradient descent method, the accelerated methods are not descent methods. In fact, they oscillate. Thus, a simple sublevel set argument does not work. Even worse, they involve extrapolations of the main iterates. Consequently, a naïve attempt to obtain an invariant set leads to an impasse: to control the extrapolations, one wants to use the Lipschitz condition, but under the local Lipschitz condition, one cannot use it before proving that they are in a fixed bounded set. Again, to the best of our knowledge, our work is the first that addresses and resolves this issue.
3. We provide a detailed discrete analysis for a nonlinear PDE. All existing works mentioned before either did not discuss numerical examples or did not explain how concrete numerical examples fit the abstract framework, and they omitted whether their numerical examples satisfy the assumptions that they imposed. In contrast, we show that our examples satisfy all the necessary assumptions.

4. We provide an intuitive explanation for the acceleration mechanism behind AGD for strongly convex and locally Lipschitz smooth objectives. Inspired by [46], we view AGD as a discretization of a certain second order ordinary differential equation (ODE) — we present how to discretize this ODE to obtain AGD — and show that the solution to this ODE converges exponentially fast to its stationary point, which is the minimizer of G . We also provide an energy based proof of the exponential convergence rate of PAGD. This proof mimics the analysis of the continuous counterpart that is previously developed and shows what dissipation mechanisms are at play to achieve the aforementioned acceleration. We also show that the rates of convergence of the ODE model and AGD match and those rates at the continuous and discrete level are both the best known rates.
5. We build a preconditioner into the problem itself (even at the continuous time level) to analyze the scheme with a preconditioner in an explicit way. This seems deceptively simple. After all, preconditioning is nothing but using a different norm, hence a numerical analysis in one norm implicitly suggests the possibility of a similar analysis in another norm. On the other hand, determining an effective preconditioning strategy is a nontrivial matter. Furthermore, it is our observation that preconditioning in the classical optimization setting, especially for problems related to data analysis and machine learning, is underutilized and is a potential growth area in the future. Likewise, preconditioning strategies related to spectral collocation methods applied to nonlinear PDE are uncommon, but, as we shall see, are effective and efficient solver tools.

1.5 Organization

This thesis is organized as follows. In Chapter 2, we summarize the notation, assumptions, and main tools that we will use. In Chapter 3, we introduce several numerical schemes that are closely related to our discussion and summarize their convergence rates. In Chapter 4, we explore the connection between PAGD and a second order ODE and how this connection can help understand the acceleration behind PAGD intuitively. In Chapter 5, we prove the existence of an invariant set for the PAGD scheme and its exponential convergence. We take an ODE inspired approach, whose intuition lies in the developments of Chapter 4. In Chapter

6, we illustrate the application of the PAGD method to the solution of some numerical PDEs. These numerical experiments show the improvement in convergence by both acceleration and preconditioning. Finally, in the Appendix, a brief literature comparison can be found.

Chapter 2

Preliminaries

Let us begin by introducing the setting, assumptions, and some basic properties of the objects that we are interested in. By $\mathbb{H}$, we denote a real and separable Hilbert space with inner product $(\cdot, \cdot)_{\mathbb{H}}$ and associated norm $\|\cdot\|_{\mathbb{H}}$. Since we will use other inner products and norms on $\mathbb{H}$, for clarity, we will refer to $(\cdot, \cdot)_{\mathbb{H}}$ and $\|\cdot\|_{\mathbb{H}}$ as the *canonical inner product* and *canonical norm*, respectively. The dual of $\mathbb{H}$ is denoted by $\mathbb{H}'$. Its canonical operator norm is denoted by $\|\cdot\|_{\mathbb{H}'}$. For $x \in \mathbb{H}$ and $f \in \mathbb{H}'$, the symbol $\langle f, x \rangle$ represents their duality pairing, that is, $\langle f, x \rangle = f(x) \in \mathbb{R}$.

2.1 Preconditioning

A *preconditioner* is defined by a linear operator $\mathcal{L} : \mathbb{H} \rightarrow \mathbb{H}'$. Such an operator induces a bilinear form: for $x, y \in \mathbb{H}$,

$$(x, y)_{\mathcal{L}} = \langle \mathcal{L}x, y \rangle = \mathcal{L}[x](y). \quad (2.1.1)$$

We further assume that the bilinear form defined in (2.1.1) satisfies the following properties: there exist $C_1, C_2 > 0$ such that, for any $x, y \in \mathbb{H}$,

$$(L1) \quad (x, y)_{\mathcal{L}} = (y, x)_{\mathcal{L}}, \quad (2.1.2)$$

$$(L2) \quad (x, y)_{\mathcal{L}} \leq C_2 \|x\|_{\mathbb{H}} \|y\|_{\mathbb{H}}, \quad (2.1.3)$$

$$(L3) \quad C_1 \|x\|_{\mathbb{H}}^2 \leq (x, x)_{\mathcal{L}}. \quad (2.1.4)$$

Let us state some immediate, but important consequences.

Proposition 2.1.1 (properties of $\mathcal{L}$). *Let $\mathbb{H}$ be a Hilbert space with an underlying inner product $(\cdot, \cdot)_{\mathbb{H}}$, and let $\mathcal{L} : \mathbb{H} \rightarrow \mathbb{H}'$ be a linear mapping satisfying (L1)–(L3). Then, $\mathcal{L}$ is a bounded linear mapping that is one-to-one and onto, and it has a continuous inverse. Furthermore, $\mathcal{L}$ induces an inner product on $\mathbb{H}$ via (2.1.1) and another inner product on $\mathbb{H}'$ via*

$$(f, g)_{\mathcal{L}^{-1}} := \langle f, \mathcal{L}^{-1}g \rangle,$$

which we call $\mathcal{L}$ –inner product and $\mathcal{L}^{-1}$ –inner product, respectively. Their associated norms are defined in the usual way:

$$\begin{aligned} \|x\|_{\mathcal{L}} &:= \sqrt{(x, x)_{\mathcal{L}}} = \sqrt{\langle \mathcal{L}x, x \rangle} \\ \|f\|_{\mathcal{L}^{-1}} &:= \sqrt{(f, f)_{\mathcal{L}^{-1}}} = \sqrt{\langle f, \mathcal{L}^{-1}f \rangle}, \end{aligned} \quad (2.1.5)$$

which we call $\mathcal{L}$ –norm and $\mathcal{L}^{-1}$ –norm, respectively. In addition, the $\mathcal{L}$ –norm is equivalent to the canonical norm, and $\mathcal{L}^{-1}$ –norm is an operator norm on $\mathbb{H}'$ with respect to $\mathcal{L}$ –norm in the sense that

$$\|f\|_{\mathcal{L}^{-1}} = \|f\|_{\mathcal{L}^*} := \sup_{x \in \mathbb{H}} \frac{\langle f, x \rangle}{\|x\|_{\mathcal{L}}}. \quad (2.1.6)$$

Finally, $\mathcal{L}$ is an isometry with respect to $\mathcal{L}$ –induced norms in the sense that

$$\|\mathcal{L}^{-1}f\|_{\mathcal{L}} = \|f\|_{\mathcal{L}^{-1}} \text{ and } \|\mathcal{L}u\|_{\mathcal{L}^{-1}} = \|u\|_{\mathcal{L}}. \quad (2.1.7)$$

In other words, $\mathcal{L}^{-1} : \mathbb{H} \rightarrow \mathbb{H}'$ is the Riesz map with respect to $(\cdot, \cdot)_{\mathcal{L}}$.

Proof. (i) (L1) implies that the bilinear form $(\cdot, \cdot)_{\mathcal{L}}$ is symmetric. And from (L2) and (L3), we see

$$\sqrt{C_1}\|x\|_{\mathbb{H}} \leq \|x\|_{\mathcal{L}} \leq \sqrt{C_2}\|x\|_{\mathbb{H}} \quad \text{for all } x \in \mathbb{H}. \quad (2.1.8)$$

Consequently, (2.1.1) defines an inner product whose induced norm is equivalent to the underlying norm.

(ii) $\mathcal{L}$ is bounded since, by (L2), we have, for $x \in \mathbb{H}$,

$$\|\mathcal{L}x\|_{\mathbb{H}'} := \sup_{y \in \mathbb{H}} \frac{\langle \mathcal{L}x, y \rangle}{\|y\|_{\mathbb{H}}} \leq C_2\|x\|_{\mathbb{H}}. \quad (2.1.9)$$

$\mathcal{L}$ is one-to-one since it is linear and it satisfies

$$\mathcal{L}x = 0 \implies \langle \mathcal{L}x, x \rangle = 0 \implies \|x\|_{\mathbb{H}} = 0 \implies x = 0, \quad (2.1.10)$$

where we used (L3) in the second implication.

$\mathcal{L}$ is onto since the Riesz representation theorem with respect to $\mathcal{L}$ -inner product implies that, for any $f \in \mathbb{H}'$, there exists a (unique) $x_f \in \mathbb{H}$ such that

$$(x_f, y)_{\mathcal{L}} = \langle f, y \rangle \quad \forall y \in \mathbb{H}. \quad (2.1.11)$$

By the definition of the bilinear form $(\cdot, \cdot)_{\mathcal{L}}$,

$$\langle \mathcal{L}x_f, y \rangle = \langle f, y \rangle \quad \forall y \in \mathbb{H}, \quad (2.1.12)$$

which implies $\mathcal{L}x_f = f$. As a result, the assignment $x \mapsto \mathcal{L}x$ is onto. Thus, the open mapping theorem implies that $\mathcal{L}$ has a continuous inverse.

(iii) The linear operator $\mathcal{L}$ also induces an inner product and a norm on the dual space, $\mathbb{H}'$, i.e. $\mathcal{L}^{-1}$ -inner product and $\mathcal{L}^{-1}$ -norm via (2.1.5). The bilinear form $(\cdot, \cdot)_{\mathcal{L}^{-1}}$ is also

symmetric as inherited from $\mathcal{L}$ -inner product:

$$\begin{aligned}(f, g)_{\mathcal{L}^{-1}} &= \langle f, \mathcal{L}^{-1}g \rangle = \langle \mathcal{L}\mathcal{L}^{-1}f, \mathcal{L}^{-1}g \rangle = (\mathcal{L}^{-1}f, \mathcal{L}^{-1}g)_{\mathcal{L}} \\ &= (\mathcal{L}^{-1}g, \mathcal{L}^{-1}f)_{\mathcal{L}} = \langle \mathcal{L}\mathcal{L}^{-1}g, \mathcal{L}^{-1}f \rangle = \langle g, \mathcal{L}^{-1}f \rangle = (g, f)_{\mathcal{L}^{-1}}.\end{aligned}$$

It also satisfies the positive definiteness: for $f \in \mathbb{H}'$,

$$(f, f)_{\mathcal{L}^{-1}} = \langle f, \mathcal{L}^{-1}f \rangle = \langle \mathcal{L}\mathcal{L}^{-1}f, \mathcal{L}^{-1}f \rangle = (\mathcal{L}^{-1}f, \mathcal{L}^{-1}f)_{\mathcal{L}} \geq C_1 \|\mathcal{L}^{-1}f\|_{\mathbb{H}}^2 \geq 0,$$

and this inequality implies that if $(f, f)_{\mathcal{L}^{-1}} = 0$, then $\mathcal{L}^{-1}f = 0$, which in turn, shows $f = 0$ upon taking $\mathcal{L}$ on both sides.

The two norms are related in the following way:

$$\|\mathcal{L}^{-1}f\|_{\mathcal{L}} = \|f\|_{\mathcal{L}^{-1}},$$

which can be seen from: for $f \in \mathbb{H}'$,

$$\|\mathcal{L}^{-1}f\|_{\mathcal{L}}^2 = (\mathcal{L}^{-1}f, \mathcal{L}^{-1}f)_{\mathcal{L}} = \langle f, \mathcal{L}^{-1}f \rangle = \|f\|_{\mathcal{L}^{-1}}^2.$$

Similarly, for $x \in \mathbb{H}$,

$$\|\mathcal{L}x\|_{\mathcal{L}^{-1}} = \|x\|_{\mathcal{L}},$$

since

$$\|\mathcal{L}x\|_{\mathcal{L}^{-1}}^2 = (\mathcal{L}x, \mathcal{L}x)_{\mathcal{L}^{-1}} = \langle \mathcal{L}x, x \rangle = \|x\|_{\mathcal{L}}^2.$$

(iv) Finally, $\|f\|_{\mathcal{L}^{-1}}$ is an operator norm of $f \in \mathbb{H}'$ with respect to $\mathcal{L}$ -norm since, for $f \in \mathbb{H}'$ and for a nonzero element $x \in \mathbb{H}$,

$$\langle f, x \rangle = \langle \mathcal{L}\mathcal{L}^{-1}f, x \rangle = (\mathcal{L}^{-1}f, x)_{\mathcal{L}} \leq \|\mathcal{L}^{-1}f\|_{\mathcal{L}} \|x\|_{\mathcal{L}} = \|f\|_{\mathcal{L}^{-1}} \|x\|_{\mathcal{L}}.$$

This implies $\|f\|_{\mathcal{L}^{-1}} \geq \|f\|_{\mathcal{L}^*}$, where $\|f\|_{\mathcal{L}^*}$ denotes the standard operator norm of f as defined in (2.1.6). On the other hand, taking $x = \mathcal{L}^{-1}f$ in the definition of the operator norm, $\|f\|_{\mathcal{L}^*}$,

$$\|f\|_{\mathcal{L}^*} \geq \frac{\langle f, \mathcal{L}^{-1}f \rangle}{\|\mathcal{L}^{-1}f\|_{\mathcal{L}}} = \frac{\|f\|_{\mathcal{L}^{-1}}^2}{\|f\|_{\mathcal{L}^{-1}}} = \|f\|_{\mathcal{L}^{-1}}.$$

□

Remark 2.1.2 (no preconditioning). By setting $\mathcal{L} = \mathfrak{R}_{\mathbb{H}}^{-1}$, we can remove the preconditioning, where $\mathfrak{R}_{\mathbb{H}} : \mathbb{H}' \rightarrow \mathbb{H}$ is the canonical Riesz map. ■

Our objective $G : \mathbb{H} \rightarrow \mathbb{R}$ will be assumed to be Fréchet differentiable at every point in $\mathbb{H}$. We denote by $G'(x) \in \mathbb{H}'$ the Fréchet derivative of G at the point $x \in \mathbb{H}$. Since the definition of Fréchet differentiability involves a norm, the actual derivative is possibly norm dependent. The following result shows that, actually, the definition is invariant as long as the norms are equivalent.

Proposition 2.1.3 (equivalent norms). *Let $\mathbb{H}$ be a real and separable Hilbert space with norm $\|\cdot\|_{\mathbb{H}}$, and $G : \mathbb{H} \rightarrow \mathbb{R}$ be Fréchet differentiable at $x \in \mathbb{H}$. Assume that $\|\|\cdot\|\|_{\mathbb{H}}$ is another norm on $\mathbb{H}$. If $\|\|\cdot\|\|_{\mathbb{H}}$ is equivalent to $\|\cdot\|_{\mathbb{H}}$, then G is also Fréchet differentiable at x with respect to $\|\|\cdot\|\|_{\mathbb{H}}$. Furthermore, the derivatives coincide.*

Proof. Since the two norms are equivalent, there exist $C'_1, C'_2 > 0$ such that

$$C'_1 \|\|\cdot\|\|_{\mathbb{H}} \leq \|\cdot\|_{\mathbb{H}} \leq C'_2 \|\|\cdot\|\|_{\mathbb{H}} \quad \text{for all } x \in \mathbb{H}.$$

By definition of the Fréchet derivative, for $x \in \mathbb{H}$, there exists $A = A(x) \in \mathbb{H}'$ such that $G(x+y) - G(x) - \langle A(x), y \rangle \rightarrow 0$ as $\|y\|_{\mathbb{H}} \rightarrow 0$. Since $\|\|\cdot\|\|_{\mathbb{H}} \rightarrow 0$ implies $\|y\|_{\mathbb{H}} \rightarrow 0$, the whole expression tends to zero as $\|\|\cdot\|\|_{\mathbb{H}} \rightarrow 0$ as well with the same choice of $A \in \mathbb{H}'$. □

Notice that nothing is said about continuity in the previous statement. For convex functions, the continuity of the derivatives is automatic once the Fréchet differentiability is guaranteed (see [34, p. 20 Corollary]).

Proposition 2.1.4 (continuity). *Let $\mathbb{H}$ be a real and separable Hilbert space and $D \subset \mathbb{H}$ be open and convex. If $G : D \rightarrow \mathbb{R}$ is convex and Fréchet differentiable, then $x \mapsto G'(x)$ is norm continuous on D .*

The following two definitions provide a framework to describe the geometry of the graph of our objective functional.

Definition 2.1.5 (Lipschitz smoothness). Let $\mathbb{H}$ be a real and separable Hilbert space, and $G : \mathbb{H} \rightarrow \mathbb{R}$ be Fréchet differentiable at every point. We say that G is *locally Lipschitz smooth* (with respect to $\mathcal{L}$ -norm) iff, for every bounded, convex set $B \subset \mathbb{H}$, there exists a constant $L_B > 0$ such that

$$\langle G'(x) - G'(y), x - y \rangle \leq L_B \|y - x\|_{\mathcal{L}}^2 \quad \forall x, y \in B. \quad (2.1.13)$$

For brevity, we say that G is L_B -smooth on B . If the constant $L_B = L > 0$ can be chosen to be independent of B , then we say that G is *globally Lipschitz smooth with a constant L* , or simply L -smooth.

Remark 2.1.6 (terminology). The above definition is a weaker notion than the local Lipschitz continuity of the Fréchet derivative of G , which is given by

$$\|G'(x) - G'(y)\|_{\mathcal{L}^{-1}} \leq L_B \|x - y\|_{\mathcal{L}} \quad \forall x, y \in B, \quad (2.1.14)$$

for some $L_B > 0$. Of course, this implies the local Lipschitz smoothness of G (2.1.13). In this thesis, to avoid confusion, whenever (2.1.14) holds, we will say that G is *locally Lipschitz smooth in the strong sense* or that G' is *locally Lipschitz in the strong sense*. We need this stronger condition when we conduct the continuous level analysis (Chapter 4). Note, however, for convex functions, the global versions of the two definitions are equivalent. That is, if $B = \mathbb{H}$, (2.1.13) implies (2.1.14) (see [33, Theorem 2.1.5 (2.1.8)]). ■

Definition 2.1.7 (strong convexity). Let $G : \mathbb{H} \rightarrow \mathbb{R}$ be Fréchet differentiable. We say that G is μ -strongly convex (with respect to $\mathcal{L}$ -norm) iff there exists a constant $\mu > 0$ such that

$$\langle G'(x) - G'(y), x - y \rangle \geq \mu \|y - x\|_{\mathcal{L}}^2 \quad \forall x, y \in \mathbb{H}. \quad (2.1.15)$$

We now state an equivalent characterization of these notions.

Theorem 2.1.8 (equivalence). *Let $\mathbb{H}$ be a real and separable Hilbert space, and $G : \mathbb{H} \rightarrow \mathbb{R}$ be Fréchet differentiable. G is L_B -smooth on the bounded convex set $B \subset \mathbb{H}$ if and only if*

$$G(y) - G(x) - \langle G'(x), y - x \rangle \leq \frac{L_B}{2} \|y - x\|_{\mathcal{L}}^2 \quad \forall x, y \in B. \quad (2.1.16)$$

Similarly, G is μ -strongly convex if and only if

$$G(y) - G(x) - \langle G'(x), y - x \rangle \geq \frac{\mu}{2} \|y - x\|_{\mathcal{L}}^2 \quad \forall x, y \in \mathbb{H}. \quad (2.1.17)$$

Proof. These results follow from Taylor's Theorem with integral remainder. See also [33, Theorem 2.1.5, Theorem 2.1.9]. $\square$

Among the two characterizations of Lipschitz smoothness and strong convexity stated above, we will call (2.1.16) and (2.1.17) the upper and the lower *quadratic trap* of G , respectively. The constant $\frac{L_B}{\mu}$ is called the *(local) condition number* of the objective functional G with respect to the $\mathcal{L}$ -norm. In what follows, we will use its reciprocal, denoted by $\rho = \frac{\mu}{L_B} \in (0, 1]$, to quantify rates of convergence. Note that the condition number crucially depends on the norm that is used to describe the geometry of the graph of G . Choosing a good preconditioner, $\mathcal{L}$, is at the heart of much of scientific computing.

We conclude this chapter by stating a pair of well-known identities which we will use frequently. For any $A, B \in \mathbb{H}$,

$$(A, B)_{\mathcal{L}} = \frac{1}{2} \|A\|_{\mathcal{L}}^2 + \frac{1}{2} \|B\|_{\mathcal{L}}^2 - \frac{1}{2} \|A - B\|_{\mathcal{L}}^2 \quad (2.1.18)$$

$$= \frac{1}{2} \|A + B\|_{\mathcal{L}}^2 - \frac{1}{2} \|A\|_{\mathcal{L}}^2 - \frac{1}{2} \|B\|_{\mathcal{L}}^2. \quad (2.1.19)$$

Chapter 3

Optimization Schemes

Here we briefly review several algorithms that are closely related to our main algorithm of interest. To focus on the main differences between the schemes of interest, we will not pay attention to choices of step size and stopping criteria of the algorithms. For those readers who are interested in these details, we refer, for instance, to [33, 10, 6, 14].

3.1 Preconditioned gradient descent method (PGD)

Algorithm 1: Preconditioned gradient descent method (PGD)

Data: G : The objective

Data: $s > 0$: The step size

Data: $x_0 \in \mathbb{H}$: The initial guess

Result: The sequence $\{x_k\}_{k \geq 1}$ that approximates x^* , the minimizer of G

for $k \geq 0$ **do**

$x_{k+1} = x_k - s\mathcal{L}^{-1}G'(x_k);$

end

We begin by presenting the PGD scheme in Algorithm 1 and describing its convergence properties. To do so, we introduce

$$B = \{x \in \mathbb{H} \mid G(x) \leq G(x_0)\},$$

which is a bounded, convex set containing the minimizer. Then, assuming that G is L_B -smooth on B and μ -strongly convex, and that the step size satisfies $s \in (0, 2/(L_B + \mu)]$, it is possible to show that $x_k \in B$ for all $k \geq 0$. Moreover, in this setting, the scheme converges exponentially fast to the minimizer (see [20, 14, 33]). In particular, if $s = 2/(L_B + \mu)$, then

$$\|x_k - x^*\|_{\mathcal{L}} \leq \left(\frac{1-\rho}{1+\rho}\right)^k \|x_0 - x^*\|_{\mathcal{L}}. \quad (3.1.1)$$

3.2 Preconditioned Nesterov's accelerated gradient descent method (PAGD)

Algorithm 2: Preconditioned accelerated gradient descent method (PAGD)

Data: G : The objective

Data: $\eta > 0$: The friction coefficient

Data: $s > 0$: The step size

Data: $x_0 \in \mathbb{H}$: The initial guess

Result: The sequence $\{x_k\}_{k \geq 1}$ that approximates x^* , the minimizer of G

Define: $\theta = \eta\sqrt{s}$ and $\lambda = \frac{1-\theta}{1+\theta}$;

Set: $x_{-1} = v_0 = x_0 \in \mathbb{H}$;

for $k \geq 0$ **do**

|

$$y_k = x_k + \lambda(x_k - x_{k-1}), \quad (3.2.1)$$

$$x_{k+1} = y_k - s\mathcal{L}^{-1}G'(y_k), \quad (3.2.2)$$

$$v_{k+1} = x_k + \frac{1}{\theta}(x_{k+1} - x_k). \quad (3.2.3)$$

end

To improve on the convergence of GD (Algorithm 1 with $\mathcal{L} = \mathfrak{R}_{\mathbb{H}}^{-1}$), Nesterov [32] devised an algorithm, which “accelerates” the rate of convergence of GD. The improved algorithm is commonly known as *Nesterov's accelerated gradient descent method* (AGD). The preconditioned version of this scheme, PAGD, is presented in Algorithm 2. Roughly speaking, it computes an extrapolation, (3.2.1), takes a gradient step there, (3.2.2), and repeats the same process. Notice that an actual implementation does not need to compute the sequence $\{v_k\}_{k \geq 0}$. We need it for the theoretical analysis. As we will see in Chapter 5,

for convergence, the algorithm must satisfy the condition $s \leq 1/L_B$ and $\eta \leq \sqrt{\mu}$ where $L_B > 0$ is the (local) Lipschitz smoothness constant of G with respect to a bounded convex neighborhood of the minimizer, B , and μ is the strong convexity constant.

It must be noted that PAGD, as presented in Algorithm 2, is practical only if the objective functional is μ -strongly convex ($\mu > 0$). Otherwise, a convergence result may not be available. There exists a more general scheme, which one may call *accelerated gradient descent method with variable weights* (see [33, p. 78]), that is applicable to merely convex objectives. We do not discuss this case here.

Let us now compare the performances of GD and AGD (Algorithm 1 and Algorithm 2 with $\mathcal{L} = \mathfrak{R}_{\mathbb{H}}^{-1}$ respectively) by comparing $G(x_k) - G^*$, where k is the number of iterations and $G^* = G(x^*)$ is the minimum of G . To the best of our knowledge, the existing results on AGD are established under the assumption that the objective is *globally* Lipschitz smooth. Thus, for the rest of the summary of this chapter, the objective G is assumed to be (globally) L -smooth. If GD is applied to a (merely) convex, L -smooth objective functional with a step size condition $0 < s \leq 1/L$, then we have a first order convergence in the objective functional, i.e., $G(x_k) - G^* \leq O(1/k)$ as $k \rightarrow \infty$ (see [33, Corollary 2.1.2]). On the other hand, AGD with variable weights (the more general version mentioned above) provides a second order convergence, that is, $G(x_k) - G^* \leq O(1/k^2)$ as $k \rightarrow \infty$. If the objective is, in addition, μ -strongly convex, the convergence rates of the two schemes become exponential. Specifically, estimate (3.1.1) and the quadratic traps show that the convergence rate of GD is $G(x_k) - G^* \leq O((\frac{1-\rho}{1+\rho})^{2k})$ as $k \rightarrow \infty$, where we recall that $\rho = \mu/L$. This is in contrast to AGD, which converges with a rate of $G(x_k) - G^* \leq O((1 - \sqrt{\rho})^k)$ as $k \rightarrow \infty$; see [33, Theorem 2.2.3]. If $\rho \ll 1$, this acceleration can be significant. As we will see later, PAGD achieves the same rate of exponential convergence even if the objective is locally Lipschitz smooth instead of the Lipschitz smoothness being imposed globally.

Chapter 4

ODE Model

We study a continuous time analogue of PAGD, a second order ODE, inspired by [46]. As we will see, the discussion in this chapter turns out to be informative. It not only provides an intuitive understanding of Nesterov’s acceleration, but also guides us to important results at the discrete level. As mentioned in the introduction, there are recent works that arrive at the same conclusion for some parts of our results using similar ideas. However, our unique contributions rely on our specific layout of various quantities and calculations. Thus, we include such details in a condensed manner while referring to existing work otherwise.

4.1 The limiting ODE

To streamline the discussion, we start by directly introducing the initial value problem (IVP) whose certain discretization leads to PAGD:

$$\begin{cases} \ddot{X}(t) + 2\eta\dot{X}(t) + \mathcal{L}^{-1}G'(X(t)) = 0, & (t > 0) \\ X(0) = x_0, \\ \dot{X}(0) = 0. \end{cases} \quad (4.1.1)$$

Interestingly, this is the same system as what Poljak [36] had in mind when he proposed the *heavy ball* method. See Section 4.4 for its derivation. Conversely, PAGD can be viewed as a discretization of this IVP although not every choice in the process can be seen as natural

or intuitive. It is given in Section 4.6. Note that it involves some ingredients appearing in Section 4.3.

Remark 4.1.1 (physical interpretation). This IVP (4.1.1) describes the motion of a ball of unit mass in the potential G with friction coefficient 2η which starts from the initial position x_0 at rest; see Figure 4.1. Our physical intuition suggests that the ball will converge to its minimal point as it exhausts the initial energy under the action of friction. If the friction, quantified by η , is too small it will oscillate much as it reaches the minimal point and will converge only after a long travel. On the other hand, if the friction is too large, it will not move sufficiently rapidly, and this, in turn, will also lead to a slow convergence.

Let us compare this with another physical system. We can interpret the physics of the gradient flow as a limiting case of the same dynamics. The gradient flow

$$\dot{X}(t) = -\mathcal{L}^{-1}G'(X(t))$$

can be viewed, up to a constant factor 2η , as a massless limit of the IVP (4.1.1). That is, a physical thought experiment suggests that the surroundings hold the particle back as soon as it gets accelerated since it is so light. A real life example of this kind is a very viscous fluid, such as honey, flowing down a bowl. Our physical experience suggests that it will not oscillate and will flow along the steepest descent direction every moment. However, it will reach the bottom slower than the rolling ball will if the friction is appropriately strong. ■

4.2 Analysis of the IVP

As one can expect from the fact that the IVP (4.1.1) describes a concrete physical situation, its solution possesses good properties. In this and the following chapter, however, we need a slightly stronger Lipschitz condition on G' than in the discrete level discussion. Part of the following proof can be found in [4, Theorem 3.1 and Proposition 4.2] in a slightly more general setting. We give our own proof, however, for the sake of completeness.

Lemma 4.2.1 (existence and uniqueness). *Suppose that $G : \mathbb{H} \rightarrow \mathbb{R}$ is μ -strongly convex and locally Lipschitz smooth in the strong sense, i.e., (2.1.14) holds. Then, for any $T > 0$,*




Figure 4.1: A rolling ball system. The IVP (4.1.1) describes a ball of unit mass rolling down a bowl-shaped potential landscape with a constant friction coefficient.

there exists a unique solution $X \in C^2(0, T; \mathbb{H})$ to the initial value problem (4.1.1) and the solution obeys the following energy identity

$$\frac{1}{2} \|\dot{X}(t)\|_{\mathcal{L}}^2 + G(X(t)) - G^* = G(x_0) - G^* - 2\eta \int_0^t \|\dot{X}(\tau)\|_{\mathcal{L}}^2 d\tau \quad \forall t \geq 0. \quad (4.2.1)$$

Consequently, the solution exists for all $t \in [0, \infty)$ and it is twice continuously differentiable.

Proof. By introducing an auxiliary variable $W : [0, \infty) \rightarrow \mathbb{H}$, the IVP (4.1.1) can be equivalently written as a system of first order ODEs

$$\begin{cases} \dot{X}(t) = W(t), & t > 0, \\ \dot{W}(t) = -2\eta W(t) - \mathcal{L}^{-1}G'(X(t)), & t > 0, \\ X(0) = x_0, \\ W(0) = 0. \end{cases} \quad (4.2.2)$$

System (4.2.2) can be written as

$$\dot{Z}(t) = F(t, Z), \quad t > 0, \quad Z(0) = (x_0, 0),$$

where $Z = (X, W) : [0, \infty) \rightarrow \mathbb{H}^2$ and $F : (0, \infty) \times \mathbb{H}^2 \rightarrow \mathbb{H}^2$ is defined via

$$F(t, Z) = (W, -2\eta W - \mathcal{L}^{-1}G'(X)).$$

It is immediate to see that F is continuous and locally Lipschitz in the strong sense in the second argument. Thus, owing to the Picard-Lindelöf theorem on Banach spaces [13, p. 110 Theorem 1.8.1], we have local existence and uniqueness of a classical solution, say on $[0, T)$ for some $T > 0$. In particular, $W = \dot{X} \in C^1(0, T; \mathbb{H})$ implies that $X \in C^2(0, T; \mathbb{H})$. Moreover, this classical solution must coincide with the mild solution represented by

$$Z(t) = (x_0, 0) + \int_0^t F(\tau, Z(\tau)) d\tau, \quad t \in [0, T). \quad (4.2.3)$$

We now show the energy law (4.2.1), which will lead to the global existence of the solution. Let us take the $\mathcal{L}$ -inner product of the first equation of (4.1.1) with $\dot{X}$ and integrate over

time $\tau \in [0, t]$. We obtain

$$\int_0^t (\ddot{X}(t), \dot{X}(t))_{\mathcal{L}} d\tau + 2\eta \int_0^t \|\dot{X}(t)\|_{\mathcal{L}}^2 d\tau + \int_0^t \langle G'(X(\tau)), \dot{X}(\tau) \rangle d\tau = 0.$$

Noting that $(\ddot{X}(t), \dot{X}(t))_{\mathcal{L}} = \frac{1}{2} \frac{d}{dt} \|\dot{X}(t)\|_{\mathcal{L}}^2$ and $\langle G'(X(t)), \dot{X}(t) \rangle = \frac{d}{dt} G(X(t))$, rearranging, and subtracting the minimum G^* from both sides, we obtain the energy identity (4.2.1) upon plugging in the initial conditions.

By discarding the velocity term and the integral term from (4.2.1), we obtain an upper bound on the distance between $X(t)$ and the minimizer x^* . That is, in conjunction with the strong convexity of G , we have $\frac{\mu}{2} \|X(t) - x^*\|_{\mathcal{L}}^2 \leq G(X(t)) - G^* \leq G(x_0) - G^*$, which, in turn, provides

$$\|X(t) - x^*\|_{\mathcal{L}} \leq \sqrt{\frac{2}{\mu} (G(x_0) - G^*)}. \quad (4.2.4)$$

Similarly, by removing different terms from (4.2.1), it follows that $\frac{1}{2} \|\dot{X}(t)\|_{\mathcal{L}}^2 \leq G(x_0) - G^*$, from which we obtain

$$\|\dot{X}(t)\|_{\mathcal{L}} = \|W(t)\|_{\mathcal{L}} \leq \sqrt{2(G(x_0) - G^*)}. \quad (4.2.5)$$

Estimates (4.2.4) and (4.2.5) ensure that the integrand involved in the definition of the mild solution (4.2.3) is uniformly bounded in t , hence $\lim_{t \uparrow T} (X(t), W(t))$ exists. Therefore, the local solution continues to exist for $[0, T + \varepsilon)$ for some $\varepsilon > 0$; see [13, p. 111]. The global existence of the solution follows by repeating this argument. $\square$

Remark 4.2.2 (smoothness of the solution). The fact that $X \in C^2((0, \infty); \mathbb{H}) \cap C([0, \infty); \mathbb{H})$ justifies the manipulations we will carry out when we derive the IVP (4.1.1) in Section 4.4.

■

4.3 Convergence to equilibrium

We now wish to prove that the solution to the IVP (4.1.1) with G being locally Lipschitz smooth and μ -strongly convex converges to its attractive steady state solution as $t \rightarrow \infty$, which is the minimal point in this case, at a matching rate with that of PAGD. This is one

of the highlights of this work. To this end, we introduce an auxiliary variable

$$V(t) = X(t) - x^* + \frac{1}{\eta} \dot{X}(t)$$

so that the first equation of the IVP (4.1.1) can be rewritten

$$\eta \dot{V}(t) + \eta \dot{X}(t) + \mathcal{L}^{-1} G'(X(t)) = 0. \quad (4.3.1)$$

We also introduce an energy

$$E(X, V) = \frac{\eta}{2} \|V\|_{\mathcal{L}}^2 + \frac{1}{\eta} (G(X) - G^*), \quad (4.3.2)$$

where we recall $G^* = G(x^*) = \min_{x \in \mathbb{H}} G(x)$. We will show that E is a *Lyapunov* function for the IVP (4.1.1). For notational convenience, set

$$E_0 := E(x_0, x_0) = \frac{1}{\eta} (G(x_0) - G^*) + \frac{\eta}{2} \|x_0 - x^*\|_{\mathcal{L}}^2. \quad (4.3.3)$$

Theorem 4.3.1 (exponential decay). *Let $G : \mathbb{H} \rightarrow \mathbb{R}$ be locally Lipschitz smooth in the strong sense and μ -strongly convex. Denote by X the unique solution to the IVP (4.1.1). If $\eta^2 \leq \mu$, the exponentially inflated energy $\mathcal{E}(t) = e^{\eta t} E(X(t), V(t))$ is nonincreasing. Consequently, the Lyapunov function (4.3.2) decays to zero at an exponential rate:*

$$E(X(t), V(t)) = \frac{\eta}{2} \|V(t)\|_{\mathcal{L}}^2 + \frac{1}{\eta} (G(X(t)) - G^*) \leq e^{-\eta t} E_0. \quad (4.3.4)$$

Proof. Existence and uniqueness of X is guaranteed by Lemma 4.2.1. Let us now prove the estimate (4.3.4). Taking the inner product of (4.3.1) with $V(t)$, and using the identity

(2.1.18), we obtain, suppressing the time variable,

$$\begin{aligned}
0 &= \eta(V, \dot{V})_{\mathcal{L}} + \eta(\dot{X}, X - x^* + \frac{1}{\eta}\dot{X})_{\mathcal{L}} + (\mathcal{L}^{-1}G'(X), X - x^*)_{\mathcal{L}} \\
&\quad + \frac{1}{\eta}(\mathcal{L}^{-1}G'(X), \dot{X})_{\mathcal{L}} \\
&= \eta(V, \dot{V})_{\mathcal{L}} + \frac{1}{2}\|\dot{X}\|_{\mathcal{L}}^2 + \frac{\eta^2}{2}\|V\|_{\mathcal{L}}^2 - \frac{\eta^2}{2}\|X - x^*\|_{\mathcal{L}}^2 + \langle G'(X), X - x^* \rangle \\
&\quad + \frac{1}{\eta}\langle G'(X), \dot{X} \rangle.
\end{aligned} \tag{4.3.5}$$

The lower quadratic trap, (2.1.17), implies

$$G(X) - G^* - \langle G'(X), X - x^* \rangle \leq -\frac{\mu}{2}\|X - x^*\|_{\mathcal{L}}^2. \tag{4.3.6}$$

Substituting (4.3.5) into the time derivative of the inflated energy and then using the above estimate (4.3.6), we have

$$\begin{aligned}
\dot{\mathcal{E}}(t) &= e^{\eta t} \left[\frac{\eta^2}{2}\|V\|_{\mathcal{L}}^2 + \eta(V, \dot{V})_{\mathcal{L}} + (G(X) - G^*) + \frac{1}{\eta}\langle G'(X), \dot{X} \rangle \right] \\
&= e^{\eta t} \left[-\frac{1}{2}\|\dot{X}\|_{\mathcal{L}}^2 + \frac{\eta^2}{2}\|X - x^*\|_{\mathcal{L}}^2 - \langle G'(X), X - x^* \rangle + G(X) - G^* \right] \\
&\leq -\frac{1}{2}e^{\eta t} \left[\|\dot{X}\|_{\mathcal{L}}^2 + \frac{\mu - \eta^2}{2}\|X - x^*\|_{\mathcal{L}}^2 \right].
\end{aligned} \tag{4.3.7}$$

The last term is always nonpositive provided $\eta^2 \leq \mu$, and this implies $\mathcal{E}(t) \leq \mathcal{E}(0) = E_0$.

This completes the proof. $\square$

Remark 4.3.2 (physical interpretation). We can rigorously explain the physical intuition given in Remark 4.1.1 through Theorem 4.3.1 and its proof. If the friction, quantified by η , is too small the decay to the attraction point is slow as η governs the decay rate $e^{-\eta t}$. On the other hand, If the friction is too large, say $\eta > \sqrt{\mu}$, then we cannot guarantee the boundedness of $\mathcal{E}(t)$. $\blacksquare$

4.4 Derivation of the ODE

While the derivation of the ODE, (4.1.1), is not an essential part of our development from the logical point of view, it provides us with inspiration and guidance. Let us start with the same approach as in [46]. We assume, as an *ansatz*, that PAGD is a discretization of an ODE, which has a solution $X : [0, \infty) \rightarrow \mathbb{H}$, which we often call a *trajectory*. We also assume that X is smooth enough, e.g., twice continuously differentiable in time. For a fixed $t \in (0, \infty)$, the assumed smoothness on X , together with the identification $t = \sqrt{s}k$ and Taylor's formula in a normed vector space (e.g., [16, Theorem 7.9-1]) implies:

$$\begin{aligned}\frac{x_{k+1} - x_k}{\sqrt{s}} &= \dot{X}(t) + \frac{1}{2}\ddot{X}(t)\sqrt{s} + o(\sqrt{s}) \quad \text{as } s \rightarrow 0, \\ \frac{x_k - x_{k-1}}{\sqrt{s}} &= \dot{X}(t) - \frac{1}{2}\ddot{X}(t)\sqrt{s} + o(\sqrt{s}) \quad \text{as } s \rightarrow 0, \\ \sqrt{s}\mathcal{L}^{-1}G'(y_k) &= \sqrt{s}\mathcal{L}^{-1}G'(X(t)) + o(\sqrt{s}) \quad \text{as } s \rightarrow 0.\end{aligned}\tag{4.4.1}$$

The last identity follows from the continuity of G' , that of $\mathcal{L}^{-1}$, and (3.2.1), from which we can deduce $y_k \rightarrow X(t)$ as $s \rightarrow 0$. Plugging (3.2.1) into (3.2.2) and dividing by $\sqrt{s}$, we have

$$\frac{x_{k+1} - x_k}{\sqrt{s}} - \lambda \frac{x_k - x_{k-1}}{\sqrt{s}} + \sqrt{s}\mathcal{L}^{-1}G'(y_k) = 0$$

Substituting the above Taylor expansions, and then rearranging, we arrive at

$$\frac{1}{2}(1 + \lambda)\ddot{X}(t) + \frac{1 - \lambda}{\sqrt{s}}\dot{X}(t) + \mathcal{L}^{-1}G'(X(t)) + o(1) = 0 \quad \text{as } s \rightarrow 0.\tag{4.4.2}$$

To make this estimate consistent, interpret λ as a function of s and further assume that $(1 - \lambda)/\sqrt{s} \rightarrow 2\eta$ as $\sqrt{s} \rightarrow 0$ for some $\eta \in (0, \infty)$, which yields

$$\ddot{X}(t) + 2\eta\dot{X}(t) + \mathcal{L}^{-1}G'(X(t)) = 0.\tag{4.4.3}$$

4.5 Derivation of the initial conditions

The initialization $y_0 = x_0$ and (3.2.2) with $k = 0$ imply

$$\frac{x_1 - x_0}{\sqrt{s}} = \sqrt{s}\mathcal{L}^{-1}G'(x_0).$$

Take the limit $s \rightarrow 0$ and conclude $\dot{X}(0) = 0$ since G' and $\dot{X}$ are assumed to be continuous. Therefore, we arrive at the desired IVP (4.1.1).

Remark 4.5.1 (momentum method). A similar procedure can be carried out far more easily for the so-called *momentum method* (MM). To see this, we recall that

$$\ddot{X}(t) \approx \frac{x_{k+1} - 2x_k + x_{k-1}}{s}, \quad \dot{X}(t) \approx \frac{x_k - x_{k-1}}{\sqrt{s}}, \quad G'(X(t)) \approx G'(x_k).$$

Then, the discrete version of the ODE (4.1.1) becomes

$$x_{k+1} = x_k - sG'(x_k) + (1 - 2\eta\sqrt{s})(x_k - x_{k-1}),$$

which is MM with the weight $1 - 2\eta\sqrt{s}$; see [36, p. 12 (9)]. This weight is close to λ :

$$\lambda = \frac{1 - \eta\sqrt{s}}{1 + \eta\sqrt{s}} = 1 - \frac{2\eta\sqrt{s}}{1 + \eta\sqrt{s}} \approx 1 - 2\eta\sqrt{s}.$$

In this sense, MM seems more natural and amenable for analysis than AGD. ■

The limiting behavior of MM can also be explained by the IVP (4.1.1). Observe that the only essential difference between MM and PAGD is where G' is evaluated, that is, x_k and y_k respectively. And in the limit $s \rightarrow 0$, x_k and y_k are not distinguishable in this setting. However, PAGD exhibits less oscillation than MM since evaluating G' at y_k serves as “foreseeing” the uphill of the objective functional, if exists, along the trajectory and “steering” to avoid unnecessary oscillating behaviors. Recently, it was shown that a higher order Taylor expansion helps differentiate their performances (see [44]).

4.6 PAGD as a discretization of the IVP

Similarly to the derivation of the ODE from the PAGD scheme, the converse direction, i.e. showing a specific discretization of the ODE leading to the scheme, is not essential in terms of logic: the theory is still meaningful as long as it provides convergence results. However, this helps us see the whole connection between the discrete and continuous levels. Let us label the step size $\sqrt{s}$, rather than s , in order to make the setting more in line with the PAGD algorithm. Again, it is helpful to have in mind the correspondence: time $t \longleftrightarrow k\sqrt{s}$ ($k = 0, 1, 2, \dots$) and position $X(t) \longleftrightarrow x_k$. First, we will see y_k corresponds to a “drifted” position without the potential landscape over $[t, t + \sqrt{s}]$. This can be modeled by

$$\ddot{X}(t) + 2\eta\dot{X}(t) = 0,$$

which leads to another energy law

$$\frac{1}{2}\|\dot{X}(t + \sqrt{s})\|_{\mathcal{L}}^2 = \frac{1}{2}\|\dot{X}(t)\|_{\mathcal{L}}^2 - 2\eta \int_t^{t+\sqrt{s}} \|\dot{X}(\tau)\|_{\mathcal{L}}^2 d\tau.$$

Approximate the speed in the integrand by the average $\frac{1}{2}(\|\dot{X}(t + \sqrt{s})\|_{\mathcal{L}} + \|\dot{X}(t)\|_{\mathcal{L}})$, then after a short calculation, one obtains

$$\|\dot{X}(t + \sqrt{s})\|_{\mathcal{L}} = \lambda \|\dot{X}(t)\|_{\mathcal{L}}.$$

Since the dynamics takes place in a single direction, this implies

$$\dot{X}(t + \sqrt{s}) = \lambda \dot{X}(t).$$

The approximations $\dot{X}(t) \approx \frac{x_k - x_{k-1}}{\sqrt{s}}$ and $\dot{X}(t + \sqrt{s}) \approx \frac{y_k - x_k}{\sqrt{s}}$ lead us to (3.2.1).

Next, we discretize the vector $V(t)$. Since we do not know the minimizer in practice, we remove it from the definition of v_k and discretize $V(t) + x^* = X(t) + \frac{1}{\eta}\dot{X}(t)$. The approximations $X(t) \approx y_k$ and $\dot{X}(t) \approx \frac{y_k - x_k}{\sqrt{s}}$ suggest

$$v_k = y_k + \frac{1}{\theta}(y_k - x_k), \tag{4.6.1}$$

which leads to the definition of $\{v_k\}_{k \geq 1}$ (3.2.3) upon combining with the definition of $\{y_k\}$.

Finally, to get the main iterates, $\{x_k\}_{k \geq 1}$, we discretize (4.3.1) using the approximations $\dot{V}(t) \approx \frac{v_{k+1} - v_k}{\sqrt{s}}$, $\dot{X}(t) \approx \frac{y_k - x_k}{\sqrt{s}}$, and the evaluation of G' at y_k , then it follows

$$\eta \frac{v_{k+1} - v_k}{\sqrt{s}} + \eta \frac{y_k - x_k}{\sqrt{s}} + \mathcal{L}^{-1} G'(y_k) = 0.$$

Plugging in (3.2.3) and (4.6.1), one obtains (3.2.2), the definition of $\{x_k\}_{k \geq 1}$.

Chapter 5

An Energy Approach to Convergence of PAGD

In this chapter, we prove the existence of an invariant set of PAGD and its exponential convergence in the objective as well as in the residual when it is applied to a strongly convex, locally Lipschitz smooth objective. We follow the ODE arguments developed in Chapter 4. Throughout this chapter, we assume $\eta = \sqrt{\mu}$, the optimal choice for the friction coefficient in view of Theorem 4.3.1. Note that this does not undermine generality. If $\tilde{\mu}$ is the largest strong convexity constant of G , that is, the supremum of μ 's that satisfies the strong convexity, (2.1.15), then any $\mu \in (0, \tilde{\mu}]$ can be taken as a (non-optimal) strong convexity constant. Thus, the general case $\eta^2 \leq \tilde{\mu}$ corresponds to $\eta^2 = \mu \leq \tilde{\mu}$, the optimal friction coefficient associated with a non-optimal strong convexity constant.

5.1 Existence of an invariant set

The first goal of this chapter is to show that the assumption of the local Lipschitz smoothness is sufficient for our analysis, as the iterates lie within a bounded set. To this end, we first show that, for every $k \geq 0$, the y_k iterate of PAGD lies in the segment between x_k and v_k . This is used frequently in the convergence proof.

Lemma 5.1.1 (convex hull). *For every $k \geq 0$, the iterates constructed in PAGD, described in Algorithm 2, satisfy $y_k \in \overline{x_k v_k}$. Specifically, they satisfy the following four equivalent*

equations:

$$\begin{cases} y_k = \frac{1}{1+\theta}x_k + \frac{\theta}{1+\theta}v_k, & x_k = (1+\theta)y_k - \theta v_k, \\ v_k = \left(1 + \frac{1}{\theta}\right)y_k - \frac{1}{\theta}x_k, & x_k - y_k = \theta(y_k - v_k). \end{cases} \quad (5.1.1)$$

Proof. If $k = 0$ this is trivial since $x_0 = y_0 = v_0$. For $k \geq 1$, we eliminate x_{k-1} from (3.2.1) and (3.2.3) with the index being $k - 1$ to get

$$\left(1 - \frac{1}{\theta}\right)y_k + \lambda v_k = \left((1+\lambda)\left(1 - \frac{1}{\theta}\right) + \frac{\lambda}{\theta}\right)x_k = -\frac{\lambda}{\theta}x_k.$$

Rearranging terms and using $\lambda = \frac{1-\theta}{1+\theta}$, we obtain the equalities that are listed above. $\square$

We now show that there is an invariant set for the iterates of PAGD.

Lemma 5.1.2 (invariant set). *Assume that the objective $G : \mathbb{H} \rightarrow \mathbb{R}$ is μ -strongly convex and locally Lipschitz smooth. Define*

$$\mathfrak{B} = \{x \in \mathbb{H} \mid \|x - x^*\|_{\mathcal{L}} \leq R\}, \quad (5.1.2)$$

where $R = R_1 + \frac{1}{\eta}R_2$, $R_1 = \sqrt{\frac{2}{\mu}(G(x_0) - G^*)}$, $R_2 = \sqrt{2r(G(x_0) - G^*)}$, and $r > 1$. Let PAGD, as described in Algorithm 2, be implemented with a step size rule

$$s \in \left(0, \min \left\{ L_B^{-1}, \left(\frac{r-1}{r+1}\right)^2 \mu^{-1} \right\} \right], \quad (5.1.3)$$

where L_B the local Lipschitz smoothness constant of G associated to the set $\mathfrak{B}$. Then, for all $k \geq 0$, we have that $\|x_k - x^*\|_{\mathcal{L}} \leq R_1$, hence $x_k \in \mathfrak{B}$, and $y_k, v_k \in \mathfrak{B}$.

Proof. The outline of this proof is simple although it is long. We mimic the energy law developed in Chapter 4.2 to obtain a bound on the distance between the main iterates and the minimizer and that on the speed. Once we get the bounds, it is easy to prescribe an appropriate ball, which will be our invariant set.

We will prove the statement by induction. For $k = 0$, the statement is trivial since $x_0 = y_0 = v_0$ and the strong convexity implies $\|x_0 - x^*\|_{\mathcal{L}} \leq R_1$. Suppose that $\|x_k - x^*\|_{\mathcal{L}} \leq R_1$ (hence $x_k \in \mathfrak{B}$) and $y_k, v_k \in \mathfrak{B}$ are true for $k = 0, 1, 2, \dots, N$. We need to show

that $\|x_{N+1} - x^*\|_{\mathcal{L}} \leq R_1$ (hence $x_{N+1} \in \mathfrak{B}$) and $v_{N+1} \in \mathfrak{B}$, then Lemma 5.1.1 implies $y_{N+1} \in \overline{x_{N+1}v_{N+1}} \subset \mathfrak{B}$ since $\mathfrak{B}$ is a convex set as a sublevel set of a convex function.

Note that the condition $s \leq (\frac{r-1}{r+1})^2 \mu^{-1}$, which is implied by (5.1.3), ensures λ^{-1} to be bounded above since

$$\frac{1}{\lambda} = \frac{1 + \sqrt{s\mu}}{1 - \sqrt{s\mu}} \leq r. \quad (5.1.4)$$

First, a similar argument to [14, Proposition 4.6] shows that the x_{N+1} update from y_N is a descent step in terms of G . That is, the section of G across the line $\overrightarrow{y_N x_{N+1}}$ also inherits the strong convexity and the local Lipschitz smoothness with the same constants on the one-dimensional affine subset

$$B_{N+1} = \{x = y_N - \tau \mathcal{L}^{-1} G'(y_N) \in \mathbb{H} \mid \tau \in \mathbb{R}\}.$$

Let $S(\tau) = G(y_N - \tau \mathcal{L}^{-1} G'(y_N))$ denote the section.

Since we know that $y_N \in \mathfrak{B}$, we can bound S in a neighborhood of $\tau = 0$ using the upper quadratic trap

$$\begin{aligned} U(\tau) &:= G(y_N) + \langle G'(y_N), -\tau \mathcal{L}^{-1} G'(y_N) \rangle + \frac{L_B}{2} \|\tau \mathcal{L}^{-1} G'(y_N)\|_{\mathcal{L}}^2 \\ &= G(y_N) - \tau \|G'(y_N)\|_{\mathcal{L}^{-1}}^2 + \frac{L_B \tau^2}{2} \|G'(y_N)\|_{\mathcal{L}^{-1}}^2. \end{aligned}$$

Observe that $S(0) = U(0) = G(y_N)$, that $U(\tau)$ is decreasing around $\tau = 0$ since $dU/d\tau(0) = -\|G'(y_N)\|_{\mathcal{L}^{-1}}^2 \leq 0$, and that the optimal step size to minimize U is $1/L_B$ since $dU/d\tau(1/L_B) = 0$. This implies that $S(s) \leq U(s) \leq U(0)$ for any $s \in [0, 2/L_B]$. Moreover, for $s \in [0, 1/L_B]$, we have

$$\begin{aligned} G(x_{N+1}) &= S(s) \leq U(s) = G(y_N) - s \|G'(y_N)\|_{\mathcal{L}^{-1}}^2 + \frac{L_B s^2}{2} \|G'(y_N)\|_{\mathcal{L}^{-1}}^2 \\ &\leq G(y_N) - \frac{s}{2} \|G'(y_N)\|_{\mathcal{L}^{-1}}^2, \end{aligned} \quad (5.1.5)$$

which is the desired descent property in G from y_N to x_{N+1} .

Now, we want to mimic the energy argument that we carried out in Chapter 4.2. Substitute (3.2.1) into (3.2.2), and add and subtract $x_k - x_{k-1}$, to obtain the discrete

counterpart of (4.1.1)

$$x_{k+1} - 2x_k + x_{k-1} + (1 - \lambda)(x_k - x_{k-1}) + s\mathcal{L}^{-1}G'(y_k) = 0. \quad (5.1.6)$$

Note that defining $x_{-1} := x_0$ allows us to extend this equality to the case $k = 0$. Take the $\mathcal{L}$ -inner product of this identity with $x_k - x_{k-1}$ and add for $0 \leq k \leq N$. Then, using (2.1.19), the first term telescopes to simplify

$$\begin{aligned} & \sum_{k=0}^N (x_{k+1} - 2x_k + x_{k-1}, x_k - x_{k-1})_{\mathcal{L}} \\ &= \frac{1}{2} \sum_{k=0}^N \left(\|x_{k+1} - x_k\|_{\mathcal{L}}^2 - \|x_{k+1} - 2x_k + x_{k-1}\|_{\mathcal{L}}^2 - \|x_k - x_{k-1}\|_{\mathcal{L}}^2 \right) \\ &= \frac{1}{2} \|x_{N+1} - x_N\|_{\mathcal{L}}^2 - \frac{1}{2} \sum_{k=0}^N \|x_{k+1} - 2x_k + x_{k-1}\|_{\mathcal{L}}^2. \end{aligned}$$

We leave the second term as it is. For the third term, using (3.2.1), $G(y_k) - G(x_k) \leq \langle G'(y_k), y_k - x_k \rangle$ from convexity, and (5.1.5), it follows

$$\begin{aligned} & s \sum_{k=0}^N \langle G'(y_k), x_k - x_{k-1} \rangle = \frac{s}{\lambda} \sum_{k=0}^N \langle G'(y_k), y_k - x_k \rangle \\ & \geq \frac{s}{\lambda} \sum_{k=0}^N (G(y_k) - G(x_k)) \geq \frac{s}{\lambda} \sum_{k=0}^N \left(G(x_{k+1}) - G(x_k) + \frac{s}{2} \|G'(y_k)\|_{\mathcal{L}^{-1}}^2 \right) \\ & = \frac{s}{\lambda} G(x_{N+1}) - \frac{s}{\lambda} G(x_0) + \frac{s^2}{2\lambda} \sum_{k=0}^N \|G'(y_k)\|_{\mathcal{L}^{-1}}^2. \end{aligned} \quad (5.1.7)$$

Gathering all the three terms together and rearranging, we get

$$\begin{aligned} \frac{1}{2} \|x_{N+1} - x_N\|_{\mathcal{L}}^2 + \frac{s}{\lambda} G(x_{N+1}) & \leq \frac{s}{\lambda} G(x_0) + \frac{1}{2} \sum_{k=0}^N \|x_{k+1} - 2x_k + x_{k-1}\|_{\mathcal{L}}^2 \\ & \quad - \frac{s^2}{2\lambda} \sum_{k=0}^N \|G'(y_k)\|_{\mathcal{L}^{-1}}^2 - (1 - \lambda) \sum_{k=0}^N \|x_k - x_{k-1}\|_{\mathcal{L}}^2. \end{aligned} \quad (5.1.8)$$

Similarly, take the $\mathcal{L}$ -inner product of (5.1.6) with $x_{k+1} - x_k$ and sum over $0 \leq k \leq N$. This time, use (2.1.18) for the first term to get

$$\begin{aligned} & \sum_{k=0}^N (x_{k+1} - 2x_k + x_{k-1}, x_{k+1} - x_k)_{\mathcal{L}} \\ &= \frac{1}{2} \|x_{N+1} - x_N\|_{\mathcal{L}}^2 + \frac{1}{2} \sum_{k=0}^N \|x_{k+1} - 2x_k + x_{k-1}\|_{\mathcal{L}}^2. \end{aligned}$$

For the second term, using Cauchy-Schwarz and Young's inequality, we have

$$\begin{aligned} & (1 - \lambda) \sum_{k=0}^N (x_k - x_{k-1}, x_{k+1} - x_k)_{\mathcal{L}} \\ & \geq -\frac{1 - \lambda}{2} \sum_{k=0}^N \left(\|x_k - x_{k-1}\|_{\mathcal{L}}^2 + \|x_{k+1} - x_k\|_{\mathcal{L}}^2 \right). \end{aligned}$$

For the third term, use (3.2.2) and argue as in (5.1.7) to get

$$\begin{aligned} & s \sum_{k=0}^N \langle G'(y_k), x_{k+1} - x_k \rangle = s \sum_{k=0}^N \langle G'(y_k), y_k - s\mathcal{L}^{-1}G'(y_k) - x_k \rangle \\ & \geq -s^2 \sum_{k=0}^N \|G'(y_k)\|_{\mathcal{L}^{-1}}^2 + s \sum_{k=0}^N \langle G'(y_k), y_k - x_k \rangle \\ & \geq sG(x_{N+1}) - sG(x_0) - \frac{s^2}{2} \sum_{k=0}^N \|G'(y_k)\|_{\mathcal{L}^{-1}}^2. \end{aligned}$$

Gathering all these estimates we get

$$\begin{aligned} & \frac{1}{2} \|x_{N+1} - x_N\|_{\mathcal{L}}^2 + sG(x_{N+1}) \leq sG(x_0) - \frac{1}{2} \sum_{k=0}^N \|x_{k+1} - 2x_k + x_{k-1}\|_{\mathcal{L}}^2 \\ & + \frac{s^2}{2} \sum_{k=0}^N \|G'(y_k)\|_{\mathcal{L}^{-1}}^2 + \frac{1 - \lambda}{2} \sum_{k=0}^N \left(\|x_k - x_{k-1}\|_{\mathcal{L}}^2 + \|x_{k+1} - x_k\|_{\mathcal{L}}^2 \right). \end{aligned} \tag{5.1.9}$$

Add (5.1.8) and (5.1.9) and rearrange, then after some cancelations, it follows

$$\begin{aligned} & \frac{1+\lambda}{2} \|x_{N+1} - x_N\|_{\mathcal{L}}^2 + s \left(1 + \frac{1}{\lambda}\right) (G(x_{N+1}) - G^*) \\ & \leq s \left(1 + \frac{1}{\lambda}\right) (G(x_0) - G^*) - \frac{s^2}{2} \left(\frac{1}{\lambda} - 1\right) \sum_{k=0}^N \|G'(y_k)\|_{\mathcal{L}^{-1}}^2 \end{aligned} \quad (5.1.10)$$

$$\leq s \left(1 + \frac{1}{\lambda}\right) (G(x_0) - G^*), \quad (5.1.11)$$

since $0 < \lambda < 1$. By removing the kinetic term from this estimate, strong convexity leads to

$$G(x_0) - G^* \geq G(x_{N+1}) - G^* \geq \frac{\mu}{2} \|x_{N+1} - x^*\|_{\mathcal{L}}^2,$$

which implies

$$\|x_{N+1} - x^*\|_{\mathcal{L}} \leq R_1, \quad (5.1.12)$$

which, in turn, proves $x_{N+1} \in \mathfrak{B}$. Similarly, discarding the potential term, dividing through $\frac{s(1+\lambda)}{2}$, and using (5.1.4), we obtain

$$\left\| \frac{x_{N+1} - x_N}{\sqrt{s}} \right\|_{\mathcal{L}} \leq \sqrt{\frac{2}{\lambda} (G(x_0) - G^*)} \leq R_2.$$

Then, from the definition of v_{N+1} (3.2.3),

$$\|v_{N+1} - x^*\|_{\mathcal{L}} \leq \|x_N - x^*\|_{\mathcal{L}} + \frac{1}{\eta} \left\| \frac{x_{N+1} - x_N}{\sqrt{s}} \right\|_{\mathcal{L}} \leq R_1 + \frac{1}{\eta} R_2 = R,$$

which implies $v_{N+1} \in \mathfrak{B}$. This completes the proof. $\square$

Remark 5.1.3 (step size restriction). The additional condition $s \leq (\frac{r-1}{r+1})^2 \mu^{-1}$ on the step size is not restrictive at all in practice. For example, if $r = 3$, we require that $s \leq 1/4\mu^{-1}$. The purpose of this condition is to bound λ^{-1} as explained in the proof. However, λ^{-1} becomes unbounded when $s\mu$ is close to 1. If we set $s = 1/L_B$, $s\mu$ is the (inverse) condition number and the (inverse) condition number being close to 1 makes the problem more amenable because it means that G is almost quadratic. Moreover, even from a theoretical point of view, as r increases, the invariant set $\mathfrak{B}$ gets larger, which means L_B^{-1} gets smaller, while

$(\frac{r-1}{r+1})^2 \mu^{-1}$ approaches μ^{-1} . Since $L_B^{-1} < \mu^{-1}$ (unless G is perfectly quadratic), the second argument of the minimum in (5.1.3) eventually becomes of no effect. $\blacksquare$

Corollary 5.1.4 (convergence of residuals). *Assume that $G : \mathbb{H} \rightarrow \mathbb{R}$ is μ -strongly convex and locally Lipschitz smooth. Suppose that PAGD, as described in Algorithm 2, is implemented with a step size that obeys condition (5.1.3), where $r > 1$, $\mathfrak{B}$ is the invariant set given by (5.1.2), and L_B is the Lipschitz smoothness constant associated with $\mathfrak{B}$. In this setting, the residuals $\{G'(y_k)\}_{k \geq 0}$ converge to zero in the $\mathcal{L}^{-1}$ -norm at least ℓ^2 -fast. In other words,*

$$\sum_{k=0}^{\infty} \|G'(y_k)\|_{\mathcal{L}^{-1}}^2 < \infty.$$

Proof. Moving the summation term of (5.1.10) to the left hand side and dropping the other nonnegative terms, we have

$$\sum_{k=0}^N \|G'(y_k)\|_{\mathcal{L}^{-1}}^2 \leq \frac{2(1+\lambda)}{s(1-\lambda)} (G(x_0) - G^*).$$

Letting $N \rightarrow \infty$ completes the proof. $\square$

Of course, this result is far from optimal. An exponential convergence of the residuals in the $\mathcal{L}^{-1}$ -norm will be proved in Corollary 5.2.5.

5.2 Convergence of PAGD

We can now begin the proof of convergence *per se*. We begin with an estimate for the discrete time derivative of the potential energy.

Lemma 5.2.1 (discrete derivative of potential energy). *Let the objective $G : \mathbb{H} \rightarrow \mathbb{R}$ be μ -strongly convex and locally Lipschitz smooth. Suppose that the step size in Algorithm 2 satisfies (5.1.3), where $r > 1$, $\mathfrak{B}$ is the invariant set given by (5.1.2), and L_B is the Lipschitz smoothness constant associated with $\mathfrak{B}$. Then, we have that*

$$\frac{1}{\eta} \frac{G(x_{k+1}) - G(x_k)}{\sqrt{s}} \leq -\frac{\sqrt{s}}{2\eta} \|G'(y_k)\|_{\mathcal{L}^{-1}}^2 + \frac{1}{\theta} \langle G'(y_k), y_k - x_k \rangle - \frac{\eta}{2\sqrt{s}} \|x_k - y_k\|_{\mathcal{L}}^2.$$

Proof. Since we have an invariant set, $\mathfrak{B}$, we can utilize the local Lipschitz smoothness with respect to it. Combined with $sL_B \leq 1$, it leads to

$$G(x_{k+1}) = G(y_k - s\mathcal{L}^{-1}G'(y_k)) \leq G(y_k) - \frac{s}{2}\|G'(y_k)\|_{\mathcal{L}^{-1}}^2. \quad (5.2.1)$$

Using this estimate, the strong convexity of G yields

$$\begin{aligned} G(x_k) &\geq G(y_k) + \langle G'(y_k), x_k - y_k \rangle + \frac{\mu}{2}\|x_k - y_k\|_{\mathcal{L}}^2 \\ &\geq G(x_{k+1}) + \frac{s}{2}\|G'(y_k)\|_{\mathcal{L}^{-1}}^2 + \langle G'(y_k), x_k - y_k \rangle + \frac{\mu}{2}\|x_k - y_k\|_{\mathcal{L}}^2. \end{aligned}$$

Rearranging the last estimate, multiplying through by $1/\theta$, and recalling $\theta = \sqrt{s\mu}$, $\eta = \sqrt{\mu}$, we obtain the desired result. $\square$

We also need an analogue of (4.3.5), a certain relation derived from the scheme.

Lemma 5.2.2 (discrete analogue of (4.3.5)). *The iterates constructed by PAGD, as described in Algorithm 2, satisfy*

$$\begin{aligned} &\frac{\eta}{2\sqrt{s}} \left(\|v_{k+1} - x^*\|_{\mathcal{L}}^2 - \|v_k - x^*\|_{\mathcal{L}}^2 \right) + \frac{1}{\theta} \langle G'(y_k), y_k - x_k \rangle + \frac{\eta^2}{2} \|v_k - x^*\|_{\mathcal{L}}^2 \\ &- \frac{\eta}{2\sqrt{s}} \|v_{k+1} - v_k\|_{\mathcal{L}}^2 + \frac{1}{2s} \|y_k - x_k\|_{\mathcal{L}}^2 - \frac{\eta^2}{2} \|y_k - x^*\|_{\mathcal{L}}^2 + \langle G'(y_k), y_k - x^* \rangle = 0. \end{aligned} \quad (5.2.2)$$

Proof. Substituting (3.2.2) in (3.2.3), and using the relations (5.1.1), we have

$$v_{k+1} = x_k + \frac{1}{\theta}(y_k - x_k) - \frac{\sqrt{s}}{\eta}\mathcal{L}^{-1}G'(y_k) = x_k + v_k - y_k - \frac{\sqrt{s}}{\eta}\mathcal{L}^{-1}G'(y_k).$$

Rearranging, and multiplying through by $\frac{\eta}{\sqrt{s}}$, we obtain the discrete analogue of the ODE (4.3.1)

$$\eta \frac{v_{k+1} - v_k}{\sqrt{s}} + \eta \frac{y_k - x_k}{\sqrt{s}} + \mathcal{L}^{-1}G'(y_k) = 0. \quad (5.2.3)$$

The discrete analogue of $V(t)$ is $v_k - x^*$, so following the proof of Theorem 4.3.1, we now take the $\mathcal{L}$ -inner product of (5.2.3) with $v_k - x^*$ to obtain

$$\frac{\eta}{\sqrt{s}} (v_{k+1} - v_k, v_k - x^*)_{\mathcal{L}} + \frac{\eta}{\sqrt{s}} (y_k - x_k, v_k - x^*)_{\mathcal{L}} + \langle G'(y_k), v_k - x^* \rangle = 0. \quad (5.2.4)$$

Using (2.1.19), the first term can be rewritten as

$$\frac{\eta}{\sqrt{s}} (v_{k+1} - v_k, v_k - x^*)_{\mathcal{L}} = \frac{\eta}{2\sqrt{s}} \left(\|v_{k+1} - x^*\|_{\mathcal{L}}^2 - \|v_k - x^*\|_{\mathcal{L}}^2 \right) - \frac{\eta}{2\sqrt{s}} \|v_{k+1} - v_k\|_{\mathcal{L}}^2.$$

For the second term in (5.2.4), we use relations (5.1.1), and then the identity (2.1.18) to get

$$\begin{aligned} \frac{\eta}{\sqrt{s}} (y_k - x_k, v_k - x^*)_{\mathcal{L}} &= \frac{\eta}{\sqrt{s}} \frac{1}{\theta} (y_k - x_k, \theta v_k - \theta x^*)_{\mathcal{L}} \\ &= \frac{1}{s} (y_k - x_k, y_k - x_k + \theta(y_k - x^*))_{\mathcal{L}} \\ &= \frac{1}{2s} \|y_k - x_k\|_{\mathcal{L}}^2 + \frac{\eta^2}{2} \|v_k - x^*\|_{\mathcal{L}}^2 - \frac{\eta^2}{2} \|y_k - x^*\|_{\mathcal{L}}^2. \end{aligned}$$

Finally, for the third term of (5.2.4), we use (5.1.1) similarly to the above, then it follows

$$\begin{aligned} \langle G'(y_k), v_k - x^* \rangle &= \frac{1}{\theta} \langle G'(y_k), \theta v_k - \theta x^* \rangle \\ &= \frac{1}{\theta} \langle G'(y_k), y_k - x_k \rangle + \langle G'(y_k), y_k - x^* \rangle. \end{aligned}$$

Then, the desired result follows upon combining the last three identities. $\square$

We need one more relation between the iterates.

Lemma 5.2.3 (relation between iterates). *The iterates constructed by PAGD, as described in Algorithm 2, satisfy*

$$\frac{\eta}{2\sqrt{s}} \|v_{k+1} - v_k\|_{\mathcal{L}}^2 = \frac{\eta}{2\sqrt{s}} \|x_k - y_k\|_{\mathcal{L}}^2 + \frac{\sqrt{s}}{2\eta} \|G'(y_k)\|_{\mathcal{L}^{-1}}^2 + \langle G'(y_k), y_k - x_k \rangle.$$

Proof. Combine (3.2.3) and the relations (5.1.1), and then use (3.2.2) to obtain

$$\begin{aligned} v_{k+1} - v_k &= x_k + \frac{1}{\theta}(x_{k+1} - x_k) + \frac{1}{\theta}x_k - \left(1 + \frac{1}{\theta}\right)y_k \\ &= x_k - y_k + \frac{1}{\theta}(x_{k+1} - y_k) = x_k - y_k - \frac{\sqrt{s}}{\eta}\mathcal{L}^{-1}G'(y_k). \end{aligned}$$

Take $\mathcal{L}$ -norm square on both sides and then multiply by $\frac{\eta}{2\sqrt{s}}$. □

We are now in a position to prove the main result of this chapter, the exponential convergence of PAGD using energy arguments. The following result and its consequences are another of the main contributions of this work. To state it, we recall that the Lyapunov function of (4.1.1) $E : \mathbb{H}^2 \rightarrow \mathbb{R}$ is defined in (4.3.2) and that E_0 is its value at the initial state as mentioned in Theorem 4.3.1.

Theorem 5.2.4 (exponential decay). *Let the objective $G : \mathbb{H} \rightarrow \mathbb{R}$ be locally Lipschitz smooth and μ -strongly convex. If PAGD, as described in Algorithm 2, is applied to approximate $x^* = \operatorname{argmin}_{x \in \mathbb{H}} G(x)$ with a step size satisfying (5.1.3), where $r > 1$, $\mathfrak{B}$ is the invariant set given by (5.1.2), and L_B is the Lipschitz smoothness constant associated with $\mathfrak{B}$, then the Lyapunov function (4.3.2) decays exponentially along the iterates $\{x_k\}_{k \geq 0}$. More specifically, for $k \geq 0$, we have*

$$E(x_{k+1}, v_{k+1} - x^*) \leq (1 - \theta)E(x_k, v_k - x^*), \quad E(x_k, v_k - x^*) \leq (1 - \theta)^k E_0. \quad (5.2.5)$$

Proof. Define, for $k \geq 0$, $\mathcal{E}_k = (1 - \theta)^{-k}E(x_k, v_k - x^*)$, which is the discrete analogue of the exponentially inflated energy in the the proof of Theorem 4.3.1. To simplify notation, we set $C_{\theta,k} = (1 - \theta)^{-(k+1)} > 0$. Then, similarly to the ODE case, one can show the discrete time derivative of $\mathcal{E}_k$ is nonpositive as follows. First, we simply use the forward difference time

derivative, rearrange, and use Lemma 5.2.1 to get

$$\begin{aligned}
\frac{\mathcal{E}_{k+1} - \mathcal{E}_k}{\sqrt{s}} &= \frac{1}{\sqrt{s}} \left[(1 - \theta)^{-(k+1)} \left(\frac{1}{\eta} (G(x_{k+1}) - G^*) + \frac{\eta}{2} \|v_{k+1} - x^*\|_{\mathcal{L}}^2 \right) \right. \\
&\quad \left. - (1 - \theta)^{-k} \left(\frac{1}{\eta} (G(x_k) - G^*) + \frac{\eta}{2} \|v_k - x^*\|_{\mathcal{L}}^2 \right) \right] \\
&= C_{\theta,k} \left[\frac{1}{\eta} \frac{G(x_{k+1}) - G(x_k)}{\sqrt{s}} + (G(x_k) - G^*) \right. \\
&\quad \left. + \frac{\eta}{2\sqrt{s}} (\|v_{k+1} - x^*\|_{\mathcal{L}}^2 - \|v_k - x^*\|_{\mathcal{L}}^2) + \frac{\eta^2}{2} \|v_k - x^*\|_{\mathcal{L}}^2 \right] \\
&\leq C_{\theta,k} \left[-\frac{\sqrt{s}}{2\eta} \|G'(y_k)\|_{\mathcal{L}^{-1}}^2 - \frac{\eta}{2\sqrt{s}} \|x_k - y_k\|_{\mathcal{L}}^2 + (G(x_k) - G^*) \right. \\
&\quad \left. + \frac{1}{\theta} \langle G'(y_k), y_k - x_k \rangle + \frac{\eta}{2\sqrt{s}} (\|v_{k+1} - x^*\|_{\mathcal{L}}^2 - \|v_k - x^*\|_{\mathcal{L}}^2) \right. \\
&\quad \left. + \frac{\eta^2}{2} \|v_k - x^*\|_{\mathcal{L}}^2 \right].
\end{aligned}$$

We continue by using Lemma 5.2.2 and then Lemma 5.2.3, then it follows

$$\begin{aligned}
\frac{\mathcal{E}_{k+1} - \mathcal{E}_k}{\sqrt{s}} &\leq C_{\theta,k} \left[-\frac{\sqrt{s}}{2\eta} \|G'(y_k)\|_{\mathcal{L}^{-1}}^2 - \left(\frac{\eta}{2\sqrt{s}} + \frac{1}{2s} \right) \|x_k - y_k\|_{\mathcal{L}}^2 + (G(x_k) - G^*) \right. \\
&\quad \left. + \frac{\eta}{2\sqrt{s}} \|v_{k+1} - v_k\|_{\mathcal{L}}^2 + \frac{\eta^2}{2} \|y_k - x^*\|_{\mathcal{L}}^2 + \langle G'(y_k), x^* - y_k \rangle \right] \\
&= C_{\theta,k} \left[-\frac{1}{2s} \|y_k - x_k\|_{\mathcal{L}}^2 + (G(x_k) - G^*) + \langle G'(y_k), y_k - x_k \rangle \right. \\
&\quad \left. + \langle G'(y_k), x^* - y_k \rangle + \frac{\eta^2}{2} \|y_k - x^*\|_{\mathcal{L}}^2 \right].
\end{aligned}$$

Finally, add and subtract $G(y_k)$ from the last expression, and use the following estimates, which are simple rearrangements of the lower and upper quadratic traps,

$$\begin{aligned}
G(y_k) - G^* + \langle G'(y_k), x^* - y_k \rangle &\leq -\frac{\mu}{2} \|y_k - x^*\|_{\mathcal{L}}^2 \\
G(x_k) - G(y_k) + \langle G'(y_k), y_k - x_k \rangle &\leq \frac{L_B}{2} \|x_k - y_k\|_{\mathcal{L}}^2,
\end{aligned}$$

then we arrive at

$$\begin{aligned}
\frac{\mathcal{E}_{k+1} - \mathcal{E}_k}{\sqrt{s}} &\leq C_{\theta,k} \left[G(x_k) - G(y_k) + \langle G'(y_k), y_k - x_k \rangle + G(y_k) - G^* \right. \\
&\quad \left. + \langle G'(y_k), x^* - y_k \rangle - \frac{1}{2s} \|y_k - x_k\|_{\mathcal{L}}^2 + \frac{\eta^2}{2} \|y_k - x^*\|_{\mathcal{L}}^2 \right] \\
&\leq C_{\theta,k} \left[\frac{1}{2} \left(L_B - \frac{1}{s} \right) \|y_k - x_k\|_{\mathcal{L}}^2 \right].
\end{aligned}$$

The step size condition forces the last term to be nonpositive. Therefore, we conclude that $\{\mathcal{E}_k\}_{k \geq 0}$ is nonincreasing, from which we obtain (5.2.5). $\square$

The following estimates are evident.

Corollary 5.2.5 (rate of convergence). *In the setting of Theorem 5.2.4, we have that the iterates of PAGD, as described in Algorithm 2, converge to x^* , the minimizer of G , at an exponential rate. More specifically, for a suitable $r > 1$ the step size can be set $s = 1/L_B$ and, in this case, for $k \geq 0$,*

$$\frac{1}{\eta} (G(x_k) - G^*) + \frac{\eta}{2} \|x_k - x^*\|_{\mathcal{L}}^2 \leq (1 - \sqrt{\rho})^k E_0, \quad (5.2.6)$$

which implies

$$G(x_k) - G^* \leq (1 - \sqrt{\rho})^k \eta E_0, \quad \|x_k - x^*\|_{\mathcal{L}} \leq (1 - \sqrt{\rho})^{\frac{k}{2}} \sqrt{\frac{2E_0}{\eta}}. \quad (5.2.7)$$

Furthermore, we have exponential convergence in the $\mathcal{L}^{-1}$ -norm of the residuals: for $k \geq 0$,

$$\|G'(y_k)\|_{\mathcal{L}^{-1}} \leq 3L_B \sqrt{\frac{2E_0}{\eta}} (1 - \sqrt{\rho})^{\frac{k-1}{2}}. \quad (5.2.8)$$

Proof. We can choose an appropriate $r > 1$ so that the step size condition (5.1.3) reduces to $s \in (0, L_B^{-1}]$; see Remark 5.1.3. Estimate (5.2.6) and the first estimate of (5.2.7) follow from (5.2.5) upon setting $s = 1/L_B$. The second estimate of (5.2.7) follows by applying strong convexity of G to the first estimate of (5.2.7).

Next, from the estimate (5.2.1), one obtains

$$G^* \leq G(x_{k+1}) \leq G(y_k) - \frac{1}{2L_B} \|G'(y_k)\|_{\mathcal{L}^{-1}}^2, \quad (5.2.9)$$

from which, one obtains the following by rearranging and then using the upper quadratic trap

$$\|G'(y_k)\|_{\mathcal{L}^{-1}} \leq \sqrt{2L_B(G(y_k) - G^*)} \leq L_B \|y_k - x^*\|_{\mathcal{L}}.$$

In addition, we also have, from the definition of y_k and $0 < \lambda < 1$,

$$\|y_k - x^*\|_{\mathcal{L}} = \|x_k - x^* + \lambda(x_k - x_{k-1} \pm x^*)\|_{\mathcal{L}} \leq 2\|x_k - x^*\|_{\mathcal{L}} + \|x_{k-1} - x^*\|_{\mathcal{L}}.$$

Combining the last two estimates and using (5.2.7), we obtain (5.2.8). $\square$

Remark 5.2.6 (total energy). The exponential decrease of the “total energy” at every step, given in (5.2.6), does not imply that the “potential energy” $G(x_{k+1}) - G^*$ or the “kinetic energy” $\frac{\mu}{2} \|v_{k+1} - x^*\|_{\mathcal{L}}^2$ decay monotonically by themselves. Corollary 5.2.5 only asserts exponential bounds. The same is true for the decay of the $\mathcal{L}^{-1}$ -norm of the residuals $\|G'(y_k)\|_{\mathcal{L}^{-1}}$. In fact, the numerical illustrations of Chapter 6 show that these quantities may oscillate. $\blacksquare$

Remark 5.2.7 (matching convergence rates). As discussed in Chapter 3, in the case of G being locally Lipschitz smooth, μ -strongly convex, the (best) contraction factor for PGD is $(\frac{1-\rho}{1+\rho})^2$ while we have $1 - \sqrt{\rho}$ for PAGD (see Theorem 5.2.4), where we recall $\rho = \mu/L_B$ and $L_B > 0$ is the Lipschitz smoothness constant on some appropriate invariant set $\mathfrak{B}$. It must be pointed out that this rate for PGD is achieved by choosing a “particularly good” step size that is only available to PGD: $s = \frac{2}{L_B + \mu}$ (see [33, Theorem 2.1.15]). More specifically, we have a contraction factor $1 - s \frac{2\mu L_B}{L_B + \mu}$ for PGD provided $0 < s \leq \frac{2}{L_B + \mu}$. If one uses the step size $s = 1/L_B$, then the contraction factor for PGD turns out to be $\frac{1-\rho}{1+\rho}$. This choice makes it easier to see the rate match the continuous time model. Setting $p = 2$ in [48, SI (Supplement Information) Theorem H.2] we see that the gradient flow $\dot{X} = -\mathcal{L}^{-1}G'(X)$ has convergence rate $G(X(t)) - G^* \leq (G(X(0)) - G^*)e^{-\mu t}$. However, using an estimate available

to μ -strongly convex functions (see [33, Theorem 2.1.10 (2.1.19)]), we can do better to get $G(X(t)) - G^* \leq (G(X(0)) - G^*)e^{-2\mu t}$. Then, we see that setting $t = sk$ and $s = 1/L_B$ for the gradient flow, and assuming $\rho \ll 1$, the contraction factor can be approximated by

$$e^{-2\mu s} \approx 1 - 2\mu s = 1 - 2\rho,$$

which is close to $\frac{1-\rho}{1+\rho}$. Similarly, setting $t = \sqrt{s}k$ and $s = 1/L_B$ in (4.3.4) and referring to Corollary 5.2.5, we have the contraction factors

$$e^{-\sqrt{\mu s}} \approx 1 - \sqrt{\mu s} = 1 - \sqrt{\rho}$$

for the IVP (4.1.1), which matches that of PAGD. ■

Chapter 6

Numerical Experiments

In this chapter, we carry out a series of numerical experiments aimed at illustrating the theory that we have developed. In all our examples, we approximate the solution to the nonlinear PDE (6.1.1) by iteratively minimizing an energy related to this PDE. The approximate solution is computed using a pseudo-spectral method (see [12, 42]), which was implemented in an in-house Matlab R2016a© code. This pseudo-spectral code heavily uses the built-in `fft` and `ifft` Matlab internal routines to invert preconditioners and apply residuals.

The energy minimization is carried out with GD, AGD, PGD, or PAGD, where the algorithm terminates if one of the following cases is true:

- (a) the ∞ -norm (when the true solution is unknown) or the $\mathcal{L}_N$ -norm (when the true solution is known) of the search direction is smaller than a certain tolerance, which we will call *convergence*;
- (b) the norm being measured is larger than a certain upper tolerance, which we will call *blow up*;
- (c) the number of iterations reaches a certain number, which we will call *no convergence*.

In the conditions above, we mean by “search directions” the residual if the scheme does not involve a preconditioner. If the scheme involves a (discrete) preconditioner $\mathcal{L}_N$ (see (6.2.8) for definition), the search direction is the solution to $\mathcal{L}_N s = r$, where r is the residual. In all implementations, the initial guess is always zero.

6.1 The continuous problem

We approximate the solution to the following “nonlocal” PDE:

$$(-\Delta)^\alpha u + |u|^{p-2}u + tu = f \quad \text{in } \Omega = (0, 1)^2 \subset \mathbb{R}^2, \quad (6.1.1)$$

supplemented with periodic boundary conditions, where $\alpha > 0$, $p \geq 2$, and $t > 0$. Here and in what follows, all functions are real-valued except for the exponential functions appearing in Fourier series and Fourier coefficients. The nonlocal operator $(-\Delta)^\alpha$ is the *spectral* fractional Laplacian, which is defined via Fourier series as follows. For every $v \in L^2_{\text{per}}(\Omega)$, we have that

$$v(\mathbf{x}) = \sum_{\mathbf{m} \in \mathbb{Z}^2} \hat{v}_{\mathbf{m}} e^{2\pi i \mathbf{m} \cdot \mathbf{x}},$$

where the equality is in the $L^2(\Omega)$ -sense, $\mathbf{x} = (x, y) \in \bar{\Omega}$, $\mathbf{i} = \sqrt{-1}$, and

$$\hat{v}_{\mathbf{m}} = \int_{\Omega} v(\mathbf{x}) e^{-2\pi i \mathbf{m} \cdot \mathbf{x}} d\mathbf{x}, \quad \mathbf{m} \in \mathbb{Z}^2.$$

Thus, we define

$$(-\Delta)^\alpha v(\mathbf{x}) = \sum_{\mathbf{m} \in \mathbb{Z}^2} (4\pi^2 |\mathbf{m}|^2)^\alpha \hat{v}_{\mathbf{m}} e^{2\pi i \mathbf{m} \cdot \mathbf{x}},$$

provided that the sum is finite.

With this definition at hand, it is not difficult to see that (6.1.1) in its weak form, can be seen as the Euler-Lagrange equation for the functional

$$G(u) = \int_{\Omega} \left(\frac{1}{2} |(-\Delta)^{\frac{\alpha}{2}} u|^2 + \frac{1}{p} |u|^p + \frac{t}{2} |u|^2 - fu \right) d\mathbf{x}, \quad (6.1.2)$$

over the space $\mathbb{H} = H^\alpha_{\text{per}}(\Omega) \cap L^p(\Omega)$. It is well known that $\{e^{2\pi i \mathbf{m} \cdot \mathbf{x}}\}_{\mathbf{m} \in \mathbb{Z}^2}$ is an orthonormal basis of $L^2_{\text{per}}(\Omega)$. Then, $\mathbb{H}$ can be equivalently defined via

$$\mathbb{H} = \left\{ v \in L^p_{\text{per}}(\Omega) \left| \sum_{\mathbf{m} \in \mathbb{Z}^2} |\mathbf{m}|^\alpha |\hat{v}_{\mathbf{m}}|^2 < \infty \right. \right\}.$$

The existence and uniqueness of a weak solution to (6.1.1) is guaranteed for any $f \in L^{p'}_{\text{per}}(\Omega)$, where $1/p + 1/p' = 1$, since, in this case, the energy is well-defined, strictly convex, and coercive.

For the space $\mathbb{H}$ to possess a Hilbert structure, a restriction on p must be imposed depending on α . For ease of notation, let $(\cdot, \cdot)$ and $\|\cdot\|$ denote the $L^2(\Omega)$ -inner product and $L^2(\Omega)$ -norm respectively. A natural inner product on $H^\alpha_{\text{per}}(\Omega)$ is given by

$$(v, w)_{H^\alpha_{\text{per}}(\Omega)} = ((-\Delta)^{\frac{\alpha}{2}} v, (-\Delta)^{\frac{\alpha}{2}} w) + (v, w),$$

and its associated norm by

$$\|v\|_{H^\alpha_{\text{per}}(\Omega)} = \sqrt{(v, v)_{H^\alpha_{\text{per}}(\Omega)}}.$$

The following is a standard Sobolev embedding result. For a proof, see, e.g., [1, Theorem 7.34].

Proposition 6.1.1 (Sobolev embedding). *Let $\alpha \in (0, 1]$. For all $p \in [2, p^*]$ with $p^* = \frac{2}{1-\alpha}$ if $\alpha < 1$ or $p \in [2, \infty)$ if $\alpha = 1$, there exists $C_{\text{emb}} = C_{\text{emb}}(p, \alpha) > 0$ such that, for all $v \in H^\alpha_{\text{per}}(\Omega)$,*

$$\|v\|_{L^p(\Omega)} \leq C_{\text{emb}} \|v\|_{H^\alpha_{\text{per}}(\Omega)}. \quad (6.1.3)$$

We introduce the preconditioner $\mathcal{L}u = (-\Delta)^\alpha u + \nu u$,

where $\nu \geq 0$ is a free parameter, which induces an inner product

$$\langle \mathcal{L}u, v \rangle = \int_{\Omega} ((-\Delta)^{\frac{\alpha}{2}} u (-\Delta)^{\frac{\alpha}{2}} v + \nu uv) \, d\mathbf{x}. \quad (6.1.4)$$

Remark 6.1.2 (notation). As it is clear from its definition, the Lipschitz constant of G' depends on the norm being used. Thus, we will make a difference between the case with preconditioner and without it. $\hat{L}$ denotes the Lipschitz constant with respect to the preconditioner-induced norm $\|\cdot\|_{\mathcal{L}}$, while L is the constant with respect to the original norm $\|\cdot\|_{\mathbb{H}}$. ■

We investigate the properties of G in the following result.

Proposition 6.1.3 (properties of G). *Let G be given by (6.1.2) and the preconditioner $\mathcal{L}$ by (6.1.4). Then, G is strongly convex with respect to $\mathcal{L}$ -norm. If, in addition, p satisfies the conditions of Proposition 6.1.1, then G is locally Lipschitz smooth with respect to the $\mathcal{L}$ -norm.*

Proof. First, the action of G' is characterized by the following: for $v, w \in \mathbb{H}$,

$$\langle G'(v), w \rangle = ((-\Delta)^{\frac{\alpha}{2}} v, (-\Delta)^{\frac{\alpha}{2}} w) + t(v, w) + (|v|^{p-2} v, w) - (f, w).$$

Note also that the following estimates hold, which are a special case of [5, Lemma 2.1]: for $p > 1$, there exist $C_{p1}, C_{p2} > 0$, which depend only on p , such that for all $\xi, \eta \in \mathbb{R}$,

$$||\xi|^{p-2} \xi - |\eta|^{p-2} \eta| \leq C_{p1} |\xi - \eta| (|\xi| + |\eta|)^{p-2}, \quad (6.1.5)$$

$$(|\xi|^{p-2} \xi - |\eta|^{p-2} \eta)(\xi - \eta) \geq C_{p2} |\xi - \eta|^2 (|\xi| + |\eta|)^{p-2}. \quad (6.1.6)$$

Thus, using (6.1.6)

$$\langle G'(v) - G'(w), v - w \rangle \geq \|(-\Delta)^{\frac{\alpha}{2}}(v - w)\|^2 + t\|v - w\|^2 \geq \hat{\mu}\|v - w\|_{\mathcal{L}}^2,$$

where

$$\hat{\mu} = \min\{1, t/\nu\}. \quad (6.1.7)$$

Observe that this holds without referring to Sobolev embedding. Note also that this implies the coercivity of G with respect to $\mathcal{L}$ -norm, that is, $\lim_{\|v\|_{\mathcal{L}} \rightarrow \infty} G(v) = \infty$.

Next, thanks to coercivity, for any bounded, convex set $B \subset \mathbb{H}$ there exists $M_B \in \mathbb{R}$ such that

$$B \subset \{x \in \mathbb{H} \mid G(x) \leq M_B\}$$

Hence, for each $v \in B$, using Cauchy-Schwarz inequality and Young's inequality, it follows that there exists $\varepsilon > 0$ such that

$$\begin{aligned} M_B &\geq G(v) = \frac{1}{2} \|(-\Delta)^{\frac{\alpha}{2}} v\|^2 + \frac{1}{p} \|v\|_{L^p(\Omega)}^p + \frac{t}{2} \|v\|^2 - (f, v) \\ &\geq \frac{1}{2} \|(-\Delta)^{\frac{\alpha}{2}} v\|^2 + \frac{1}{p} \|v\|_{L^p(\Omega)}^p + \frac{t}{4} \|v\|^2 - \frac{1}{2\varepsilon} \|f\|^2. \end{aligned} \quad (6.1.8)$$

Rearranging this, we see that there exists $C_{f,t,p,B} > 0$ such that

$$\|v\|_{L^p(\Omega)} \leq C_{f,t,p,B} \quad \forall v \in B. \quad (6.1.9)$$

On the other hand, using (6.1.5), Hölder's inequality, and (6.1.9), we have, for all $v, w \in B$,

$$\begin{aligned} &\langle G'(v) - G'(w), v - w \rangle \\ &= \|(-\Delta)^{\frac{\alpha}{2}}(v - w)\|^2 + t\|v - w\|^2 + (|v|^{p-2}v - |w|^{p-2}w, v - w) \\ &\leq \|(-\Delta)^{\frac{\alpha}{2}}(v - w)\|^2 + t\|v - w\|^2 + C_{p1} \int_{\Omega} |v - w|^2 (|v| + |w|)^{p-2} d\mathbf{x} \\ &\leq \|(-\Delta)^{\frac{\alpha}{2}}(v - w)\|^2 + t\|v - w\|^2 + C_{p3} \|v - w\|_{L^p(\Omega)}^2 \left(\|v\|_{L^p(\Omega)}^{p-2} + \|w\|_{L^p(\Omega)}^{p-2} \right) \\ &\leq \|(-\Delta)^{\frac{\alpha}{2}}(v - w)\|^2 + t\|v - w\|^2 + 2C_{f,t,p,B}^{p-2} C_{p3} \|v - w\|_{L^p(\Omega)}^2, \end{aligned}$$

where $C_{p3} > 0$ is a constant reflecting the equivalence between $(|v| + |w|)^{p-2}$ and $|v|^{p-2} + |w|^{p-2}$.

Finally, owing to the restriction on p , Proposition 6.1.1 guarantees that

$$\|v - w\|_{L^p(\Omega)}^2 \leq C_{emb}^2 \|v - w\|_{H_{\text{per}}^{\alpha}(\Omega)}^2,$$

so that

$$\langle G'(v) - G'(w), v - w \rangle \leq \hat{L}_B \|(v - w)\|_{\mathcal{L}}^2,$$

with $\hat{L}_B = \max\{1, t/\nu, 2C_{f,t,p,B}^{p-2} C_{p3} C_{emb}^2\}$. □

Remark 6.1.4 (strong Lipschitz smoothness). The proof of Proposition 6.1.3 can be easily modified to show that G is locally Lipschitz smooth in the strong sense, i.e., (2.1.14) holds. ■

6.2 Discretization

We discretize the model problem (6.1.1) by introducing a uniform grid of points. To simplify the presentation, we choose $N \in \mathbb{N}$ with $N = 2K + 1$ for some integer $K \geq 1$. (The details for the case that N is even are only slightly more complicated.) Define $h = 1/N$, and introduce the grid domain

$$\Omega_N = \{(x_\ell, y_m) \in [0, 1]^2 \mid x_\ell = \ell h, y_m = mh, 0 \leq \ell, m \leq N\}.$$

For ease of notation, let us introduce

$$\begin{aligned} \mathbb{N}_N^2 &= \{\mathbf{m} = (m_1, m_2) \in \mathbb{Z}^2 \mid 1 \leq m_1, m_2 \leq N\} \\ \mathbb{Z}_K^2 &= \{\mathbf{r} = (r_1, r_2) \in \mathbb{Z}^2 \mid -K \leq r_1, r_2 \leq K\} \end{aligned}$$

Then, for $\mathbf{m} \in \mathbb{N}_N^2$, we can denote $\mathbf{x}_\mathbf{m} = (x_{m_1}, y_{m_2}) \in \Omega_N$. This notation must not be confused with that of the iterates of PAGD. Define the space of periodic grid functions

$$\begin{aligned} \mathbb{H}_N = \{v_N : \Omega_N \rightarrow \mathbb{R} \mid v_N(0, hm) = v_N(hN, hm), v_N(h\ell, 0) = v_N(h\ell, hN), \\ 0 \leq m, \ell \leq N\}, \end{aligned} \quad (6.2.1)$$

endowed with the L_N^2 -inner product

$$(v_N, w_N)_N = h^2 \sum_{\mathbf{m} \in \mathbb{N}_N^2} v_N(\mathbf{x}_\mathbf{m}) w_N(\mathbf{x}_\mathbf{m}). \quad (6.2.2)$$

More generally, for $p \geq 1$, define

$$\|w_N\|_{N,p} = \left(h^2 \sum_{\mathbf{m} \in \mathbb{N}_N^2} |w_N(\mathbf{x}_\mathbf{m})|^p \right)^{\frac{1}{p}}.$$

Given $w_N \in \mathbb{H}_N$, its *discrete Fourier transform* (DFT) is

$$\hat{w}_K(\mathbf{r}) = h^2 \sum_{\mathbf{s} \in \mathbb{N}_N^2} w_N(\mathbf{x}_\mathbf{s}) e^{-2\pi i \mathbf{r} \cdot \mathbf{x}_\mathbf{s}}, \quad \mathbf{r} \in \mathbb{Z}_K^2.$$

The *discrete fractional Laplacian* $(-\Delta_N)^\alpha : \mathbb{H}_N \rightarrow \mathbb{H}_N$ is defined by

$$[(-\Delta_N)^\alpha w_N](\mathbf{x}_\mathbf{m}) = \sum_{\mathbf{r} \in \mathbb{Z}_K^2} (4\pi^2 |\mathbf{r}|^2)^\alpha \hat{w}_K(\mathbf{r}) e^{2\pi i \mathbf{r} \cdot \mathbf{x}_\mathbf{m}}. \quad (6.2.3)$$

Finally, for $v_N, w_N \in \mathbb{H}_N$, the H_N^α -inner product is given by

$$(v_N, w_N)_{H_N^\alpha} = (v_N, w_N)_N + ((-\Delta_N)^{\frac{\alpha}{2}} v_N, (-\Delta_N)^{\frac{\alpha}{2}} w_N)_N, \quad (6.2.4)$$

and $\|w_N\|_{H_N^\alpha} = \sqrt{(w_N, w_N)_{H_N^\alpha}}$.

We comment that there are, at least, three different natural choices for the underlying inner product for $\mathbb{H}_N$: the L_N^2 -inner product, the H_N^α -inner product, and the $\mathcal{L}_N$ -inner product, (6.2.8). We will choose the first option, i.e., the L_N^2 -inner product for several reasons. To begin with, this is the way numerical experiments are usually done if no special distinction is made between representers of the residual with respect to multiple inner products. In addition, this illustrates the effect of preconditioning more vividly. For example, if we adopt H_N^α -inner product, this leads to a preconditioned scheme in disguise: finding a representer of the residual with respect to this inner product is equivalent to using $\mathcal{L}_N$ -inner product (6.2.8) with $\nu_N = 1$. On the other hand, L_N^2 -inner product leads to truly *non-preconditioned* schemes such as GD or AGD.

We need a fractional version of Parseval's theorem to complete our numerical analysis.

Proposition 6.2.1 (Parseval). *Let $\alpha, \beta \geq 0$. For grid functions $v_N, w_N \in \mathbb{H}_N$, let $v = I_K(v_N), w = I_K(w_N) \in \mathcal{P}_K$, where $\mathcal{P}_K$ is the space of Ω -periodic trigonometric polynomials*

of degree at most K and $I_K : \mathbb{H}_N \rightarrow \mathcal{P}_K$ maps a grid function, v_N , to a (unique) trigonometric polynomial, $v = I_K(v_N)$, that interpolates v_N , i.e., $v_N(\mathbf{x}_m) = v(\mathbf{x}_m)$ for all $\mathbf{x}_m \in \Omega_N$. Then,

$$((-\Delta_N)^\alpha v_N, (-\Delta_N)^\beta w_N)_N = ((-\Delta)^\alpha v, (-\Delta)^\beta w).$$

Consequently,

$$\|v_N\|_{H_N^\alpha} = \|I_K(v_N)\|_{H_{\text{per}}^\alpha(\Omega)}. \quad (6.2.5)$$

Proof. This is a direct consequence of the orthonormality of $\{e^{2\pi i \mathbf{n} \cdot \mathbf{x}_m}\}_{\mathbf{n} \in \mathbb{Z}_K^2} \subset \mathbb{H}_N$ with respect to $(\cdot, \cdot)_N$ and that of $\{e^{2\pi i \mathbf{n} \cdot \mathbf{x}}\}_{\mathbf{n} \in \mathbb{Z}_K^2} \subset \mathcal{P}_K$ with respect to $(\cdot, \cdot)$ once we realize that $v(\mathbf{x}) = I_K(v_N)(\mathbf{x}) = \sum_{\mathbf{r} \in \mathbb{Z}_K^2} \hat{v}_K(\mathbf{r}) e^{2\pi i \mathbf{r} \cdot \mathbf{x}}$ and that the same relation holds between w and w_N : in words, the linear combination of $\{e^{2\pi i \mathbf{n} \cdot \mathbf{x}}\}_{\mathbf{n} \in \mathbb{Z}_K^2}$ with the DFT of v_N yields the unique trigonometric polynomial that interpolates it. $\square$

After having introduced all this notation, we can write our discrete problem as: given $f_N \in \mathbb{H}_N$, find $u_N \in \mathbb{H}_N$ such that

$$(-\Delta_N)^\alpha u_N + |u_N|^{p-2} u_N + t u_N = f_N. \quad (6.2.6)$$

In this problem, $f_N \in \mathbb{H}_N$ is some approximation of the problem data f . For example, if f is continuous, $f_N(\mathbf{x}_m) = f(\mathbf{x}_m)$ is a natural option, and if f is only an $L^2(\Omega)$ -function, then the sampling at the nodes of the $L^2(\Omega)$ -projection of f onto $\mathcal{P}_K$, the trigonometric polynomial of degree at most K , is natural although these two may not agree even if one starts with the same continuous function f . In fact, the difference between these two possibilities is very small if f is smooth and its derivatives are periodic (see [12, pp. 44–45]).

Our discrete problem has a similar energy structure to the continuous problem. It is the Euler-Lagrange equation of the following functional

$$G_N(v_N) = \frac{1}{2} \|(-\Delta_N)^{\frac{\alpha}{2}} v_N\|_N^2 + \frac{1}{p} \|v_N\|_{N,p}^2 + \frac{t}{2} \|v_N\|_N^2 - (f_N, v_N)_N. \quad (6.2.7)$$

We introduce a (discrete) preconditioner

$$\mathcal{L}_N = (-\Delta_N)^\alpha + \nu_N \text{id}_N, \quad (6.2.8)$$

where $\nu_N > 0$ and $\text{id}_N : \mathbb{H}_N \rightarrow \mathbb{H}_N$ is the identity map. The parameter $\nu_N > 0$ will be determined later. This preconditioner induces an inner product on $\mathbb{H}_N$ given by

$$(v_N, w_N)_{\mathcal{L}_N} = \nu_N (v_N, w_N)_N + \left((-\Delta_N)^{\frac{\alpha}{2}} v_N, (-\Delta_N)^{\frac{\alpha}{2}} w_N \right)_N, \quad (6.2.9)$$

and an associated norm $\|v_N\|_{\mathcal{L}_N} = \sqrt{(v_N, v_N)_{\mathcal{L}_N}}$. It is desirable that the convergence of our scheme does not deteriorate as we refine the grid points. We can ensure this under a certain restriction on p . The following proposition provides an important tool for that purpose.

Proposition 6.2.2 (discrete Sobolev embedding). *Let $\alpha \in (0, 1]$. For all $p \in [2, p^*]$ with $p^* = \frac{2}{1-\alpha}$ if $\alpha < 1$ or for all $p \in [2, \infty)$ if $\alpha = 1$, there exists a constant $C_{p,\alpha} > 0$ such that, for all $v_N \in \mathbb{H}_N$,*

$$\|v_N\|_{N,p} \leq C_{p,\alpha} \|v_N\|_{H_N^\alpha}. \quad (6.2.10)$$

$C_{p,\alpha}$ is independent of v_N and N .

Proof. It is well-known that, in our setting, the stability of the norm $\|\cdot\|_{N,p}$, is guaranteed. That is, there exists a constant C , which depends only on the dimension of Ω , such that $\|v_N\|_{N,p} \leq C \|v\|_{L^p(\Omega)}$ for all $v_N \in \mathbb{H}_N$; see, for instance, [29, Lemma 2.48], where $v = I_K(v_N)$ as introduced in Proposition 6.2.1. In conjunction with the Sobolev embedding at the continuous level, (6.1.3), and the Parseval's identity, (6.2.5), we have, for any $v_N \in \mathbb{H}_N$,

$$\|v_N\|_{N,p} \leq C \|v\|_{L^p(\Omega)} \leq C C_{p,\alpha} \|v\|_{H_{\text{per}}^\alpha(\Omega)} = C C_{p,\alpha} \|v_N\|_{H_N^\alpha}. \quad \square$$

The following result addresses dimension-independence of the (inverse) condition number as well as the structure of G_N that is needed to apply the theory we have developed in Chapter 5. Note that, in the following statement, the sublevel sets of G_N provide a

compatible way to describe bounded, convex sets when we consider multiple resolutions since, strictly speaking, for different values of N , functions in $\mathbb{H}_N$ may not be directly comparable.

Theorem 6.2.3 (properties of G_N). *Let the space of grid functions $\mathbb{H}_N$ be given by (6.2.1) and the preconditioner $\mathcal{L}_N$ by (6.2.8). Then, the energy functional $G_N : \mathbb{H}_N \rightarrow \mathbb{R}$ defined by (6.2.7) is strongly convex and locally Lipschitz smooth with respect to the $\mathcal{L}_N$ -norm. Moreover, the strong convexity constant $\hat{\mu}_N$ is independent of N . Suppose, in addition, that p satisfies the conditions of Proposition 6.2.2 and that f_N is defined in a stable manner when we pose the discrete problem (6.2.6), i.e., there exists $C > 0$ independent of N such that*

$$\|f_N\|_{N,2} \leq C \|f\|_{L^2(\Omega)}.$$

Then, the local Lipschitz smoothness constant $\hat{L}_N$ is also independent of N in the sense that, for each $M \in \mathbb{R}$, G_N is $\hat{L}_N$ -Lipschitz smooth on the sublevel set

$$\{v_N \in \mathbb{H}_N \mid G(v_N) \leq M\}$$

with $\hat{L}_N$ independent of N . Consequently, the (inverse) condition number $\hat{\mu}_N/\hat{L}_N$ with respect to the $\mathcal{L}_N$ -norm stays away from 0 as $N \rightarrow \infty$.

Proof. The proof of the strong convexity is parallel to that of Proposition 6.1.3. The proof of local Lipschitz smoothness is also parallel, but we need the assumed stability of $\|f_N\|_{N,2}$ to proceed from (6.1.8) to (6.1.9). Finally, to complete the proof, we simply replace the embedding constant C_{emb} with its discrete counterpart $C_{p,\alpha}$ as given in Proposition 6.2.2. $\square$

Since the (inverse) condition number $\hat{\mu}_N/\hat{L}_N$ governs the rate of convergence of PAGD (Corollary 5.2.5), the previous theorem guarantees that one can achieve the same rate of convergence even if we refine the number of grid points $N \rightarrow \infty$. However, this is true in terms of the number of iterations, but the wall clock time will take longer as the refinement is conducted.

Being in finite dimensions, G_N is also strongly convex, and locally Lipschitz smooth with respect to *any norm*, for instance $\|\cdot\|_N$. The constants in this case, however, are different and depend on the dimension of $\mathbb{H}_N$, which obviously depends on the number of grid points, and

thus on N . We label them μ_N and L_N to distinguish them from the dimension-independent constants $\hat{\mu}_N$ and $\hat{L}_N$ respectively.

6.3 A problem with a manufactured solution


In this first experiment, we solve (6.2.6) by minimizing the energy (6.2.7). To compute the errors and energies, the following manufactured solution is used

$$u_N(\mathbf{x}_m) = \exp \left(\sin 2\pi \left(x_{m_1} - \frac{1}{4} \right) + \sin 4\pi \left(y_{m_2} - \frac{3}{8} \right) \right).$$


We set $\alpha = 0.5$, $p = 4$, $t = 1$, $N = 64$, and found, experimentally, that the values $\nu_N = 1.2$, $\mu_N = 1$ are optimal, while we set $\hat{\mu}_N = 5/6 = \min\{1, t/\nu_N\}$ in view of (6.1.7). To specify step sizes, recall the step size rules that theoretically guarantee convergence (see Chapter 3): $s = 2/(L_N + \mu_N)$ for GD, $s = 1/L_N$ for AGD, $s = 2/(\hat{L}_N + \hat{\mu}_N)$ for PGD, and $s = 1/\hat{L}_N$ for PAGD. Step sizes are set by these relations with $L_N = 500$ and $\hat{L}_N = 20$, which are also experimentally proved to be optimal. However, it must be noted that this is just a way of setting step sizes. We do not really know neither whether the values for L_N or $\hat{L}_N$ are the Lipschitz constants of the corresponding energy functionals nor whether the aforementioned step size rules give the optimal results even if we knew the Lipschitz constants. In fact, our last experiment suggests that larger step sizes than what is theoretically proven seem to work.

Figure 6.1 shows the performance of GD, AGD, PGD, and PAGD when used to solve (6.2.6) by minimizing (6.2.7), where the data is as described above. The stopping criteria take the following parameters: the tolerance is 10^{-8} , the upper tolerance is 10^{10} , and the maximum number of iterations is 200.

Figure 6.1 (a) shows the decay of the objective, $G_N(x_k) - G_N(u_N)$ which is, up to a constant, the same as the decay of the potential energy, for all four schemes. Here k is the number of iterations. Figure 6.1 (b) shows the decay of the $\mathcal{L}_N$ -norm of the errors. Notice that PAGD performs significantly better than all the other methods.



(a) Objective, $G_N(x_k) - G_N(u_N)$, plot of GD, AGD, PGD, and PAGD. (b) $\mathcal{L}_N$ -norm of errors that are generated by GD, AGD, PGD, and PAGD.



(c) Potential, kinetic, and total energy plot of GD and AGD.

Figure 6.1: Objective, error, and energy decay plots for GD, AGD, PGD, and PAGD. They are implemented to solve (6.2.6) by minimizing (6.2.7) ($\alpha = 0.5$, $p = 4$, $t = 1$, $N = 64$, $\nu_N = 1.2$, $L_N = 500$, $\mu_N = 1$, $\hat{L}_N = 20$, $\hat{\mu}_N = 5/6 = \min\{1, t/\nu_N\}$, and step sizes are set via $s = 2/(L_N + \mu_N)$ for GD, $s = 1/L_N$ for AGD, $s = 2/(\hat{L}_N + \hat{\mu}_N)$ for PGD, and $s = 1/\hat{L}_N$ for PAGD). The vertical axes (logarithmic scale) show the value of the objective, $\mathcal{L}_N$ -norm of errors, or various energies while the horizontal axis (linear scale) shows the number of iterations.

Figure 6.1 (c) shows the performance of GD and AGD. Since these schemes do not involve a preconditioner, the corresponding total energy is defined by

$$E_N(x_k, v_k) = \frac{1}{\eta_N}(G_N(x_k) - G_N(u_N)) + \frac{\eta_N}{2}\|v_k - u_N\|_N^2, \quad (6.3.1)$$

where k is the number of iterations and $\eta_N = \sqrt{\mu_N}$. The first and second terms can be understood as potential and kinetic energy respectively. Figure 6.1 (c) shows the decay of various energies for nonpreconditioned schemes. This figure better illustrates our analysis of the previous chapter than the preconditioned ones since they converge slower. As expected, AGD performs substantially better than GD. The total energy of AGD decreases steadily and exponentially fast. Notice that the vertical axis is in logarithmic scale. This matches what is predicted by the theory in Theorem 5.2.4. Observe also that the potential and kinetic energies of AGD, by themselves, oscillate; see Remark 5.2.6. The physical analogy for AGD described in Remark 4.1.1 is clear from this picture. A fraction of the potential energy is converted to kinetic energy and they fluctuate as the mechanical system converges to equilibrium.

6.4 A problem where the solution is unknown

In this second experiment we, again, solve (6.2.6) by minimizing the energy (6.2.7). The discrete right hand side f_N is given by


$$f_N(\mathbf{x}_m) = \exp(\sin 2\pi(x_{m_1} - 0.25) + \sin 2\pi(y_{m_2} - 0.25)). \quad (6.4.1)$$

The parameters of the PDE are set to $\alpha = 0.5$, $p = 10$, and $t = 1$. Observe that for these values of α and p we do not have that $H_{\text{per}}^\alpha(\Omega) \hookrightarrow L^p(\Omega)$. We found, experimentally, that the choice $\nu_N = 0.9$ is optimal for the preconditioner. We also set $\mu_N = 1$ and $\hat{\mu}_N = 1 = \min\{1, t/\nu_N\}$ in view of (6.1.7) as before. Step sizes are set in the same way as in the previous experiment with $L_N = 300$ or 3000 , and $\hat{L}_N = 9$. The values of μ_N , L_N , and $\hat{L}_N$ were experimentally found to be optimal except for $L_N = 3000$. That is, they yield the best


convergence rate with all other parameters being fixed. A more detailed explanation about $L_N = 3000$ will follow. Two different degrees of resolution are used to show the dimension dependence of nonpreconditioned schemes. The stopping criterion parameters are as before.

Figure 6.2 shows the ∞ -norm of the search directions for GD, AGD, PGD, and PAGD with varying degrees of resolution and with two different step sizes for GD and AGD, which are determined by the same step size rules as in the previous experiment with $L_N \in \{300, 3000\}$. In Figure 6.2 (a), we observe a similar performance as in Figure 6.1. Recall that we do not have Sobolev embedding. Thus, one can expect the Lipschitz constant L_N , hence the step size, to depend on the number of grid points. In fact, theory predicts that even $\hat{L}_N$ depends on it. However, for $\hat{L}_N$, such dependence is not observed within the range of N that we have chosen. We see that the step size for convergence indeed depends on N in Figure 6.2 (b). As we increase the resolution of the grid from $N = 64$ to $N = 512$, nonpreconditioned schemes become unstable. Figure 6.2 (c) shows that the stability of GD and AGD is recovered after L_N is increased from 300 to 3000, which amounts to decreasing the step size to roughly a tenth of the old one. ($L_N = 3000$ is not optimally chosen).

Figure 6.3 shows the dependence of L_N , hence the step size, on the number of grid points with the same experiment. However, here we use different tolerances and a different maximum number of iterations to best illustrate the dependence. For $N \in \{16, 32, 64, 128, 256, 512\}$, Figure 6.3 records the number of iterations for ∞ -norm of the search direction generated by each scheme to reach a tolerance 10^{-3} (“convergence”) or the maximum number of iterations, which is set to be 1000, if it does not reach the tolerance (“no convergence”). If the ∞ -norm of the search direction reaches an upper tolerance 10^8 , the algorithm records the number of iteration taken as 1100, which indicates “blowing up.” Figure 6.3 (a) shows when the step sizes of the nonpreconditioned schemes correspond to $L_N = 300$ and those of the preconditioned ones correspond to $\hat{L}_N = 9$. GD and AGD converge until $N = 64$. However, they become unstable for $N \geq 128$. Figure 6.3 (b) shows the same experiment with smaller step sizes, which correspond to $L_N = 3000$. In this case, we recover the stability of GD and AGD.




(a) ∞ -norm of search direction ($N = 64, L_N = 300$). (b) ∞ -norm of search direction ($N = 512, L_N = 300$).




(c) ∞ -norm of search direction ($N = 512, L_N = 3000$).

Figure 6.2: ∞ -norm plots of the search directions for GD, AGD, PGD, and PAGD. They are implemented to solve (6.2.6) by minimizing (6.2.7) with varying resolutions $N \in \{64, 512\}$ and varying step sizes for GD and AGD; $s = 2/(L_N + \mu_N)$ for GD and $s = 1/L_N$ for AGD with $L_N \in \{300, 3000\}$. The other parameters are set to $\alpha = 0.5$, $p = 10$, $t = 1$, $\nu_N = 0.9$, $\mu_N = 1$, $\hat{\mu}_N = 1 = \min\{1, t/\nu_N\}$, $s = 2/(\hat{L}_N + \hat{\mu}_N)$ for PGD, and $s = 1/\hat{L}_N$ for PAGD with $\hat{L}_N = 9$. The horizontal axis (linear scale) represents the number of iterations. The vertical axis (logarithmic scale) represents ∞ -norm of the search directions.



(a) Number of iteration to reach a tolerance ($L_N = 300$, $\hat{L}_N = 9$).



(b) Number of iteration to reach a tolerance ($L_N = 3000$, $\hat{L}_N = 9$).

Figure 6.3: Number of iterations for ∞ -norm of the search directions to reach the tolerance 10^{-3} for GD, AGD, PGD, and PAGD. They are implemented to solve (6.2.6) by minimizing (6.2.7) with varying resolutions $N = 16, 32, 64, 128, 256, 512$ and varying $L_N = 300, 3000$ ($\alpha = 0.5$, $p = 10$, $t = 1$, $\nu_N = 0.9$, $\mu_N = 1$, L_N as indicated in the subfigures, $\hat{\mu}_N = 1 = \min\{1, t/\nu_N\}$, $\hat{L}_N = 9$). The horizontal axis represents the degrees of resolution, N . The vertical axis represents the minimum of the number of iterations for the ∞ -norm of the search directions to reach the tolerance 10^{-3} (convergence) or 1000 iterations. The number of iterations being 1100 means that the ∞ -norm of the search directions have reached the upper tolerance 10^8 (blow up).

Table 6.1: The minimal number of iterations needed for the ∞ -norm of the search direction of PGD and PAGD to reach a tolerance of 10^{-9} and the values of ν_N and s (step size) that led to the minimum iterations for a range of values of α . They are implemented to solve (6.2.6) by minimizing the energy (6.2.7). $N = 64$, $\alpha \in \{0.1j \mid j = 1, 2, 3, \dots, 10\} \cup \{1.5, 2.0, 2.5, 3.0\} \subset (0, 3]$, $p = 6$, $t = 1$, $\hat{\mu}_N = \min\{1, t/\nu_N\}$, and f_N is given by (6.4.1). For each value of α , we consider $\nu_N \in \{0.1j \mid j = 1, 2, 3, \dots, 100\} \subset (0, 10]$ and $s \in \{0.01j \mid j = 1, 2, 3, \dots, 200\} \subset (0, 2]$. Among the possible 20,000 possible combinations of ν_N and s , we display the values that give the minimal number of iterations.

α	PGD			PAGD		
	# iterations	ν_N	step size	# iterations	ν_N	step size
0.1	64	1.0	0.20	38	0.9	0.14
0.2	50	1.1	0.25	32	1.0	0.18
0.3	39	1.2	0.31	29	1.1	0.22
0.4	29	2.6	0.57	26	1.2	0.26
0.5	22	2.8	0.66	24	1.3	0.30
0.6	16	4.1	0.97	20	5.5	0.83
0.7	13	3.4	0.90	17	5.2	0.91
0.8	11	4.6	1.04	15	4.2	0.88
0.9	12	3.8	0.89	12	5.0	0.96
1.0	10	4.0	0.95	12	4.3	0.92
1.5	9	4.5	0.97	11	4.5	0.97
2.0	8	4.8	1.03	10	4.5	0.96
2.5	8	4.1	0.88	9	4.2	0.90
3.0	8	4.1	0.88	9	4.2	0.90

6.5 A comparison between PGD and PAGD

In this final collection of experiments, we aim at comparing the performance of PGD and PAGD in different scenarios. To do so, we solve the discrete problem (6.2.6) by minimizing the energy (6.2.7) with the right hand side given by (6.4.1) as before. The problem parameters are set as $\alpha \in \{0.1j \mid j = 1, 2, 3, \dots, 10\} \cup \{1.5, 2.0, 2.5, 3.0\} \subset (0, 3]$, $p = 6$, and $t = 1$. We set $N = 64$ and $\hat{\mu}_N = \min\{1, t/\nu_N\}$. Then, for each value of α (column 1 of Table 6.1), PGD and PAGD are applied with $\nu_N \in \{0.1j \mid j = 1, 2, 3, \dots, 100\} \subset (0, 10]$ and the step size $s \in \{0.01j \mid j = 1, 2, 3, \dots, 200\} \subset (0, 2]$. Observe that neither Algorithm 1 nor Algorithm 2 require knowledge of $\hat{L}_N$ a priori. Instead, we directly set the step size in this last experiment. Among these 20,000 possible values of ν_N and s , the minimal number of iterations for the ∞ -norm of the search direction generated by PGD and PAGD to reach a tolerance of 10^{-9} (convergence) is recorded (column 2 and column 5 of Table 6.1, respectively). A pair of values, ν_N and s , that led to the minimal number of iterations is also recorded (columns 3 and 4 of Table 6.1 for PGD and columns 6 and 7 of Table 6.1 for PAGD). There can be multiple such pairs. If this is the case, the pair (ν_N, s) that comes the first in the lexicographic order is recorded.

As we can see from Table 6.1, for the nonlocal PDE (6.2.6) with small α ($\alpha = 0.1, 0.2, 0.3, 0.4$), PAGD performs better than PGD when they are implemented with their own best pair of parameters ν_N and s among those pairs that were considered. In particular, in the cases of $\alpha = 0.1, 0.2, 0.3$, PAGD outperforms PGD while the best values of ν_N for the two schemes are similar, hence directly comparing their performances roughly make sense. An interesting thing, however, is that one cannot say that PAGD is *always* better than PGD. In fact, for the remaining values of α , PGD takes fewer iterations to converge in the aforementioned sense than PAGD provided they are equipped with their “best” parameters for each method. It must be noted that this result does not contradict our theory. The theory only tells us some upper bounds about the rate of convergence of the two schemes within a certain range of step size when they involve the same preconditioner. It does not explain what happens outside of that. The result provided here perhaps illustrates the latter

case. In any case, we can see an improvement in the convergence of PAGD compared to PGD for “harder” problems ((6.2.6) with small α), where a stronger nonlocality is involved.

Bibliography

- [1] Adams, R. A. and Fournier, J. J. F. (2003). *Sobolev spaces*, volume 140 of *Pure and Applied Mathematics (Amsterdam)*. Elsevier/Academic Press, Amsterdam, second edition. [48](#)
- [2] Al-Baali, M. (1985). Descent property and global convergence of the Fletcher-Reeves method with inexact line search. *IMA J. Numer. Anal.*, 5(1):121–124. [7](#)
- [3] Allen-Zhu, Z. and Orecchia, L. (2017). Linear coupling: an ultimate unification of gradient and mirror descent. In *8th Innovations in Theoretical Computer Science Conference*, volume 67 of *LIPICs. Leibniz Int. Proc. Inform.*, pages Art. No. 3, 22. Schloss Dagstuhl. Leibniz-Zent. Inform., Wadern. [8](#)
- [4] Attouch, H., Goudou, X., and Redont, P. (2000). The heavy ball with friction method. I. The continuous dynamical system: global exploration of the local minima of a real-valued function by asymptotic analysis of a dissipative dynamical system. *Commun. Contemp. Math.*, 2(1):1–34. [4](#), [8](#), [22](#), [73](#)
- [5] Barrett, J. W. and Liu, W. B. (1993). Finite element approximation of the p -Laplacian. *Math. Comp.*, 61(204):523–537. [2](#), [49](#)
- [6] Beck, A. (2017). *First-order methods in optimization*, volume 25 of *MOS-SIAM Series on Optimization*. Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA; Mathematical Optimization Society, Philadelphia, PA. [18](#)
- [7] Benyamin, M., Calder, J., Sundaramoorthi, G., and Yezzi, A. (2020). Accelerated variational PDEs for efficient solution of regularized inversion problems. *J. Math. Imaging Vision*, 62(1):10–36. [6](#)
- [8] Bertsekas, D. P. (1999). *Nonlinear programming*. Athena Scientific Optimization and Computation Series. Athena Scientific, Belmont, MA, second edition. [2](#)
- [9] Bhaya, A. and Kaszkurewicz, E. (2004). Steepest descent with momentum for quadratic functions is a version of the conjugate gradient method. *Neural Networks*, 17(1):65–71. [6](#)

- [10] Boyd, S. and Vandenberghe, L. (2004). *Convex optimization*. Cambridge University Press, Cambridge. [18](#)
- [11] Calder, J. and Yezzi, A. (2019). PDE acceleration: a convergence rate analysis and applications to obstacle problems. *Res. Math. Sci.*, 6(4):Paper No. 35, 30. [6](#), [73](#)
- [12] Canuto, C., Hussaini, M. Y., Quarteroni, A., and Zang, T. A. (2006). *Spectral methods*. Scientific Computation. Springer-Verlag, Berlin. Fundamentals in single domains. [46](#), [53](#)
- [13] Cartan, H., Moore, J., Husemoller, D., and Maestro, K. (2017). *Differential Calculus on Normed Spaces: A Course in Analysis*. CreateSpace Independent Publishing Platform. [24](#), [25](#)
- [14] Chen, L., Hu, X., and Wise, S. M. (2020). Convergence analysis of the fast subspace descent method for convex optimization problems. *Math. Comp.*, 89(325):2249–2282. [18](#), [19](#), [34](#)
- [15] Ciarlet, P. G. (1989). *Introduction to numerical linear algebra and optimisation*. Cambridge Texts in Applied Mathematics. Cambridge University Press, Cambridge. With the assistance of Bernadette Miara and Jean-Marie Thomas, Translated from the French by A. Buttigieg. [1](#)
- [16] Ciarlet, P. G. (2013). *Linear and nonlinear functional analysis with applications*. Society for Industrial and Applied Mathematics, Philadelphia, PA. [28](#)
- [17] Dai, Y. H. and Yuan, Y. (2001). An efficient hybrid conjugate gradient method for unconstrained optimization. volume 103, pages 33–47. Optimization and numerical algebra (Nanjing, 1999). [7](#)
- [18] Evans, L. C. (2010). *Partial differential equations*, volume 19 of *Graduate Studies in Mathematics*. American Mathematical Society, Providence, RI, second edition. [2](#)
- [19] Feng, W., Guan, Z., Lowengrub, J., Wang, C., Wise, S. M., and Chen, Y. (2018). A uniquely solvable, energy stable numerical scheme for the functionalized Cahn-Hilliard equation and its convergence analysis. *J. Sci. Comput.*, 76(3):1938–1967. [2](#)

- [20] Feng, W., Salgado, A. J., Wang, C., and Wise, S. M. (2017). Preconditioned steepest descent methods for some nonlinear elliptic equations involving p-Laplacian terms. *J. Comput. Phys.*, 334:45–67. [4](#), [19](#)
- [21] Fletcher, R. (1987). *Practical methods of optimization*. A Wiley-Interscience Publication. John Wiley & Sons, Ltd., Chichester, second edition. [7](#)
- [22] Fletcher, R. and Reeves, C. M. (1964). Function minimization by conjugate gradients. *Comput. J.*, 7:149–154. [7](#)
- [23] Fox, L., Huskey, H. D., and Wilkinson, J. H. (1948). Notes on the solution of algebraic linear simultaneous equations. *Quart. J. Mech. Appl. Math.*, 1:149–173. [6](#)
- [24] Goudou, X. and Munier, J. (2009). The gradient and heavy ball with friction dynamical systems: the quasiconvex case. *Math. Program.*, 116(1-2, Ser. B):173–191. [4](#), [8](#), [73](#)
- [25] Grippo, L. and Lucidi, S. (1997). A globally convergent version of the Polak-Ribière conjugate gradient method. *Math. Programming*, 78(3, Ser. A):375–391. [7](#)
- [26] Hager, W. W. and Zhang, H. (2005). A new conjugate gradient method with guaranteed descent and an efficient line search. *SIAM J. Optim.*, 16(1):170–192. [7](#)
- [27] Hager, W. W. and Zhang, H. (2006). A survey of nonlinear conjugate gradient methods. *Pac. J. Optim.*, 2(1):35–58. [7](#)
- [28] Hestenes, M. R. and Stiefel, E. (1952). Methods of conjugate gradients for solving linear systems. *J. Research Nat. Bur. Standards*, 49:409–436 (1953). [6](#)
- [29] Jovanović, B. S. and Süli, E. (2014). *Analysis of finite difference schemes*, volume 46 of *Springer Series in Computational Mathematics*. Springer, London. For linear partial differential equations with generalized solutions. [54](#)
- [30] Laborde, M. and Oberman, A. (2020). A Lyapunov analysis for accelerated gradient methods: from deterministic to stochastic case. In Chiappa, S. and Calandra, R., editors, *Proceedings of the Twenty Third International Conference on Artificial Intelligence and*

- Statistics*, volume 108 of *Proceedings of Machine Learning Research*, pages 602–612, Online. PMLR. [6](#), [8](#), [73](#)
- [31] Luo, H. and Chen, L. (2020). From differential equation solvers to accelerated first-order methods for convex optimization. arXiv:1909.03145. [5](#), [8](#), [73](#)
- [32] Nesterov, Y. E. (1983). A method for solving the convex programming problem with convergence rate $O(1/k^2)$. *Dokl. Akad. Nauk SSSR*, 269(3):543–547. [4](#), [19](#)
- [33] Nesterov, Y. E. (2004). *Introductory lectures on convex optimization*, volume 87 of *Applied Optimization*. Kluwer Academic Publishers, Boston, MA. A basic course. [3](#), [4](#), [8](#), [16](#), [17](#), [18](#), [19](#), [20](#), [44](#), [45](#)
- [34] Phelps, R. R. (1993). *Convex functions, monotone operators and differentiability*, volume 1364 of *Lecture Notes in Mathematics*. Springer-Verlag, Berlin, second edition. [15](#)
- [35] Polak, E. and Ribiere, G. (1969). Note sur la convergence de méthodes de directions conjuguées. *ESAIM: Mathematical Modelling and Numerical Analysis-Modélisation Mathématique et Analyse Numérique*, 3(R1):35–43. [7](#)
- [36] Poljak, B. T. (1964). Some methods of speeding up the convergence of iterative methods. *Ž. Vyčisl. Mat i Mat. Fiz.*, 4:791–803. [21](#), [29](#)
- [37] Powell, M. J. D. (1977). Restart procedures for the conjugate gradient method. *Mathematical programming*, 12(1):241–254. [7](#)
- [38] Powell, M. J. D. (1984). Nonconvex minimization calculations and the conjugate gradient method. In *Numerical analysis (Dundee, 1983)*, volume 1066 of *Lecture Notes in Math.*, pages 122–141. Springer, Berlin. [7](#)
- [39] Reid, J. K. (1971). On the method of conjugate gradients for the solution of large sparse systems of linear equations. In *Large sparse sets of linear equations (Proc. Conf., St. Catherine’s Coll., Oxford, 1970)*, pages 231–254. [6](#)

- [40] Schaeffer, H. and Hou, T. Y. (2016). An accelerated method for nonlinear elliptic PDE. *J. Sci. Comput.*, 69(2):556–580. [6](#), [8](#), [73](#)
- [41] Schmidt, E. (1908). Über die auflösung linearer gleichungen mit unendlich vielen unbekannten. *Rendiconti del Circolo Matematico di Palermo (1884-1940)*, 25(1):53–77. [6](#)
- [42] Shen, J., Tang, T., and Wang, L.-L. (2011). *Spectral methods*, volume 41 of *Springer Series in Computational Mathematics*. Springer, Heidelberg. Algorithms, analysis and applications. [46](#)
- [43] Shewchuk, J. R. et al. (1994). An introduction to the conjugate gradient method without the agonizing pain. [6](#)
- [44] Shi, B., Du, S. S., Jordan, M. I., and Su, W. J. (2018). Understanding the acceleration phenomenon via high-resolution differential equations. arXiv:1810.08907. [5](#), [8](#), [29](#), [73](#)
- [45] Siegel, J. W. (2019). Accelerated first-order methods: Differential equations and Lyapunov functions. arXiv:1903.05671. [5](#), [8](#), [73](#)
- [46] Su, W., Boyd, S., and Candes, E. (2014). A differential equation for modeling Nesterov’s accelerated gradient method: Theory and insights. In Ghahramani, Z., Welling, M., Cortes, C., Lawrence, N. D., and Weinberger, K. Q., editors, *Advances in Neural Information Processing Systems 27*, pages 2510–2518. Curran Associates, Inc. [5](#), [8](#), [9](#), [21](#), [28](#)
- [47] Touati-Ahmed, D. and Storey, C. (1990). Efficient hybrid conjugate gradient techniques. *J. Optim. Theory Appl.*, 64(2):379–397. [7](#)
- [48] Wibisono, A., Wilson, A. C., and Jordan, M. I. (2016). A variational perspective on accelerated methods in optimization. *Proc. Natl. Acad. Sci. USA*, 113(47):E7351–E7358. [2](#), [5](#), [6](#), [8](#), [44](#)
- [49] Wilson, A. C., Recht, B., and Jordan, M. I. (2018). A Lyapunov analysis of momentum methods in optimization. arXiv:611.02635. [5](#), [8](#), [73](#)

- [50] Yuan*, Y.-X. (1993). Analysis on the conjugate gradient method. *Dynamical Systems*, 2(1):19–29. [7](#)
- [51] Zoutendijk, G. (1970). Nonlinear programming, computational methods. In *Integer and nonlinear programming*, pages 37–86. [7](#)

Appendix

A Literature comparison

We summarize our discussion on the existing literature works, and contrast them with our contributions, in Table 2.

Table 2: A comparison of recent works from a numerical PDE point of view. All works that provide convergence rates either in the continuous (column $s = 0$) or discrete (column $s > 0$) case assume the global Lipschitz condition (column L). No work addresses invariant sets (column $\mathfrak{B}$), incorporates preconditioning explicitly (column $\mathcal{L}$), nor it explains how concrete numerical examples fit the abstract theory (column Numerical Analysis).

Ref.	$s = 0$	$s > 0$	L	$\mathfrak{B}$	$\mathcal{L}$	Numerics	Numerical Analysis
[49]	Opt.	Opt.	Glob.	×	Bregman	×	×
[4]	×	×	Loc.	×	×	✓	×
[30]	Opt.	Opt.	Glob.	×	×	×	×
[40]	×	×	Glob.	×	×	✓	×
[24]	×	×	Loc.	×	×	×	×
[11]	Sub.	×	Glob.	×	×	✓	×
[31]	Opt.	Opt.	Glob.	×	×	×	×
[45]	Opt.	Opt.	Glob.	×	×	×	×
[44]	Sub.	Sub.	Glob.	×	×	✓	×
Ours	Opt.	Opt.	Loc.	✓	✓	✓	✓

Vita

Jea Hyun Park was born and grew up in Korea. He earned his Bachelor of Science degree in Mathematics Education at Kongju National University (Korea), and Master of Education degree at Korea National University of Education (Korea). He also worked as a math teacher in Korea for several years. Then, he decided to pursue a Doctor of Philosophy degree in Mathematics at University of Tennessee, Knoxville. His research interests include optimization and numerical partial differential equations. He loved the years of his graduate study at Knoxville very much.